COMPSTAT'2008

Tutorial on
Computational Methods in Finance

James E. Gentle

Physica-Verlag
A Springer Company

Computational Methods in Finance

James E. Gentle

Department of Computational and Data Sciences, George Mason University
Fairfax, Virginia 22030, U.S.A., *gentle@gmu.edu*

Abstract. This tutorial will cover a variety of methods of computational statistics in the analysis of financial data. Because I do not assume a background in finance, I begin with an overview of the properties of asset prices and the models used in their analysis. Many of the standard models of financial data rely on simplifying assumptions about the distribution of random components. Stochastic models of the price of a stock may view the price as a random variable that depends on previous prices and some characteristic parameters of the particular stock. Under the assumption of a Gaussian distribution of rates of return, simple formulas for fair prices of derivative assets can be developed. We will discuss this approach, some of the problems the Gaussian model, and then some variations of the stochastic models.

We will address the use of Monte Carlo methods in the pricing of derivatives and in other financial applications. We will also discuss the use of clustering, classification, and pattern recognition in studying classes of assets and asset prices. The emphasis will be on the statistical methods and on the computations, rather than on topics in the domain of finance.

None of the material presented in this tutorial represents original research; rather, it is a review of some of the important results and a discussion of some open problems and future directions. Modeling of financial data presents interesting challenges to any statistician, and many problems on the frontier of financial modeling can be addressed by persons with only a brief introduction to the field of finance.

Keywords: Black-Scholes, derivatives pricing, GARCH models, geometric Brownian motion, Monte Carlo methods, volatility

1 Introduction: Financial Data

The study of financial data probably began with the emergence of capitalism (even before it was called that). There are many reasons for studying financial data. One, of course, is the profit motive. Investors, speculators, and operators seek an advantage over others in the trading of financial assets. Academics often find a different motive for studying financial data: the challenges of developing models for price movements. (In fact many academic researchers believe in models of price changes that do not allow one trader to have an advantage over another "informed" trader.) Finally, government regulators and others are motivated by an interest in maintaining a fair and orderly market. With more and more retirees depending on equity investments for their livelihood, it becomes very important to understand and control the risk in portfolios.

Many of the early applications of statistics involved financial data. Various simple statistics such as means, variances, and correlations were computed for given assets over fixed time intervals. Methods for analysis of time series were applied, and indeed many advances in time series analysis were motivated by study of financial data. Methods for statistical graphics were also developed to aid in the analysis of financial data.

Data mining of asset prices also has a long history. Traders have always looked for patterns and other indicators that they think may help to predict stock price movements. (I use the term "data mining" here to refer to exploratory data analysis, especially of large or amorphous data sets.) In data mining the analyst is not just fitting the coefficients in a regression model or estimating the autoregressive order in a time series. Rather, although formal methods of statistical inference may be part of the process, in data mining the goals are somewhat less focused. The hope is to discover the unexpected. A common characteristic of data mining is the use of data from multiple sources, in different formats, and collected for different purposes.

Traders have sifted through massive amounts of data from disparate data sets looking for meaningful relationships. A somewhat frivolous example of this was the discovery several years ago of an empirical relationship between American football wins and losses and an index of stock prices. It was discovered that when the winner of the Super Bowl was a team from the old American Football Conference, the market (measured by the Dow Jones Industrial Index) went up between the date of the game and the end of the year.

No one would have expected such a relationship. It could have been discovered by mining of large and disparate datasets. We must categorize this as knowledge discovery. It actually happened. It is an interesting fact, but it is worthless. Data mining and knowledge discovery must be kept in context.

1.1 The Nature of Financial Data

Whether a statistician has any interest in financial applications or not, financial data have many fascinating properties that would be of interest to statisticians. Providing further appeal to the statistician or data analyst is the ready availability of interesting data.

Sources of Financial Data Some financial data are difficult to obtain; some have been collected at great expense, and may be expected to provide some monetary advantage to their possessors. To satisfy general academic interests, however, there is a wealth of easily available data. Statistics collected by government agencies are usually available at the agencies' websites. Most governments require that companies whose stock is publicly traded file periodic reports of particular financial characteristics of those companies. Different governments require different levels of detail and hold different standards for the reported data, but for publicly traded companies there is generally a large and reasonably current body of financial data available. Finally, the prices at which securities have traded, at least on a daily basis, are available easily readily. These securities include the basic shares of stock as well as various derivatives of these shares.

There are also various proprietary indexes of stock prices, such as the Dow-Jones Industrial Average and the Standard & Poor's 500. Some of these indexes are tracked by exchange traded funds, and some can be traded on futures markets.

One of the easiest places to get current or historical stock price data for companies traded in American markets is `http://finance.yahoo.com/`

In the following, I will occasionally use data from this source.

Diversity of Financial Data There are many different types of financial data including balance sheets, earnings statements and their various details, stock prices, and so on. Taking the view that "data" means any kind of information, we include in financial data the names and backgrounds of directors and other company officers, news items relating to the company, general economic news, and so on. Two major components of financial data are the opinions of financial analysts and the chatter of the large army of people with a computer and a connection to the internet where they can post touts. This data must be included in the broad class of financial data because they actually have an effect on other financial data such as stock prices. In the following, I will use the term "financial data" in this very broad sense.

Three of the general types of financial data are numerical. One type of numerical data relates to trading of financial assets, and for present purposes I will restrict this to trading of stock equities and options on those equities. This kind of data is objective and highly reliable. Two other types of numerical data have to do with the general financial state of an individual company and of the national and global economy. These kinds of data, while ostensibly objective, depend on the subjectivity of the definitions of the terms. ("Unemployment rate", for example, is not a measure of how many employable persons are not employed, as the term would seem to imply.) Finally, a fourth type of financial data is the text data that are relevant to the trading of stock equities. This includes such objective information as the names of the company officers and products or services of the company. It also includes statements and predictions by anyone with the ability to publicize anything. This kind of data varies from very objective to very subjective, and from very accurate to completely erroneous.

Data Quality Data analysts often deal with data that have different levels of accuracy. When it is possible to assign relative variances to different subsets of numerical data, the data can easily be combined using weights inversely proportional to the variance. The differences in the nature of financial data, however, make it difficult to integrate the data in any meaningful way. The totality of relevant financial data is difficult to define. The provenance of financial data is almost impossible to track. Rumors easily become data. Even "hard data", that is, numerical data on book values, earnings, and so on, are not always reliable. A PE ratio for example may be based on "actual earnings" in some trailing period or "estimated earnings" in some period that includes some future time. Unfortunately, "earnings", even "actual", is a rather subjective quantity, subject to methods of booking. (Even if the fundamental operations of the market were efficient, as is assumed in most academic models, the basic premise of efficiency, that is, that everyone has the same information, is not satisfied. Analysts and traders with the ability to obtain really good data have an advantage.) The integration of data of various types is perhaps the primary challenge in mining financial data. A major obstacle to the relevance of any financial data is the fact that the phenomena measured by or described by the data are changing over time.

The primary challenges in the mining of financial data arise from the vast diversity of types and sources of data and from the variation of the data in time. Understanding of the temporal variability can be facilitated by mathematical models. The models contain parameters, such as mean drift and the standard deviation of the drift in a diffusion process. These model parameters are not directly observable; rather, they must be derived from observable data. Although the models are still very approximate, in some ways the modeling has advanced faster than the statistical methods for estimating the parameters in the models. An example of this is the indirect estimation of the standard deviation using price data for derivatives and the model that relates the price to the standard deviation.

Seeking the Unexpected in Disparate Datasets An interesting discovery that resulted from mining of financial data is called the "January effect". Several years ago, it was discovered that there are anomalies in security prices during the first few days of January. There are various details in the differences of January stock prices and the prices in other months (see, for example, Cataldo and Savage, 2000), but the most relevant general fact is that for a period of over 80 years the average rate of return of the major indexes of stock prices during January was more than double the average rate of return for any other month.

While this discovery comes from just an obvious and straightforward statistical computation, and thus is a rather trivial example in data mining, it serves to illustrate a characteristic of financial data. I call it the "uncertainty principle", in analogy to Heisenberg's uncertainty principle that states that making a measurement affects the state of the thing being measured. After the January effect became common knowledge, the effect seemed to occur earlier (the "Santa Claus rally"). This, of course, is exactly what one might expect. Carrying such expectations to the second order, the Santa Claus rally may occur in November, and then perhaps in October.

1.2 Stylized Properties of Financial Data

A basic type of financial data is the price of an asset, and a basic property of the price is that it fluctuates. There is a certain amount of jargon used for other aspects of the price. Both "risk" and "volatility" refer to the standard deviation of the price. (I am not being precise here; "standard deviation" of course has a very precise meaning in probability theory, but my use of the term must be interpreted more loosely.) The term "stochastic volatility" refers to a situation in which the volatility ("standard deviation") is changing randomly in time.

