CURRENT TRENDS
IN SCIENCE

PLATINUM JUBILEE SPECIAL

A tutorial introduction to financial engineering

M VIDYASAGAR

Cecil & Ida Green (II) Professor, Erik Jonsson School of Engineering & Computer Science, University of Texas at Dallas, 800 W. Campbell Road, Richardson, TX 75080, USA
e-mail: M.Vidyasagar@utdallas.edu

In this paper we give a brief, elementary introduction to various aspects of financial engineering. Specific topics discussed include the determination of a fair price for an option or other derivative instrument, hedging strategies, etc. We also discuss what role, if any, financial engineering has played in the current financial crisis.

1. Introduction

Buying and selling assets as a human activity has been going on since the dawn of civilization. Buying and selling (and thus speculating in) *derivative instruments* based on these assets is almost as old an activity. While banking as a formal, regulated activity is just a few hundred years old, there can be no doubt even ancient societies had their version of money lenders. Borrowing against assets, or 'leveraged buying' is neither new nor disreputable. A home-owner who buys a house with only part down payment and borrows the rest to be paid in monthly instalments is in effect speculating that his future earnings will be adequate to cover the payments; simultaneously, he may also be betting that the value of his home will go up before the mortage matures. At the same time, the lender (holder of the mortgage) is also betting on the same thing.

For the past several decades, the financial world has become more and more mathematical. By now practically everyone has heard of the Black–Scholes formula, even if he might not actually know what it is (or even what it is for). The objective of this article is to give a very brief, and very elementary, tutorial-level introduction to various topics in financial engineering. In all the problems discussed here, the emphasis is on *giving the flavour* of the problem, rather than on stating theorems with the strongest possible conclusions in the most general setting possible. The bibliography at the end of the article contains several authoritative books that could be consulted by readers interested in pursuing the topic further.

The article is aimed at both mathematical as well as non-mathematical readers. A non-mathematical reader can read only sections 5., 8., 9. and 10., which do not have a single equation in them. A person who knows a little bit of mathematics can *also* read sections 2. and 6., except for subsection 6.4. Section 2. discusses an extremely simplified problem that nevertheless brings out many of the salient aspects, while section 6. discusses slightly more general situations. Mathematically inclined readers, of course, can read the entire article.

2. One-period model

We begin with a very simple model that illustrates many of the important ideas, known popularly as the 'one-period model.' Suppose there are two assets in which we can invest, referred to respectively as a 'bond' and as a 'stock.' The bond offers an absolutely guaranteed rate of return, whereas the stock price can go up or down. Specifically, let $B(0)$ denote the value of one unit of the bond at the present time $T = 0$. Then at the next period (say after one month), the value of one unit of the bond is

$$B(1) = (1 + r)B(0),$$

Keywords. Financial engineering; quantitative finance; Black–Scholes theory; option pricing; derivatives; hedging.

where r is the rate of return over one period. Under normal circumstances $r > 0$. The behaviour of the stock's value, in contrast, is assumed to be stochastic. Let $S(0)$ denote the value of unit of the stock at the present time, and let $S(1)$ denote the value of one unit of the stock at the next time instant. Then the stochastic model describing the stock price is

$$S(1) = \begin{cases} S(0)u & \text{with probability } p, \\ S(0)d & \text{with probability } 1-p. \end{cases}$$

Here the symbols u and d are meant to suggest 'up' and 'down.' Also, for reasons of convenience, u and d represent the rate of return on the stock price, as opposed to the absolute change in the stock price. Clearly, both u and d are positive, and $d < u$. But it is possible that $d < 1$; that is, it is possible for the stock price to decline in absolute terms. In any case, it is assumed that

$$d < 1 + r < u.$$

In other words, the guaranteed rate of return $1+r$ is bracketed by the best and worst rates of return on the risky investment. Clearly this is a reasonable assumption. If $1 + r < d$, then no one would have any reason to buy bonds, whereas if $u < 1+r$, then no one would have any reason to buy stocks.

An 'option' on a stock is an instrument that gives the buyer the right, but not the obligation, to buy the stock at a prespecified price known as the 'strike price.' Suppose the strike price is K. The option gives the buyer the right (but not the obligation) to buy one unit of the stock at the price K at time $T = 1$, irrespective of what the prevailing market price might be. So if $S(1) > K$, then the buyer of the option will exercise the option and make an instant profit of $S(1) - K$ (because the market value of the stock is $S(1)$ whereas the option buyer pays only K). On the other hand, if $S(1) < K$ then the buyer does not exercise the option, and its value is zero. So if $X(1)$ denotes the value of the option at time $T = 1$, then

$$X(1) = \max\{S(1) - K, 0\} = [S(1) - K]_+.$$

One can think of $X(1)$ as the outgo of the person selling the option, at time $T = 1$. Since $S(1)$ is a random variable, so is X. Moreover, since the value of $X(1)$ is *derived* from that of $S(1)$, the option is called a 'derivative' instrument, or simply a 'derivative.'

The theory of option pricing deals with the following question: Suppose we are provided with the stochastic model of the stock price and the strike price K. What is the 'correct' price at time $T = 0$ for the option X? In other words, what price should a buyer be prepared to pay for the option, and what price should the seller of the option be prepared to accept for the option? Most interesting perhaps, are the two prices the same?

One's first impulse might be to just compute the discounted expected value of the option at time $T = 1$. Clearly the option makes sense only if

$$S(0)d < K < S(0)u,$$

that is, only if the strike price is bracketed by the maximum and minimum stock prices at time $T = 1$; so let us assume this. Therefore it follows that

$$X(1) = \begin{cases} S(0)u - K & \text{with probability } p \\ 0 & \text{with probability } 1-p. \end{cases}$$

So the expected value of $X(1)$ with the distribution $(p, 1-p)$, discounted by the factor $1 + r$, is

$$(1 + r)^{-1} p[S(0)u - K].$$

One may think that this number is the correct price for the option at time $T = 0$.

The reasoning behind this suggestion is as follows: Suppose the seller of the option receives an amount v (as yet unspecified) at time $T = 0$, for an option at time $T = 1$ with the strike price K. Suppose the seller were to invest *all* of this amount v in the 'safe' asset, namely the bond. Then at time $T = 1$, his investment would be worth $(1 + r)v$. If the stock price goes up so that $S(1) = S(0)u > K$ then the seller of the option would receive K, but would then have to procure the stock by paying the prevailing price of $S(0)u$ to deliver to the buyer. So in this case the seller of the option would incur a loss, and his payout (out of pocket expense) would be

$$S(0)u - K - v(1 + r).$$

This happens with probability p. If the stock price goes down, then the option is not exercised, and the seller of the option does not incur any additional expenditure. This happens with probability $1 - p$. So to be neutral, the value v would have to satisfy

$$[v(1 + r) + K - S(0)u]p + v(1 + r)(1 - p) = 0,$$

which leads to

$$v = (1 + r)^{-1}[S(0)u - K].$$

But this intuition is *wrong*. The reason why it is wrong is that the seller of the option can *hedge* against future changes in the price of the stock. Specifically, the seller of the option need not put *all* of the money received into the safe investment. Instead, he can invest some of it in the stock itself. Suppose he were to buy a units of the stock and b units of the bond at time $T = 0$. Let us treat the quantum of investments in both the bond and the stock as being infinitely divisible, so that a and b be treated as real numbers. Note that we explicitly permit both a and b to be negative. Negative a corresponds to short-selling the stock, whereas negative b corresponds to borrowing money at the fixed rate of interest r. Now at time $T = 1$, suppose the stock price goes up so that the option is exercised by the buyer. Then the seller is obliged to deliver one unit of stock. But he is already in possession of a units of the stock; so he has to procure only $1 - a$ units of stock at the prevailing (high) market price. This is the logic behind hedging. Now let us see how the seller of the option should split his initial investment between the stock and the bond.

Let us return to the stock price $S(1)$, which is a random variable. Let us define

$$u' = \frac{u}{1+r}, \quad d' = \frac{d}{1+r}.$$

Then u', d' are respectively the *discounted* returns of the stock price. Note that

$$d' < 1 < u'$$

by assumption. Now suppose Y is a random variable representing the return on the stock investment. In the 'real world' $Y = u'$ with probability p and $Y = d'$ with probability $1-p$. Let us disregard this, and instead choose an 'artificial' probability q such that $Y = u'$ with probability q, $Y = d'$ with probability $1 - q$, and most important, Y *is risk-neutral under this probability distribution.* In other words, we would like the expected value of Y to equal one, so that the expected value of the share price $S(1)$ *under this artificial distribution* is equal to $S(0)(1 + r)$. Since Y has only two possible values, there is a unique choice of q that achieves this property, namely the solution of

$$qu' + (1 - q)d' = 1,$$

which is

$$q = \frac{1 - d'}{u' - d'}, \quad 1 - q = \frac{u' - 1}{u' - d'}. \tag{2.1}$$

What is the significance of this number q? Suppose a random variable $\mathtt{S}(1)$ equals $S(0)u$ with probability q (not p), and equals $S(0)d$ with probability $1 - q$. Thus $\mathtt{S}(1)$ assumes the same values as the original random variable $S(1)$, but not necessarily with the same probabilities. Then

$$E[\mathtt{S}(1)] = (1 + r)S(0). \tag{2.2}$$

This is our first encounter with what is known as a 'martingale measure'. We shall meet it again.

Now it is claimed that the 'correct' value for the option at time $T = 0$ is the discounted expected value of the payoff function $X(1) = \max\{S(1)-K\}$ under the distribution $(q, 1-q)$. In other words, we claim that the 'correct' price for the option with strike price K is

$$\begin{aligned} c &= (1 + r)^{-1}[S(0)u - K]q \\ &= (1 + r)^{-1}[S(0)u - K]\frac{(1 + r) - d}{u - d}. \end{aligned} \tag{2.3}$$

Up to now we have discussed very simple 'options.' However, the theory goes through readily to the case where the payout need not simply equal $\max\{S(0)u - K, 0\}$. Suppose an individual is selling a 'contingent claim' as follows: At time $T = 1$, if the stock price $S(1)$ equals $S(0)u$, then he will pay an amount of X_u to the buyer of the claim, whereas if the stock price $S(1)$ equals $S(0)d$, then he will pay an amount of X_d to the buyer. Since the amount paid out is 'contingent' on the price of the stock, this kind of instrument is called a 'contingent claim.' In the case of a simple option, we have $X_u = S(0)u - K$ and $X_d = 0$. In the general case, we will show that the 'correct' price for the contingent claim is

$$c = (1 + r)^{-1}[X_u q + X_d(1 - q)], \tag{2.4}$$

where $q = (u' - 1)/(u' - d')$ as defined earlier. Thus the correct price is the expected value of the contingent claim X under the distribution $(q, 1-q)$, which is then discounted by the factor $1 + r$. The reason for the discounting is that the expected value is in some sense computed at time $T = 1$, whereas the seller of the option is receiving the money at time $T = 0$.

Why is this the 'correct' price? We show next that if the price is some other number, then there exists an opportunity for 'arbitrage.' In the present setting, 'arbitrage' refers to a trading strategy that results in profit without risk. Specifically, arbitrage refers to a trading strategy wherein the profit to the investor is a random variable that (i) assumes only nonnegative values, and (ii) has a positive expected value. Since the profit assumes only nonnegative values, the investor can never make a loss, and since the expected value is positive, he will make a profit with positive probability. This situation is unsustainable in an 'efficient' market, where it is assumed that everyone has access to the same amount of information. Hence the price of any trading instrument will settle at a value where no one, neither buyer nor seller, has an opportunity for arbitrage.

To justify why the above value is the only one that results in no arbitrage, we introduce the notion of a 'replicating strategy.' The idea is that the seller of the contingent claim will take a certain amount of money at time $T = 0$ and split it between stocks and bonds in such a way that, irrespective of whether the stock goes up or down, the value of his portfolio at time $T = 1$ exactly equals the value of the contingent claim. Let a, b denote our holdings in stocks and bonds at time $T = 0$. The idea is to choose these numbers in such a way that

$$aS(0)u + bB(0)(1+r) = X_u,$$
$$aS(0)d + bB(0)(1+r) = X_d.$$

This is a simple set of simultaneous equations that can be written as

$$[a \;\; b] \begin{bmatrix} S(0)u & S(0)d \\ B(0)(1+r) & B(0)(1+r) \end{bmatrix} = [X_u \;\; X_d].$$

Clearly the unique solution of these equations is

$$[a \;\; b] = [X_u \;\; X_d] \begin{bmatrix} S(0)u & S(0)d \\ B(0)(1+r) & B(0)(1+r) \end{bmatrix}^{-1}. \tag{2.5}$$

Let us make the formula in (2.5) a little more explicit. Note that

$$\begin{bmatrix} S(0)u & S(0)d \\ B(0)(1+r) & B(0)(1+r) \end{bmatrix} = (1+r) \begin{bmatrix} S(0) & 0 \\ 0 & B(0) \end{bmatrix} \begin{bmatrix} u' & d' \\ 1 & 1 \end{bmatrix}.$$

Hence

$$[a \;\; b] = \frac{1}{(1+r)(u'-d')} [X_u \;\; X_d] \begin{bmatrix} 1 & -d' \\ -1 & u' \end{bmatrix} \times \begin{bmatrix} 1/S(0) & 0 \\ 0 & 1/B(0) \end{bmatrix}$$
$$= \frac{1}{(1+r)(u'-d')} \begin{bmatrix} \dfrac{X_u - X_d}{S(0)} & \dfrac{u'X_d - d'X_u}{B(0)} \end{bmatrix}. \tag{2.6}$$

Note that, depending on the relative values of X_u, X_d, u', d', the initial investment in bonds could be negative. In the specific case of options, we have $X_d = 0$, so that the value of b is always negative. This means that the seller of the claim has to borrow money at the interest rate of r to finance the purchase of the replicating portfolio.

