CHAPTER 1

Introduction to the Techniques of Derivative Modeling

1.1. INTRODUCTION

"How do I model derivatives?"

For the desk quantitative analyst ("desk-quant"), the quantitative-programmer, the quantitative-trader or risk analyst—at a firm actively trading, risk managing, or even auditing books of derivatives—who needs to know the basic answer to this question, it is contained in a toolkit of well-established mathematical models and techniques. These professionals might be prepared to forgo mathematical rigor and the security against subtle errors that a thorough foundation could perhaps provide. They might be prepared to take a few shortcuts to get at these techniques relying on another team member or risk-group reviewer who has the thorough mathematical background to provide a safety net against subtle errors and misunderstandings. But the desk quant, the programmer, and the trader do not have to forgo everything. This book is aimed at such readers who want a good quick grasp of these techniques.

This first chapter reviews in coarse outline the two most typical techniques for theoretically modeling and pricing derivatives and takes the first motivational step of mentioning the model of the process for stock that underpins the first technique and the simplest implementations of the second. The remainder of the book will take the reader through the mathematical tools that underpin these two and many other techniques.

1.2. MODELS

1.2.1. What Is a Derivative?

The archetypal example of a *derivative* is a stock option. An equity call option is a financial contract. It is often an exchange traded security much like the stock itself. It is the right, but not the obligation, to buy one share of stock on a given date called the *expiration date* and at a price, the so-called *strike price,* contractually determined on the *trade date*. A put option is the right, but not the obligation, to sell at a preset *strike*. The action of buying or selling stock under the terms of a put or call option is called *exercise. American-style options* are exercisable on any day up to their expiration date—and *European options* are exercisable *only* on the expiration date.

Figure 1.1 shows the value of a call option on the expiration date, the *intrinsic value,* as a function of the market price of the stock (which we obviously don't know before expiration). Clearly, if the stock is trading above the strike, we can make a profit by exercising the option, that is, buying the stock for the strike price and then selling it in the marketplace for a profit of the difference between the two. If the stock price is below the strike on expiration date, the option is worthless. This graph of intrinsic value is obviously the fair price for the option at the moment of expiration. The fair price is generally different at all times previous to this. For example, it will have a small but nonzero fair value below strike at all times previous to expiration, representing the real, albeit small probability that the stock might end up above the strike at expiration.

Generally, a financial derivative or, better, a *contingent claim* is a loosely defined term meaning a security whose price is dependent on the price of another security that itself is called the derivative's *underlying* security. There may be more than one underlying. In the general case of an option, because there is some kind of choice involved, we can generally say that there clearly must be at least two underlying securities.

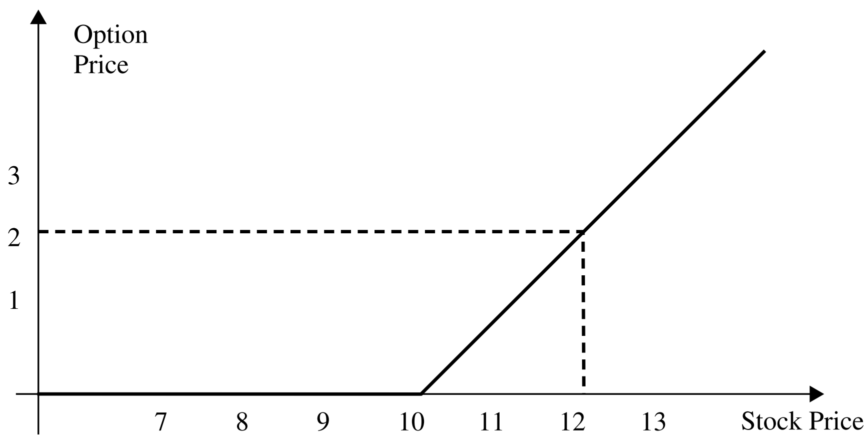

FIGURE 1.1 The value of a call option struck at 10 on maturity date as a function of the underlying stock price at maturity.

