1

Introduction to Energy Derivatives and Fundamentals of Modelling and Pricing

1.1 Introduction to Energy Derivatives

Energy markets around the world are under going rapid deregulation, leading to more competition, increased volatility in energy prices, and exposing participants to potentially much greater risks. Deregulation impacts both consumers and producers and has lead to a heightened awareness of the need for risk management and the use of derivatives for controlling exposure to energy prices. However, this is not the only source of the development of risk management products. Investment banks are being drawn into the area as they look for new markets in which to operate. There is also an increasing number of power marketers entering the market and companies like Enron[1] are establishing themselves in a role which might be described as an 'energy investment bank'. This combination of the two different sides of the market, along with the sheer size of the market at the sales level, has the potential to make energy derivatives one of the fastest growing of all derivatives markets[2].

For many market participants, energy derivatives appear to be a new phenomenon. Although it is true that traded derivatives are a relatively new concept in the energy markets, the structures have been around for centuries and contracts with derivative characteristics have existed in energy markets for decades. For example, options have been embedded in supply and purchasing agreements which have traditionally offered a high degree of flexibility in terms of price, volume, timing and location of delivery. Although there is now a realisation that these contracts should be priced to reflect the optionality in such agreements, they have been trading for many years.

There are many contracts that enable the user to manage their exposure to energy prices, with derivatives often providing the simplest and most flexible solutions for precise risk management.

A derivative security can be defined as a security whose payoff depends on the value of

[1] http://www.enron.com

[2] Readers interested in the market growth and the development of competitive electricity markets are referred to Kaminski (1997) and Masson (1999).

other more basic variables[3]. The simplest types of derivatives are forward and futures contracts.

Futures and Forward Contracts

A futures contract is an agreement to buy or sell the underlying asset in the spot market (often called the spot asset) at a predetermined time in the future for a certain price, which is agreed today. Futures contracts are standardised, in terms of the future date, amount traded, etc. and can be retraded through time on a futures exchange. Forward contracts are also agreements to transact on fixed terms at a future date, but these are direct agreements between two parties. Although forwards and futures are similar contracts involving an agreement to buy or sell on a certain date for a certain price, important differences exist. Firstly, as we have just seen, futures are exchange standardised contracts, whereas forward contracts trade between individual institutions. Secondly, the cash flows of the two contracts occur at different times – futures are daily marked to market with cashflows passing between the long and the short position to reflect the daily futures price change, whereas forwards are settled once at maturity[4]. Despite these differences, if future interest rates are known with certainty then futures and forwards can be treated as the same for pricing purposes and we will, for the most part, use the terms interchangeably.

There are two sides to every forward contract. The party who agrees to buy the asset is said to hold a long forward position, whilst the seller is said to hold a short forward position. At the maturity of the contract (the 'forward date') the short position delivers the asset to the long position in return for the cash amount agreed in the contract – which is often called the delivery price.

If T represents the contract maturity date, then mathematically this long forward payoff can be expressed as $S_T - K$ where S_T represents the asset price at time T, and K represents the agreed delivery price. Figure 1.1 shows the profit and loss profile to the long forward position at the maturity of the contract. The payoff can obviously be positive or negative, depending on the relative values of S_T and K. The short position, by definition, has the opposite payoff to the long position (i.e. $-S_T + K$) as every time the long position makes a profit, a loss is suffered by the short, and vice versa. Since the holder of a long forward contract is guaranteed to pay a known fixed price for the spot asset, futures and forwards can be seen as insurance contracts providing protection against the price uncertainty in the spot markets.

A straight-forward arbitrage relationship means that the forward price must be equal to the cost of financing the purchase of the spot asset today and holding it until the forward maturity date[5]. Let F represent the price of a forward contract on the spot asset

[3] Derivatives are often referred to as contingent claims as the payout to the security (often referred to as the maturity payoff), and hence value, is contingent on other events.
[4] For credit purposes, some forward contracts are also marked to market on a regular basis.
[5] See chapter 4 for more details.