So how much money is needed at time $T = 0$ to implement this replicating strategy? The answer is

$$c = [a \;\; b] \begin{bmatrix} S(0) \\ B(0) \end{bmatrix}$$
$$= [X_u \;\; X_d] \begin{bmatrix} S(0)u & S(0)d \\ B(0)(1+r) & B(0)(1+r) \end{bmatrix}^{-1} \begin{bmatrix} S(0) \\ B(0) \end{bmatrix}$$
$$= (1+r)^{-1} [X_u \;\; X_d] \begin{bmatrix} q_u \\ q_d \end{bmatrix}, \tag{2.7}$$

where

$$\begin{bmatrix} q_u \\ q_d \end{bmatrix} = \begin{bmatrix} u' & d' \\ 1 & 1 \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ 1 \end{bmatrix}. \tag{2.8}$$

Now it is easy to see that $(q_u, q_d) = (q, 1-q)$, which is the probability distribution defined earlier. Thus we have established two facts relatively painlessly. First, there exists an investment strategy that perfectly replicates the contingent claim at time $T = 1$. Second, the initial investment c needed to implement this replicating strategy equals the discounted expected value of the contingent claim under a very special probability distribution. A corollary of the second statement is that the initial investment c depends only on the two (discounted) returns u' and d', and does not at all depend on the *probability* with which these two events occur. This counter-intuitive result can be explained by noting that since the replicating strategy perfectly replicates the contingent claim in each of the two possible outcomes, the probability with which each outcome occurs is of no importance.

Now let us return to the 'correctness' of c as the price for the contingent claim. Suppose that at time $T = 0$, a buyer is willing to pay a price $c' > c$ for the contingent claim. Then the seller of the contingent claim can collect c', use c of that to implement the replicating strategy, and pocket the difference $c' - c$. At time $T = 1$, the replicating strategy will enable the seller of the option to settle the claim by using the current value of the portfolio, irrespective of whether the stock price goes up or down. So the seller of the option can make a risk-free profit of $c' - c$. In other words, if the price is greater than c, then the seller of the contingent claim can make a risk-free profit, and thus has an arbitrage opportunity. By reversing signs, it is clear that if the contingent claim is offered at a price $c' < c$, then the buyer of the claim has an arbitrage opportunity.

Thus the conclusions of the analysis of the one-period model can be summarized as follows: The seller of a contingent claim with payouts X_u, X_d at time $T = 1$ collects an amount equal to c in (2.7) at time $T = 0$. He then invests this amount in stocks and bonds in the ratio indicated by (2.5). Since this portfolio is replicating, the value of the portfolio exactly equals the value of the claim, irrespective of whether the stock price goes up or goes down. In this way, the seller of the claim has perfectly 'hedged' the claim.

3. Multiple time periods: The binomial model

In this section, we build upon the analysis of section 2. to study a somewhat more realistic model of stock price movement, known as the 'binomial model.' For this model, we can once again compute the 'correct' price for a class of contingent claims known as 'European' claims. By taking the limit as the time interval between successive price movements approaches zero, it is possible to obtain the celebrated Black–Scholes formula for option pricing. That is done in section 4..

As before, we consider two assets: A 'bond' whose price movement is deterministic, and a 'stock' whose value is stochastic. We study these assets over N time instants, after the initial time 0. Initially the bond is worth B_0 and the stock is worth S_0. At the n-time instant, the bond gets an assured return of r_n, so that

$$B_{n+1} = (1 + r_n)B_n.$$

As for the stock, if its price is S_n at time n, then at time $n + 1$ its price can have one of two possible values, namely $S_n u_n$ and $S_n d_n$. Therefore u_n and d_n represent respectively, the higher and lower returns on the stock during the n-th time interval. We could in principle associate 'probabilities' with these two movements, but as in the single period case, it is only the values $S_n u_n$ and $S_n d_n$ that matter, and not any associated probabilities. We could of course assume that the return r_n and the returns u_n, d_n are the same for all time, but this does not significantly simplify the analysis.

From this description, it is clear that at time N there are 2^N possible sample paths. These sample paths can be identified in an obvious way with the 2^N strings in the set $\{u, d\}^N$. For instance, if $N = 5$ and we choose the string $uudud$, then

$$S_1 = S_0u_0, \quad S_2 = S_1u_1 = S_0u_0u_1, \ldots,$$
$$S_5 = S_0u_0u_1d_2u_3d_4.$$

Thus there are 2^N possible values S_N that the stock can have at time N. Depending on the values of u_n and d_n, these 2^N values may not be distinct. For instance, if $u_n = u, d_n = d$ for all n, then there are only $N+1$ possible values for S_N, namely $S_0u^kd^{N-k}$ as k varies from 0 to N. Let $\mathbf{h} \in \{u, d\}^N$ denote some string of length N over $\{u, d\}$, and let $X_\mathbf{h}$ denote the payout at time N if the stock price movement follows the sequence in $\mathbf{h}$. This implies in particular that the payout can be 'path-dependent.' For instance, we explicitly permit that the payout in the case of a sequence du can be different from that for the sequence ud even if the final price of the stock may be the same in both cases. Thus the seller of the contingent claim needs to pay $X_\mathbf{h}$ to the buyer of the claim at time N if the stock price movement follows the sequence $\mathbf{h}$.

In option pricing theory, a 'European Contingent Claim (ECC)' is an instrument that can be exercised only *at the end* of a fixed duration N. In contrast, an 'American Contingent Claim' is an instrument that can be exercised *at any time* during the fixed duration N. Hence the pricing and hedging of ACCs is much more difficult than for ECCs. For this reason, the discussion in the present section is restricted to ECCs.

Now we adapt the reasoning of section 2. to compute the 'correct' or arbitrage-free price of a ECC with a payout of $X_\mathbf{h}$ corresponding to a trajectory $\mathbf{h} \in \{u, d\}^N$. For this purpose we again introduce an artificial probability distribution, under the very reasonable assumption that

$$d_n < 1 + r_n < u_n \ \forall n, \quad \text{or} \quad d'_n < 1 < u'_n \ \forall n,$$

where $d'_n = d_n/(1+r_n), u'_n = u_n/(1+r_n)$. We define

$$q_{u,n} = \frac{1 - d'_n}{u'_n - d'_n}, \quad q_{d,n} = \frac{u'_n - 1}{u'_n - d'_n;}. \tag{3.1}$$

Now we define a stochastic process $\mathtt{S}_n$ as follows: $\mathtt{S}_0 = S_0$ is determininstic, and for $n = 0, \ldots, N-1$, we have

$$\mathtt{S}_{n+1} = \begin{cases} \mathtt{S}_n u'_n & \text{with probability } q_{u,n}, \\ \mathtt{S}_n d'_n & \text{with probability } q_{d,n}. \end{cases} \tag{3.2}$$

Thus the process $\{\mathtt{S}_n\}$ assumes the same values as the original stock price $\{S_n\}$, but with possibly different probabilities. Because of the way in which the quantities $q_{u,n}$, $q_{d,n}$ are defined, the process $\{\mathtt{S}_n\}$ has the very special property that

$$E\{\mathtt{S}_{n+1}|\mathtt{S}_n, \mathtt{S}_{n-1}, \ldots, \mathtt{S}_0\} = (1 + r_n)\mathtt{S}_n, \quad \text{for } n = 0, \ldots, N-1. \tag{3.3}$$

In other words, the discounted process

$$\left\{ \prod_{i=0}^{n-1} (1 + r_i)^{-1} \mathtt{S}_n \right\},$$

where the empty product is taken as one, is a martingale. The distribution defined by the q's is the martingale measure. When there are only two possible transitions at each time step, it is easy to see that the above choice of the q's is unique; in other words, this is the only way to choose the q's to satisfy (3.3). If there are more than two possible values at any stage, then the choice of the martingale measure is not unique. This situation is discussed in section 6.4.

Next, we show that the 'replicating' strategy introduced in section 2. for a one-period model can be extended to the binomial model. Some good notation makes the task simpler. For each N-tuple $\mathbf{h} \in \{u, d\}^N$, the probability $q_\mathbf{h}$ is defined as

$$q_\mathbf{h} = \prod_{n=0}^{N-1} q_{h,n}, \tag{3.4}$$

where

$$q_{h,n} = \begin{cases} q_{u,n} & \text{if } h_n = u, \\ q_{d,n} & \text{if } h_n = d. \end{cases}$$

Then it is easy to see that $\{q_{\mathbf{h}}\}$ defines a probability distribution on $\{u,d\}^N$. Let $X_{\mathbf{h}}$ denote the payout if the stock price follows the set of transitions corresponding to $\mathbf{h}$. Now define

$$c_0 := \left[\prod_{n=0}^{N-1} (1+r_n) \right]^{-1} \sum_{\mathbf{h} \in \{u,d\}^N} X_{\mathbf{h}} q_{\mathbf{h}}. \qquad (3.5)$$

Thus c_0 is the expected value of the random payout X with the distribution

$$\Pr\{X = X_{\mathbf{h}}\} = q_{\mathbf{h}}, \ \forall \mathbf{h} \in \{u,d\}^N,$$

which is then discounted by the factor $\left[\prod_{n=0}^{N-1} \times (1+r_n) \right]^{-1}$ to make it comparable to the risk-free return on the bond. We now show that c_0 is the only arbitrage-free price for the contingent claim by showing that there exists a replicating strategy.

For this purpose, suppose $1 \leq n < N$ and that $\mathbf{j} \in \{u,d\}^n$ is a string of length n. Then we define

$$c_{\mathbf{j}} := \left[\prod_{i=n}^{N-1} (1+r_i) \right]^{-1} \sum_{\mathbf{k} \in \{u,d\}^{N-n}} X_{\mathbf{jk}} q_{\mathbf{jk}}. \qquad (3.6)$$

The quantity $c_{\mathbf{j}}$ is a kind of intermediate discounted expected value, after the transitions corresponding to $\mathbf{j}$ have already taken place. Then it is easy to see that $c_{\mathbf{j}}$ satisfies the following recursive relationship:

$$c_{\mathbf{j}} = (1+r_{n-1})^{-1} \left(c_{\mathbf{j}u} q_{u,n} + c_{\mathbf{j}d} q_{d,n} \right). \qquad (3.7)$$

The above recursive relationship follows readily from the product nature of the distribution $q_{\mathbf{h}}$ and the definition (3.6). Finally, we can interpret c_0 as $c_\emptyset$, corresponding to the empty string and $n = 0$.

The replicating strategy is now easily described. As time 0, the seller of the claim receives an amount of money equal to c_0. He then invests this in a_0 stocks and b_0 bonds as per the formula

$$[a_0 \ \ b_0] = [c_u \ \ c_d] \begin{bmatrix} S_0 u_0 & S_0 d_0 \\ B_0(1+r_0) & B_0(1+r_0) \end{bmatrix}^{-1}.$$

This formula is analogous to (2.5) with c_u interpreted as the payout if the stock price goes up at time 0, and c_d interpreted as the payout if the stock price goes down at time 0. Since this is a one-period replicating strategy, the cost of this investment is precisely $c_0 = c_\emptyset$. Moreover, the value of the portfolio is c_u if the stock price goes up at time 0 and is c_d if the stock price goes down at time 0. This can be expressed as: The value of the portfolio at time 1 equals c_{i_0} irrespective of whether $i_0 = u$ or $i_0 = d$. Now at time 1, we again implement a one-period replicating strategy as per the formula

$$[a_1 \ \ b_1] = [c_{i_0 u} \ \ c_{i_0 d}] \begin{bmatrix} S_1 u_1 & S_1 d_1 \\ B_1(1+r_1) & B_1(1+r_1) \end{bmatrix}^{-1},$$

where S_1 is the *actual* price of the stock at time 1, that is, after the transition $i_0 = u$ or d at time 0. Then, irrespective of whether the stock price goes up or down at time 1, it follows that the value of the portfolio equals $c_{i_0 i_1}$, where $i_1 = u$ or d. This argument can be repeated all way until $n = N-1$, to show that the strategy *completely replicates* the payout $X_{\mathbf{h}}$ for all 2^N possible sample paths. Moreover, this strategy is *self-financing* in the sense that, once the seller receives the initial price of c at time 0, all subsequent investments can be financed with the value of the current portfolio. Note that there is no analog of 'self-financing' in the one-period model.