NOTE ON DERIVATIVE

The reader should beware of confusing the unrelated concept from mathematics, the derivative of a function, which means the rate of change of the given function with respect to one of its variables. Derivative in this mathematical sense can naturally arise in a discussion of financial derivatives, possibly even, and most confusingly, in the same sentence!

1.2.2. What Is a Model?

The typical mathematical pricing formula for options is the *Black-Scholes formula* (Figure 1.2). It outputs the theoretical fair price that a European option on a dividendless stock should trade for given the following set of inputs: call or put; current stock price; option strike price; time to expiration; average future stock volatility (for example, standard deviation of daily stock returns) from today to expiration, and current market value of the "riskless" interest rate from today to maturity (as determined by prices of treasury bonds or maybe interest rate swaps offered by big banks).

Deriving this formula requires a set of assumptions that are very important to bear in mind when using the formula. They are discussed in detail throughout this book. Essentially these assumptions are the *model*.

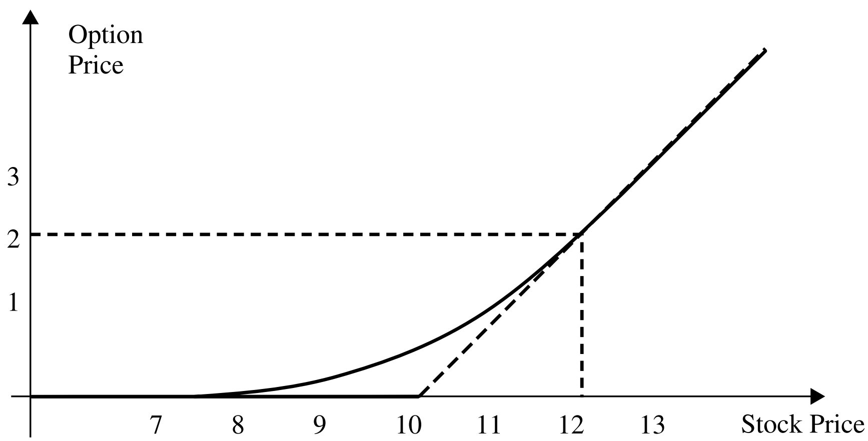

FIGURE 1.2 The value of a call option struck at 10 before maturity date as a function of the underlying stock price today.

Different sets of assumptions are different models. The model that underpins the Black-Scholes formula includes the following assumptions:

- Stock price changes are statistical (or stochastic), meaning the changes are random but they have a well-defined distribution of values.
- Stock price changes are continuous. The shorter the time over which you measure the average variation of size of changes, the smaller the average.
- The distribution of stock price changes is only a function of today's stock price and other values that are determined today and these changes have no dependence on the history of the stock or any other variables (models with this property are referred to as *Markovian* models)
- Stock price returns (i.e., changes expressed as percentages) have a normal (i.e., Gaussian) or bell-curve distribution around a mean that is the stock percentage growth rate.

1.2.3. Two Initial Methods for Modeling Derivatives

Many options and financial optionality can be approximately, quantitatively, and certainly qualitatively understood using simple modifications of the Black-Scholes formula. If this technique is lacking, the general machinery used to derive the formula provides a framework for more complex and thus more realistic analysis.

This means that an initial understanding of stock options and the Black-Scholes formula immediately opens up the world of derivatives and not just stock options themselves. The severe restriction is that the modeler can only study options under the assumption that the underlying security price has (continuous Markovian) normally distributed returns. Further study, leading to an understanding of the derivation of the Black-Scholes formula, opens that world far wider to practitioners, allowing them outside of the assumptions of a lognormal distribution of security changes to other perhaps more realistic distributions.

Mastering these two tools gives practitioners a strong suite of skills to analyze the trading and risks of derivatives. The derivatives accessible with the study of stock options include, among many others:

- Bond options and interest rate options such as swaptions
- Callable and putable bonds
- Credit default swaps (also called *credit-default put options*)

Indeed, any financial contract with an embedded clause that the issuer, or holder has the right but not the obligation to exercise can potentially be analyzed to some degree, even if only qualitatively, using one of the two methods mentioned: (1) the Black-Scholes formula with modifications or (2) the more general ideas behind the derivation of Black-Scholes formula.