A substantial amount of statistical theory and methodology has been developed for the Gaussian or normal distribution. The normal distribution appears to be a very good model for a wide range of phenomena. Much financial data, however, do not fit a normal distribution well.

Some of the characteristics of financial data that make it interesting and challenging to analyze are the following.

- Heavy tails. The frequency distribution of rates of return decrease more slowly than $\exp(-x^2)$.
- Asymmetry in rates of return. Rates of return are slightly negatively skewed. (Because traders react more strongly to negative information than to positive information.)
- Asymmetry in lagged correlations. Coarse volatility predicts fine volatility better than the other way around.
- Custering of volatility.
- Aggregational normality.
- Quasi long range dependence.
- Seasonality.

1.3 Models of Asset Prices

Anything that is openly traded has a market price that may be more or less than some "fair" price. In financial studies, a general objective is to measure "fair price", "value", or "worth". For shares of stock, the fair price is likely to be some complicated function of intrinsic (or "book") current value of identifiable assets owned by the company, expected rate of growth, future dividends, and other factors.

The price, either the market price or the fair price, varies over time. We often assume discrete time, $t_0, t_1, t_2, \ldots$ or $t, t+1, t+2, \ldots$.

The prices of individual securities, even if they follow similar models, behave in a way peculiar to the security. There are more security-specific extraordinary events that affect the price of a given security, than there are extraordinary events that affect the overall market. For that reason, at least as a beginning, we study the prices of some index of stock prices. The S&P 500 is a commonly-used index. A graph of the daily closes of the S&P 500 Index from July 1, 1987, to June 30, 2008, is shown in Figure 1.

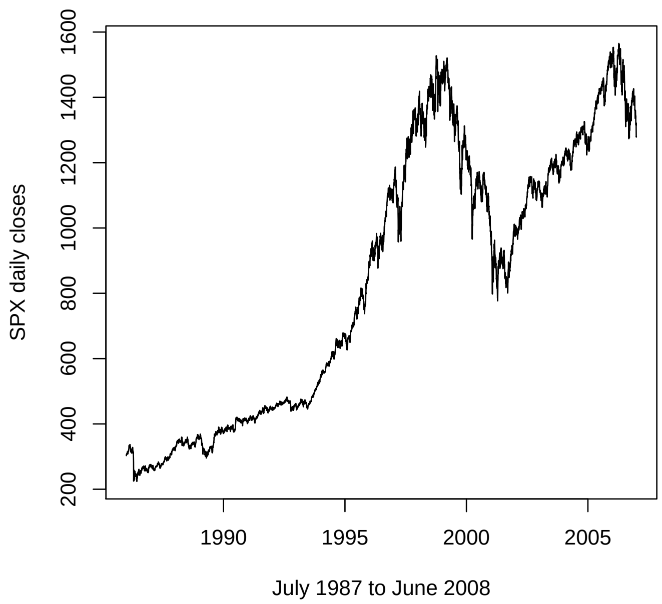

Fig. 1. S&P 500 Daily Closes (Tick marks are at beginning of the year).

A stochastic model of the price of a stock may view the price as a random variable that depends on previous prices and some characteristic parameters of the particular stock. For example, in discrete time:

$$S_{t+1} = f(S_t, \mu, \sigma) \tag{1}$$

where t indexes time, μ and σ are parameters, and f is some function that contains a random component. The randomness in f may be assumed to reflect all variation in the price that is not accounted for in the model.

There are two basic types of models of asset prices. One type is a stochastic diffusion differential equation and the other is an autoregressive model. We will briefly discuss these below.

Diffusion Models of Relative Changes of Stock Prices In the absence of exogenous forces, the movement of stock prices is usually assumed to be some kind of random walk. A simple random walk would have step sizes independent of location and furthermore could take the prices negative. Also, it seems intuitive that the random walk should have a mean step size that is proportional to the magnitude of the price. The proportional rate of change, $(S_{t+1} - S_t)/S_{t+1}$, therefore, is more interesting than the prices themselves, and is more amenable to fitting to a probability model.

The proportional rate of change, $(S_{t+1} - S_t)/S_{t+1}$, shown in Figure 2, is more interesting than the prices. The event of October 19, 1987, clearly stands out.

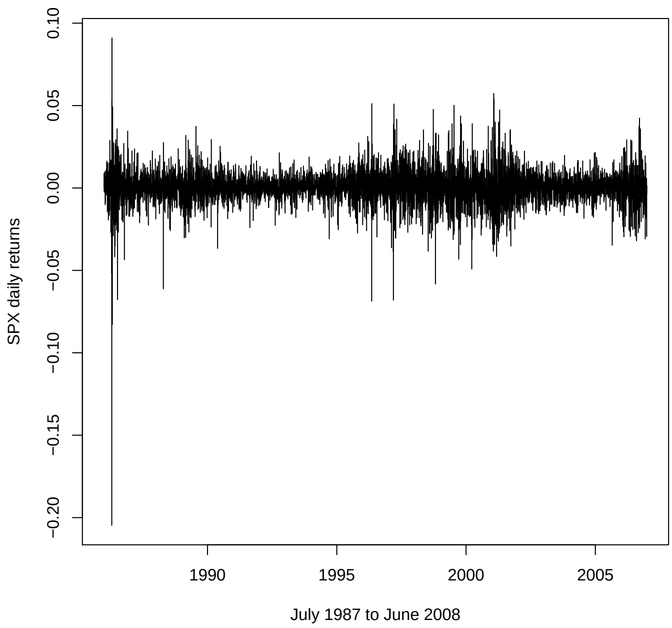

Fig. 2. S&P 500.

A good, general model of a random walk is a Brownian motion, or a Wiener process; hence, we may write the model as a drift and diffusion,

$$\frac{\mathrm{d}S(t)}{S(t)} = \mu(S(t), t)\mathrm{d}t + \sigma(S(t), t)\mathrm{d}B,$$

where $\mathrm{d}B$ is a Brownian motion, which has a standard deviation of 1. The standard deviation of the rate of change is therefore $\sigma(S(t), t)$. In this equation both the mean drift and the standard deviation may depend on both the magnitude of the price $S(t)$ and also on the time t itself.

If we assume that $\mu(\cdot)$ and $\sigma(\cdot)$ do not depend on the value of the state and that they are constant in time, we have the model

$$\frac{\mathrm{d}S(t)}{S(t)} = \mu \mathrm{d}t + \sigma \mathrm{d}B,$$

which is called a geometric Brownian motion. In this form μ is called the drift and the diffusion component σ is called the volatility. We can estimate μ and σ from the historical rates of return; the mean and the standard deviation respectively.

Although it is easy to understand the model, it is not so obvious how we would estimate the parameters. First of all, obviously, we cannot work with $\mathrm{d}t$. We must use finite differences based on Δt, but the question is how long is Δt? We first note an interesting property of Brownian motion. Its variation seems to depend on the frequency at which it is observed; it is infinitely "wiggly". (Technically, the first variation of Brownian motion is infinite.)

We note that as a model for the rate of return, $\mathrm{d}S(t)/S(t)$ geometric Brownian motion is similar to other common statistical models:

$$\frac{\mathrm{d}S(t)}{S(t)} \quad = \quad \mu \mathrm{d}t \quad + \quad \sigma \mathrm{d}B(t)$$

or

response = systematic component + random error.

Also, note that without the stochastic component, the differential equation has the simple solution

$$S(t) = c\mathrm{e}^{\mu t},$$

from which we get the formula for continuous compounding for a rate μ.

The rate of growth we expect for S just from the systematic component in the geometric Brownian motion model is μ. Because the expected value of the random component is 0, we might think that the overall expected rate of growth is just μ. Closer analysis, however, in which we consider the rate of change σ being equally likely to be positive or negative and the effect on a given quantity if there is an uptick of σ followed by a downtick of equal magnitude yields a net result of $-\sigma^2$ for the two periods. The average over the two periods therefore is $-\sigma^2/2$. The stochastic component reduces the expected rate of μ by $-\sigma^2/2$. This is the price of risk.

The geometric Brownian motion model is the simplest model for stock prices that is somewhat realistic.

Autoregressive Models of Relative Changes of Stock Prices The most widely-used time series models in the time domain are those that incorporate either a moving average or an autoregression term. These ARMA models are not even close approximations to observational data. The main problem seems to be that the assumption of a constant variance does correspond to reality. Variations, such as ARCH or GARCH models, models seem to fit the data better.

Fitting the Models Returning now to the problem of estimating a parameter of a continuous-time process, we consider the similar model for continuous compounding. If an amount A is invested for n periods at a per-period rate R that is compounded m times per period, the terminal value is

$$A\left(1+\frac{R}{m}\right)^{nm}.$$

The limit of the terminal value as $m \to \infty$ is

$$A\mathrm{e}^{Rn}.$$

This suggests that Δt can be chosen arbitrarily, and, since rates are usually quoted on an annualized basis, we chose Δt to be one year. Using the formula for compounded interest, we first transform the closing price data $S_0, S_1, S_2 \ldots$ to $r_i = \log(S_i/S_{i-1})$. Now, an obvious estimator of the annualized volatility, σ, based on N periods each of length Δt (measured in years) is

$$\tilde{\sigma} = \frac{1}{\sqrt{\Delta t}}\sqrt{\frac{1}{N}\sum_{i=1}^{N}(r_i - \bar{r})^2}.$$

We are now in a position to use the geometric Brownian motion drift-diffusion model (with the simplifying assumptions of constant drift and diffusion parameters).