Now it is not difficult to see that c_0 is the only price that can be charged in such a way that neither the buyer of the claim nor the seller of the claim can have an arbitrage opportunity. As in the one-period case, if a buyer is willing to pay a price $c' > c_0$ for the contingent claim, the seller of the claim can use c_0 of that to implement a replicating strategy, and thus make a risk-free profit of $c' - c_0$, leading to arbitrage. Similarly, if a contingent claim is available for less than c_0, then the buyer can make a risk-free profit.

In order to keep implementing the self-financing and fully replicating strategy, the seller of the claim has to adjust his portfolio at every time instant. After n time instants, suppose the set of price movements that have occured is given by $\mathbf{i} \in \{u,d\}^n$. Then the holding at time n is adjusted in accordance with the formula

$$[a_n \ \ b_n] = [c_{\mathbf{i}u} \ \ c_{\mathbf{i}d}] \begin{bmatrix} S_n u_n & S_n d_n \\ B_n(1+r_n) & B_n(1+r_n) \end{bmatrix}^{-1}, \qquad (3.8)$$

where S_n is the actual price of the stock at time n. Thus the implementation of the replicating strategy assumes that *there are no transaction costs*. If stock is bought or sold, the transaction does not attract any commission. Similarly, at some times the bond holding B_n may be negative, which corresponds to borrowing money from the bank at the rate of interest r_n. If on the other hand the bond holding B_n is positive, then it attracts interest at the same rate r_n. In other words, there is no difference between the rates of interest when borrowing money and saving money (in the form of bonds). Both of these assumptions (no transaction costs and no difference in rates of interest for borrowing versus saving money) are clearly unrealistic. Nevertheless, they allow quite a deep analysis of the problem of selling and hedging options. In addition, we have also assumed that both bonds and stocks are infinitely divisible. Markets that satisfy these conditions are said to be 'frictionless.'

There is another, more serious, issue to ponder. Suppose an agent sells a contingent claim at time 0 and collects the arbitrage-free price c_0. Then at each time instant starting at time 0, he keeps on adjusting his portfolio in accordance with (3.8). Since the investment strategy is replicating and self-financing, at time N his portfolio is worth precisely $X_{\mathbf{h}}$ where $\mathbf{h}$ is the set of transitions that actually occurs. The seller then settles his claim by paying out the amount $X_{\mathbf{h}}$ and is thus left with ... precisely nothing. On the other hand another agent who has not bothered to 'play the game' at all is also left with nothing at time N. So why should anyone bother to sell an option at all? This question is revisited in section 5., after we have presented the Black–Scholes theory.

Finally, it should be obvious that there is nothing special about having *just one* stock. The above theory applies equally well, with obvious changes in notation, to the case where there are finitely many stocks, but each stock can just move up or down by a fixed amount at each time instant. In such a case too, it is possible to construct a replicating strategy, and thus prove that the correct price for a contingent claim is the expected value of the claim under the martingale measure, discounted to time $T = 0$.

4. Black–Scholes theory

In this section, we present the famous Black–Scholes formula for the pricing of options when the underlying stock price follows a stochastic model known as the 'geometric Brownian motion' model. This section is quite abstract and may be difficult to follow. Then in the next section we discuss the *implications* of the model, how it is actually used, and so on. This discussion is not mathematical at all, and is far easier to follow.

4.1 *Review of the basics of probability and random variables*

The objective of this subsection is to present a formal definition of Brownian motion. To set the stage for that, we first define a random variable in a very formal setting. The discussion below is a summary of the basic definitions and facts needed to define Brownian motion, but is not a substitute for a proper study of probability theory. The reader is referred to any standard text for further details on this topic.

Suppose Z is some set, and Ω is a collection of subsets of Z. Then Ω is said to be a **σ-algebra** if it satisfies the following properties:

1. Both the complete set Z and the empty set $\emptyset$ belong to Ω.
2. If $A \subseteq Z$ belongs to Ω, so does its complement A^c.
3. If A_i, $i = 1, 2, \ldots$ is a countable collection of sets such that A_i belongs to Ω for each i, then the union $\cup_i A_i$ also belongs to Ω.

The pair (Z, Ω) is often referred to as a **measurable space**.

A **probability measure** P on the measurable space (Z, Ω) is a map from Ω into the interval $[0, 1]$ such that the following axioms hold:

1. $P(Z) = 1, P(\emptyset) = 0$.
2. If $A_i, i = 1, 2, \ldots$ is a countable collection of pairwise disjoint sets such that A_i belongs to Ω for each i, then

$$P\left(\bigcup_i^{\infty} A_i\right) = \sum_{i=1}^{\infty} P(A_i).$$

The triple (Z, Ω, P) is called a **probability space**. The probability space is said to be **complete** if, for every set $A \in \Omega$ such that $P(A) = 0$, it is true that every subset $B \subseteq A$ is also measurable (and of course $P(B) = 0$). Any probability space can be 'completed' using a standard procedure, and dealing with complete probability spaces often avoids a lot of technicalities.

Given two probability measures P, Q on a common measurable space (Z, Ω), the **total variation metric** between them is defined as

$$\rho(P, Q) = \sup_{A \in \Omega} |P(A) - Q(A)|.$$

In the common case where Z is a finite set, say $Z = \{z_1, \ldots, z_n\}$, and $(p_1, \ldots, p_n), (q_1, \ldots, q_n)$ are two probability distributions on Z, it is easy to show that

$$\rho(P, Q) = \frac{1}{2} \sum_{i=1}^{n} |p_i - q_i|.$$

Given any collection of subsets $\mathcal{C}$ of Z, the smallest σ-algebra that contains every set in $\mathcal{C}$ is referred to as the σ-algebra generated by $\mathcal{C}$. In particular, if $Z = \mathbb{R}$, the set of real numbers, then the σ-algebra generated by the collection of all closed sets is referred to as the **Borel σ-algebra** and is denoted by $\mathcal{B}$. A random variable X defined on a measurable space (Z, Ω) is a function $X : Z \to \mathbb{R}$ such that whenever $A \subseteq \mathbb{R}$ belongs to the Borel σ-algebra, the preimage $X^{-1}(A)$ belongs to Ω. In other words,

X is a **measurable function** from Z into $\mathbb{R}$ where Z is equipped with the σ-algebra Ω and the real set $\mathbb{R}$ is equipped with the Borel σ-algebra. Quite often, we take $Z = \mathbb{R}$ and $\Omega = \mathcal{B}$, the Borel σ-algebra. This is known as the 'canonical representation' of a random variable. For such a random variable, the function $\Psi : \mathbb{R} \to \mathbb{R}$ defined by

$$\Psi(c) := P\{X^{-1}(-\infty, c]\} = \Pr\{X \le c\}$$

is called the **cumulative distribution function** of the random variable X. It is always well-defined since the semi-infinite interval $(-\infty, c]$ belongs to the Borel σ-algebra $\mathcal{B}$. Clearly, it is also a non-decreasing function. Moreover, the function $\Psi(\cdot)$ is continuous from the right, and has limits from the left. In other words,

$$\lim_{c \to c_0^+} \Psi(c) = \Psi(c_0), \lim_{c \to c_0^-} \Psi(c) \text{ exists and is } \le \Psi(c_0).$$

Among the most useful distribution functions is the **normal distribution** defined by

$$\Phi(c) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{c} e^{-u^2/2} du. \tag{4.1}$$

Hereafter we use Φ exclusively to denote the normal distribution.

Given a probability space (Z, Ω, P), two sets $A, B \in \Omega$ are said to be **independent** if $P(A \cap B) = P(A)P(B)$. Two random variables X, Y defined on a common probability space (Z, Ω, P) are said to be independent if the sets $X^{-1}(A), Y^{-1}(B)$ are independent for each pair of sets A, B in $\mathcal{B}$. If we use the canonical representation, the above definition is equivalent to

$$\Pr\{X \le c \& Y \le d\} = \Pr\{X \le c\} \cdot \Pr\{Y \le d\}.$$

Suppose $\{X_l\}$ is a sequence of random variables defined on a common probability space (Z, Ω, P), and that X_* is another random variable defined on (Z, Ω, P). Then the sequence $\{X_l\}$ is said to **converge in probability** to X_* if, for every $\epsilon > 0$, it is the case that

$$P\{\omega \in Z : |X_l(\omega) - X_*(\omega)| > \epsilon\} \to 0 \text{ as } l \to \infty.$$

Now suppose $\{X_t\}_{t \ge 0}$ is a set of random variables indexed by the quantity $t \ge 0$. One can think of the index t as 'time' but strictly speaking this is not necessary. We can equip the set $[0, \infty)$ with the Borel σ-algebra and define any reasonable probability measure on the half-line. Then the indexed family $\{X_t\}$ is said to be a **stochastic process** if X_t is right-continuous with respect to t and has left-limits with respect to t, where convergence is with respect to probability. What this means is that (i) $\lim_{t \to t_0^+} X_t = X_{t_0}$ for almost all t_0, where the convergence of X_t to X_{t_0} is in probability, as defined above, and (ii) $\lim_{t \to t_0^-} X_t$ exists, but may not necessarily equal X_{t_0}.

A stochastic process $\{W_t\}_{t \ge 0}$ is called a **Brownian motion** if the following properties hold:

1. Each W_t is a random variable defined on a common measurable space $(\mathbb{R}, \mathcal{B})$.
2. $W_0(\omega) = 0$ almost surely with respect to ω. In other words, $W_0(\omega) = 0$ except for possibly on a set of measure zero.
3. The process has 'independent increments.' Thus, whenever $s < t \le u < v$, the random variables $W_t - W_s$ and $W_v - W_u$ are independent.
4. The process has 'stationary increments.' Thus the distribution of $W_t - W_s$ is a function only of the difference $t - s$ and not of t and s separately.
5. The process has 'normal increments.' Thus, for $s < t$, the random variable $W_t - W_s$ has the normal distribution with mean zero and variance $t - s$. Thus

$$\Psi_{W_t - W_s}(c) = \Phi[c/(t-s)].$$

4.2 *The Black–Scholes formula and partial differential equation*

Now we are in a position to state the Black–Scholes formula and the Black–Scholes partial differential equation.

In the Black–Scholes formalism, the price of the stock S_t is a *continuous-time* stochastic process described by

$$S_t = S_0 \exp\left[\left(\mu - \frac{1}{2}\sigma^2\right) t + \sigma W_t\right], \quad t \in [0, T], \tag{4.2}$$

where W_t is a Brownian motion. In the above formula, the mean μ is often called the 'drift term' of the geometric Brownian motion, whereas σ is called the 'volatility.' Since we can also write (4.2) as

$$\log \frac{S_t}{S_0} = \left(\mu - \frac{1}{2}\sigma^2\right) t + \sigma W_t, \quad t \in [0, T],$$

the formula (4.2) is often referred to as the 'log-normal' or 'geometric Brownian motion' model for the stock price. The bond price B_t is a deterministic function of time, determined by

$$B_t = B_0 e^{rt}, \tag{4.3}$$

where r is the risk-free rate of return. Note that, to simplify notation, we have assumed that all the three terms μ, σ, r are all constant with respect to time. However, this assumption can be dispensed with, at the expense of more cumbersome notation.

We can think of the geometric Brownian motion model as a limiting case of the binomial model described in section 3.. The phrase 'limiting case' should be taken with a pinch of salt, since a rigorous analysis of continuous stochastic processes requires far more machinery than that of discrete-time stochastic processes. Thus in the present section we only motivate the results and give plausibility arguments, whereas in the previous section the proofs are completely rigorous.

Let us think of the binomial model where the up and down returns are equal and uniform with time, and also equally probable. In other words,

$$S_{n+1} = \begin{cases} S_n e^{\lambda} & \text{with probability } 0.5, \\ S_n e^{-\lambda} & \text{with probability } 0.5. \end{cases}$$

Suppose that the time interval between successive increments is Δ, and that we wish the variance of the quantity $\log(S_{n+1}/S_n)$ to equal Δ. Then it is easy to verify that the quantum of the jump λ has to equal $\Delta^{1/2}$. Thus the 'slope' or 'rate' of the jump equals $\Delta^{-1/2}$, which approaches infinity as $\Delta \to 0$. Now one can think of Brownian motion as the limiting case of a random walk with equal probabilities of moving left or right, where the time interval between successive movements approaches zero; moreover, the successive jumps are independent, and the variance of the total movement over a time interval of width T is equal to T. Because of the variance requirement, the slope of the individual jumps approaches infinity as $\Delta \to 0$, which is why the sample paths of Brownian motion are nowhere differentiable, and have unbounded variation over any nonzero interval. These are the features of Brownian motion that require the use of very advanced mathematics. The stock price itself can be thought of as the exponential of Brownian motion with a drift term.

