10

Financial Engineering: Some Tools of the Trade

The expansion of derivatives markets has given birth to the new profession of financial engineering. Financial engineers are the specialists who deal with the quantitative aspects of the derivatives business and in this chapter we discuss the evolution of this profession. We describe some of the basic numerical tools that are used by financial engineers to price derivatives and in risk management. We also discuss two challenging problems that financial engineers have worked on and have now solved: the valuation of Asian options and the pricing of complex American style options using Monte Carlo simulation.

We saw in the last two chapters that the increased use of derivatives has caused profound changes in financial practice. There has been an expanded use of derivatives both by financial and non-financial corporations. Derivative instruments have become more complicated and more sophisticated. The technology has been extended to new areas of application such as credit, power and weather. In Chapter 6 we saw that modern risk management often involves complex derivative strategies. The overall management of risk is now of central importance to financial institutions and non-financial corporations. These developments have, in turn, created a demand for individuals with strong analytical and quantitative skills who can handle the technical aspects of derivatives and risk management.

Individuals who work in this area have backgrounds in quantitative disciplines such as mathematics, engineering, physics or economics. In particular, the employment opportunities in this field are often attractive to physicists. As many physicists made the transition to Wall Street, the term “rocket scientists” was coined in the 1980s to describe them. This term has now become somewhat *passé* and the more prosaic job title of quantitative analyst (“quant” for short) is now widely used. Less flattering terms such as “derivatives geek” are also used. The term “financial engineer” is now the most popular to describe individuals who work as quantitative analysts in the derivatives and risk management fields.

Initially the demand for strong quantitative skills came mainly from Wall Street investment banks. Traders needed advice on how to price and hedge different types of derivatives. Nowadays, financial engineers work in many different types of organisations across the world and their tasks can range from constructing models of electricity markets to implementing a risk management system for a pension fund. They could also give advice to accountants who may not otherwise be capable of auditing a derivatives book. In recent years, financial engineers have been involved increasingly in risk management, building the underlying models and creating the necessary software.

Financial engineers use a wide range of computational tools in their trade. These methods were already well known to the scientific community before they were first used in financial applications because finance has only emerged as a quantitative discipline within the last 50 years. These tools have become very important in financial engineering as the applications have become more complex and, because of increased computer power, they can now solve very large-scale problems.

In this chapter, we describe some very basic numerical approaches that are used in pricing derivatives and in risk management. These approaches were initially used to obtain prices for simple contracts that were extensions of the basic European call and put contracts but now they have much broader applications. The spread of option pricing was greatly facilitated by the introduction of a number of numerical methods. These methods not only made it easier to understand the basic model: they also made it easier to value non-standard options and compute the items required to set up the replicating portfolio. The stimuli for the development of numerical methods in finance were the Black, Scholes and Merton papers. When these papers were published in 1973, they were inaccessible to most practitioners and finance professors, who did not have the mathematical background to understand them or use the results. In this connection, the binomial method played a valuable role in translating the esoteric mysteries of the Itô calculus into a simple and intuitive numerical method that could be understood by traders and implemented by MBA students. It provides a flexible method of obtaining prices for some basic derivatives contracts.

THE BINOMIAL METHOD

The idea of approximating a continuous distribution with a simpler discrete distribution has a long history in physics and mathematics and was used by Bachelier (1900) in his thesis. The first person to suggest using the binomial model as a method to price options seems to have been Bill Sharpe. He had the idea of using this model to capture both the stock price movements as well as the essence of the hedging argument.[1] The binomial method was developed more fully by John Cox, Steve Ross and Mark Rubinstein and published in 1979. Rubinstein summarised the method as follows:[2]

> It showed in a very simple way the basic economics that underlay option-pricing theory in a mathematically unadorned fashion.

The binomial method became much better known to the financial community through the publication of an influential book by Cox and Rubinstein (1985).