Solution of the Stochastic Differential Equation The solution of a differential equation is obtained by integrating both sides and allowing for constant terms. Constant terms are evaluated by satisfying known boundary conditions, or initial values. In a stochastic differential equation (SDE), we must be careful in how the integration is performed, although different interpretations may be equally appropriate.

The SDE defines a simple Ito process with constant coefficients,

$$\mathrm{d}S(t) = \mu S(t)\mathrm{d}t + \sigma S(t)\mathrm{d}B(t).$$

We solve this using Ito's formula. Integrating to time $t = T$, we have

$$S(T) = S(t_0)\exp\left(\left(\mu - \frac{1}{2}\sigma^2\right)\Delta t + \sigma\Delta B\right).$$

Numerical methods for the solution of SDEs are not as well-developed as those for deterministic differential equations. We will not consider these numerical methods here. The interested reader is referred to Seydel (2006) or Kloeden et al. (1994).

Modeling Prices of Derivative Assets For any asset whose price varies over time, there may be a desire to insure against loss, or there may be interest in putting up a small amount of capital now so as to share in future gains. The insurance or the stake that depends on an "underlying asset" is itself an asset that can be traded. (In the following, we will refer to an underlying asset as an "underlying".)

There are many ways the insurance or stake can be structured. It most certainly would have an expiration date. It could be a contract or it could be an option. It could be exercisable only at a fixed date or at anytime before expiration. Its price depends on its nature and on the price of the underlying, so it is a "derivative asset". (In the following, we will refer to a derivative asset as a "derivative".) Determination of the fair price of derivatives is one of the major motivations for developing pricing models for stock prices as discussed in the previous section. This section is intended to provide a quick review of the derivative pricing models. More complete discussions are available in many texts, for example, Chriss (1997) or Hull (2005).

Common derivatives are "puts" and "calls" on stocks or stock indexes. A put is the right to sell; a call is the right to buy. They are "options".

There are various ways that an option agreement can be structured, and so there are different types of options. "American" style options, for example, carry the right to exercise anytime between the time of acquisition of the right and the expiration date.

How to price a derivative is a difficult question. A model for the "value" of an option may be expressed as an equation in the form

$$V(t) = g(S(t), \mu(t), \sigma(t)),$$

where $V(t)$ is "value", that is, "correct price" of the option at time t (relative to the expiration date), $S(t)$ is market price of the underlying at time t, $\mu(t)$ and $\sigma(t)$ are parametric characteristics of the underlying, and g is some function that contains a random component. As usual, we assume frictionless trading; that is, we ignore transaction costs.

The randomness in g may be assumed to reflect all variation in the price that is not accounted for in the model.

The price of the underlying will fluctuate, and so the price of the derivative is related to the expected value of the underlying at expiration.

The Price of a European Call Option A *European call option* is a contract that gives the owner the right to buy a specified amount of an underlying for a fixed *strike price*, K on the *expiration* or *maturity* date T. The owner of the option does not have any obligations in the contract.

The *payoff*, h, of the option at time T is either 0 or the excess of the price of the underlying $S(T)$ over the strike price K. Once the parameters K and

T are set, the payoff is a function of $S(T)$:

$$h(S(T)) = \begin{cases} S(T) - K & \text{if } S(T) > K \\ 0 & \text{otherwise} \end{cases}$$

The price of the option at any time is a function of the time t, and the price of the underlying s. We denote it as $P(t, s)$.

We wish to determine the fair price at time $t = 0$.

It seems natural that the price of the European call option should be the expected value of the payoff of the option at expiration, discounted back to $t = 0$:

$$P(0, s) = \mathrm{e}^{-rT}\mathrm{E}(h(S(T))).$$

The basic approach in the Black-Scholes pricing scheme is to seek a portfolio with zero expected value, that consists of short and/or long positions in the option, the underlying, and a risk-free bond.

There are two key ideas in developing pricing formulas for derivatives:

- no-arbitrage principle
- replicating, or hedging, portfolio

An *arbitrage* is a trading strategy with a guaranteed rate of return that exceeds the riskless rate of return. In financial analysis, we assume that arbitrages do not exist. This follows from an assumption that the market is "efficient"; that is, the assumption that all market participants receive and act on all of the relevant information as soon as it becomes available. One does not need any deep knowledge of the market to see that this assumption does not hold, but without the assumption it would not be possible to develop a general model. Every model would have to provide for input from different levels of information for different participants; hence, the model would necessarily apply to a given set of participants. While the hypothesis of an efficient market clearly cannot hold, we can develop useful models under that assumption. (The situation is as described in the quote often attributed to George Box: "All models are wrong, but some are useful.")

There are two essentially equivalent approaches to determining the fair price of a derivative, use of delta hedging and use of a replicating portfolio. In the following, we will briefly describe replicating portfolios.

The replication approach is to determine a portfolio and an associated trading strategy that will provide a payout that is identical to that of the underlying. This portfolio and trading strategy *replicates* the derivative. A replicating strategy involves both long and short positions. If every derivative can be replicated by positions in the underlying (and cash), the economy or market is said to be *complete*. We will generally assume complete markets.

The Black-Scholes approach leads to the idea of a self-financing replicating hedging strategy. The approach yields the interesting fact that the price of the call does not depend on the expected value of the underlying. It does depend on its volatility, however.

Expected Rate of Return on Stock Assume that XYZ is selling at $S(t_0)$ and pays no dividends. Its expected value at time $T > t_0$ is merely the forward price for what it could be bought now, where the forward price is calculated as $\mathrm{e}^{r(T-t_0)}S(t_0)$, where r is the risk-free rate of return, $S(t_0)$ is the spot price, and $T - t_0$ is the time interval.

This is an application of the no-arbitrage principle.

The holder of the forward contract (long position) on XYZ must buy stock at time T for $\mathrm{e}^{r(T-t_0)}S(t_0)$, and the holder of a call option buys stock only if $S(T) > K$.

Now we must consider the role of the volatility. For a holder of forward contract, volatility is not good, but for a call option holder volatility is good, that is, it enhances the value of the option.

Under the assumptions above, the volatility of the underlying affects the value of an option, but the expected rate of return of the underlying does not.

A simple model of the market assumes two assets:

- a *riskless asset* with price at time t of β_t
- a *risky asset* with price at time t of $S(t)$.

The price of a derivative can be determined based on trading strategies involving these two assets.

The price of the riskless asset follows the deterministic ordinary differential equation

$$\mathrm{d}\beta_t = r\beta_t \mathrm{d}t,$$

where r is the instantaneous riskfree interest rate.

The price of the risky asset follows the stochastic differential equation

$$\mathrm{d}S(t) = \mu S(t)\mathrm{d}t + \sigma S(t)\mathrm{d}B_t.$$

We have then for the price of the call

$$C(t) = \varDelta_t S(t) - \mathrm{e}^{r(T-t)}R(t),$$

where $R(t)$ is the current value of a riskless bond.

We speak of a portfolio as a vector p whose elements sum to 1. The length of the vector is the number of assets in our universe. By the no-arbitrage principle, there does not exist a p such that for some $t > 0$, either

- $p^{\mathrm{T}}s_0 < 0$ and $p^{\mathrm{T}}S(t)(\omega) \geq 0$ for all ω,

or

- $p^{\mathrm{T}}s_0 \leq 0$ and $p^{\mathrm{T}}S(t)(\omega) \geq 0$ for all ω, and $p^{\mathrm{T}}S(t)(\omega) > 0$ for some ω.

A derivative D is said to be *attainable* (over a universe of assets $S = (S^{(1)}, S^{(2)}, \ldots, S^{(k)})$) if there exists a portfolio p such that for all ω and t,

$$D_t(\omega) = p^{\mathrm{T}} S(t)(\omega).$$

Not all derivatives are attainable. The replicating portfolio approach to pricing derivatives applies only to those that are attainable.

The value of a derivative changes in time and as a function of the value of the underlying; therefore, a replicating portfolio must be changing in time or "dynamic". (Note that transaction costs are ignored.) The replicating portfolio is self-financing; that is, once the portfolio is initiated, no further capital is required. Every purchase is financed by a sale.

If the portfolio is self-financing

$$\mathrm{d}(a_t S(t) + b_t e^{rt}) = a_t \mathrm{d}S(t) + r b_t e^{rt} \mathrm{d}t.$$

The Black-Scholes Differential Equation Consider the fair value V of a European call option at time $t < T$. At any time this is a function of both t and the price of the underlying S_t. We would like to construct a dynamic, self-financing portfolio (a_t, b_t) that will replicate the derivative at maturity. If we can, then the no-arbitrage principle requires that

$$a_t S_t + b_t \mathrm{e}^{rt} = V(t, S_t),$$

for $t < T$.