Now suppose the stock price follows the formula (4.2), the bond price follows (4.3), and that a person wishes to sell a European option on the stock with the strike price K. Thus, after time T, the buyer of the claim has the right, but not the obligation, to pay an amount of K to receive one unit of the stock. What is the arbitrage-free or correct price for such an option?

The famous paper of Black and Scholes shows that the correct price is given by a formula that involves the following parameters: The final time T, the discounted strike price $K^* = Ke^{-rT}$, and the volatility σ. Interestingly, the drift parameter *μ does not affect the price*. The formula is

$$C_0 = S_0 \Phi\left(\frac{\log(S_0/K^*)}{\sigma\sqrt{T}} + \frac{1}{2}\sigma\sqrt{T}\right) - K^* \Phi\left(\frac{\log(S_0/K^*)}{\sigma\sqrt{T}} - \frac{1}{2}\sigma\sqrt{T}\right), \quad (4.4)$$

where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal variable defined in (4.1).

Equation (4.4) is the **Black–Scholes formula** for the price of an option with strike price K at time T. More generally, suppose the seller of the European contingent claim is obliged to pay an amount of $e^{rT}\psi(e^{-rT}x)$ to the buyer if $S_T = x$, where $\psi : \mathbb{R}_+ \to \mathbb{R}_+$. The various exponentials in the definition of the payout function are meant to discount both the payout as well as the price at time T to their values at time $T = 0$, using the risk-free rate of return. In the case of an option with strike price of K, we can simply take $\psi(x) = (x - K^*)_+$. The **Black–Scholes partial differential equation** states that the arbitrage-free price for the contingent claim is given in terms of the solution to the partial differential equation

$$\frac{\partial f}{\partial t} + \frac{1}{2}\sigma^2 x^2 \frac{\partial^2 f}{\partial x^2} = 0, \ \forall (t,x) \in (0,T) \times (0,\infty), \quad (4.5)$$

with the boundary condition

$$f(T,x) = \psi(x). \quad (4.6)$$

The arbitrage-free price of the claim is

$$C_0 = f(0, S_0). \quad (4.7)$$

Thus the Black–Scholes theory gives an 'explicit' expression for the arbitrage-free price of a European contingent claim in terms of the solution of a *fixed* partial differential equation, which is a diffusion equation. Only the boundary condition changes, depending on the payout function $\psi(\cdot)$. The Black–Scholes partial differential equation does not always have a closed-form solution. In the important case of option-pricing, where $\psi(x) = (x - K^*)_+$ where K is the strike price of the option, there is indeed a closed-form solution, and this leads to the expression (4.4) for the arbitrage-free price C_0. In a general situation however, no closed-form solution may be available and one would have to solve the partial differential equation numerically.

It is also possible to define a fully replicating self-financing strategy in terms of the solution $f(t,x)$ of the Black–Scholes partial differential equation (whether or not a closed-form formula is available for f). Define

$$C_t^* = C_0 + \int_0^t f_x(s, S_s^*) dS_s^*, \quad t \in (0,T), \quad (4.8)$$

where $S_t^* = e^{-rt}S_t$ is the discounted stock price, f_x denotes the partial derivative of f with respect to x, and the integral is a stochastic integral (since S_t^* is a stochastic process). At each time t, the seller of the contingent claim should hold α_t units of stock and β_t units of the bond, where

$$\alpha_t = f_x(t, S_t^*), \quad \beta_t = C_t^* - \alpha_t S_t^*. \tag{4.9}$$

Thus far we have studied the case where there is just one stock. However, Black–Scholes theory can be readily extended to the case where there are *multiple* stocks $S_t^{(i)}, i = 1, \ldots, d$ of the form

$$S_t^{(i)} = S_0^{(i)} \exp\left[\left(\mu^{(i)} - \frac{1}{2}[\sigma^{(i)}]^2\right)t + \sigma W_t^{(i)}\right], t \in [0, T], \tag{4.10}$$

where $W_t^{(i)}, i = 1, \ldots, d$ are Brownian motions that are possibly correlated. In this case, there is analog of the Black–Scholes partial differential equation (4.5). However, the p.d.e. does not always have a nice closed form solution.

5. Implications and usage of the Black–Scholes formula

While the mathematics in the previous section may be very difficult to follow for the uninitiated, there are a few points that emerge, which can be stated in purely non-mathematical terms.

1. For the case where the stock price follows a geometric Brownian motion model, there is only one arbitrage-free price, and it is given in terms of the solution of a partial differential equation. A closed-form formula is available for this arbitrage-free price for some situations, such as when the contingent claim is an option. This formula involves only the strike price of the option, the final time, and the volatility of the stock, but *not* the drift term of the Brownian motion. In general, if the payout is some complicated function of the final price, then no closed-form solution may be available, and one has to solve the Black–Scholes partial differential equation numerically to obtain the arbitrage-free price.
2. When the seller of a contingent claims receives this arbitrage-free price, he can then implement a fully replicating, self-financing trading strategy whereby, irrespective of the movement of the stock price, he will be in a position to meet the claim at the final time T.
3. This fully replicating, self-financing strategy requires continuous trading. This is not surprising because, in the binomial model, the holding of stocks and bonds is adjusted at each time instant. As a consequence, in the limiting case where the time interval goes to zero, the trading has to be continuous.
4. The implementation of the trading strategy assumes that the seller does not incur any *transaction costs*. Similarly, when his bond holdings go negative, meaning that he has to follow money to buy bonds, he is able to borrow money at the same rate of return as the bond itself. Thus there is no differential between the interest rates for borrowing and saving. Moreover, it is assumed that both bonds and stocks are available as real numbers (and not as multiples of some integer), in unlimited quantities, and also as both positive as well as negative numbers. All of these assumptions are captured by saying that the market is 'frictionless.'

The geometric Brownian motion model is unrealistic in practice. Data from several decades around the world shows that actual asset prices have 'fat tails.' In other words, the likelihood of stock prices straying beyond $\pm 3\sigma$ from the mean value are well in excess of that predicted by the normal (Gaussian) distribution. Moreover, there is a definite skew in the probability distribution, in that the tail on the high side is noticeably fatter than on the low side. In other words, the likelihood of a stock price straying far from the mean is higher than that predicted by the normal distribution; but this effect is more pronounced when the stock price rises than when it falls (though even a steep fall is also more likely than the normal distribution predicts). To do a better job of capturing the non-normality of asset price returns, researchers have introduced more general models of asset prices known as geometric Lévy processes. A Lévy process is more general than a Wiener process and can thus model fatter tails than the normal distribution. In subsequent sections we briefly discuss the problem of pricing claims when the asset price follows a geometric Lévy process model. Note that a Lévy process is the most general process with independent increments. This brings us squarely face to face with the notion of 'independent increments.' The assumption implies that whether today's return on a stock is positive or negative is completely independent of what happened yesterday or the day before. This is clearly not true! In the geometric Brownian motion model, the returns on the stock over nonoverlapping intervals are assumed to be independent random variables. Unfortunately, so far the research community has not come up with any tractable models of asset price movement that do away with the assumption of independent increments in the returns. Hence this is clearly an important area for future research.

Let us now take up a different issue, namely the actual application of the Black–Scholes theory 'in practice.' Suppose we wish to sell a contingent claim on a particular stock, for which historical data is available for many years if not many decades. Ignoring for the moment the fact that historical data rarely follows a geometrical Brownian motion model, we need to confront another important point, namely: historical data is rarely a good guide to the future! One could use the variation in the price of a stock to compute its volatility over a period of a day or thirty days. However, experience has shown that the volatility over the past thirty days (say) is rarely matched by the volatility over the *next* thirty days (say). Instead, the practitioners of Black–Scholes theory use a quantity called 'implied volatility.' In other words, they assume that 'the market is always right.' Since the actual players in the options market define the price of various options on popular and widely traded stocks, the analysts simply *retrofit* the option price on the stock to the formula (4.4), and from the formula compute σ. The quantity so computed is the 'implied volatility' of the stock, because the marketplace is pricing options on the stock assuming that this is the volatility. Since option prices reflect what the market thinks will happen in future, the implied volatility is believed to be a better guide to the future volatility of the stock than any quantity computed on the basis of historical prices.

The computation of the implied volatility throws up a few practical issues. First, on every major stock there are options sold at a variety of strike prices and at a variety of maturity dates, known as 'terms.' Thus, if a stock trades at \$ 100, there may be options with strike prices ranging from \$ 80 to \$ 120, in intervals of \$ 10. Similarly, there may be options with terms of thirty, sixty, or ninety days into the future. (Normally options are not sold with terms longer than ninety days.) Taking all of these into account, there may be dozens of options on any one stock! Now, for each option that is sold in the marketplace, it is possible to compute an associated implied volatility of the stock. Unfortunately, these implied volatilities do not always match. An option is said to be 'at the money' if its strike price is (more or less) equal to the current price. If the current price is less than the strike price, then it is said to be 'out of the money,' whereas if the current price is more than the strike price, then the option is said to be 'in the money.' Data and analysis of thousands of options on hundreds of stocks shows that, as a rule, the volatility implied by the 'at the money' options is less than the volatility implied by the 'in the money' or 'out of the money' stocks. This phenomenon is usually referred to as the *volatility smile*, because a plot of the implied volatility against the strike price shows a minimum at the current price and turns up both above and below the current price. Thus the plot resembles a parabola or a smile. Similarly, for the same strike price, as the maturity period is farther in the future, the implied volatility becomes higher. Thus a three-dimensional plot of the implied volatility against the strike price and the term resembles a sloping gutter pipe (or what the Americans would call an eavestrough), with the cross-section for a fixed term resembling a smile, and the cross-section for a fixed strike price showing an increasing function of the term length.

Now let us return to an issue first raised in section 3., namely that the replicating strategy essentially enables the seller of the claim to settle the outstanding claim at the end of the term and be left with precisely nothing. Moreover, this ideal outcome takes place only under two unrealistic assumptions, namely: **(i)** no transaction costs, and **(ii)** interest on borrowings and savings being equal. So what is the benefit of selling a claim and then hedging it?

There are in effect two kinds of traders in the marketplace. First, there are those who hold a definite view about whether a stock will go up or go down. If a person believes that a stock is undervalued and will go up, he can either buy the stock itself, or buy an option on the stock. Buying an option gives 'leverage,' in the sense that for the same quantum of investment, the investor can increase his potential profit significantly. To illustrate, suppose an investor has \$ 10,000 to invest, and believes that a stock currently priced at \$ 50 will rise to \$ 60 in the next sixty days. If he were to buy 200 units of the stock, and if the stock does indeed go up to \$ 60, then he would make a profit of \$ 2,000. On other hand, suppose an option with a strike price of \$ 60 and a term of sixty days is selling for \$ 2.50. Then the investor would be able to purchase 4,000 options. Suppose the stock price after sixty days is \$ 60. When the option matures, its value would become equal to \$ 10, which is the difference between the actual price and the strike price. Thus the investor would be able to sell the options at \$ 10, realize \$ 40,000, and make a profit of \$ 30,000. Thus by buying options instead of the underlying stock, the investor can make significantly more profit. On the other hand, the world is always uncertain, and buying options also carries greater downside risk. Suppose that the stock price drops to \$ 40 after sixty days. Then the strategy of buying 200 units of stock would leave the investor with a loss of \$ 2,000, while the strategy of buying options would entail a loss of the entire amount invested (because the options would be worthless at maturity). So an investor who sincerely believes that the stock will go up would buy the underlying stock, but limit his downside risk by simultaneously buying a put at a strike lower price than the current stock price.[1] Alternatively, the investor could buy a call option at a strike price higher than the current stock price, and simultaneously a put option at a strike price lower than the current stock price.

The second kind of investor usually holds no opinion on whether a stock will go up or down, but is only looking for situations where some option is priced differently from the trend. The underlying premise of such investors is that any deviations from trends will reverse themselves and align with the trend. Suppose that on some stock, there are a variety of options offered, each of which carries its own implied volatility. Now suppose that most of these implied volatilities follow the trend of the volatility smile, but *one option* deviates significantly from the trend. Then the investor would bet that the implied volatility would shortly revert to trend. If the implied volatility is lower than it should be, compared to other options, then the investor would bet that the implied volatility would increase, assuming all other factors remain the same. Since greater volatility implies greater option price, the investor would buy the option. If the implied volatility is higher than it should be, the investor would sell the option, or (more or less) equivalently, buy a put option. This is a highly simplified summary. In the 'real world,' investors use very sophisticated regression (or other statistical) methods to make their investment decisions.