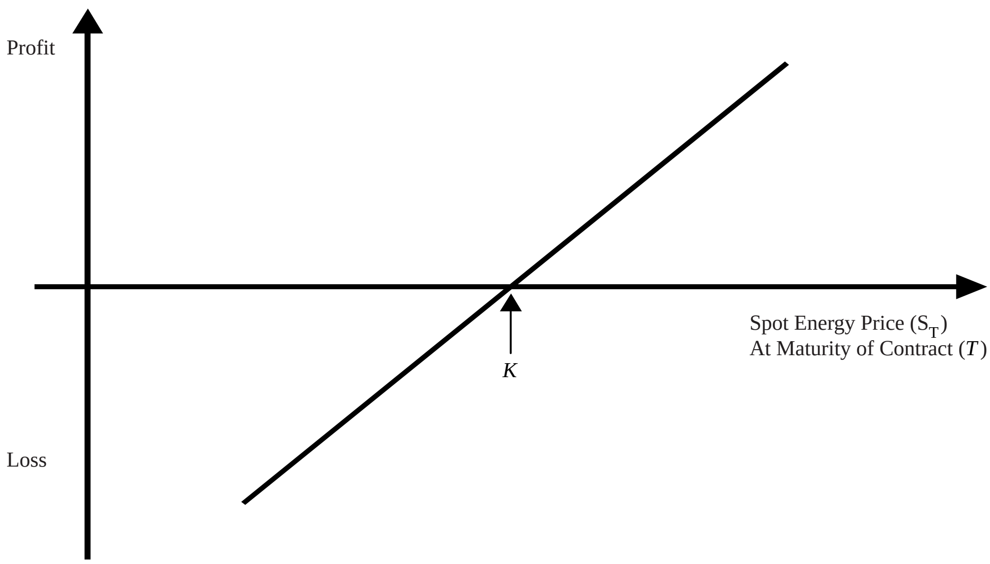

FIGURE 1.1 Payoff to Long Forward Position

that is currently trading at S. If the maturity date of the contract is T, c represents the cost of holding the spot asset (which includes the borrowing costs for the initial purchase and any storage costs), and δ the continuous dividend yield paid out by the underlying asset, then the price of a forward contract, at current time t, and the spot instrument on which it is written are related via the 'cost of carry' formula[6];

$$F = Se^{(c-\delta)(T-t)} \tag{1.1}$$

The continuous dividend yield can be interpreted as the yield on an index for index futures, as the foreign interest rate in foreign exchange futures contracts, and as the convenience yield for various energy contracts[7].

Options Contracts

Options contracts are the second cornerstone to the derivatives market. There are two basic types of options. A call option gives the holder the right, but not the obligation, to buy the spot asset on or before a predetermined date (the maturity date) at a certain price (the strike price), which is agreed today. Figure 1.2 graphs the payoff to the holder of such an option.

Options differ from forward and futures contracts in that a payment, usually at the time the contract is entered into, must be made by the buyer – this is the option price or

[6] $T - t$ is measured in years.
[7] See Chapter 6 for a discussion of convenience yields in energy and commodity markets.

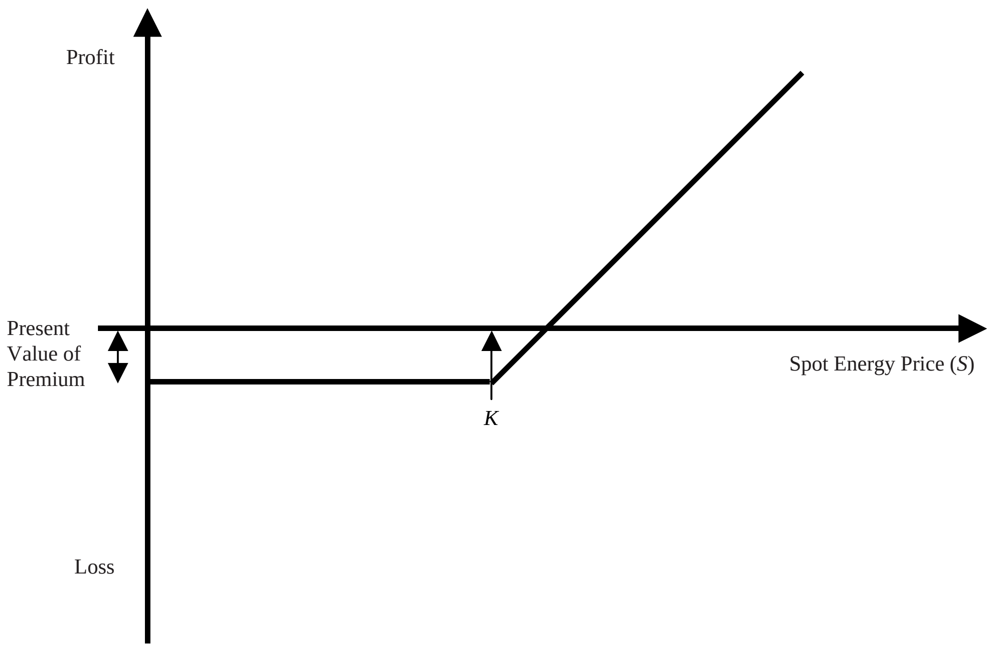

FIGURE 1.2 Payoff to Call Option

premium. At the maturity date, for spot asset prices below the agreed strike price (denoted by K), the holder lets the option expire worthless, forfeits the premium, and buys the asset in the spot market. For asset prices greater than K, the holder exercises the option, buying the asset at K and has the ability to immediately make a profit equal to the difference between the two prices (less the initial premium). Therefore, the holder of the call option essentially has the same positive payoff as the long forward contract, but without the downside, resulting in a so-called 'dog leg' payoff profile. The payoff to a call option can be described mathematically as follows:

$$\max(0, S - K) \tag{1.2}$$

The second basic type of option, a put option, gives the holder the right, but not the obligation, to sell the asset on or before the maturity date at the strike price. Figure 1.3 shows the payoff profile to the holder of a put option. Mathematically, the payoff for a put option can be written:

$$\max(0, K - S) \tag{1.3}$$

One of the difficulties experienced by newcomers to the derivatives market is the amount of terminology involved. As we have already seen, the date specified in the contract is known as the maturity date, but it is also known as the expiration, exercise, or strike date. The strike price is often referred to as the exercise price. Options are also classified