We introduced the binomial method in Chapter 4 and can summarise it as follows. We divide the time period into discrete steps and assume that in a single step the asset price can move either up or down. The size of the up movement and the down movement remains fixed. This framework enables us to model the uncertainty in the underlying asset's price in a convenient way. At any vertex there are just two possibilities: the asset price either goes up or down. We saw that if we had another asset that was risk-free, then at each step we could match the value of a derivative security to that of a portfolio, which had the right investments in the underlying asset and the risk-free bond. As we then have a portfolio that replicates the derivative's value one time step ahead, we can use the no-arbitrage principle (from Chapter 3) to find the current price of the derivative. Under the binomial approach we work backwards, one step at a time, until we obtain the price of the derivative at the current time.

The binomial method has a number of advantages:

- ❑ It is a very useful way to obtain the price of a number of common derivative contracts. For example, it can be used to price an American option because the early exercise feature can be modelled at each time step by testing if it is better to exercise the option or hold on to it.
- ❑ It has the simplicity and visual clarity of a spreadsheet: one can see directly how the method works and, just as with a spreadsheet, it is very easy to handle on a computer.
- ❑ It has an elegant economic interpretation because the construction of the replicating portfolio, which is an economic concept, ties in directly with the structure of the binomial tree.

It is common in science for the same discovery to be made almost simultaneously by different people and the binomial tree model for the pricing of stock options is a case in point. Within weeks of the publication of the Cox–Ross–Rubinstein paper, Richard Rendleman and Brit Bartter published a paper on the very same topic. The Cox–Ross–Rubinstein paper was the lead paper in the September 1979 issue of the *Journal of Financial Economics*[3] and the Rendleman–Bartter paper was the lead paper in the December, 1979 issue of the *Journal of Finance*. Bartter and Rendleman collaborated on this project when both were on the faculty at Northwestern University in the late 1970s.

We next turn to a discussion of two other methods that are also widely used to obtain prices of options and other derivatives. Both these methods were first applied to price derivatives in the early 1970s.

The Canadian connection

We start at the University of British Columbia (UBC) in Canada. This university became an important centre for option research in the 1970s. In particular, it became an active research centre for option pricing. Michael Brennan and Eduardo Schwartz (a doctoral student of Brennan's) made pioneering contributions to derivatives research. David Emanuel invented the Asian option while he was an assistant professor of finance at UBC. Another faculty member, Phelim Boyle, wrote the first paper that applied the Monte Carlo method to finance problems. Other academics who would later make contributions to the field also spent time at UBC. These included Stuart Turnbull, who later worked with Robert Jarrow on the development of credit risk models, and John Hull (1999), who was active in the options area in mid-1980s and wrote a well-known textbook, was also a visiting professor at UBC.

Under Brennan's guidance, Schwartz became interested in the problem of valuing American warrants in the Black–Scholes–Merton (BSM) framework. American warrants pay dividends and their exercise prices can change but their price still obeys the BSM differential equation. At this time, the Cox–Ross binomial method had not been published so there was no simple way to price them.

THE FINITE DIFFERENCE METHOD

We mentioned in Chapter 5 that the Black–Scholes equation for the price of an option is a differential equation. Merton had shown that any type of derivative contract written on a stock satisfied a similar type of equation. The contractual provisions could be translated into mathematical conditions known as boundary conditions. Until the advent of the BSM model such equations were not widely used in finance. However, they had been used for a long time in mathematics, physics, engineering and chemistry. In a few exceptional cases these equations have closed-form solutions. Otherwise they can be solved using *finite difference methods*.

Merton had set out the problem as a partial differential equation so it was natural to use methods from this field. To get warrant prices this equation would have to be solved numerically. Schwartz discussed this problem with Phelim Boyle, who put Schwartz in touch with Alvin Fowler. Fowler had a background in nuclear engineering and was an expert in computer programming.[4] He was well used to solving partial differential equations and had employed them before in physics and fluid dynamics. With Fowler's help, Schwartz wrote a Fortran program that was able to provide numerical solutions to the BSM equation for American options. This method was the finite difference method and we describe it in more detail later.