1.2.4. Price Processes

The central idea behind the financial models that price stock options is the future distribution of stock prices. Furthermore, a random walk has characteristics similar to a stock or bond price path and has a future distribution of path end points. The walk may tend in a general direction, a *nonstochastic* drift, such as making a general observation that "on average" prices increase by \$1 every year or each stock makes a \$1/365 step up per day. But the walk also includes a *stochastic* component, referring to the fact that any particular price path has an additional random step either up or down.

There are many possible choices for this distribution of price changes and empirical data from the markets can often be inconclusive (principally because it seems the real world process itself is not stationary). More confusingly, the mathematical form of this distribution might look very different on different time scales. However, the normal distribution is a very useful tool for *first approximation* because it turns out that if the distribution of one-day price changes (i.e., steps) is normally distributed and has a standard deviation of, say, 5¢ with a mean of 0.27¢, then the distribution of annual changes (after many daily steps are added up) is also a normal distribution just with different mean and standard deviation. In fact, the normal distribution of annual price changes has a standard deviation of $\$0.05 \times \sqrt{365} \approx \$1$ and a mean of $\$0.0027 \times 365 \approx \$1$. The square root appeared because for normal distributions, the variance, the square of the standard deviation, grows linearly with time rather than the standard deviation itself.

This is an example of a model for the price process that starts at today's stock price; drifts by an average of \$1 every year (and thus 0.27¢ every day); and has a standard deviation of annual changes of \$1 (equivalent to a standard deviation of daily changes of 5¢).

A normal distribution of stock price changes with a drift could be used to find derivative prices with qualitatively correct features. The serious drawback, however, is that no matter where the stock starts, future stock prices have a potentially very large probability of being negative. This is a bad feature and will ruin even some qualitative observations, let alone quantitative ones. We need to refine this model slightly to get the simplest qualitatively correct model and hence make it (possibly) useful quantitatively even if only approximately.

1.2.5. The Archetypal Security Process: Normal Returns

This problem—of the normal distribution of future stock prices giving a probability for negative prices in the future—is easily corrected by modeling the price process with a normal distribution of returns on stock rather than the stock price changes. This is because allowing all possible percentage changes never results in negative stock prices if the path starts at a positive stock price. Equivalently, we may say that the natural log of stock prices is the normally distributed *state variable* over which the random walk occurs. This is a simple variable change and results in a much more compelling model qualitatively.

A normal distribution for changes of the natural log of stock price is equivalently described as a *lognormal* distribution of stock price changes (note this is a definition).

So one of the simplest models of the stock price process is given by a process in which the natural log of stock prices follows a random path, whose expected mean $\langle S \rangle$, grows with time (i.e., drifts) according to

$$\langle S \rangle = S_0 \exp(\mu(t - t_0)),$$

where the path begins at stock price S_0, at time t_0, and grows exponentially with time t according to some constant μ.

The path's size of statistical fluctuations is measured approximately by the standard deviation of the stock price percentage returns or, more accurately, by the standard deviation of the changes of the natural log of stock prices. This calculated result for an actual stock (or indeed any security) price history is common in derivatives finance and is called the *historical volatility* of the security.

$$\text{Historical Volatility} = \sqrt{\frac{\sum_{j=1}^{j=N} \left[\ln\left(\frac{S_j}{S_{j-1}}\right) - \text{mean}\left(\ln\left(\frac{S_j}{S_{j-1}}\right)\right)\right]^2}{N-1}}.$$

Note that $\ln(x)$ denotes the natural log of x, meaning that if $y = \ln(x)$ then $x = e^y$.

To bring the model we are constructing into contact with reality, imagine looking at a stock price that follows a path with a distribution of price changes that is *perfectly* lognormal.