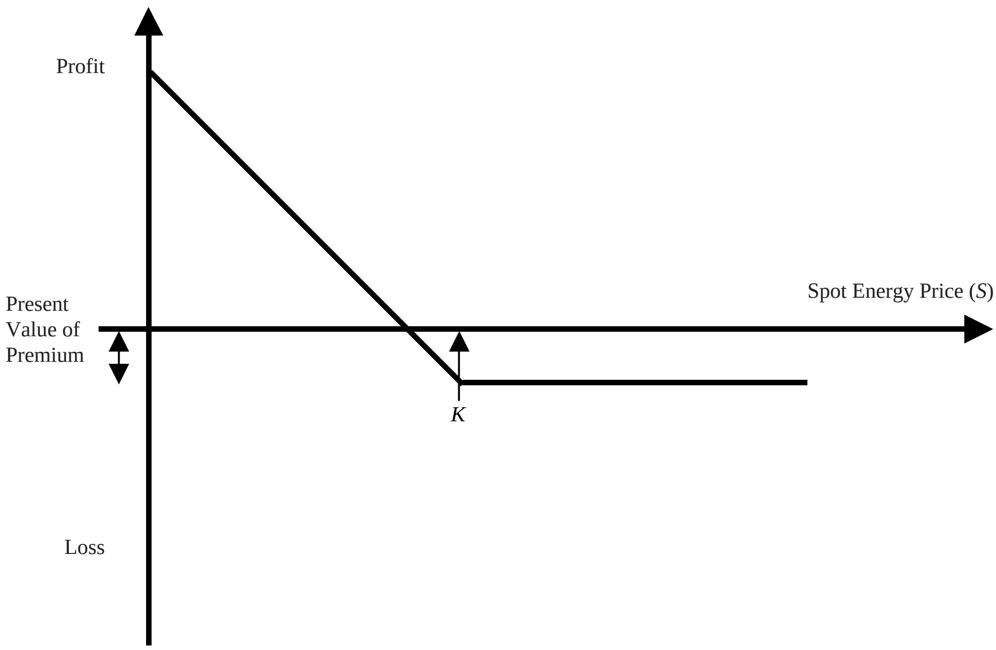

FIGURE 1.3 Payoff to Put Option

with respect to their exercise conventions. European options can only be exercised on the maturity date itself, whereas American style options can be exercised at any time up to and including the expiration date.

As with forwards, there are two sides to every option contract. One side has bought the option and has the long position, whilst the other side has sold (or written) the option and has taken a short position. Figure 1.4 shows the four possible combinations of terminal payoffs for long and short positions in European call and put options with maturity date T.

The futures and options of this section describe the basic building blocks of all derivative securities and the principals are consistent across all underlying markets. However, derivative structures in energy markets exhibit a number of important differences from other underlying markets. The differences arise because of the complex contract types that exist in the energy industry as well as the complex characteristics of energy prices. Both the type of derivatives that trade, and the modelling needed to capture the evolution of prices, reflect these differences. For example, many contracts in the energy industry are based on averages (often weekly or monthly in the oil industry and hourly or less in the electricity markets) of prices and this has led to the wide acceptance of so-called Asian or average price options[8]. Basis risk, widely defined to

[8] See chapter 5 for further details on Asian options.

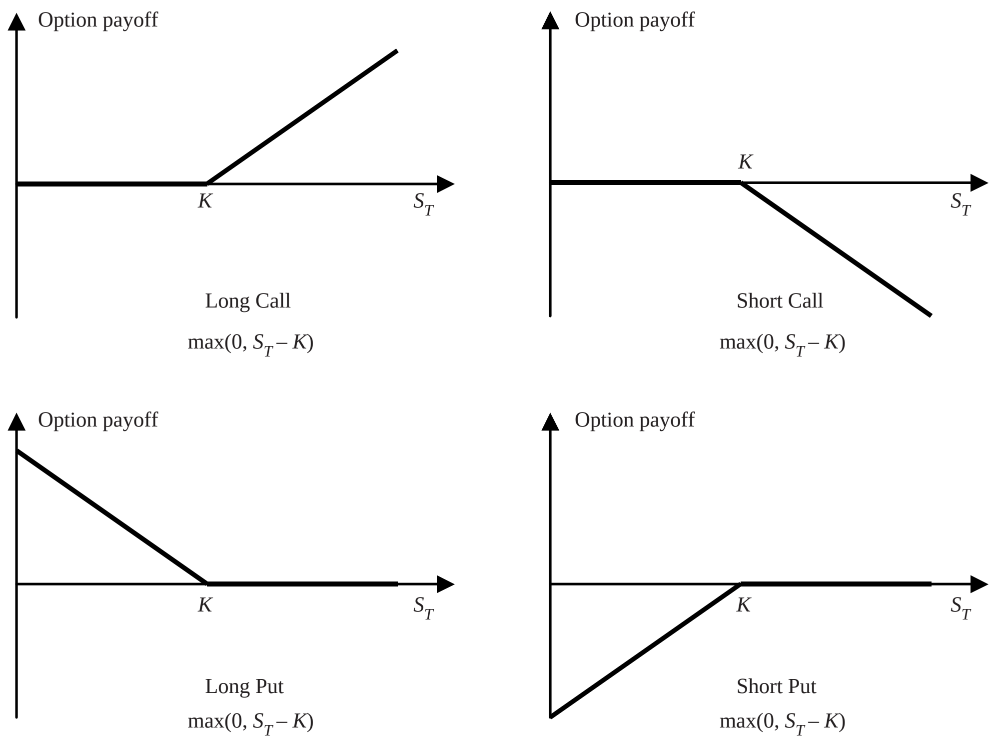

FIGURE 1.4 Terminal Payoffs for European Options

mean the difference between two different prices, is very important in energy markets as production processes often involve the conversion of one energy (say natural gas) into another (electricity) thus exposing the company to the price differential. This has lead to the development of a wide variety of spread options. The complications extend to modelling of the price dynamics. For example, large variations in the cost of electricity generation and high demand variability contribute to high price volatility and jumps in prices. High levels of seasonality exhibited by energies are also important to capture[9].