The idea behind the finite difference method is to start at maturity where the solution is known and then find the solution at regular time intervals all the way back to the present. The future is mapped into a regular grid of stock prices and times to maturity. Figure 10.1 illustrates this method. The vertical lines correspond to a fixed time and the horizontal lines correspond to fixed stock prices. At maturity, we know the value of the call option for each stock price so we can fill in all the maturity option prices. We know how the call price evolves according to the BSM equation. By using this equation to handle discrete time steps we can connect the call values at two successive time points. (Wilmot, Dewynne and Howison (1993) explain in detail how to use the finite approach to value derivatives.) This gives a large set of equations for the call prices one period earlier for each stock price on the grid. These equations can be solved on a computer to give the individual option prices at each grid point one small time step from maturity. We then repeat the process moving backwards, one step at a time, until eventually we arrive at the current time. Special features like dividend payments can be accommodated in the program and it can also be modified to handle the early exercise feature of American options.

Eduardo Schwartz used the finite difference method in two different applications: the valuation of AT&T warrants, which were more complicated than standard options, and the valuation of the guarantees embedded in certain types of life insurance contracts.

Under these insurance contracts, the premiums were invested in a stock portfolio and when the policy matured the policyholder would receive the

Figure 10.1 Finite difference method

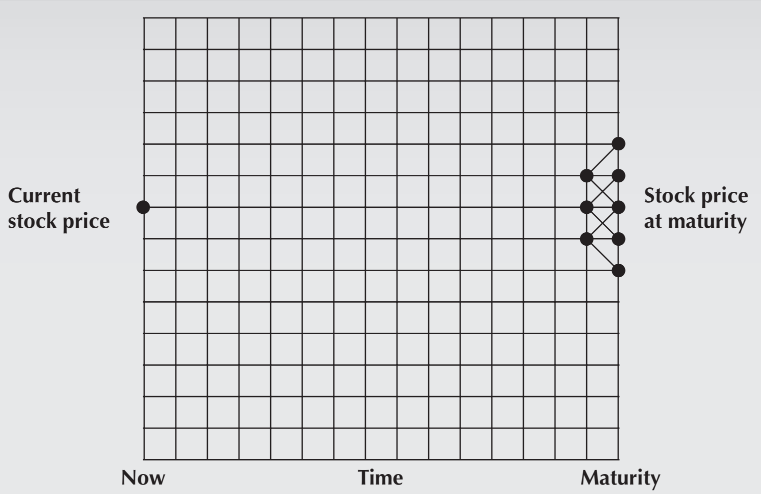

market value of the portfolio. However, there was also a guaranteed minimum floor in case the stock market did poorly. These guarantees were popular in the UK and the market fall in 1974 provided a vivid reminder of their value. Traditional actuarial methods were not really suitable for dealing with financial guarantees of this nature. Michael Brennan noted that these guarantees corresponded to long-term put options on equity portfolios. In his thesis, Schwartz used the finite difference method to obtain prices for these guarantees and also AT&T warrants. In subsequent work, Brennan and Schwartz wielded this weapon with considerable success.

Brennan (1999) has noted the importance of this numerical approach in the introduction to a volume of his collected papers:

> Armed with numerical skills, we discovered that the solution to a whole range of problems was within our reach. We valued American put options using over the counter data from Myron Scholes and found that before the Black–Scholes era there were big differences between the Black–Scholes prices and the market prices. Contemporaneously with Oldrich Vasicek, we began to apply the same principles to interest rate contingent claims. (Brennan and Schwartz (1977)); our inspiration was the humble savings bond which gave the investor the right to redeem early and which at that time played a major role in Canadian government finance.

The finite difference approach continues to be a useful tool for the computation of numerical values of the prices of derivatives. This approach can handle the early exercise feature of American options. However, if the derivative is based on the value of several assets or variables it is generally more efficient to use another method called the Monte Carlo method.