We assume no-arbitrage and we assume that a risk-free return is available.

Assuming $V(t, S_t)$ is continuously twice-differentiable, we differentiate both sides of the equation that represents a replicating portfolio with no arbitrage:

$$\begin{aligned} a_t \mathrm{d}S_t + r b_t e^{rt} \mathrm{d}t = & \left(\frac{\partial V}{\partial S_t} \mu S_t + \frac{\partial V}{\partial t} + \frac{1}{2} \frac{\partial^2 V}{\partial S_t^2} \sigma^2 S_t^2 \right) \mathrm{d}t \\ & + \frac{\partial V}{\partial S_t} (\sigma S_t) \mathrm{d}B_t \end{aligned}$$

By the market model for $\mathrm{d}S_t$ the left-hand side is

$$(a_t \mu S_t + r b_t e^{rt}) \mathrm{d}t + a_t \sigma S_t \mathrm{d}B_t.$$

Equating the coefficients of $\mathrm{d}B_t$, we have

$$a_t = \frac{\partial V}{\partial S_t}.$$

From our equation for the replicating portfolio we have

$$b_t = (V(t, S_t) - a_t S_t) \, e^{-rt}.$$

Now, equating coefficients of dt and substituting for a_t and b_t, we have the Black-Scholes differential equation,

$$r\left(V - S_t\frac{\partial V}{\partial S_t}\right) = \frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S_t^2\frac{\partial^2 V}{\partial S_t^2}$$

Notice that μ is not in the equation.

Instead of European calls, we can consider European puts, and proceed in the same way using a replicating portfolio, and we arrive at the same the Black-Scholes differential equation.

The Black-Scholes Formula The solution of the differential equation depends on the boundary conditions. In the case of European options, these are simple. For calls, they are

$$V_{\mathrm{c}}(T, S_t) = (S_t - K)^+,$$

and for puts, they are

$$V_{\mathrm{p}}(T, S_t) = (K - S_t)^+.$$

With these boundary conditions, there are closed form solutions to the Black-Scholes differential equation. For the call, for example, it is

$$C_{\mathrm{BS}}(t, S_t) = S_t\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2),$$

where

$$d_1 = \frac{\log(S_t/K) + (r + \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}},$$

$$d_2 = d_1 - \sigma\sqrt{T-t},$$

and

$$\Phi(s) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{s} e^{-y^2/2}\mathrm{d}y.$$

We recall the assumptions of the Black-Scholes model:

- differentiability of stock prices with respect to time
- a dynamic replicating portfolio can be maintained without transaction costs
- returns are
  - independent
  - normal
  - mean stationary
  - variance stationary

1.4 Assessment of Models for Price Movements

Our focus is on rates of return, because that is the fundamental quantity in our pricing models. Rates of return are not directly observable, and as we indicated before because of the continuous time in the geometric Brownian motion model, there are various ways we may evaluate these derived observations.

However they are measured, the data do not appear to meet the assumptions of the model. Just from the plot in Figure 2, without any additional analyses, we can notice three things the violate the assumptions:

- the data have several outliers
- the data are asymmetric
- the data values seem to cluster

Because of the outliers, the data do not appear to be from a normal distribution. The departure from normality can be seen more easily in a normal qq plot. The plot in Figure 3 is dramatic evidence of the lack of normality.

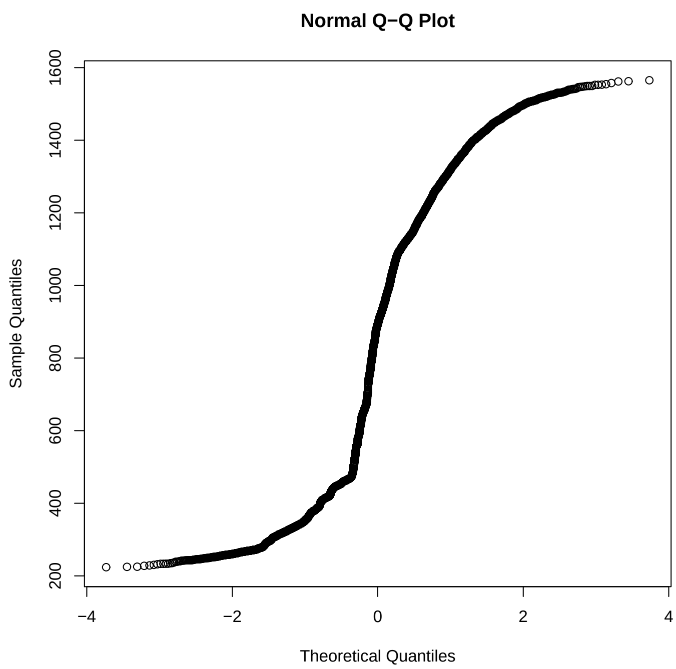

Fig. 3. S&P 500 Daily Closes (Tick marks are at beginning of the year).

Although the data seem to be centered about 0, and there seem to be roughly as many positive values as negative values, there appear to be larger extreme negative values than extreme positive values. The extreme values are more than what we would expect in a random sample from a normal distribution. The crash of October 19, 1987, is an extreme outlier; but there are others. Finally, the data do not seem to be independent; specifically, the extreme values seem to be clustered. The large drop on October 19, 1987, is followed by a large gain.

All of these empirical facts seem to bring into question the whole idea that the market is efficient. Either basic facts are changing rapidly, or there are irrational motives at play.

The returns in this example are for an index of 500 large stocks. We would expect this index to behave more in line with any model based on assumptions of independent normality than the price of some individual stock. These kinds of violations of the assumptions, however, can be observed in other indexes as well as in individual stock prices. Jondeau et al. (2007) and Rachev et al. (2005) discuss many of the issues of developing models similar to those we discuss above, except without the assumption of normality.

Both standard stochastic differential equations and simple autoregressive processes have a constant variance, but such models do not correspond well with empirical observations. This means that neither a single SDE nor a stationary autoregressive model will be adequate. There are various approaches for modifying the simple models. Instead of an SDE with a single stochastic differential component, we may incorporate a second stochastic differential component that represents Poisson jumps (see Bjursell and Gentle (2008) and Cont and Tankov (2004)), or else we can use a couple system of SDEs, one for the standard deviation of the other one. There are various ways of modifying the basic autoregressive model. One modification is generalized autoregressive conditional heteroskedasticity (GARCH).

2 Estimation of Model Parameters

In addition to the problem of choosing a correct models, in order to use the model we are still faced with the problem of estimating the parameters in the model. This is not as straightforward as it may seem. The most important parameter is th volatility; that is, the standard deviation.

2.1 Volatility

As in all areas of science, our understanding of a phenomenon begins by identifying the quantifiable aspects of the phenomenon and then is limited by our ability to measure those quantifiable aspects.

Most models of prices of financial assets have a parameter for volatility. For a model to be useful, of course, we must have some way of supplying a value for the parameter, either by direct measurement or by estimation. The volatility parameter presents special challenges for the analyst. (Recall "volatility" is defined as the standard deviation of the rate of return; and, ideally, here we mean the instantaneous rate of return.)

Although often the model assumes that the volatility is constant, volatility, like most parameters in financial models, varies over time. There may be many reasons for this variation in volatility including arrival of news. When the variation is not explained by a specific event, the condition is known as "stochastic volatility". Empirical evidence supports the notion that the volatility is stochastic.

Volatility itself is a measure of propensity for change in time. This results in the insurmountable problem of providing a value for a parameter that depends on changes of other values in time, but which itself changes in time.

There are some interesting stylized facts about volatility (see McQueen and Vorkink (2004)):

- volatility is serially correlated in time
- both positive news and negative news lead to higher levels of volatility
- negative news tends to increase future volatility more than positive news
- there are two distinct components to the effect of news on volatility, one with a rapid decay and one with a slow decay
- volatility has an effect on the risk premium

The Black-Scholes model and the resulting Black-Scholes formula include a volatility parameter, σ. One use of the Black-Scholes formula obviously is to provide an approximate fair price for an option. This was the motivation for its development. Another formula for the fair price of anything for which there is an active market, however, is the market price. For a given option (underlying, type, strike price, and expiry) and given the price of the underlying and the riskfree rate, the Black-Scholes formula relates the option price to the volatility of the underlying; that is, the volatility determines the model option price. The actual price at which the option trades can be observed, however. If this price is plugged into the Black-Scholes formula, the volatility can be computed. This is called the "implied volatility".

The implied volatility may be affected by either the absence of "fairness" in the market price, or by the market price being correlated to Black-Scholes formulaic prices.

In the case of thinly traded assets, the market price may be strongly affected by idiosyncrasies of the traders.