1.2 Fundamentals of Modelling and Pricing

The modern theory of option pricing is possibly one of the most important contributions to the whole area of financial economics. The breakthrough came in the early 1970's with work by Fisher Black, Myron Scholes and Robert Merton (Black and Scholes (1973),

[9] See Chapter 4 for a description of types of contract provision that cause embedded options to take on complicating characteristics and their impact on pricing. See also Kaminski (1997).

Merton (1973)). The Black-Scholes-Merton (BSM) modelling approach proved not only important for providing a computationally efficient and relatively easy way of pricing the then recently developed exchange traded equity options in Chicago, but also for demonstrating the principal of no-arbitrage, risk neutral, valuation. Their analysis showed that the payoff to an option could be perfectly replicated with a continuously adjusted holding in the underlying asset and the risk free bond. Since the risk of writing an option can be completely eliminated the risk appetites or preferences of market participants are irrelevant to the valuation problem, and we can assume they are risk neutral. In such a world, all assets earn the riskless rate of interest, thus the actual expected return on the asset does not appear in the Black-Scholes formula.

In many energy markets the concept of being able to perfectly replicate options by continuously trading the underlying asset is unrealistic. For example, spot electricity cannot be easily stored[10] and therefore a continuously adjusted position is not possible. Similar arguments can be applied, albeit in a less extreme sense to many other spot energies. However, many energy derivatives actually depend on futures prices rather than the spot price and futures can be used to replicate options positions allowing the application of the risk neutral pricing approach. In cases where it is not reasonable to apply risk neutral pricing we can argue that the risk neutral price provides a good reference with which to compare other pricing methods. An alternative and useful 'insurance' based approach is to calculate the expected payoff of the option under a model for the actual behaviour or real-world measure (as opposed to the risk neutral measure) of the market prices. The methods we describe in this book can also be used in this way. Finally, as a writer of options we may wish to price them on the basis of our expected cost of hedging the option given our access to hedging markets and management systems. This can also be done using the models and methods which we describe.

In this section we look at a number of different methodologies that have been developed for pricing options. We start with approaches that have been developed under the BSM assumptions of costless trading in continuous time, infinite divisibility of the underlying asset, a non-dividend paying asset, constant interest rates and constant volatility. From the perspective of this chapter, however, the most important assumption in the BSM model is the mathematical description of how asset prices evolve through time. This is the well-known Geometric Brownian Motion (GBM) assumption where proportional changes in the asset price, denoted by S, are assumed to have constant instantaneous drift, μ, and volatility, σ. The mathematical description of this property is given by the following stochastic differential equation[11];

$$dS = \mu S dt + \sigma S dz \tag{1.4}$$

[10] Electricity can be stored by hydroelectric schemes by using it to pump water into the reservoir, the electricity can then be recovered by releasing the water through the turbines. Electricity can also be indirectly stored by generators in the form of the fuel used to generate it.

[11] Most models of asset price behaviour for pricing derivatives are formulated in a continuous time framework by assuming a stochastic differential equation describing the stochastic process followed by the asset price.

Here dS represents the increment in the asset price process during a (infinitesimally) small interval of time dt, and dz is the underlying uncertainty driving the model and represents an increment in a Weiner process during dt. The risk-neutral assumption implies that the drift can be replaced by the riskless rate of interest (i.e. $\mu = r$). Any process describing the stochastic behaviour of the asset price will lead to a characterisation of the distribution of future asset values and the assumption in equation (1.4) implies that future asset prices are lognormally distributed, or alternatively, that returns are to the asset are normally distributed.

Let C represent the value of any derivative security (a call or a put, a forward, or any of the other more complicated derivatives we look at throughout this book). The arguments of BSM allow for the derivation of the following partial differential equation describing the evolution of the derivative price through time,

$$\frac{\partial C}{\partial t} + rS\frac{\partial C}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} = rC \tag{1.5}$$

The value of any derivative whose payoff is contingent on the level of the asset price following equation (1.4), and time, must satisfy this equation. In order to evaluate the prices of specific options (for example, European call and put options) this equation must be solved with the appropriate boundary conditions – given as the option maturity payoff (i.e. $C_T = \max(0, S_T - K)$ for a European call and $C_T = \max(0, K - S_T)$ for a put). For a European option equation (1.5) can be solved in a variety of ways, yielding the familiar Black-Scholes formula (here for a call evaluated at time t),

$$C(t) = SN(d_1) - Ke^{-r(T-t)}N(d_2) \tag{1.6}$$

where

$$d_1 = \frac{\ln(S/K) + (r + \frac{1}{2}\sigma^2)(T - t)}{\sigma\sqrt{T - t}}$$
$$d_2 = d_1 - \sigma\sqrt{T - t}$$

and where the parameters S, K, r, t, T, and σ have been previously defined, ln(.) is the natural logarithm and N(.) is the standard cumulative normal distribution function. The use of the risk free rate for discounting is based on the notion of risk neutral valuation. One of the qualities that has led to the enduring success of the Black-Scholes model is its simplicity. The inputs of the model are defined by the contract being priced or are directly observable from the market. The only exception to this is the volatility parameter and there is now a vast amount of published material in the finance literature for deriving estimates of this figure either from historical data or as implied by the market prices of options.