THE MONTE CARLO METHOD

We start with the story of how the Monte Carlo method was first used to value options. While Eduardo Schwartz was testing his first programs to value the European put options embedded in the insurance contracts, he had frequent discussions with Phelim Boyle. Boyle wanted a quick way to obtain values that would verify Schwartz's numbers. He was motivated by reading a working paper by Cox and Ross (1976), which showed how an option could be valued by pretending that the stock's average return was equal to the risk-free return and discounting the expected value of the option payoff under this assumption. The Monte Carlo method provides a simple way to compute an average, so Boyle used this method and was able to verify Schwartz's results.

The name "Monte Carlo" comes from the city of the same name in Monaco because the method is based on the use of so-called *random* numbers, which can be generated by a roulette wheel. The first large-scale applications of the Monte Carlo method were in physics and arose from work on the Manhattan Project in the 1940s. Truly random numbers are unpredictable. For example, if you throw a six-sided die then it will land on any one of the numbers from one to six; this is one way of generating a random number between one and six.

Here is an example that illustrates how the Monte Carlo method can be used to value a security using random numbers to compute the average value. Suppose there is a security that will pay either 10, 20, 30, 40, 50 or 60 and that each of the six possible payoffs is equally likely. We could simulate the situation by throwing a die. If the die shows a one, the security pays 10; if the die shows a 2 the security pays 20 and so on. To estimate the average payoff of the security using the Monte Carlo method, we would throw the die a large number of times and find the average value of the payoff. Instead of throwing the die, we can generate the outcomes on a computer. The technical term used to describe the generation of a possible outcome on the computer is a *simulation trial*. For example, we used 100 simulation trials and found that the average payoff was 36.6. If we increase the number of throws, we will obtain a more accurate estimate. For example, based on 100,000 simulation trials we obtained an estimate of 35.06.

For this example, we can compute the accurate value by other methods and it works out to be 35. The Monte Carlo method has the property that, as we increase the number of simulation trials, the estimate will converge to the true value (35 in this case). The estimate that we obtain, however, contains some error. Nonetheless, our estimate of the average value itself has a distribution around the true value. In fact, it will have a normal distribution and we can estimate our error because we can estimate the standard deviation of this distribution. This means that when we use the Monte Carlo method we obtain not only an estimate of the answer but information on how accurate the results are.

The Monte Carlo method can be used to value derivatives and we illustrate the procedure for the case of a standard European call option. First, we generate a possible stock price at the maturity of the option. This can be easily carried out on a computer by selecting a random outcome from the stock price distribution. In the BSM case, this distribution is lognormal and we only require its expected value and the standard deviation to generate its distribution. Second, we compute the option payoff by comparing this stock price with the option's strike price. If the stock price exceeds the strike price, the call payoff will be equal to the difference. If the stock price is below the strike price the call payoff will be zero. Third, we repeat this process many times thus obtaining the values of the call option at maturity for the different stock prices. Fourth, we compute the average of these option payoffs. Finally, we convert this average payoff at option maturity to its current value using the risk-free interest rate. This provides an estimate of the price of the call option.

We now give an example that will help to explain the method. Assume we want to value a European call option on a stock whose current price is US$100; the strike price of the option is also US$100 and it will mature in one month. We assume that this stock does not pay any dividend during the next three months. The standard deviation of the return on the stock is 25% and the risk-free interest rate is 6% per annum. Table 10.1 shows the steps in the Monte Carlo method, assuming we just use 10 trials.

In the first trial, the computer generated a stock price at maturity of US$126.81 and in this case the payoff on the call option was US$26.11. In the third trial, the computer generated a stock price at maturity of US$93.88 and the payoff on the call option was zero. The average of the 10 possible payoffs in the third column is 6.96. To obtain an estimate of the current option price, we discount it for three months at 6% obtaining 6.86. The Monte Carlo estimate of the call price is US$6.86. The accurate price in this case, from the BSM formula, is 5.73. The Monte Carlo price differs considerably from the accurate price because we just used 10 simulation trials. If we had used more trials the Monte Carlo estimate would have been closer to the accurate value. The accuracy of the Monte Carlo method is proportional to the square root of the number of simulation trials. This

Table 10.1

Trial number	Stock price at maturity (US$)	Option payoff (US$)
1	126.81	26.81
2	122.77	22.77
3	93.88	0
4	96.57	0
5	88.66	0
6	92.28	0
7	89.47	0
8	115.94	15.94
9	104.11	4.11
10	92.26	0

means that if you want to increase the accuracy by a factor of 10, you have to increase the number of trials by a factor of 100.