Bjursell and Gentle (2008) consider the effect of a superimposed jump process of the volatility, and discuss method of testing for the present of jumps

2.2 Implied Volatility

Let c be the observed price of the call. Now, set $C_{\mathrm{BS}}(t, S_t) = c$, and

$$f(\sigma) = S_t \Phi(d_1) - K e^{-r(T-t)} \Phi(d_2),$$

We have

$$c = f(\sigma).$$

Given a value for c, that is, the observed price of the option, we can solve for σ. There is no closed-form solution, so we solve iteratively. Beginning with $\sigma^{(0)}$, we can use the Newton updates,

$$\sigma^{(k+1)} = \sigma^{(k)} - (f(\sigma^{(k)}) - c)/f'(\sigma^{(k)}).$$

We have

$$\begin{aligned} f'(\sigma) &= S_t \frac{\mathrm{d}\Phi(d_1)}{\mathrm{d}\sigma} - K e^{-r(T-t)} \frac{\mathrm{d}\Phi(d_2)}{\mathrm{d}\sigma} \\ &= S_t \phi(d_1) \frac{\mathrm{d}d_1}{\mathrm{d}\sigma} - K e^{-r(T-t)} \phi(d_2) \frac{\mathrm{d}d_2}{\mathrm{d}\sigma}, \end{aligned}$$

where

$$\phi(y) = \frac{1}{\sqrt{2\pi}} e^{-y^2/2},$$

$$\frac{\mathrm{d}d_1}{\mathrm{d}\sigma} = \frac{\sigma^2(T-t) - \log(S_t/K) - (r + \frac{1}{2}\sigma^2)(T-t)}{\sigma^2\sqrt{T-t}},$$

and

$$\frac{\mathrm{d}d_2}{\mathrm{d}\sigma} = \frac{\mathrm{d}d_1}{\mathrm{d}\sigma} - \sqrt{T-t}.$$

As we have mentioned, one of the problems of the Black-Scholes formula is its assumption that the volatility σ is constant. Because of this, obviously if we substitute the observed market price of a particular option for the Black-Scholes price for that option, and do the same for a different option on the same underlying, we are likely to get different values for the implied volatility. In any event, computing an implied volatility is not straightforward (see Hentschel, 2003).

The implied volatility from the Black-Scholes model should be the same at all points, but it is not. The implied volatility, for given T and S_t, depends on the strike price, K.

In general, the implied volatility is greater than the empirical volatility, but the implied volatility is even greater for far out-of-the money calls. It also generally increases for deep in-the-money calls.

This variation in implied volatility is called the “smile curve”, or the “volatility smile”, as shown in Figure 4. The curve is computed for a given security and the traded derivatives on that security with a common expiry.

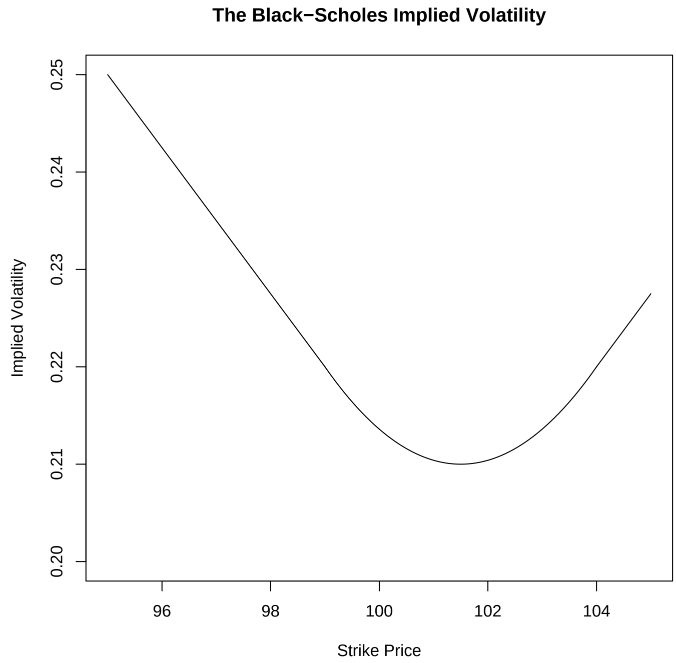

Fig. 4. The Volatility Smile.

The empirical price for each derivative at that expiry is used in the Black-Scholes formula and, upon inversion of the formula, yields a value for the volatility. Tavella (2002) and Fouque et al. (2000) provide extensive discussion of this variation in the apparent volatility volatility.

The available strike prices are not continuous, of course. This curve is a smoothed (and idealized) fit of the observed points.

The smile curve is not well-understood, although we have a lot of empirical observations on it. Interestingly, prior to the 1987 crash, the minimum of the smile curve was at or near the market price S_t. Since then it is generally at a point larger than the market price.

In addition to the variation in implied volatility at a given expiry, there is variation in the implied volatility at a given strike at different expiries. This is a term structure in volatility similar to the term structure of interest rates. The reasons for the term structure of volatility are probably related to the reasons for term structure in interest rates, but this is also not well-understood. Gatheral (2006) provides an extensive discussion of the implied volatility function from an empirical perspective.

Variation in Volatility over Time The volatility also varies in time. (This is not the term structure referred to above, which is a variation at a fixed time for events scheduled at different future times.) Volatility varies in time, with periods of high volatility and other periods of low volatility. This is called "volatility clustering".

The volatility of an index is somewhat similar to that of an individual stock. The volatility of an index is a reflection of market sentiment. (There are various ways of interpreting this!)

In general, a declining market is viewed as "more risky" than a rising market, and hence, it is generally true that the volatility in a declining market is higher.

Contrarians believe high volatility is bullish because it lags market trends.

A standard measure of the overall volatility of the market is the CBOE Volatility Index, VIX, which CBOE introduced in 1993 as a weighted average of the Black-Scholes-implied volatilities of the S&P 100 Index from at-the-money near-term call and put options. ("At-the-money" is defined as the strike price with the smallest difference between the call price and the put price.)

In 2004, futures on the VIX began trading on the CBOE Futures Exchange (CFE), and in 2006, CBOE listed European-style calls and puts on the VIX.

Another measure of the overall market volatility is the CBOE Nasdaq Volatility Index, VXN, which CBOE computes from the Nasdaq-100 Index, NDX, similarly to the VIX. (Note that the more widely-watched Nasdaq Index is the Composite, IXIC.)

The VIX initially was computed from the Black-Scholes formula. Now the empirical prices are used to fit an implied probability distribution, from which an implied volatility is computed. In 2006, CBOE changed the way the VIX is computed. It is now based on the volatilities of the S&P 500 Index implied by several call and put options, not just those at the money, and it uses near-term and next-term options (where "near-term" is the earliest expiry more than 8 days away).

The CBOE in computing the VIX uses the prices of calls with strikes above the current price of the underlying, starting with the first out-of-the money call and sequentially including all with higher strikes until two consecutive such calls have no bids. It uses the prices of puts with strikes below the current price of the underlying in a similar manner.

The price of an option is the "mid-quote" price, i.e. the average of the bid and ask prices.

Let $K_1 = K_2 < K_3 < \cdots < K_{n-1} < K_n = K_{n+1}$ be the strike prices of the options that are to be used.

The VIX is defined as $100 \times \sigma$, where

$$\sigma^2 = \frac{2\mathrm{e}^{rT}}{T}\left(\sum_{i=2;i\neq j}^{n} \frac{\Delta K_i}{K_i^2} Q(K_i) + \frac{\Delta K_j}{K_j^2}\left(Q(K_j\text{put}) + Q(K_j\text{call})\right)/2\right)$$
$$-\frac{1}{T}\left(\frac{F}{K_j} - 1\right)^2,$$

where T is the time to expiry (in our usual notation, we would use $T-t$, but we can let $t=0$), F, called the "forward index level", is the at-the-money strike plus e^{rT} times the difference in the call and put prices for that strike, K_i is the strike price of the i^{th} out-of-the-money strike price (that is, of a put if $K_i < F$ and of a call if $F < K_i$), $\Delta K_i = (K_{i+1} - K_{i-1})/2$, $Q(K_i)$ is the mid-quote price of the option, r is the risk-free interest rate, and K_j is the largest strike price less than F. Hentschel (2003), discusses different methods of computing an implied volatility estimation, and compares some with the VIX computations prior to 2003 (when the VIX was based on Black-Scholes).

Figure 5 shows the VIX for the period January 1, 1990, to December 31, 2007, together with the absolute value of the lograte returns of the S&P 500, on a different scale. The peaks of the two measures correspond, and in general, the VIX and the lograte returns of the S&P 500 seem to have similar distributions.