All of the preceding discussion can be briefly summarized as follows: The Black–Scholes theory *should not* be viewed as a valid mathematical model of how 'real' markets work, in the same way that the Navier–Stokes equation accurately models aerodynamics. Because the Navier–Stokes equation is indeed a highly accurate model of aerodynamics, one can carry out very detailed and very deep analysis of the equation, comfortable in the knowledge that any predictions made by the theory would be borne out in practice. For instance, one can optimize the shape of an aircraft wing purely using computational fluid dynamics, and use wind tunnel tests solely as a verification of the theoretical predictions (as opposed to a design methodology as in decades past). Clearly, the Black–Scholes theory, or the more general theories to be advanced in subsequent sections, do not and cannot aspire to this level of realism or accuracy. Instead, the Black–Scholes theory gives a kind of gross approximation to the behavior of option prices, which individual investors can use to calibrate their own behavior, or to determine their own investment strategies.

[1]A put option is the reverse of the call option, and gives the buyer the right, but not the obligation, to *sell* the stock at a prespecified price at (or before) the maturity date.

6. Incomplete markets

Until now, the analysis has been facilitated by the fact that there exists a replicating strategy. In other words, there exists an investment strategy with the property that, whatever sample path the stock price follows, at the final time the value of the portfolio equals the payout. Moreover, the replicating strategy is also self-financing, in the sense that once the seller of the European contingent claim receives a certain amount at time $T = 0$, the value of the portfolio at any future time is sufficient to finance all subsequent adjustments in the portfolio. Finally, the replicating strategy offers a perfect hedge of the contingent claim because the final value of the portfolio equals the value of the claim. Such a market is said to be **complete**, and if there does not exist a replicating strategy, then the market is said to be **incomplete**.

The binomial model discussed in section 3. is an example of a complete market. One need not look very hard to find an example of an incomplete market. The simplest example that was studied in section 2., namely the one-period model, becomes incomplete if the number of possible stock prices is merely increased from two to three! To reformulate the problem, suppose there is a stock S and a bond B. At time 0, the bond has price $B(0)$, which increases in a deterministic fashion to $(1+r)B(0)$ at time $T = 1$. The stock price is $S(0)$ at time 0, and can assume one of three values, as follows:

$$S(1) = \begin{cases} S(0)u & \text{with probability } p_u, \\ S(0)m & \text{with probability } p_m, \\ S(0)d & \text{with probability } p_d, \end{cases}$$

where the numbers are arranged such that $d < m < u$ (and the notation is intended to suggest down, median, and up respectively). Of course $p_u + p_m + p_d = 1$. Suppose there is a contingent claim whose payout value equals X_u, X_m, X_d according as $S(1)$ equals $S(0)u, S(0)m, S(0)d$, respectively. Suppose we wish to find a pair of numbers (a, b) that replicate the movement of the stock price, and ensure that the value of the portfolio at time $T = 1$ exactly equals the payout. Thus we wish the pair of numbers (a, b) to satisfy

$$[a \;\; b] \begin{bmatrix} S(0)u & S(0)m & S(0)d \\ (1+r)B(0) & (1+r)B(0) & (1+r)B(0) \end{bmatrix} = [X_u \;\; X_m \;\; X_d].$$

Note that there are only two parameters that can be adjusted, namely a and b, while there are three equations to be satisfied. Thus in general these equations do not have a solution. As a result, no replicating strategy exists in general, and the market is incomplete.

The incompleteness of the market also has implications on the martingale measure defined in (2.1). Suppose that, as before, we wish to replace the original random variable $S(1)$ with another random variable $\mathtt{S}(1)$ that assumes the same three values os $S(1)$, but with probabilities q_u, q_m, q_d respectively. Moreover, these three probabilities should be chosen in such a way that (2.2) is satisfied. This requires that two equations be satisfied simultaneously, namely

$$1 + r = q_u u + q_m m + q_d d, \quad q_u + q_m + q_d = 1.$$

Clearly these equations do not have a unique solution because there are more parameters than equations. In contrast, if there are only two possible values as in section 2. (meaning that there is no middle price), then these equations have a unique solution given by

$$q_u = \frac{1 - d'}{u' - d'}, \quad q_d = \frac{u' - 1}{u' - d'},$$

which we have already seen. And finally, unless $d \leq 1 + r \leq u$, the above equations do not have a *nonnegative solution* for the q's, irrespective of whether the number of stock prices at time $T = 1$ is two or three. Thus the solution of the above equations would not be a *probability measure* unless $d \leq 1 + r \leq u$. The violation of the condition $d \leq 1+r \leq u$ leads to arbitrage in an obvious fashion. If $1 + r < d$, then one borrows money at the rate of interest r and invests it in stocks, whereas if $u < 1+r$, then one short-sells the stock and invests the proceeds at the guaranteed rate of return r. What this means is that unless the guaranteed rate of return on the bond is bracketed by the best and worst returns on the stock, no martingale measure exists.

One can now ask: How general are these conclusions? There are two tentative conclusions here. First, if it is possible to have arbitrage, then no martingale measure exists. Second, unless there is a replicating strategy, the martingale measure is not unique. It turns out that the first conclusion is perfectly general. There exists at least one martingale measure if and only if it is not possible to have arbitrage. As for the second, a slightly modified version of the conclusion is true. There is a replicating strategy if and only if the expected value of the payoff random variable is exactly the same under all martingale measures. In the next few paragraphs, we will make these two statements precise. Note that, in order to make the exposition simple, we do not always study the most general possible situation.

6.1 *Existence of a Martingale measure*

Suppose the marketplace offers a bond with a risk-free rate of return, as well as d different stocks whose prices are random variables. To make the notation simple, it is assumed that the price of the bond at time $T = 0$ is 1. Also, future prices of the stocks are discounted by the guaranteed rate of return on the bond, so that (in constant currency), the bond price is equal to 1 at all times n. At the beginning of the investment period, each stock $S^{(i)}$ has a deterministic price denoted by $S_0^{(i)}$. At the end of day n, the price of the i-th stock changes from $S_n^{(i)}$ to $S_{n+1}^{(i)}$. The investor, at the beginning of day n, chooses a set of weights $\theta_n^{(i)}, i = 0, \ldots, d$, where θ_0 is the quantum of investment in the bond.[2] Thus at the start of day n, the total wealth of the investor is $\bar{\theta}_n \cdot \bar{S}_n$, where

$$\bar{S}_n = [1 \;\; S_n^{(1)} \ldots S_n^{(d)}], \quad \bar{\theta}_n = [\theta_n^{(0)} \;\; \theta_n^{(1)} \ldots \theta_n^{(d)}].$$

At the end of day n, the price of the i-th stock moves from $S_n^{(i)}$ to $S_{n+1}^{(i)}$, and the investor's portfolio is worth $\bar{\theta}_n \cdot \bar{S}_{n+1}$. He then chooses the weights $\theta_{n+1}^{(i)}$ for the next day. To ensure that the portfolio is self-financing, the constraint

$$\bar{\theta}_{n+1} \cdot \bar{S}_{n+1} = \bar{\theta}_n \cdot \bar{S}_{n+1}$$

must be satisfied. Moreover, the vector $\bar{\theta}_n$ must be a measurable function of $\bar{S}_0, \ldots, \bar{S}_n$. The final set of weights to be chosen is $\bar{\theta}_{N-1}$, and at the end of the period, the portfolio is worth $\bar{\theta}_{N-1} \cdot \bar{S}_N$.

In this setting, we say that **there exists an arbitrage opportunity** if there exists a way of choosing the θ's such that

$$\bar{\theta}_{N-1} \cdot \bar{S}_N \geq 0 \;\; \tilde{P}\text{-a.s., and } \tilde{P}\{\bar{\theta}_{N-1} \cdot \bar{S}_N > 0\} > 0,$$

where $\tilde{P}$ is the joint probability law of the d stock prices over N time periods. Note that we make no assumptions whatsoever about the nature of the stock price random variables, such as independent increments, Markovian transitions, etc. Thus an arbitrage opportunity is just a way of choosing the portfolio in such a way that one never incurs a loss, and can make a profit with positive probability. In this very general setting, the following theorem holds (see the paper by Rogers [1]):

Theorem. The following statements are equivalent:

1. There exists a probability measure $\tilde{Q}$ on $\mathbb{R}^{dN}$ that is equivalent to $\tilde{P}$ such that, under $\tilde{Q}$, the process $\{S_n\}_{n=0}^N$ is a Martingale.

[2] Note that our notation differs from that of Rogers [1]; our θ_n is his θ_{n+1}.

2. There does not exist an arbitrage opportunity.

Moreover, if either of these equivalent conditions holds, then it is possible to choose the measure $\tilde{Q}$ in such a way that the Radon–Nikodym derivative $(d\tilde{Q}/d\tilde{P})$ is bounded.

In the above theorem, the measure $\tilde{Q}$ is said to be equivalent to $\tilde{P}$ if $\tilde{Q}(A) = 0$ if and only if $\tilde{P}(A) = 0$. In the case of the one-period, three-value model, the distribution (q_d, q_m, q_u) is equivalent to the original distribution (p_d, p_m, p_u) if and only if every component of q is positive.

The above theorem is popularly known in the literature as the Dalang–Morton–Willinger theorem [2]. However, the DMW theorem builds on a long list of earlier results due to various persons such as Harrison, Pliska and others, which are mentioned in the paper by Rogers [1]. Leaving aside all the technicalities, one can say roughly that there exists an equivalent martingale measure if and only if there is no opportunity for arbitrage.

6.2 *Existence of a replicating strategy*

Now let us get back to the one-period, three-price model. Then a distribution (q_d, q_m, q_u) is a martingale measure if and only if it satisfies the two conditions stated earlier, namely

$$1 + r = q_u u + q_m m + q_d d, \quad q_u + q_m + q_d = 1.$$

The distribution (q_d, q_m, q_u) is an *equivalent* martingale measure if and only if, in addition, each component of q is positive. Let $\mathcal{M}$ denote the set of all martingale measures. Then it is easy to show that $\mathcal{M}$ is a convex and closed subset of the set of all probability measures (in the topology induced by the total variation metric). Moreover, in many problems of interest $\mathcal{M}$ is a polyhedral set, and is thus the convex hull of a finite number of extremal measures. However, the set of *equivalent* martingale measures is the interior of the set $\mathcal{M}$, and is thus not a closed set (though it is convex).

What is so important about the set of martingale measures? If we replace the original probability distribution $\tilde{P}$ by any martingale measure $\tilde{Q}$, then *under the modified distribution $\tilde{Q}$ the stock prices form a risk-neutral process.* Thus, for a Europen contingent claim whose payout function X depends only on the final stock price vector $\bar{S}_N$, it could be argued that a 'fair' price for the claim at time $T = 0$ is the expected value

$$c_0 = E[X, \tilde{Q}], \tag{6.1}$$

where E denotes the expected value. In the above formula, we are using 'constant currency' so that there is no discount factor of the form $(1 + r)^{-N}$. Now the potential difficulty with the above formula is that the martingale measure $\tilde{Q}$ is not unique. In the binomial model, there is only one martingale measure, so the formula (6.1) gives an unambiguous fair price for the contingent claim. In general, however, one would have to define two quantities

$$V_-(X) := \min_{\tilde{Q} \in \mathcal{M}} E[X, \tilde{Q}], V_+(X) := \max_{\tilde{Q} \in \mathcal{M}} E[X, \tilde{Q}]. \tag{6.2}$$

Any and all numbers in the interval $[V_-(X), V_+(X)]$ are 'fair prices.' The seller of the claim could, quite reasonably, insist on being paid the amount $V_+(X)$, since that is the maximum he can expect under the condition that the stock price movements are risk-neutral. Moreover, it is also the *minimum* price under which he would be able to super-hedge against all possible stock price movements. On the other hand, the buyer of the claim could, equally reasonably, be prepared to pay only $V_-(X)$. This is also the maximum price that he would be ready to pay and still be able to make a profit under all outcomes using sub-hedges of his own. In this context, the following theorem assumes significance. See the book by Williams [3], Corollary 3.5.2:

Theorem. A European contingent claim with the payout function X is replicable if and only if $V_-(X) = V_+(X)$.

In other words, a European contingent claim is replicable if and only if there is a unique fair price. In principle there could be more than one martingale measure, but all martingale measures must lead to the same fair price. Under a few technical conditions, it can be shown that a European contingent claim is replicable if and only if there is a unique martingale measure.

From the above discussion, it is clear that any price for the claim that lies outside the interval $[V_-(X), V_+(X)]$ leads to arbitrage. If the seller of a claim were to receive an amount greater than $V_+(X)$, then he could use $V_+(X)$ of the amount to achieve a replication of the claim, and pocket the difference; this leads to arbitrage. A similar argument holds in the case where the claim is priced below $V(X)$. On the other hand, every price in the interval $[V - (X), V_+(X)]$ is a fair price, since there exists *some* martingale measure that leads to this expected value for the contingent claim. At this point, it is up to the buyer and seller of the claim to negotiate a mutually acceptable price.