The Monte Carlo method is well suited for complicated valuation problems. For example, it can be used to find the price of an equity derivative whose payoff depends on several underlying stock prices.

Another example of the application of the Monte Carlo method would be the valuation of the Asian option that we introduced in Chapter 2. The payoff on the Asian option depends on the average of the asset prices over some time. Using the Monte Carlo method we use the computer to simulate one possible price path. Along each path, we can simulate the asset price path so that we obtain a value for the asset price at each point on the path where it is needed. For instance, the contract might define the average based on prices at the end of each day, or at the end of each week. We can compute the average price of the asset along this path from these prices and this average is used to compute the option's payoff for this particular path. Then we repeat this process and obtain the Asian option's payoff for each path. We take the average value of these payoffs over all the paths. The final step is to convert this average payoff at option maturity to its current value using the risk-free interest rate. This provides an estimate of the price of the Asian call option.

The Monte Carlo method is now widely used in risk management applications. A common problem involves the estimation of the distribution of the profit-and-loss statement of a portfolio at some future date. This information may be required as an input for a value-at-risk (VAR) calculation (see Chapter 6). The future value of the portfolio can be estimated from the price movements of each of its component securities. The Monte Carlo method can be used to estimate the future value of the portfolio by estimating the market values of its individual parts.

The Monte Carlo method has two main drawbacks. For large-scale problems, a naive application of the method can waste a lot of computation time. Indeed it has been described as "The most brutish of the brute force methods".[5] However, there are tricks that can be used to make the method more efficient. The second drawback concerns the valuation of American options by Monte Carlo. This has proved to be a challenging numerical problem and, at one time, it was believed that American options could not be valued using this method. As we will see below, financial engineers have made considerable progress in solving this problem.

ASIAN OPTIONS: THE QUEST FOR SOLUTIONS

Asian options, or average options, have their payoffs computed with reference to the average price of the underlying asset or commodity. They are widely used for hedging commodity price risk and currency risk. This averaging feature means that Asian options are more difficult to value than standard options because the payoff depends on the asset price at many different times, not just at the time when the option contract matures. Financial engineers have developed a number of different ways to handle this problem. In this section, we start with a brief review of the development of Asian options and then discuss some of the approaches that have been developed to value them.

The idea of basing a contract on the average value of some variable has been around for many years. For example, in some pension plans the pension benefit is based on the plan member's average yearly salary taken over the five years prior to retirement. To our knowledge, David Emanuel was the first person to propose an option based on the average when he was an assistant professor at UBC in 1979. Emanuel also noted that if the option payoff is based on the geometric average rather than the arithmetic average, then there would be a simple expression for the price of the option.[6]

Angelien Kemna and Ton Vorst independently discovered this result in 1987. Kemna and Vorst's research was motivated by a commodity-linked bond issued in 1985 by the Dutch venture capital company Oranje Nassau. Each bond contained an embedded call option to purchase 10.5 barrels of North Sea oil. An investor who bought the bond was entitled to the appreciation (if any) in oil prices over the strike price. To pay for this feature, Oranje Nassau was able to pay a lower coupon rate on the bond than if the bond did not have the option feature. In order to protect itself against possible price manipulation just prior to the option maturity, Oranje Nassau based the settlement price of the option on the average of oil prices over the previous year. Kemna and Vorst (1990) showed how this feature of the contract could be valued.

The term Asian options was first coined by financial engineers and traders working for Bankers Trust who independently invented this concept. Bill Falloon (1995), describes the story of how they came up with the idea.