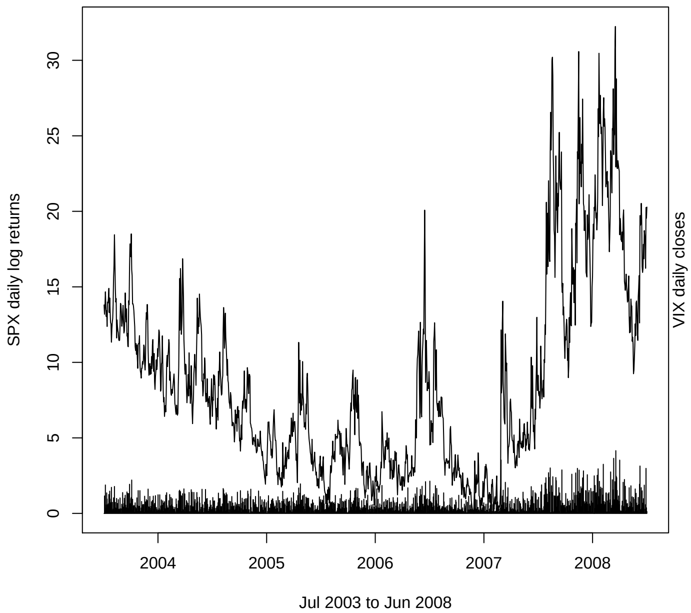

Fig. 5. S&P 500 Daily Lograte Returns and VIX Daily Closes.

3 Monte Carlo Methods

In Monte Carlo methods, we simulate observations from the model, and then perform operations on those pseudo-observations in order to make inferences on the population being studied. Monte Carlo methods require a source of random numbers and methods to transform a basic sequence of random numbers to a sequence that simulates any particular distribution. Gentle (2003) describe general Monte Carlo methods and Glasserman (2004) and Jäckel(2002) describe specific Monte Carlo methods for applications in the analysis of financial data.

3.1 Monte Carlo Estimation

In its simplest form, Monte Carlo simulation is the evaluation of a definite integral

$$\theta = \int_D f(x)\,\mathrm{d}x \tag{2}$$

by identifying a random variable Y with support on D and density $p(y)$ and a function g such that the expected value of $g(Y)$ is θ:

$$\begin{aligned}
\mathrm{E}(g(Y)) &= \int_D g(y)p(y)\,\mathrm{d}y \\
&= \int_D f(y)\,\mathrm{d}y \\
&= \theta.
\end{aligned}$$

Let us first consider the case in which D is the interval $[a, b]$, Y is taken to be a random variable with a uniform density over $[a, b]$, and g is taken to be f. In this case,

$$\theta = (b - a)\mathrm{E}(f(Y)).$$

The problem of evaluating the integral becomes the familiar statistical problem of estimating a mean, $\mathrm{E}(f(Y))$.

The statistician quite naturally takes a random sample and uses the sample mean. For a sample of size m, an estimate of θ is

$$\widehat{\theta} = (b - a)\,\frac{\sum_{i=1}^{m} f(y_i)}{m}, \tag{3}$$

where the y_i are values of a random sample from a uniform distribution over (a, b). The estimate is unbiased:

$$\begin{aligned}
\mathrm{E}(\widehat{\theta}) &= (b - a)\,\frac{\sum \mathrm{E}(f(Y_i))}{m} \\
&= (b - a)\mathrm{E}(f(Y)) \\
&= \int_a^b f(x)\,\mathrm{d}x.
\end{aligned}$$

The variance is

$$\begin{aligned}
\mathrm{V}(\widehat{\theta}) &= (b-a)^2 \, \frac{\sum \mathrm{V}(f(Y_i))}{m^2} \\
&= \frac{(b-a)^2}{m} \, \mathrm{V}(f(Y)) \\
&= \frac{(b-a)}{m} \int_a^b \left(f(x) - \int_a^b f(t)\,\mathrm{d}t \right)^2 \mathrm{d}x. \qquad (4)
\end{aligned}$$

The integral in equation (4) is a measure of the *roughness* of the function. (There are various ways of defining roughness. Most definitions involve derivatives. The more derivatives that exist, the less rough the function. Other definitions, such as the one here, are based on a norm of a function. The L_2 norm of the difference of the function from its integrated value is a very natural measure of roughness of the function. Another measure is just the L_2 norm of the function itself, which, of course, is not translation-invariant.) The method of estimating an integral just described is sometimes called "crude Monte Carlo".

Suppose that the original integral can be written as

$$\begin{aligned}
\theta &= \int_D f(x)\,\mathrm{d}x \\
&= \int_D g(x)p(x)\,\mathrm{d}x, \qquad (5)
\end{aligned}$$

where $p(x)$ is a probability density over D. As with the uniform example considered earlier, it may require some scaling to get the density to be over the interval D. (In the uniform case, $D = (a, b)$, both a and b must be finite, and $p(x) = 1/(b-a)$.) Now, suppose that we can generate m random variates y_i from the distribution with density p. Then, our estimate of θ is just

$$\widehat{\theta} = \frac{\sum g(y_i)}{m}. \qquad (6)$$

Compare this estimator with the estimator in equation (3).

The use of a probability density as a weighting function allows us to apply the Monte Carlo method to improper integrals (that is, integrals with infinite ranges of integration. The first thing to note, therefore, is that the estimator (6) applies to integrals over general domains, while the estimator (3) applies only to integrals over finite intervals. Another important difference is that the variance of the estimator in equation (6) is likely to be smaller than that of the estimator in equation (3).

Quadrature is an important topic in numerical analysis, and a number of quadrature methods are available. They are generally classified as Newton–Cotes methods, extrapolation or Romberg methods, and Gaussian quadrature. These methods involve various approximations to the integrand over various subdomains of the range of integration. The use of these methods involves consideration of error bounds, which are often stated in terms of some function of a derivative of the integrand evaluated at some unknown point. Monte Carlo quadrature differs from these numerical methods in a fundamental way: Monte Carlo methods involve random (or pseudorandom) sampling. The expressions in the Monte Carlo quadrature formulas do not involve any approximations, so questions of bounds of the error of approximation do not arise. Instead of error bounds or order of the error as some function of the integrand, we use the variance of the random estimator to indicate the extent of the uncertainty in the solution.

The square root of the variance (that is, the standard deviation of the estimator) is a good measure of the range within which different realizations of the estimator of the integral may fall. Under certain assumptions, using the standard deviation of the estimator, we can define statistical "confidence intervals" for the true value of the integral θ. Loosely speaking, a confidence interval is an interval about an estimator $\widehat{\theta}$ that in repeated sampling would include the true value θ a specified portion of the time. (The specified portion is the "level" of the confidence interval and is often chosen to be 90% or 95%. Obviously, all other things being equal, the higher the level of confidence, the wider the interval must be.)

Because of the dependence of the confidence interval on the standard deviation, the standard deviation is sometimes called a "probabilistic error bound". The word "bound" is misused here, of course, but in any event, the standard deviation does provide some measure of a sampling "error".

The important thing to note from equation (4) is the order of error in terms of the Monte Carlo sample size; it is $\mathrm{O}(m^{-\frac{1}{2}})$. This results in the usual diminished returns of ordinary statistical estimators; to halve the error, the sample size must be quadrupled.

We should be aware of a very important aspect of a discussion of error bounds for the Monte Carlo estimators. It applies to random numbers. The pseudorandom numbers that we actually use only simulate the random numbers, so "unbiasedness" and "variance" must be interpreted carefully.

3.2 Application: Pricing Options

Important questions in finance concern the fair price for various assets. Although there are many fundamental issues that must be considered in order to determine fair prices for some types of assets, there are other assets, called "derivatives", for which prices depend primarily on the prices and price movement characteristics of other assets (the underlying assets or the "underlying") but are conditionally independent of other fundamental economic conditions. Derivatives are financial instruments having values dependent on constraints on their trading (their "exercise") and on the price of other assets (the underlying) or on some measure of the state of the economy or nature.

A derivative is an agreement between two sides: a *long position* and a *short position*.

Many pricing problems in finance cannot be solved analytically. The traditional approach has been to develop overly simplified models that approximate what is believed to be a more realistic description of market behavior. Pricing of various derivative instruments is an area in finance in which Monte Carlo methods can be used to analyze more realistic models.

Pricing Forward Contracts One of the simplest kinds of derivative is a *forward contract*, which is an agreement to buy or sell an asset at a specified time at a specified price. The agreement to buy is a long position, and the agreement to sell is a short position. The agreed upon price is the *delivery price*. The consummation of the agreement is an *execution*.

Forward contracts are relatively simple to price, and their analysis helps in developing pricing methods for other derivatives. Let k be the delivery price or "strike price" at the settlement time t_s, and let X_t be the value of the underlying. (We consider all times to be measured in years.) A graph of the profit or the payoff at the settlement date is shown in Figure 6 for both long and short positions.

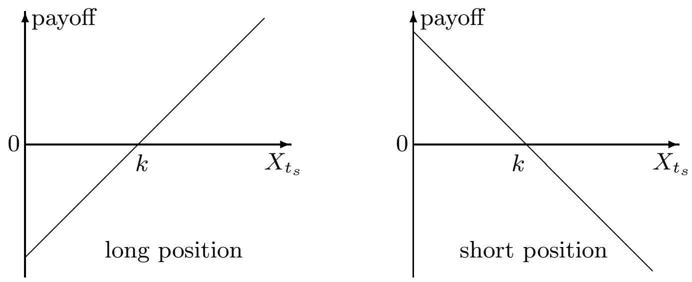

Fig. 6. Payoff of a Forward Contract.