6.3 *A dual problem formulation*

We conclude this section by demonstrating a pair of interesting convex optimization problems.

In this subsection, we directly tackle the problem of finding the maximum and minimum 'hedgable' prices for a contingent claim, and show that the dual problem formulation leads very naturally to martingale measures. In the next subsection, we address incomplete markets and ask: Out of all possible martingale measures, which one should we choose? A popular answer is to choose one that has minimum relative entropy with respect to the original probability distribution followed by the stock. It is shown that this is a fairly straight-forward convex optimization problem that can be solved iteratively. The message of this subsection is that convex optimization problems arise naturally in finance problems. To keep the exposition simple, in both cases we address only one-period problems with a single 'uncertain asset'. But the theory itself can be extended without difficulty to multi-period problems with multiple uncertain assets, with more cumbersome notation.

Let us begin by studying the problem of sub-hedging and super-hedging a random stock price over one time-period. Thus there is a bond whose value at time $T = 0$ is $B(0)$, and whose value at time $T = 1$ is $(1 + r)B(0)$, where r is the risk-free return. There is also a stock whose price at time $T = 0$ is a deterministic quantity $S(0)$, and which assumes values $u_1, \ldots, u_n$ with probabilities $p_1, \ldots, p_n$ respectively. Finally, suppose there is a contingent claim X whose value equals x_i if the stock price equals u_i. As we have already seen, if $n \geq 3$, then the market is incomplete, and there is no replicating strategy. So the seller of the claim wishes to protect himself against all possible outcomes in the stock price. Thus he wishes to choose numbers a, b such that

$$au_i + b(1 + r)B(0) \geq x_i, \quad i = 1, \ldots, n,$$

or in matrix form

$$[a \;\; b] \begin{bmatrix} u_1 & \ldots & u_n \\ (1+r)B(0) & \ldots & (1+r)B(0) \end{bmatrix} \geq [x_1 \ldots x_n].$$

The inequalities can be written in obvious vectorial notation as

$$[a \;\; b] \; M \geq \mathbf{x},$$

where

$$M = \begin{bmatrix} u_1 & \ldots & u_n \\ (1+r)B(0) & \ldots & (1+r)B(0) \end{bmatrix}.$$

Any choice of the pair (a, b) by the seller of the claim that satisfies the above inequalities can be referred to as a 'super-hedging strategy' since the choice will enable him to meet the value of the contingent claim irrespective of the outcome of the stock price at time $T = 1$.

Now let us ask the question: What is the lowest price that the seller of the claim would be willing to accept that would still permit him to super-hedge? The answer is

$$V_+(X) = \min_{a,b} [a \;\; b] \begin{bmatrix} S(0) \\ B(0) \end{bmatrix} \text{ s.t. } [a \;\; b]M \geq \mathbf{x}.$$

Since this is a linear (and thus convex) optimization problem, we can compute the value of the minimum via the dual problem formulation. Let $\boldsymbol{\phi}' \in \mathbb{R}^n$ denote the variables in the dual problem. This leads to the formulation

$$V_+(X) = \max_{\boldsymbol{\phi}'} \mathbf{x}\boldsymbol{\phi}' \text{ s.t. } M\boldsymbol{\phi}' = \begin{bmatrix} S(0) \\ B(0) \end{bmatrix}, \quad \boldsymbol{\phi}' \geq \mathbf{0}.$$

Now define $\boldsymbol{\phi} = (1 + r)\boldsymbol{\phi}'$. Then the dual problem formulation can be rewritten as

$$V_+(X) = \max_{\boldsymbol{\phi}} (1 + r)^{-1}\mathbf{x}\boldsymbol{\phi} \text{ s.t.}$$

$$(1 + r)^{-1}M\boldsymbol{\phi} = \begin{bmatrix} S(0) \\ B(0) \end{bmatrix}, \quad \boldsymbol{\phi} \geq \mathbf{0}.$$

Let us now examine the constraints on $\boldsymbol{\phi}$. In expanded form, this can be written as

$$\begin{bmatrix} u_1^* & \ldots & u_n^* \\ B(0) & \ldots & B(0) \end{bmatrix} \boldsymbol{\phi} = \begin{bmatrix} S(0) \\ B(0) \end{bmatrix}, \quad \boldsymbol{\phi} \geq \mathbf{0},$$

where $u_i^* = (1 + r)^{-1}u_i$ denotes the *discounted* i-th outcome of the stock price at time $T = 1$. The second row, together with the constraint $\boldsymbol{\phi} \geq \mathbf{0}$, shows that $\boldsymbol{\phi}$ is a probability distribution, because its components are nonnegative and add up to one. The first row works out to

$$E[S(1), \boldsymbol{\phi}] = S(0),$$

which is to say: Under the measure $\boldsymbol{\phi}$, the stock price $S(1)$ is a martingale. Now the quantity to be maximized is

$$(1 + r)^{-1}\mathbf{x}\boldsymbol{\phi} = \sum_{i=1}^{n} x_i^* \phi_i = E[X^*, \boldsymbol{\phi}],$$

which is the discounted expected value of the payout with respect to the martingale measure $\boldsymbol{\phi}$. Hence the minimum price that the seller of a contingent claim can accept while still being able to hedge all possible outcomes is given by

$$\begin{aligned} V_+(X) &= \max_{\phi \in \mathbb{S}_n} E[X^*, \boldsymbol{\phi}] \text{ s.t. } E[S(1), \boldsymbol{\phi}] = S(0) \\ &= \max_{\phi \in \mathcal{M}} E[X^*, \boldsymbol{\phi}], \end{aligned} \tag{6.3}$$

where $X^* = (1 + r)^{-1}X$ is the discounted claim function, $\mathcal{M}$ denotes the set of martingale measures, and $\mathbb{S}_n$ denotes the set of all probability distributions with n components. Thus by using the dual formulation, we readily obtain the formulas in (6.2) for the highest and lowest fair prices, as being the maximum and minimum discounted expected values of the payout function over the set of martingale measures.

In many books on mathematical finance, the optimization is carried out over the set of *equivalent* martingale measures, rather than the set of *all* martingale measures. Note that a distribution $\boldsymbol{\phi}$ is equivalent to the original distribution $\mathbf{p} = (p_1, \ldots, p_n)$ if and only if $\phi_i > 0$ for all i. Thus the set of all equivalent martingale measures is the *interior* of the set $\mathcal{M}$. Moreover, while $\mathcal{M}$ is a polyhedral set, being the convex hull of a finite number of extremal distributions, the set of all equivalent martingale measures does not have any such nice structure. Finally, if one wishes to restrict to equivalent martingale measures, then the optimization problem does not attain its extremal value. For all these reasons, the present author prefers to speak of just martingale measures than equivalent martingale measures.

The above discussion can be extended readily, albeit with more notation, to multi-period models where the number of outcomes is finite at each period. The problem of finding the lowest acceptable price that would still allow the seller of the claim to hedge all possible outcomes at time $t = 1$ can once again be formulated as a linear programming problem. The dual problem formulation once again leads to the maximization of the expected value of the discounted payoff over the set of martingale measures. There are just a couple of additional technical issues to watch out for. First, the hedging strategy must be *self-financing*, and second, the hedging strategy must be *non-anticipative*, in that the investment adjustments at time n can depend only on the stock prices until time n. Neither of these features arises in the one-period case.

6.4 *Minimum relative entropy Martingale measures*

As we have seen, in an incomplete market the set of martingale measures consists of more than one measure. It is obvious that the set of martingale measures is convex, because if a stochastic process is a martingale under two measures $\boldsymbol{\phi}$, $\boldsymbol{\psi}$, then it is also a martingale under their convex combination $\alpha\boldsymbol{\phi} + (1-\alpha)\boldsymbol{\psi}$ for each $\alpha \in (0,1)$. Thus in reality the set of martingale measures is either a singleton set (which can happen only in a complete market), or else is an infinite set. This raises the question: What is the 'right' martingale measure to use? This topic has been discussed by various authors, and in this subsection we present one such choice, known as the MREMM (minimum relative entropy Martingale measure).

To make the discussion simple, we once again consider a one-period problem. As before, there is a bond whose value at time $T = 0$ is $B(0)$, and whose value at time $T = 1$ is $(1+r)B(0)$, where r is the risk-free return. There is also a stock whose price at time $T = 0$ is a deterministic quantity $S(0)$, and which assumes values $u_1, \ldots, u_n$ at time $T = 1$ with probabilities $p_1, \ldots, p_n$, respectively. Let P denote the distribution $(p_1, \ldots, p_n)$, and let $Q = (q_1, \ldots, q_n)$ be any other distribution over the same set of n outcomes of the stock price at time $T = 1$. The quantity

$$H(Q \parallel P) = \sum_{i=1}^{n} q_i \ln\left(\frac{q_i}{p_i}\right)$$

is called the **relative entropy**, or the **Kullback–Leibler divergence** between Q and P. Therefore, given that the set of Martingale measures is not a singleton (and is thus infinite), one way to choose the 'right' Martingale measure is to choose Q as the solution of the problem

$$\min_{Q \in \mathcal{M}} H(Q \parallel P),$$

whereas before $\mathcal{M}$ is the set of Martingale measures. Naturally, such a Q is called the **MREMM (minimum relative entropy Martingale measure)**. We first show that this problem is easily solved through an iterative technique. Then we reproduce, very briefly, the arguments from Frittelli [4] on the interpretation of the MREMM.

Let us begin by reformulating the problem. Since the stock can assume values $u_1, \ldots, u_n$, the no-arbitrage condition, which is equivalent to the condition that Q should be a Martingale measure, can be written as

$$\sum_{i=1}^{n} q_i u_i = (1+r).$$

Thus the optimization problem at hand is

$$\min_{\mathbf{q}} \sum_{i=1}^{n} q_i \ln\left(\frac{q_i}{p_i}\right) \text{ s.t.}$$

$$q \geq \mathbf{0}, \mathbf{q}^t\mathbf{e}_n = 1, \quad \mathbf{q}^t\mathbf{u} = (1+r),$$

where $\mathbf{q} = (q_1, \ldots, q_n)$, $\mathbf{e}_n$ denotes a column vector with n one's, and $\mathbf{u} = [u_1 \ldots u_n]^t$. To solve the problem, we use Lagrange multipliers. From the augmented objective function

$$J = \sum_{i=1}^{n} q_i \ln\left(\frac{q_i}{p_i}\right) + \lambda\left[1 - \sum_{i=1}^{n} q_i\right] + \gamma\left[1 + r - \sum_{i=1}^{n} q_i u_i\right].$$

Then

$$\frac{\partial J}{\partial q_i} = \ln \frac{q_i}{p_i} + 1 - \lambda - \gamma u_i.$$

Setting this quantity equal to zero for all i shows that

$$\ln \frac{q_i}{p_i} - \gamma u_i = \lambda - 1$$

is independent of i. Thus $\ln(q_i/p_i) - \gamma u_i$ is a constant, independent of i. This shows that q_i/p_i is proportional to $\exp(\gamma u_i)$, or q_i is proportional to $p_i \exp(\gamma u_i)$. We can determine the constant of proportionality by using $\mathbf{q}^t \mathbf{e}_n = 1$. Thus

$$q_i = \frac{p_i e^{\gamma u_i}}{\sum_{j=1}^n p_j e^{\gamma u_j}} = \frac{p_i e^{\gamma u_i}}{E[e^{\gamma S(1)}, P]}, \tag{6.4}$$

where $S(1)$ is the stock price at time $T = 1$ and has the distribution P.

This still leaves the unknown constant γ. This constant can be determined from the no arbitrage condition $\mathbf{q}^t \mathbf{u} = 1+r$. Substituting for q_i from (6.4) leads to

$$\frac{\sum_{i=1}^n p_i u_i e^{\gamma u_i}}{E[e^{\gamma S(1)}, P]} = 1 + r,$$

or

$$\frac{E[ue^{\gamma u}, P]}{E[e^{\gamma u}, P]} = 1 + r.$$

This equation is of the form $g(\gamma) = 1 + r$, where

$$g(\gamma) := \frac{E[ue^{\gamma u}, P]}{E[e^{\gamma u}, P]}. \tag{6.5}$$

Now it is shown that

$$g'(\gamma) > 0 \ \forall \gamma, \text{ unless all } u_i\text{'s are equal.}$$

To show this, note that

$$g'(\gamma) = \frac{E[e^{\gamma u}, P] E[u^2 e^{\gamma u}, P] - \{E[ue^{\gamma u}, P]\}^2}{\{E[ue^{\gamma u}, P]\}^2}.$$

Now the denominator in the above equation is always positive, while the numerator is positive provided that

$$\{E[ue^{\gamma u}]\}^2 < E[e^{\gamma u}] E[u^2 e^{\gamma u}].$$

This last inequality follows readily from Schwarz' inequality factor

$$ue^{\gamma u} = e^{\gamma u/2} \cdot ue^{\gamma u/2}.$$

Then Schwarz' inequality implies that

$$\{E[ue^{\gamma u}]\}^2 \leq E[e^{\gamma u}] E[u^2 e^{\gamma u}].$$

with equality if and only if $e^{\gamma u/2}$ and $ue^{\gamma u/2}$ are proportional, i.e., if and only if all components of $\mathbf{u}$ are the same. It is routine to show that

$$g(\gamma) \to \min_i u_i \text{ as } \gamma \to -\infty, g(\gamma) \to \max_i u_i \text{ as } \gamma \to \infty.$$

Hence, provided

$$\min_i u_i < 1 + r < \max_i u_i, \tag{6.6}$$

there is a unique γ such that $g(\gamma) = 1 + r$. Clearly (6.6) is a necessary condition to have a meaningful problem.