> The geographical associations can be confusing. Most Asian options can only be exercised at maturity and hence they are of the European type. However some Asian option contracts can be exercised early. There is no standard name for such contracts but the meaning of the terms Asian American or American Asian is already firmly established in the language.

We now turn to a discussion of a few of the different approaches that have been developed by financial engineers to value Asian options. More precisely, we will discuss options where the payoff is based on the arithmetic average of the price of the asset and which can only be exercised at maturity. As we have mentioned earlier, this problem can be solved numerically using Monte Carlo simulation.

There is a clever trick that can be used in this case to speed up the computation time if we are using Monte Carlo simulation to solve the problem. The trick involves using information from a related problem for which we know the exact solution. In the case of the Asian option, the related problem is an option based on the geometric average. In terms of simulation, the technique is known as the *control variate procedure* and the option based on the geometric average of the prices is the control variate in our case.

We first note that the arithmetic average of a set of stock prices will be strongly correlated with the corresponding geometric average. If the arithmetic average is large, so is the geometric average and if the arithmetic average is small, so too will be the geometric average.

There is a very simple formula for the price of the option based on the geometric average used. We use the Monte Carlo method to estimate the price of the option based on the arithmetic average and the corresponding option based on the geometric average taking care to use the *same random numbers* for both calculations. Then we compare the estimate of the geometric average option from our simulations with the accurate price from the formula. This comparison tells us how biased the estimate of the geometric average option price is. It is reasonable to suppose that the estimate of options based on the arithmetic average suffers from a similar bias because it was generated using the very same random numbers. We can use this information to remove the bias from our Monte Carlo estimate of the arithmetic average option. This procedure gives excellent numerical options for short to medium-term options (up to five years).

Binomial trees

The binomial method is a very inefficient tool for pricing Asian options because it quickly leads to an enormous number of computations. This is because in the binomial tree the number of terminal asset prices increases at the same rate as the number of time steps. In a one-period tree we have two possible final asset prices. In a two-period tree, we have three possible final asset prices. (The figures in Chapter 4 illustrate this.) In a three-period tree, there are four final asset prices. In general, when the number of periods is equal to N the number of final asset prices is $(N + 1)$. We describe this pace of growth as being "linear".

To use the binomial method to price an option based on the average, we need to store information on all the different possible paths through the tree. This is because we need to compute the average asset price for all the possible price paths. The number of paths quickly becomes very large and this is the source of the problem. In a one-period tree, there are just two paths, in a two-period tree there are four paths and in a three-period tree the number of different paths through the tree is eight. For a general N-period tree, the number of different paths through the tree is 2^N. The number of paths is equal to 1,024 for a ten-period tree, over a million for a 20-period tree and over 33.5 million for a 25-period tree. If we have a one-year option with 50 weekly averaging points, then the number of different paths through the 50-period tree is $(1.25)10^{15}$ (a number with 16 digits). This number of paths is much too large to deal with on a computer, which is why the binomial method is not suitable for pricing Asian options.

It turns out that there is a closed-form solution for the price of an Asian option, which is based on the arithmetic average. This solution is based on fairly sophisticated mathematics.

At this point, it may be useful to explain why financial engineers find the quest for closed-form solutions so fascinating. Recall that the towering example of closed form solution in this field is the BSM formula for a standard European option. We described this formula in the Appendix to Chapter 5 and saw that the price of a standard call can be written in terms of five input variables. Closed-form solutions are often simpler and more intuitive than numerical solutions. They can lead to fresh insights and sometimes have an intrinsic beauty of their own. Sometimes, as in the case of the American put option, the closed-form solution does not appear to exist. In other cases such as the arithmetic Asian option case, it was not known if a solution existed or not. The intellectual challenge was therefore to find it if does exist.