First, for a large number of observations of price changes N, the average of the percentage changes of the stock price tends to a constant, the *expectation value,* driven by the model drift, reflecting that average prices in this model drift upward exponentially.

Second, the measured historical volatility will tend to the input volatility σ of the model as we calculate it using more and more stock price changes, that is, as N tends to infinity. Taking an average of many measurements n, of the historical variance over N stock price changes will have a distribution with a mean that is the input volatility squared. Alternatively, taking an average of n values for historical volatility using N market-close-price changes for the stock (*N-day historical vol*' in market parlance) will have a mean that is only approximately the input volatility, even as the number of observations n goes to infinity. However, as long as more than thirty price changes are used ($N > 30$) for each calculation of historical volatility, the difference between the expected value of the average and the input volatility is less than a few percent.

Note that an N-day historical volatility has an expected standard deviation of order volatility over $\sqrt{N}$. This means that an observed path of trailing historical volatility fluctuates around its mean with percentage difference to the mean of order $1/\sqrt{N}$. For example, a graph of trailing 30-day historical volatility on a perfectly lognormal stock price with actual volatility 0.30, will fluctuate around a mean value within a few percent of 0.30, and more than half the time will be within a band approximately given by 0.25 and 0.35. In conclusion, we *can* make observations directly comparing or contrasting real markets with this lognormal model.

These features of a lognormal distribution are already of use for some qualitative match to the real markets. For example, historical graphs of the major indices show long periods of something like exponential growth (most easily seen as straight lines on *log index* versus *time* plots) and graphs of trailing historical volatility, for a particular stock, do show a lower variance with larger samples of daily price changes. However, while a cursory look at the data shows some qualitative match, it simultaneously shows significant differences that make this model a schematic fit at best. Typical stocks can have a historical volatility graph that has a mean near 0.30 and typical range of fluctuations between 0.25 and 0.35 for a year or more and then a business announcement can cause a stock price change that makes historical volatility jump to 0.60. Such events are too frequent for the lognormal distribution and lead directly to discussions of fat tails.

Finally, if this process were to imply that there is a unique option price, then this would obviously be a very valuable qualitative tool with the potential to be quantitatively useful. This theoretical option price might not match the market exactly because the real market for the stock does not have exactly lognormally distributed stock price changes (in fact, the real market is not continuous and is influenced by the recent past, i.e. non-Markovian), but this price would still be a useful guide to very expensive or very cheap options and how to trade and risk manage them.

Many processes for stocks can indeed be used to derive option prices and the simplest model—an option on a stock with a lognormally distributed price process and no dividend—results in the famous Black-Scholes formula for the price of the option. Understanding the techniques of modeling price processes and of finding option prices consistent with a chosen underlying price process are the principal objectives of this book.

1.2.6. Book Outline

This book outlines all of the basic mathematics to understand the derivation of pricing equations for various derivatives from the starting point of assuming a process for the underlying security price. It turns out that a partial differential equation (PDE) exists for each process and which the pricing formula of all derivatives on this underlying security solve. We will derive this equation and solve it. Solving the equation can be done analytically or, more likely, numerically.

After some preliminary mathematics review in chapter 2, chapters 3 and 4 deal with a formalism to describe stochastic processes and its application to finance. Chapters 5 and 6 then derive the pricing equation based solely on stochastic price fluctuations as the source of risk. Chapters 7 and 8 repeat the development again for interest rates with a focus on the constraints on the form for the process that arise due to trying to consistently model the stochastic movements of the interest rate *curve* rather than just a single security price. Chapter 9 deals with analytic and numerical solutions to the types of equations that arise in derivatives modeling. Chapter 10 incorporates a simple binary probabilistic model of default into the framework and finally Chapter 11 combines much of the previous work into some specific examples of models of derivative securities. A few sample exercises together with solutions are provided and range from the almost inane, but actually very important, “*do I get this?*” type question to much more difficult “*I can’t do this!*” type questions. Four topics are relegated to appendices because they were either algebra intensive or outside the flow of the text but nevertheless they are important.