Although the pricing formula (1.6) was originally applied to equity markets, some of the rigid assumptions have been relaxed by later authors, extending the model to other markets. For example, Merton (1973) extended the model by firstly showing that the discounting in the model could be done in terms of a pure discount bond of the same maturity as the option, thus taking into account non-constant interest rates. A pure discount bond is defined as a bond which pays one unit of cash at its maturity date only. If we represent with, $P(t,T)$, the price at time t of a pure discount bond with maturity date T, the BSM formula can be written;

$$C(t) = SN(d_1) - KP(t, T)N(d_2) \tag{1.7}$$

Merton also showed how a non-constant, but deterministic, volatility can be handled by using the average volatility over the life of the option. Another, widely used, relaxation of the original formula takes into account assets that pay a constant proportional dividend. Assets of this kind are handled by reducing the expected growth rate of the asset by the amount of the dividend yield. If the asset pays a constant proportional dividend at a rate δ, over the life of the option, then we can use the original Black-Scholes call formula (1.6) with the adjustment that the parameter S is replaced by the term $Se^{-\delta(T-t)}$. This adjustment has been applied to value options on broad-based equity indices as well as options on foreign exchange rates – see Garman and Kohlhagen (1983) for the latter[12].

In practice it is now well established that the Black-Scholes-Merton model is used not with the constant parameter volatility assumption of equation (1.6) but in conjunction with what is termed 'implied volatility smiles'. This is the practice of adjusting the volatility which is entered into the Black-Scholes formula for options which are away from the money[13]. As we shall see in section 3.3, volatility smiles correspond to the probability distribution *implied* in option prices differing from the lognormal distribution implied by the GBM assumption of equation (1.4), resulting in options being priced at different levels of volatility. In chapter 2 we show that smiles can be introduced into models by incorporating stochastic volatility and jumps. In addition to varying volatility dependent on strike price, traders frequently adjust volatility dependent on the maturity of the option often resulting in the smile becoming less pronounced as the option maturity increases.

Although it is possible to obtain closed form solutions such as equation (1.6) for certain derivative pricing problems there are many situations when analytical solutions are not obtainable and numerical techniques need to be applied. Examples that we will see in this book include American options, and other options where there are early exercise opportunities, 'path dependent' options with discrete observation frequencies, models that incorporate jumps, and models dependent on multiple random factors. The description of two of these techniques is the subject of the next section.

1.3 Numerical Techniques

In this section we describe two numerical techniques which are most commonly used by practitioners to value derivatives in the absence of closed-form solutions. Although we restrict our attention to (trinomial) tree building and Monte Carlo simulation, other

[12] The proportional dividend is interpreted as the risk free rate on the foreign currency in options of this type.
[13] See section 3.3 for a discussion of volatility smiles in energy markets.

techniques such as finite difference schemes (see Clewlow and Strickland (1998)), numerical integration, finite element methods, and others, are also sometimes used by practitioners. However, these methods require more advanced expertise in numerical techniques. For both of the techniques we outline it is possible to price not only derivatives with complicated payoff functions dependent on the final energy price, but also derivatives whose payoff is determined also by the path the underlying price follows during its life.

Monte Carlo simulation provides a simple and flexible method for valuing complex derivatives for which analytical formulae are not possible. The method can easily deal with multiple random factors; for example, options on multiple energy prices or models with random volatility, convenience yield, or interest rates. Monte Carlo simulation can also be used to value complex path dependent options, such as average rate, barrier, and lookback options, and also allows the incorporation of more realistic energy price processes, such as jumps in prices and more realistic market conditions such as the discrete fixing of exotic path dependent options.

For many American style options early exercise can be optimal depending on the level of the underlying energy price. It is rare to find closed-form solutions for prices and risk parameters of these options, so numerical procedures must again be applied. However, using Monte Carlo simulation for pricing American style options is difficult. The problem arises because simulation methods generate trajectories of state variables forward in time, whereas a backward dynamic programming approach is required to efficiently determine optimal exercise decisions for pricing American options. Therefore, binomial and trinomial trees are usually used by practitioners for pricing American options[14].

1.3.1 The Trinomial Method

The binomial model of Cox, Ross and Rubinstein (1979) is a well-known alternative discrete time representation of the behaviour of asset prices to GBM. This model is important in several ways. Firstly, the continuous time limit of the proportional binomial process is exactly the GBM process. Second, and perhaps most importantly, the binomial model is the basis of the dynamic programming solution to the valuation of American options[15]. Although binomial trees are used by many practitioners for pricing American style options, we and many other practitioners prefer trinomial trees. The trinomial tree has a number of advantages over the binomial tree. Because there are three possible future movements of the asset price over each time period, rather than the two of the binomial approach, the trinomial tree provides a better approximation to a continuous price process than the binomial tree for the same number of time steps. Also,

[14] Tilley (1993), Li and Zhang (1996), Broadie and Glasserman (1997a, 1997b) and Clewlow and Strickland (1998), amongst others have described methods of extending Monte Carlo simulation to the valuation of options with early exercise opportunities. In chapter 8 we show how Monte Carlo simulation can be applied to American options in conjunction with a tree based approach.