The closed form solution to the Asian option involves some elegant but complicated mathematical expressions. To the best of our knowledge the first person to solve this problem was Eric Reiner while he was a doctoral student in chemical engineering. Unfortunately Reiner's solution has not been published. Independently, Marc Yor and Hélyette Geman (Yor, 1993; Geman and Yor, 1993) also developed a closed-form solution for the price of an Asian option based on the arithmetic average. This closed-form solution deepens our knowledge about the theoretical structure of Asian options. The formula is elegant from a mathematical perspective but it is hard work to obtain numerical solutions from it in practice. Other approaches such as finite difference methods and Monte Carlo methods are normally used.

VALUATION OF AMERICAN OPTIONS USING MONTE CARLO

American options are harder to value than European options because they can be exercised at any time. For some basic contracts such as an American option on one underlying asset, either the binomial tree or the finite difference approach provides a practical and efficient method of finding the price. For certain more complicated American-style derivatives, such as those based on several underlying assets, both these methods become inefficient. Normally, the numerical weapon of choice when there are many variables would be the Monte Carlo simulation.

It turns out that the valuation of an American style derivative using Monte Carlo simulation is a very hard problem to solve. Indeed, until Tilley published a paper in 1993, it was generally believed that American options could not be valued using the Monte Carlo approach. We now explain why the problem is so challenging.

The price of an American option is based on the assumption that the holder of the option exercises it optimally. The valuation procedure has to incorporate this decision problem. At each step, the decision is whether to exercise the option or continue to hold it. Usually the best way to tackle this problem is by working backwards from the option's maturity. However, in the Monte Carlo approach the future asset prices are generated from the current asset price and so we are marching along the price path. At any point on the price path, the early exercise decision requires some information about the future and, in the standard Monte Carlo approach, all we know is the price path up to this point; the future is still to unravel. This forward marching approach is in direct conflict with the requirements for the valuation of an American option because we have to use information based on the future to decide whether we should exercise the option or hold on to it.

Tilley's key insight was to adjust the Monte Carlo method to capture some of the aspects of a binomial tree. He had the idea of sorting the stock price at each time step into ordered bundles, so that the stock prices in a given bundle were close to one another. He then assumed that all the stock prices in a bundle had the same *holding value*. The holding value is the value of the option if it is not exercised. Tilley computed the holding value for each bundle by discounting the expected value of the option prices associated with the successor stock prices of the bundle one step ahead. In his own words:

> The goal of this paper is to dispel the prevailing belief that American-style options cannot be valued efficiently in a simulation model, and thus remove what has been considered a major impediment to the use of simulation models for valuing financial instruments. We present a general algorithm for estimating the value of American options on an underlying instrument or index for which the arbitrage-free probability distribution of paths through time can be simulated. The general algorithm is tested by an example for which the exact option premium can be determined.

Since Tilley's paper, other authors have developed more generalised and efficient methods to value American options by simulation, but Tilley's paper was of great importance because it showed that the problem could be solved.

CONCLUSION

In this chapter, we have discussed some of the methods used by financial engineers. We concentrated on the basic numerical methods for pricing derivatives and provided some historical context. Nowadays the most challenging numerical problems arise in the context of portfolios and risk measurement. The field of financial engineering is attracting some very gifted graduates who are well-equipped to surmount these challenges.

1 Mark Rubinstein describes discussions with Bill Sharpe on this topic at a 1975 conference in Israel. For details see Rubinstein (1999).
2 *Derivatives Strategy*, March 2000. Interview with Mark Rubinstein.
3 See Cox, Ross and Rubinstein (1979).
4 Alvin Fowler passed away on February 8, 1999. A summary of his accomplishments is contained in the website: URL: http://www.itservices.ubc.ca/newscentre/into_it/spr99/memoriam.shtmlA.
5 Oren Cheyette (1997), website: URL: http://www.barra.com/Newsletter/nl164/TNCNL164.asp
6 The simplest way to explain the geometric average is by example. The geometric average of any two numbers is the square root of their product. For example, the geometric average of 1 and 4 is 2. The arithmetic average in this case is 2.5. In the case of three numbers, the geometric average is the cube root of their product, eg, the geometric average of 1, 3 and 9 is the cube root of 27, which is 3. The arithmetic average of these last three numbers is 4.33.