Thus, to solve the equation $g(\gamma) = 1 + r$, we can use a simple bisection search method. In view of (6.6), we can first find numbers γ_-, γ_+ such that

$$g(\gamma_-) < 1 + r < g(\gamma_+).$$

Then we bifurcate the interval $[\gamma_-, \gamma_+]$ and compute $g(\cdot)$ at the midpoint. If g exceeds $1 + r$ at the midpoint, we bisect the left half-interval, whereas if g is less than $1+r$ at the midpoint, then we bisect the right half-interval.

We conclude this subsection by discussing the significance of the MREMM. As shown in (6.4), the weight q_i of the MREMM is proportional to the exponential of the stock price u_i. Thus the weights of the MREMM can be viewed as being proportional to the marginal utility of terminal wealth when we use exponential utility functions. See the paper of Frittelli [4] for an elaboration of this argument, as well as a treatment of the multi-period case.

7. Alternate models for asset prices and options

7.1 *Alternate models for asset prices*

Until now the discussion has been focused on two distinct situations: discrete-time, discrete-valued processes, and continuous-time, continuous-valued processes. In the latter case, the price of the stock has been modeled in terms of geometric Brownian motion, namely (4.2), whereas the price of the bond has been modeled as an exponential function with a constant exponent, as in (4.3). However, this model has a number of drawbacks.

1. Analysis of historical data of thousands of stock prices over decades shows that the actual price distribution is in fact *not log-normal.* As mentioned earlier, the actual price distributions

tend to have fatter tails than the normal distribution would suggest. Moreover, there is a further skew towards the higher side, meaning that the tails are fatter on the side of increasing stock prices than on the side of decreasing stock prices. Thus alternative (and presumably more realistic) distributions would be needed.
2. Even if one were to accept the models (4.2) for stock prices and (4.3) for bond prices, there is no reason to accept that the volatility σ and the interest rate r are both fixed and constant. Anyone reading the newspaper everyday would observe that the interest rates on even the safest government bonds tend to fluctuate daily due to a variety of external factors, none of which can be predicted ahead of time. Thus it would be far more realistic to treat both the risk-free interest rate and the volatility of the stock as being themselves stochastic processes.

In either of these situations, the analysis given in section 4. based on Black–Scholes theory would not apply.

To address the first shortcoming, the research community has been studying the use of geometric Lévy processes, as an alternative to geometric Wiener processes. (Recall that a Wiener process is another name for Brownian motion.) To introduce Lévy processes, we first talk about Poisson processes.

A **Poisson process** (also called a 'jump process,' which explains the notation) is a continuous-time, integer-valued process $\{J_t\}_{t\geq 0}$ with the following properties:

1. $J_0 = 0$ almost surely.
2. For $t > s \geq 0$, the distribution of $J_t - J_s$ depends only on the difference $t - s$ and not on t and s individually.
3. For Δ approaching zero, we have that $\Pr\{J_{s+\Delta} - J_s = 1\} = \lambda\Delta$, where λ is called the *rate* of the process. Along the same vein, we have that $\Pr\{J_{s+\Delta} - J_s \geq 2\} = o(\Delta)$. Thus, in an interval of width Δ, the integer-valued process $\{J_t\}$ increases by one with probability roughly equal to $\lambda\Delta$, while the probability that *two or more jumps* occur in a small interval is essentially equal to zero.

With these assumptions, it can be shown that, over an interval $[0, T]$, the number of jumps has the following distribution:

$$\Pr\{J_T = k\} = \frac{(\lambda T)^k}{k!}\exp(-\lambda T), \quad k = 0, 1, \ldots.$$

If we ignore for a moment the fact that k is an integer and maximize the right side with respect to k viewed as a real variable, it is easy to see that the maximum is attained at λT. Thus the most likely number of jumps in an interval of length T is λT, which is why λ is called the rate of the Poisson process.

Now we come to a Lévy process. We can define a Lévy process using the same set of axioms as a Wiener process, *except that* we say nothing about the distribution of the differences. Specifically, a process $\{L_t\}_{t\geq 0}$ is a **Lévy process** if

1. Each L_t is a random variable defined on a common measurable space $(\mathbb{R}, \mathcal{B})$.
2. $L_0(\omega) = 0$ almost surely with respect to ω. In other words, $L_0(\omega) = 0$ except for possibly on a set of measure zero.
3. The process has 'independent increments.' Thus, whenever $s < t \leq u < v$, the random variables $L_t - L_s$ and $L_v - L_u$ are independent.
4. The process has 'stationary increments.' Thus the distribution of $L_t - L_s$ is a function only of the difference $t-s$ and not of t and s separately.

Thus a Lévy process is the most general process with stationary increments, starting from $L_0 = 0$ almost surely. As it turns out, the seemingly simple requirement of stationary increments tremendously constrains what a Lévy process can look like. In a tutorial introduction like this one, we wish to avoid getting into too many technicalities. So to put it a bit sloppily, every Lévy process can be expressed as the sum of three processes: (i) A Wiener process with drift term, (ii) a Poisson process, and (iii) a 'pure jump' martingale that can have a countable number of jumps in every finite interval. This decomposition corresponds to the fact that every measure can be expressed as the sum of three types of measures: an absolutely continuous (with respect to the Lebesgue measure) component, a purely atomic component, and a singular component that is supported on a set of Lebesgue measure zero. A reader interested in a proper elaboration of the representation of Lévy processes is referred to section 1.4 of the book by Protter [5].

Once we have (more or less) understood the notion of a Lévy process, we can replace the geometric Brownian motion model of (4.2) with the more general representation

$$S_t = S_0 \exp(L_t),$$

where $\{L_t\}$ is a Lévy process. Thus, instead of modeling the stock price as a geometric Wiener process in (4.2), we now model it as a geometric Lévy process.

The geometric Lévy process has greater expressive power than the geometric Brownian motion model. In particular, by choosing the Lévy process appropriately, one can achieve fatter tails than the normal distribution. On the other side, the key negative factor is that *the only complete market where the stock price follows a geometric Lévy process model is one with a geometric Brownian motion model.* Thus, the greater realism achieved by using a non-normal distribution is partially negated by the fact that there is not a unique martingale measure. Instead, one has to cope with the fact that there are infinitely many martingale measures, and then try to choose some hedging strategy.

Among the more tractable versions of geometric Lévy processes are a class known as 'variance gamma' processes. These are discussed in the book by Fu *et al* [6].

7.2 *Alternate types of options*

Until now the emphasis has been on so-called European contingent claims, whose defining feature is that they can be exercised only *at a fixed point in time*, which we have been denoting by T, and been referring to as the 'term.' However, there are other, indeed more popular, types of options. Among the most popular is the so-called **American contingent claim**, which can be exercised *at any time up to and including* some fixed point in time. To highlight the difference, suppose there is a stock whose price at time t is denoted by S_t, and suppose a trader sells an option with a strike price of K and maturity term of T (taking the present as time 0). If the option is European, the buyer of the option has to hope that the price of the stock will be in excess of the strike price at time T. In other words, he has to hope that $S_T \geq K$. In contrast, if the option is American, then the buyer of the option has to hope that the price of the stock will exceed the strike price K *at some time during the interval* $[0, T]$. If there exists some intermediate time $t \in (0, T)$ such that $S_t \geq K$, then the buyer of the option is 'in the money' at that point in time. So he could exercise the option and acquire the stock by paying the strike price of K, immediately sell it for S_t, and make a profit of $(S_t - K)_+$. But then he faces a (potentially agonizing) choice of determining whether to exercise the option at once (that is to say, 'take the money and run'), or to wait in the hope that the stock price will climb still higher before time T. He could also turn around and sell his own option on the stock at a strike price higher than K.

For the seller of the option, the hedging strategies against a American claim will be different than the hedging strategies against an European claim. In the case of a European claim, the seller will implement a strategy that will ensure that, *at the end of the term*, the value of his hedging portfolio will equal or exceed that of the claim. But in the case of an American claim, the seller will have to implement a strategy to ensure that the strike price is covered *at all times up to and including time* T. This kind of strategy is called a 'super-hedging' strategy.

In general the pricing theory for American contingent claims is more complicated, and more 'incomplete' than for European contingent claims. In the special case where the stock price follows the geometric Brownian motion model of (4.2) and the payout is $(S_t - K)_+$ (in other words, the contingent claim is just an option), it can be shown that the correct price is the same as with the European option, namely (4.4). Not much is known about other situations however.

There are yet more claims besides European and American. In the case of the Bermudan option, a term T is specified, along with various intermediate times $T_1, \ldots, T_k = T$; the option can be exercised only the discrete time instants T_i, $i = 1, \ldots, k$, but not at intermediate times. Needless to say, the analysis of Bermudan options is extremely complex. Finally there is the Russian option, in which the term T equals infinity, and the right to exercise the option is triggered by some other considerations. It turns out that the analysis of Russian options is simpler than that of Bermudan options. The book by Kallianpur and Karandikar [7] is one of the few to discuss Russian options.

7.3 *Sensitivities and the 'Greeks'*

We conclude this section with a very brief discussion of the sensitivities of the value of a contingent claim with respect to various parameters, popularly known in the literature (or perhaps more accurately, amongst the trading community as opposed to the theorem-proving community) as 'the greeks.' A very thorough discussion of the greeks is found in the wikipedia article on the topic; see [8].

To make the discussion simpler, let us stick to the Black–Scholes framework as set out in section 4.2. In this case, there is a function $f(t, x)$ that satisfies the Black–Scholes partial differential equation (4.5) with the associated boundary condition (4.6), where $\psi(\cdot)$ is the payout function. Then the arbitrage-free value of a European contingent claim at time $T = 0$ is given by $C_0 = f(0, S_0)$, as stated in (4.7). With this background, some of the greek symbols are defined as follows:

$$\Delta = \frac{\partial C_0}{\partial S_0}, \quad \Gamma = \frac{\partial \Delta}{\partial S_0} = \frac{\partial^2 C_0}{\partial S_0^2}, \quad \text{Vega} = \nu = \frac{\partial C_0}{\partial \sigma},$$

$$\theta = -\frac{\partial f(t, X_t)}{\partial t}, \quad \rho = \frac{\partial f(t, X_t)}{\partial r}.$$

The greeks are used to devise fairly elaborate hedging strategies that are insensitive to various assumptions. For instance, a ‘delta-hedging’ strategy is one whose sensitivity to S_0 is zero. Thus, to a first-order approximation, the return is insensitive to the initial stock price. A ‘delta-gamma-hedging strategy’ is one that is insensitive to the initial stock price to a second-order approximation. And so on.

In case the contingent claim is a simple option, then we have an explicit formula for $C_0 = f(0, S_0)$ in (4.4). In this case it is possible to write down explicit formulas for the various greeks; this is done in [8]. However, the definition of the greeks is not restricted to the case where the contingent claim is an option. So the question arises: In case we are not able to compute the solution to the Black–Scholes partial differential equation in closed form, how can we compute the greeks? More generally, suppose we abandon the geometric Brownian motion model for a geometric Lévy process model. In this case we do not even have a unique martingale measure, nor a replicating strategy. How then can we compute the greeks corresponding to whatever claim-pricing strategy we may adopt?

When no closed-form solutions are available for the price of the claim, we can resort to Monte Carlo type of simulations. The book by Seydel [9] discusses many of these issues. In such a case, the price of the claim is itself determined through simulation, and if we wish to compute the sensitivity of the price to various parameters, we would have to resort to numerical differentiation, which is likely to be highly susceptible to numerical errors. In this connection, Malliavin calculus is an excellent tool. By using Malliavin calculus, it is possible to express all of the greeks as stochastic integrals. When no closed-form solution is available for the price of a claim, the stochastic integral representing a greek (i.e., a sensitivity) would still have to be estimated via Monte Carlo simulation. But at least one can avoid multiple simulations and numerical differentiation of the resulting values. Originally Malliavin calculus was developed for an entirely different problem, in the 1970s, and its applicability to problems of mathematical finance, especially to the computation of the greeks, was discovered in the 1990s. In the original form, Malliavin calculus is applicable only to geometric Brownian motion; see the book by Malliavin and Thalmaier [10] for an exposition. However, the theory has been extended to more general processes in stages, and finally, the theory has been extended to arbitrary geometric Lévy processes in the PhD thesis of Petrou [11].