[15] See Chapter 2 of Clewlow and Strickland (1998) for an in-depth discussion of implementing the binomial method.

the trinomial tree is easier to work with because of its more regular grid and is more flexible, allowing it to be fitted more easily to market prices of forwards and standard options, an important practical consideration.

For an asset paying a continuous dividend yield, the stochastic differential equation for the risk neutral GBM model is given by equation (1.4) where the drift is replaced by the difference between the riskless rate and the continuous yield:

$$dS = (r - \delta)Sdt + \sigma Sdz \tag{1.8}$$

In the following we will find it more convenient to work in terms of the natural logarithm of the spot price, $x = \ln(S)$, and under this transformation we have the following process for x:

$$dx = \nu dt + \sigma dz \tag{1.9}$$

where $\nu = r - \delta - \frac{1}{2}\sigma^2$. Consider a trinomial model of the asset price in which, over a small time interval Δt, the asset price can increase by Δx (the space step), stay the same or decrease by Δx, with probabilities p_u, p_m, and p_d respectively. This is depicted in terms of x in figure 1.5.

The drift and volatility parameters of the asset price are now captured in this simplified discrete process by Δx, p_u, p_m, and p_d. It can been shown that the space step cannot be chosen independently of the time step, and that a good choice is $\Delta x = \sigma\sqrt{3\Delta t}$. The relationship between the parameters of the continuous time process and the trinomial process is obtained by equating the mean and variance over the time interval Δt and requiring that the probabilities sum to one, i.e.:

$$E[\Delta x] = p_u(\Delta x) + p_m(0) + p_d(-\Delta x) = \nu\Delta t \tag{1.10}$$

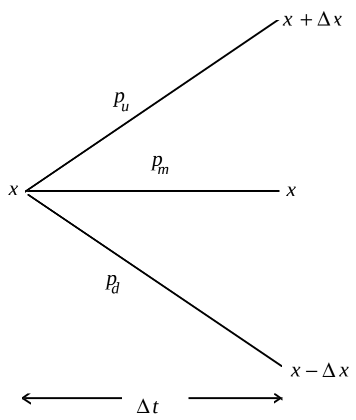

FIGURE 1.5 Trinomial Model of an Asset Price

$$E[\Delta x^2] = p_u(\Delta x^2) + p_m(0) + p_d(\Delta x^2) = \sigma^2 \Delta t + \nu^2 \Delta t^2 \tag{1.11}$$

$$p_u + p_m + p_d = 1 \tag{1.12}$$

Solving equations (1.10) to (1.12) yields the following explicit expressions for the transitional probabilities;

$$p_u = \tfrac{1}{2}\left(\frac{\sigma^2 \Delta t + \nu^2 \Delta t^2}{\Delta x^2} + \frac{\nu \Delta t}{\Delta x}\right) \tag{1.13}$$

$$p_m = 1 - \frac{\sigma^2 \Delta t + \nu^2 \Delta t^2}{\Delta x^2} \tag{1.14}$$

$$p_d = \tfrac{1}{2}\left(\frac{\sigma^2 \Delta t + \nu^2 \Delta t^2}{\Delta x^2} - \frac{\nu \Delta t}{\Delta x}\right) \tag{1.15}$$

The single period trinomial process in figure 1.5 can be extended to form a trinomial tree. Figure 1.6 depicts such a tree.