Principles and Procedures for Pricing Derivatives A basic assumption in pricing financial assets is that there exists a fixed rate of return on some available asset that is constant and "guaranteed"; that is, there exists an asset that can be purchased, and the value of that asset changes at a fixed, or "riskless", rate. The concept of "risk-free rate of return" is a financial abstraction based on the assumption that there is some absolute controller of funds that will pay a fixed rate of interest over an indefinite length of time. In current world markets, the rate of interest on a certain financial instrument issued by the United States Treasury is used as this value.

Given a riskless positive rate of return, another key assumption in economic analyses is that there are no opportunities for arbitrage. An *arbitrage* is a trading strategy with a guaranteed rate of return that exceeds the riskless rate of return. In financial analysis, we assume that arbitrages do not exist. This is the "no-arbitrage principle". This means, basically, that all markets have two liquid sides and that there are market participants who can establish positions on either side of the market.

The assumption of a riskless rate of return leads to the development of a *replicating portfolio*, having fluctuations in total valuation that can match the expected rate of return of any asset. The replication approach to pricing derivatives is to determine a portfolio and an associated trading strategy that will provide a payout that is identical to that of the underlying. This portfolio and trading strategy *replicates* the derivative.

Consider a forward contract that obligates one to pay k at t_s for the underlying. The value of the contract at expiry is $X_{t_s} - k$, but of course we do not know X_{t_s}. If we have a riskless (annual) rate of return r, we can use the no-arbitrage principle to determine the correct price of the contract.

To apply the no-arbitrage principle, consider the following strategy:

- take a long position in the forward contract;
- take a short position of the same amount in the underlying (sell the underlying short).

With this strategy, the investor immediately receives X_t for the short sale of the underlying. At the settlement time t_s, this amount can be guaranteed to be

$$X_t \mathrm{e}^{r(t_s - t)} \tag{7}$$

using the risk-free rate of return. Now, if the settlement price k is such that

$$k < X_t \mathrm{e}^{r(t_s - t)}, \tag{8}$$

a long position in the forward contract and a short position in the underlying is an arbitrage, so by the no-arbitrage principle, this is not possible. Conversely, if

$$k > X_t \mathrm{e}^{r(t_s - t)}, \tag{9}$$

a short position in the forward contract and a long position in the underlying is an arbitrage, and again, by the no-arbitrage principle, this is not possible. Therefore, under the no-arbitrage assumption, the correct value of the forward contract, or its "fair price", at time t is $X_t \mathrm{e}^{r(t_s - t)}$.

There are several modifications to the basic forward contract that involve different types of underlying, differences in when the agreements can be executed, and differences in the nature of the agreement: whether it conveys a right (that is, a "contingent claim") or an obligation.

The common types of derivatives include stock options, index options, commodity futures, and rate futures. Stock options are used by individual investors and by investment companies for leverage, hedging, and income. Index options are used by individual investors and by investment companies for hedging and speculative income. Commodity futures are used by individual investors for speculative income, by investment companies for income, and by producers and traders for hedging. Rate futures are used by individual investors for speculative income, by investment companies for income and hedging, and by traders for hedging.

Stock Options A single call option on a stock conveys to the owner the right to buy (usually) 100 shares of the underlying stock at a fixed price, the strike price, anytime before the expiration date. A single put option conveys to the owner the right to sell (usually) 100 shares of the underlying stock at a fixed price, the strike price, anytime before the expiration date. Most real-world stock options can be exercised at any time (during trading hours) prior to expiration. Such options that can be exercised at any time are called "American options". We will consider a modification, the "European option", which can only be exercised at a specified time. There are some European options that are actually traded, but they are generally for large amounts, and they are rarely traded by individuals. European options are studied because the analysis of their fair price is easier.

A major difference between stock options and forward contracts is that stock options depend on the fluctuating (and unpredictable) prices of the underlying.

Another important difference between stock options and forward contracts is that stock options are *rights*, not obligations. The payoff therefore cannot be negative. Because the payoff cannot be negative, there must be a cost to obtain a stock option. The profit is the difference between the payoff and the price paid.

Pricing of Stock Options It is difficult to determine the appropriate price of stock options because stocks are *risky assets*; that is, they are assets whose prices vary randomly (or at least unpredictably).

Pricing formulas for stock options can be developed from a simple model of the market that assumes two types of assets: the risky asset (that is, the stock) with price at time t of X_t and a *riskless asset* with price at time t of β_t. The price of a stock option can be analyzed based on trading strategies involving these two assets, as we briefly outline below. (See Hull (2005) for a much more extensive discussion.)

The price of the riskless asset, like the price of a forward contract, follows the deterministic ordinary differential equation

$$\mathrm{d}\beta_t = r\beta_t \mathrm{d}t, \tag{10}$$

where r is the instantaneous risk-free interest rate.

A useful model for the price of the risky asset is the stochastic differential equation

$$\mathrm{d}X_t = \mu X_t \mathrm{d}t + \sigma X_t \mathrm{d}B_t, \tag{11}$$

where μ is a constant mean rate of return, σ is a positive constant, called the "volatility", and B_t is a Brownian motion; that is,

1. the change ΔB_t during the small time interval Δt is
$$\Delta B_t = Z\sqrt{\Delta t},$$
where Z is a random variable with distribution $\mathrm{N}(0,1)$;
2. ΔB_{t_1} and ΔB_{t_2} are independent for $t_1 \neq t_2$.

Notice, therefore, that for $0 < t_1 < t_2$, $B_{t_2} - B_{t_1}$ has a $\mathrm{N}(0, t_2 - t_1)$ distribution. Notice also that the change in the time interval Δt is randomly proportional to $\sqrt{\Delta t}$. By convention, we set $B_0 = 0$, so B_{t_2} has a $\mathrm{N}(0, t_2)$ distribution. A process following equation (11) is a special case of an *Ornstein-Uhlenbeck process.*

Given a starting stock price, X_0, the differential equation (11) specifies a random *path* of stock prices. Any realization of X_0, $x(0)$, and any realization of B_t at $t \in [0, t_1]$, $b(t)$, yields a realized path, $x(t)$.

We should note three simplifying assumptions in this model:

- μ is constant;
- σ is constant;
- the stochastic component is a Brownian motion (that is, i.i.d. normal).

The instantaneous rate of return from equation (11) is
$$\frac{\mathrm{d}X_t}{X_t} = \mu \mathrm{d}t + \sigma \mathrm{d}B_t, \tag{12}$$
so under the assumptions of the model (11), the price of the stock itself follows a lognormal distribution.

A discrete-time version of the change in the stock price that corresponds to the stochastic differential equation (11) is the stochastic difference equation
$$\begin{aligned} \Delta X_t &= \mu X_t \Delta t + \sigma X_t \Delta B_t \\ &= \mu X_t \Delta t + \sigma X_t Z \sqrt{\Delta t}, \end{aligned} \tag{13}$$
and a discrete-time version of the rate of return is
$$R_t(\Delta t) = \mu \Delta t + \sigma Z \sqrt{\Delta t},$$
where, as before, Z is a random variable with distribution $\mathrm{N}(0,1)$, so $R_t(\Delta t)$ is $\mathrm{N}(\mu\Delta t,\ \sigma^2\Delta t)$. The quantity $\mu\Delta t$ is the expected value of the return in the time period Δt and by assumption is constant. The quantity $\sigma Z\sqrt{\Delta t}$ is the stochastic component of the return, where, by assumption, σ, the "volatility" or the "risk", is constant.

Figure 7 shows 100 simulated paths of the price of a stock for one year using the model (13) with $x(0) = 20$, $\Delta t = 0.01$, $\mu = 0.1$, and $\sigma = 0.2$.

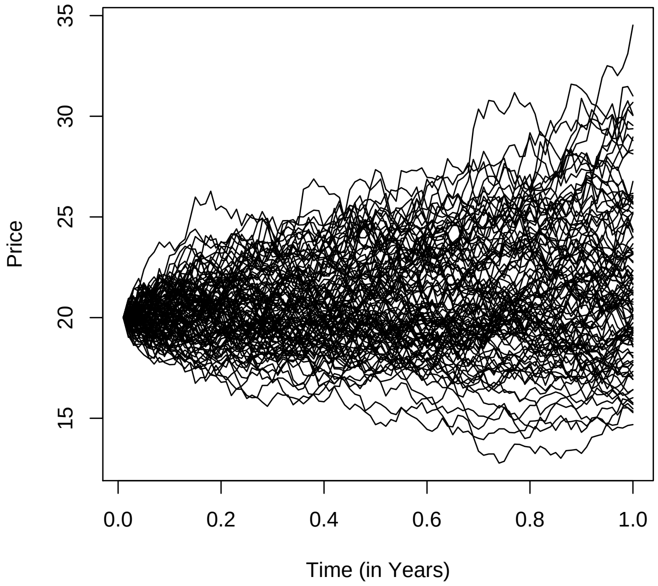

Fig. 7. 100 Simulated Paths of the Price of a Stock with $\mu = 0.1$ and $\sigma = 0.2$.