8. What went wrong?

Any discussion of financial engineering would be sterile if it did not take into account the current economic climate. Therefore I will use this section to discuss the burning question ‘What went wrong?’[3] At the time of this writing, the world economy is in a tail spin, caused primarily by various financial institutions losing enormous sums of money. The US government has had to inject a stimulus of more than *one trillion dollars* into its economy, as well as hand out what amount to blank cheques to various banks and financial institutions that are deemed to be ‘too big to fail.’ By some estimates, the accumulated losses of the financial institutions around the world during the past eighteen months are more than their total profits in all of recorded history! Thus, in one stroke, banking as an activity has become ‘net negative.’ It is therefore pertinent to ask what role if any financial engineering has had in this meltdown.

There is no denying the fact that financial engineering as it is currently practised makes heavy use of the models and results briefly reviewed in the present article. There cannot also be any doubt that several practioners of financial engineering were blindly applying various formulas, neither knowing nor caring what assumptions led to these formulas, and whether these assumptions held in their specific situation. Finally, existing option pricing theory is based (like much of economics) on assumptions of rationality, enlightened self-interest, equilibrium dynamics, etc. Attempts to incorporate behavioural factors such as the herd mentality, greed, nonequilibrium dynamics, etc. into option pricing are still in their infancy. So one can legitimately ask: Was the current financial crisis caused by various statisticians seriously underestimating the risks involved?

In my view, the current financial debacle owes very little to poor modeling of risks. Far more significant factors were a lax to nonexistent regulatory environment, which encouraged reckless risk-taking by financial institutions; a ‘one-way’ reward system for individual traders; the phenomenon of ‘capitalism for the masses, socialism for the rich,’ etc. In subsequent paragraphs, I explain each of these factors briefly.

One of the consequences of the great depression in the USA was the passage of the Glass–Steagalls Act, which clearly separated banking from investment. As a result, there was a clear dividing line between ‘investment’ banks which speculated with the money given to them by investors for that very purpose, and ‘traditional’ banks which were entrusted by the public at large with its money for safe-keeping, not speculation. As a part of

[3]Since this section contains only my personal opinions, I have chosen to write it in the first person, instead of using the usual impersonal style appropriate for a scientific paper.

this separation, traditional banks were obliged to maintain certain minimal levels of reserves, and deposits in the bank were insured by the government (up to some limit per individual). In contrast, investment banks were both exempt from many regulations that governed traditional banks, and at the same time, the depositors were not protected by government insurance. The underlying premise was that customers of investment banks were wealthy individuals who could afford the risks taken by such banks, whereas customers of traditional banks were mostly interested in safety first. The Glass–Steagalls act was repealed in the 1990's, thus permitting traditional banks also to speculate with the money given to them for safekeeping. Moreover, in the absence of proper regulations, banks and other financial institutions were allowed to create various 'special purpose vehicles,' resulting in 'off-balance sheet transactions.' These off-balance sheet transactions effectively meant that the so-called balance sheets of banks did not reflect their true financial (ill-)health.

At the level of individual traders employed by financial institutions, the reward system was 'one-way' in the following sense: The traders were rewarded by way of huge bonuses when the value of their portfolios went up, but were *not penalized* when the value went down. In this kind of 'heads I win, tails the shareholders/depositors lose' scenario, it is not surprising that individual traders took enormous risks. The situation of one-way rewards to traders was exacerbated by another factor that may be called 'real bonuses on virtual profits.' When a trader is said to have 'made money' during a particular period, all it meant was that *the paper value* of his portfolio went up; but since the portfolio was not actually liquidated at that point, the profits were virtual. Nevertheless, the bonuses paid out to the traders were quite real in the form of hard cash or shares in the parent company. When the portfolios were finally liquidated, often years after the bonuses were paid out, the virtual profits had disappeared in many cases, along with the traders themselves and their very real bonuses.

When the financial institutions got into really deep trouble due to their reckless investments, they retreated behind the argument that they were 'too big to fail' and thus had to be bailed out by the taxpayers. Thus the USA now epitomizes the credo of 'capitalism for the masses, and socialism for the rich.'

There is *one aspect* of the financial crisis that can legitimately be blamed on the modelers, in my view. In reality, a financial instrument (either a stock or a derivative written on that stock) does not have any 'intrinsic' value, and its value is whatever the marketplace collectively thinks it is. In the case of traditional options of the kind studied here, there are 'exchanges' where these are freely traded, thus allowing the marketplace as a whole to determine their price. The quest for profits spurred the trading community to come up with ever more exotic (i.e., non-standard) derivative instruments, many of which were traded 'over the counter,' and not on an exchange. Thus there was no price discovery process, and the price for each transaction was negotiated between the two parties. In such a scenario, at least a first-cut approximation of the price was generated by the statisticians (the so-called 'quants'). In such a situation, the statisticians had an *obligation* to highlight their lack of certitude about the quality of their analysis, and of the various outcomes of the analysis (such as the fair price, value at risk, probability of failure/default, etc.). I do not see sufficient evidence that this actually happened. Having said that, however, I am firmly of the view that, even if extremely precise methods had been available for determining the true values, and/or assessing the risk levels of these instruments had been available, there was simply no incentive for the investment houses to pay any attention to the statisticians, and all incentives for them to ignore the statisticians. Thus in my view the community of statisticians has had a very marginal role in the financial collapse.

9. Conclusions

In this paper we have seen rather elementary versions of some interesting approaches to finance, which is often referred to as 'financial engineering.' While there is no doubt that the underlying mathematics is very beautiful, that does not address the issue of whether financial engineering really is 'engineering' in the same sense as designing VLSI chips, or optimizing aircraft wing shapes.

The answer is a clear 'no.' In a very readable article, Emanuel Derman [12] says:

> "In finance we study how to manage funds ... Physics, because of its astonishing success at predicting the future behavior of material objects from their present state, has inspired most financial modeling. ... The method works. The laws of atomic physics are accurate to more than ten decimal places. ... Financial theory has tried hard to emulate the style and elegance of physics in order to discover its own laws. But markets are made of people, who are influenced by events, by their ephemeral feelings about events and by their expectations of other people's feelings. The truth is that there are no fundamental laws in finance."

So what *should* financial engineers be doing? It may not be inappropriate again to quote Derman [12], specifically the 'Modeler's Hippocratic oath':

- I will remember that I didn't make the world, and it doesn't satisfy my equations.
- Though I will use models boldly to estimate value, I will not be overly impressed by mathematics.
- I will never sacrifice reality for elegance without explaining why I have done so.
- Nor will I give the people who use my model false comfort about its accuracy. Instead, I will make explicit its assumptions and oversights.
- I understand that my work may have enormous effects on society and the economy, many of them beyond my comprehension.

This is an excellent philosophy to keep in mind while studying the subject called financial engineering.

10. Suggested reading

Since this article is a tutorial introduction and not a survey, the bibliography is rather short, and is not in any way meant to be exhaustive (nor exhausting).

For a general reader who is comfortable with the formalism of mathematics, the book by Williams [3] is an excellent introduction to mathematical finance. One of the nice features about her book is that, even while discussing the case of discrete-time, finite-valued stock prices, she uses the *notation* of continuous-time, real-valued stock prices, so that the transition to the latter case becomes easy.

For quite authoritative treatments that are 'for adults only' in terms of the level of mathematical background required, the reader is directed to the books by Kallianpur and Karandikar [7], and by Karatzas and Shreve [13].

The books by Benth [14] and Seydel [9] fall somewhere between that of Williams and those of Kallianpur–Karandikar and Karatzas–Shreve. Benth discusses the fact that real asset prices do not follow a log-normal distribution, and discusses alternative models. The book edited by Fu *et al* [6] contains a series of articles that include, among others, a description of the variance gamma processes that are an alternative to geometric Brownian motion. Seydel [9] has quite extensive discussion about issues such as simulation, Monte Carlo approximation, etc. that are not too common in mathematical finance texts.

All the books cited above concentrate on the mathematical aspects of pricing derivatives, and do not go very deeply into how 'real' financial markets actually work. In partial contrast, the book by Lin [15] attempts to give more details about actual financial markets.

When it comes to the use of mathematical finance methods in the 'real world,' Paul Wilmott has developed a very high reputation, and has written several books. Only one three volume set by Wilmott [16] is included here as an illustration.

One of the problems besetting the use of Black–Scholes theory in practice is that many of the quantities used in the theory (such as the volatility for example), are *not intrinsic*, but are inferred. Thus the user community is very concerned about the *sensitivity* of various quantities to these parameters. These sensitivities are known in the literature as 'the greeks,' and are just partial derivatives of the quantities with respect to the parameters. When the Black–Scholes theory leads to closed-form formulas, it is relatively straight-forward to compute 'the greeks.' However, depending on the nature of the payout function, there might not exist a closed-form formula for the price of the claim. In such a case, the price of the claim is often computed via numerical methods. When this is done, the greeks themselves need to be computed via numerical differentiation. Trying to compute the sensitivity of quantities that are themselves computed numerically is fraught with the risk of numerical instability. Malliavin calculus [10] alleviates this difficulty by expressing the sensitivities as *stochastic integrals.* There are situations when even the stochastic integrals of Malliavin calculus have to be estimated via simulation; but at the very least one avoids numerical *differentiation* and uses the more stable process of numerical *integration.* It may be mentioned that Malliavin calculus was invented in the 1970s for entirely different purposes, and its applicability to mathematical finance was discovered during the 1990s.

In its original form, Malliavin calculus is applicable only to geometric Brownian motion. Since then, the research community has been busy trying to extend the theory to arbitrary geometric Lévy processes, which are the most general processes with independent increments. Among the most comprehensive results in this direction is the PhD thesis of Petrou [11], which claims to have extended the Malliavin calculus to arbitrary geometric Lévy processes. However, the present author must confess that he has not independently verified this claim.

One of the serious limitations of Black–Scholes theory is the assumption that there are no transaction costs. Attempts to remove this limitation have not been very successful. Amongst the very few successful theories is the paper by Davis *et al* [17], which treats the case where the cost of a transaction is linearly proportional to the value of the transaction.

Finally, the paper by Rogers [1] presents a very general result on the equivalence between the existence of a Martingale measure and the absence of arbitrage, and in the process generalizes the well-known theorem of Dalang, Morton and Willinger [2].

References

[1] Rogers L C G 1994 Equivalent martingale measures and no arbitrage; *Stochastics and Stochastics Rep.* **51** 41–49.

[2] Dalang R C, Morton A and Willinger W 1990 Equivalent martingale measures and no-arbitrage in stochastic securities market models; *Stochastics and Stochastics Rep.* **29** 185–201.

[3] Williams R J 2006 *Introduction to the mathematics of finance*; Amer. Math. Soc., Providence, RI.

[4] Frittelli M 2000 The minimal entropy martingale measure and the valuation problem in incomplete markets; *Mathematical Finance* **10(1)** 39–52.

[5] Protter P E 2005 *Stochastic integration and differential equations* (Second Edition) (Berlin: Springer-Verlag).

[6] Fu M C, Jarrow R A, Yen J-Y J and Elliott R J (eds) 2007 *Advances in mathematical finance* (Boston: Birkhauser).

[7] Kallianpur G and Karandikar R L 2000 *Introduction to option pricing theory* (Boston: Birkhauser).

[8] http://en.wikipedia.org/wiki/Greeks_ (finance)

[9] Seydel R U 2006 *Tools for computational finance* (Third Edition) (Berlin: Springer-Verlag).

[10] Malliavin P and Thalmaier A 2006 *Stochastic calculus of variations in mathematical finance* (Berlin: Springer).

[11] Petrou E 2006 *Malliavin calculus in Lévy processes and applications in finance*; Ph.D. Thesis, Imperial College, UK.

[12] Derman E The modeler's Hippocratic oath; http://www.wilmott.com/blogs/eman/ index.cfm/2009/1/8/The-Financial-Modelers-Manifesto

[13] Karatzas I and Shreve S E 2001 *Methods of mathematical finance* (Berlin: Springer).

[14] Benth F E 2004 *Option theory with stochastic analysis* (Berlin: Springer).

[15] Sheldon Lin X 2006 *Introductory stochastic analysis for finance and insurance* (New York: Wiley).

[16] Wilmott P 2006 *Paul Wilmott on quantitative finance* (3 volumes) (Chichister: Wiley).

[17] Davis M H A, Panas V G and Zariphopoulou T 1993 European option pricing with transaction costs; *SIAM J. Control and Optimization* **31(2)** 470–493.