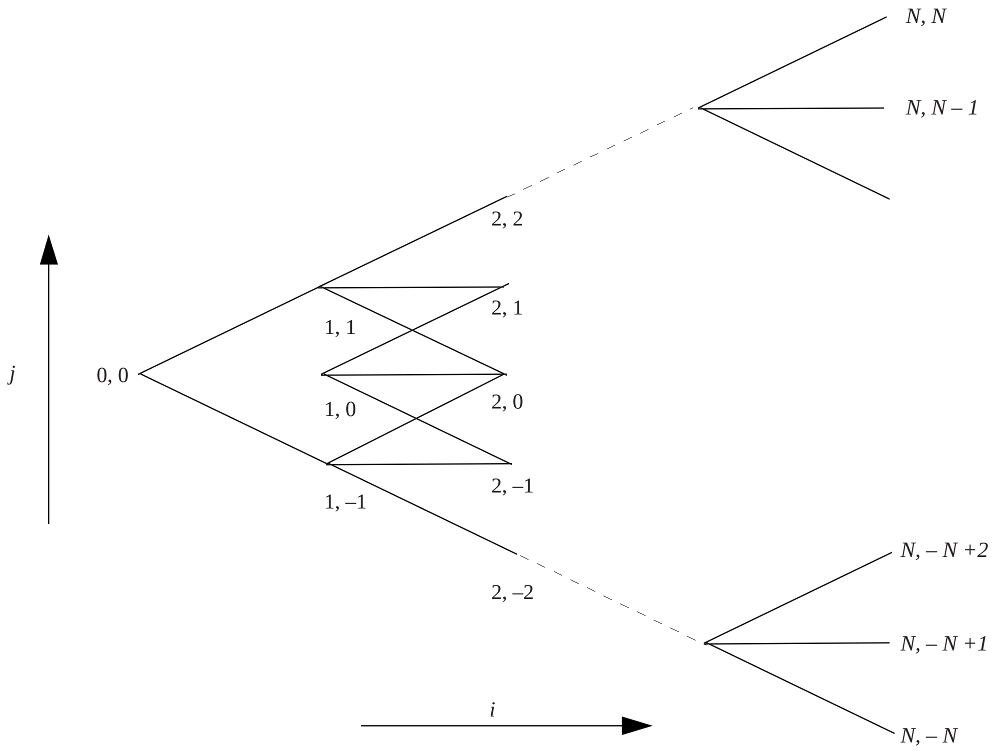

FIGURE 1.6 A Trinomial Tree Model of an Asset Price

Let i denote the number of the time step and j the level of the asset price relative to the initial asset price in the tree. If $S_{i,j}$ denotes the level of the asset price at node (i, j) then we have $t = t_i = i\Delta t$, and an asset price level of $S\exp(j\Delta x)$. Once the tree has been constructed we know the spot price at every time and every state of the world consistent with our original assumptions about its behaviour process, and we can use the tree to derive prices for a wide range of derivatives. We will illustrate the procedure with reference to pricing a European and American call option with strike price K on the spot price.

We represent the value of an option at node (i, j) by $C_{i,j}$. In order to value an option we construct the tree representing the evolution of the spot price from the current date out to the maturity date of the option – let time step N correspond to the maturity date in terms of the number of time steps in the tree, i.e. $T = N\Delta t$. The values of the option at maturity are determined by the values of the spot price in the tree at time step N and the strike price of the option;

$$C_{N,j} = \max(0, S_{N,j} - K); \; j = -N, \ldots, N \tag{1.16}$$

It can be shown that we can compute option values as discounted expectations in a risk neutral world[16], and therefore the values of the option at earlier nodes can be computed as discounted expectations of the values at the following three nodes to which the asset price could jump:

$$C_{i,j} = e^{-r\Delta t}(p_u C_{i+1,j+1} + p_m C_{i+1,j} + p_d C_{i+1,j-1}) \tag{1.17}$$

where $e^{-r\Delta t}$ is the single period discount factor. This procedure is often referred to as 'backwards induction' as it links the option value at time i to known values at time $i + 1$. The attraction of this method is the ease with which American option values can be evaluated. During the inductive stage we simply compare the immediate exercise value of the option with the value if not exercised, computed from equation (1.17). If the immediate exercise value is greater, then we store this value at the node, i.e.:

$$C_{i,j} = \max\left\{e^{-r\Delta t}(p_u C_{i+1,j+1} + p_m C_{i+1,j} + p_d C_{i+1,j-1}), S_{i,j} - K\right\} \tag{1.18}$$

This method also gives us the optimal exercise strategy for the American option since for every possible future state of the world, i.e. every node in the tree, we know whether we should exercise the option or not. The value of the option today is given by the value in the tree at node (0,0), $C_{0,0}$. Although we have so far discussed only the valuation of a simple derivative using the tree structure, it is also possible to price many path dependent options. In section 7.5 we explain how path dependent options can be priced in an energy spot price tree fitted to the observed forward curve.

The standard option hedge sensitivities: delta, gamma and theta, can be calculated straightforwardly using the tree since they can be approximated by finite difference ratios. Vega and rho can be computed by re-evaluation of the price for small changes in

[16] See Clewlow and Strickland (1998) for an in-depth discussion of the implementation of trinomial trees for derivative pricing.

the volatility and interest rate respectively (see section 7.2.5 below and section 3.8 of Clewlow and Strickland (1998)).

1.3.2 Monte Carlo Simulation[17]

Monte Carlo simulation is easy to implement, works for a wide range of path dependent options, and is suitable for handling multiple stochastic factors. This last property implies that it is straightforward to add multiple sources of uncertainty such as stochastic volatility or random jumps to the basic model as well as valuing derivatives whose payoff depends on some function of two or more energy prices[18].

As we have already seen, in general, the present value of an option is the expectation of its discounted payoff. We can obtain an estimate of this expectation by computing the average of a large number of discounted payoffs computed via Monte Carlo simulation. Originally applied to the pricing of financial instruments by Boyle (1977), the Monte Carlo technique involves simulating the possible paths that the asset price can take from today until the maturity of the option.

We can discretise the transformed GBM process represented in equation (1.8) in the following way:

$$x_{t+\Delta t} = x_t + (\nu \Delta t + \sigma(z_{t+\Delta t} - z_t)) \tag{1.19}$$

Alternatively, in terms of the original asset price we have the discrete form

$$S_{t+\Delta t} = S_t \exp(\nu \Delta t + \sigma(z_{t+\Delta t} - z_t)) \tag{1.20}$$

Equations (1.19) or (1.20) can be used to simulate the evolution of the spot price through time. The change in the random Brownian motion, $z_{t+\Delta t} - z_t$, has a mean of zero and a variance of Δt. It can therefore be simulated using random samples from a standard normal multiplied by $\sqrt{\Delta t}$, i.e. $\sqrt{\Delta t}\varepsilon$ where $\varepsilon \sim N(0, 1)$. In order to simulate the spot price we divide the time period over which we wish to simulate, $[0,T]$, into N intervals such that $\Delta t = T/N$, $t_i = i\Delta t$, $i = 1, \ldots, N$. Using, for example, equation (1.20) we have

$$S_{t_i} = S_{t_{i-1}} \exp(\nu \Delta t + \sigma \sqrt{\Delta t} \varepsilon_i) \tag{1.21}$$

It is important to note that, since the drift and volatility terms do not depend on the variables S and t, the discretisation is correct for any time step we choose. This enables us to jump straight to the maturity date of the option in a single time step, if the payoff to the derivative is only a function of the terminal asset value, and does not depend on the path taken by the asset during the life of the option. Repeating this process N times, choosing ε_i randomly each time, leads to one possible path for the spot price. Figure 1.7 illustrates the result of repeating this single path simulation one thousand times.

$$S = 100, r - \delta = 0.05, \sigma = 0.30, \Delta t = 1/(365 \times 24).$$

[17] This section is based on Chapter 4 of Clewlow and Strickland (1998).

[18] For example a crack spread that pays the difference between gas and electricity prices.

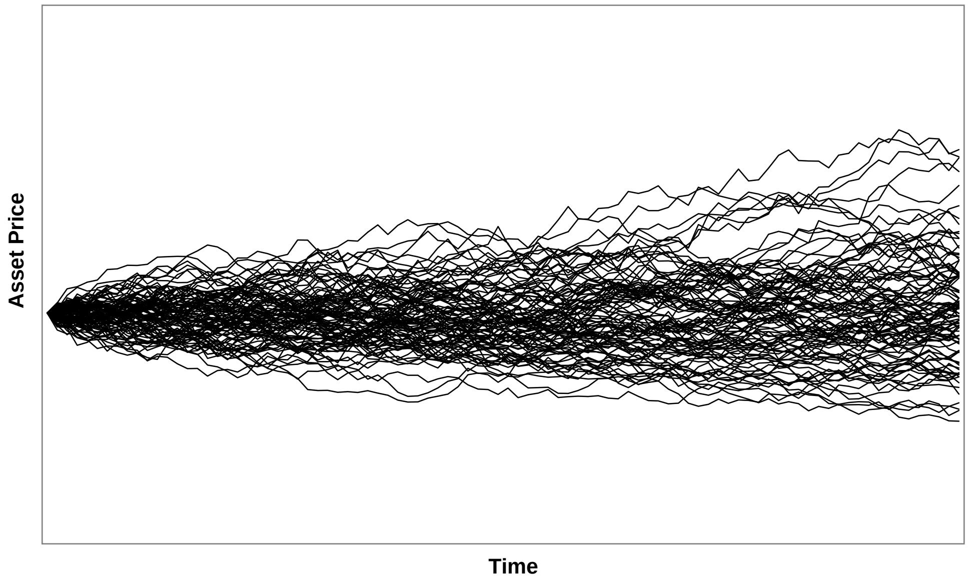

FIGURE 1.7 Illustration of 1000 Simulated GBM Paths

At the end of each simulated path the terminal value of the option (C_T) is evaluated. Let $C_{T,j}$ represent the payoff to the contingent claim under the j^{th} simulation. For example, for a standard European call option the terminal value is given by;

$$C_{T,j} = \max(0, S_{T,j} - K) \tag{1.22}$$

Each payoff is discounted using the simulated short-term interest rate sequence,

$$C_{0,j} = \exp\left(-\int_0^T r_u du\right) C_{T,j} \tag{1.23}$$

In the case of constant or deterministic interest rates equation (1.23) simplifies to

$$C_{0,j} = P(0,T)C_{T,j}$$

This value represents the value of the option along one possible path that the asset can follow. The simulations are repeated many (say M) times and the average of all the outcomes is taken to compute the expectation, and hence option price;

$$\hat{C}_0 = \frac{1}{M}\sum_{j=1}^{M} C_{0,j} \tag{1.24}$$

Therefore $\hat{C}_0$ is an estimate of the true value of the option, C_0, but with an error due to the fact that it is an average of randomly generated samples and so is itself random. In order to obtain a measure of the error we estimate the standard error *SE*(.) as the sample standard deviation, *SD*(.), of $C_{0,j}$ divided by the square root of the number of samples;

$$SE(\hat{C}_0) = \frac{SD(C_{0,j})}{\sqrt{M}} \tag{1.25}$$

where $SD(C_{0,j})$ is the standard deviation of $\hat{C}_0$;

$$SD(C_{0,j}) = \sqrt{\frac{1}{M-1}\sum_{j=1}^{M}\left(C_{0,j} - \hat{C}_0\right)^2}$$

The biggest criticism of Monte Carlo methods concerns the speed with which derivative values can be evaluated. It is not unknown for the technique to take many hours to return a price that is sufficiently accurate, due to the number of simulations that have to be performed. However, a number of authors, including Kemna and Vorst (1990), Clewlow and Carverhill (1994) and Clewlow and Strickland (1997, 1998), have proposed methods to speed up the process. These techniques are known as variance reduction techniques, as their aim is to reduce the variance of the estimate obtained via the simulations. Further variance reduction techniques involve the implementation of quasi-random sequences (see for example Joy, Boyle, and Tan (1996)). Another criticism of Monte Carlo is its perceived inability to handle American options. However, we discuss in Chapter 8 how American options can be priced using a combination of tree and simulation techniques.

1.4 Summary

In this chapter we have introduced energy derivatives, describing some simple structures, and outlined the differences between energy and other underlying markets. We have presented an overview of the fundamental pricing principals that are applied to derivative valuation in a Black-Scholes-Merton world, and described two numerical procedures often implemented by practitioners to evaluate derivative prices and risk sensitivities. In the following chapter we look at the applicability of GBM for modelling energy price processes and describe other price processes often applied to energy price movements.