We first consider the fair price of the option at the time of its expiration. The *payoff*, h, of the option at time t_s is either 0 or the excess of the price of the underlying X_{t_s} over the strike price k. Once the parameters k and t_s are fixed, h is a function of the realized value of the random variable X_{t_s}:

$$h(x_{t_s}) = \begin{cases} x_{t_s} - k & \text{if } x_{t_s} > k, \\ 0 & \text{otherwise.} \end{cases}$$

The price of the option at any time t prior to expiration t_s is a function of t and the price of the underlying x. We denote the price as $P(t, x)$. With full generality, we can set the time of interest as $t = 0$.

The price of the European call option should be the expected value of the payoff of the option at expiration discounted back to $t = 0$,

$$P(0, x) = \mathrm{e}^{-qt_s}\mathrm{E}(h(X_{t_s})), \tag{14}$$

where q is the rate of growth of an asset. Likewise, for an American option, we could maximize the expected value over all stopping times, $0 < \tau < t_s$:

$$P(0, x) = \sup_{\tau \leq t_s} \mathrm{e}^{-q\tau}\mathrm{E}(h(X_\tau)). \tag{15}$$

The problem with expressions (14) and (15), however, is the choice of the rate of growth q. If the rate of growth r of the riskless asset in equation (10) is different from the mean rate of growth μ in equation (11), then there is an opportunity for arbitrage. We must therefore consider a completely different approach. Consideration of a "replicating strategy" leads us to a fair price for options under the assumptions of the model in equation (11) and consistent with a no-arbitrage assumption.

Replicating Strategies A replicating strategy involves both long and short positions that together match the price fluctuations in the underlying, and thus in the fair price of the derivative. We will generally assume that every derivative can be replicated by positions in the underlying and a risk-free asset. (In that case, the economy or market is said to be *complete*.) We assume a finite universe of assets, all priced consistently with some pricing unit. (In general, we call the price, or the pricing unit, a *numeraire*. A more careful development of this concept rests on the idea of a *pricing kernel*.) The set of positions, both long and short, is called a *portfolio*.

The value of a derivative changes in time and as a function of the value of the underlying. Therefore, a replicating portfolio must be changing in time or "dynamic". In analyses with replicating portfolios, transaction costs are ignored. Also, the replicating portfolio must be self-financing; that is, once the portfolio is initiated, no further capital is required. Every purchase is financed by a sale.

Now, using our simple market model, with a *riskless asset* with price at time t of β_t and a *risky asset* with price at time t of X_t (with the usual assumptions on the prices of these assets), we can construct a portfolio with a value that will be the payoff of a European call option on the risky asset at time T.

At time t, the portfolio consists of a_t units of the risky asset and b_t units of the riskless asset. Therefore, the value of the portfolio is $a_tX_t + b_t\beta_t$. If we scale β_t so that $\beta_0 = 1$ and adjust b_t accordingly, the expression simplifies so that $\beta_t = \mathrm{e}^{rt}$.

The portfolio replicates the value of the option at time t_s if almost surely

$$a_{t_s}X_{t_s} + b_{t_s}\mathrm{e}^{rt_s} = h(X_{t_s}). \tag{16}$$

The portfolio is self-financing if at any time t

$$\mathrm{d}(a_tX_t + b_t\mathrm{e}^{rt}) = a_t\mathrm{d}X_t + rb_t\mathrm{e}^{rt}\mathrm{d}t. \tag{17}$$

The Black–Scholes Differential Equation Consider the price P of a European call option at time $t < T$. At any time, this is a function of both t and the price of the underlying X_t. We would like to construct a dynamic, self-financing portfolio (a_t, b_t) that will replicate the derivative at maturity. If we can, then the no-arbitrage principle requires that

$$a_tX_t + b_t\mathrm{e}^{rt} = P(t, X_t) \tag{18}$$

for $t < t_s$.

Now, differentiate both sides of this equation. If a_t is constant, the differential of the left-hand side is the left-hand side of equation (17), which must be satisfied by a self-financing portfolio. (*The assumption that a_t is constant is not correct, but the approximation does not seem to have serious consequences.*)

The derivative of $P(t, X_t)$ is rather complicated because of its dependence on X_t and the fact that $\mathrm{d}X_t$ has components of both $\mathrm{d}t$ and $\mathrm{d}B_t$ in the stochastic differential equation (11). If $P(t, X_t)$ is continuously twice-differentiable, we can use Itô's formula (see Øksendal, 1998, for example) to develop the expression for the differential of the right-hand side of equation (18),

$$\mathrm{d}P(t, X_t) = \left(\frac{\partial P}{\partial t} + \frac{\partial P}{\partial X_t}\mu X_t + \frac{1}{2}\frac{\partial^2 P}{\partial X_t^2}\sigma^2 X_t^2 \right) \mathrm{d}t + \frac{\partial P}{\partial X_t}(\sigma X_t)\mathrm{d}B_t. \qquad (19)$$

By inserting the market model (11) for $\mathrm{d}X_t$ into the differential of the left-hand side, we have

$$(a_t\mu X_t + r b_t \mathrm{e}^{rt})\mathrm{d}t + a_t \sigma X_t \mathrm{d}B_t.$$

Now, equating the coefficients of $\mathrm{d}B_t$, we have

$$a_t = \frac{\partial P}{\partial X_t}.$$

From our equation for the replicating portfolio, we have

$$b_t = (P(t, X_t) - a_t X_t)\, \mathrm{e}^{-rt}.$$

Finally, equating coefficients of $\mathrm{d}t$ and substituting for a_t and b_t, we have the Black–Scholes differential equation,

$$r\left(P - X_t \frac{\partial P}{\partial X_t}\right) = \frac{\partial P}{\partial t} + \frac{1}{2}\sigma^2 X_t^2 \frac{\partial^2 P}{\partial X_t^2}. \qquad (20)$$

Instead of European calls, we can consider European puts and proceed in the same way. We arrive at the same Black–Scholes differential equation (written in a slightly different but equivalent form from the equation above):

$$\frac{\partial P}{\partial t} + r X_t \frac{\partial P}{\partial X_t} + \frac{1}{2}\sigma^2 X_t^2 \frac{\partial^2 P}{\partial X_t^2} = rP. \qquad (21)$$

It is interesting to note that μ is not in the equations, but r has effectively replaced it.

A similar development could be used for American options. Other approaches can be used to develop these equations. One method is called "delta hedging" (see Hull (2005)).

The Black–Scholes Formula The solution of the differential equations depends on the boundary conditions. These come from the price of the option at expiration. In the case of European options, these are simple. For calls, they are

$$P_\mathrm{c}(t_s, X_{t_s}) = \max(X_{t_s} - k, 0),$$

and, for puts, they are

$$P_\mathrm{p}(t_s, X_{t_s}) = \max(k - X_{t_s}, 0),$$

where k is the strike price in either case. With these boundary conditions, the Black–Scholes differential equations have closed-form solutions. For the call, the solution is the "Black–Scholes formula",

$$P_c(t, X_t) = X_t \Phi(d_1) - k\mathrm{e}^{-r(t_s - t)}\Phi(d_2), \tag{22}$$

and, for the put, it is

$$P_p(t, X_t) = k\mathrm{e}^{-r(t_s - t)}\Phi(-d_2) - X_t\Phi(-d_1), \tag{23}$$

where $\Phi(\cdot)$ is the standard normal CDF,

$$d_1 = \frac{\log(X_t/k) + (r + \frac{1}{2}\sigma^2)(t_s - t)}{\sigma\sqrt{t_s - t}},$$

and

$$d_2 = d_1 - \sigma\sqrt{t_s - t}.$$

The Black–Scholes formulas are widely used in determining a fair price for options on risky assets. In practice, of course, σ^2 is not known. The standard procedure is to use price data over some fixed time interval, perhaps a year, compute the sample variance of the rates of return during some fixed-length subintervals, perhaps subintervals of length one day, and use the sample variance as an estimate of σ^2. A data analyst who has looked at such data will see the effects of the rather arbitrary choice of the fixed times.

The prices have systematic relationships to the prices of the underlying, as shown in Figure 8.

More Realistic Models As we have seen, several simplifying assumptions were made in the development of the Black–Scholes formulas. As in most financial analyses, we have assumed throughout that there are no costs for making transactions. In trades involving derivatives, the transaction costs (commissions) can be quite high.

One of the most troubling assumptions is that σ is constant. Because real data on stock prices, X_t, and corresponding option prices, $P(t, X_t)$, are available, the unknown value of σ can be empirically determined from equations (20) and (21) for any given value of t and any fixed value of the strike

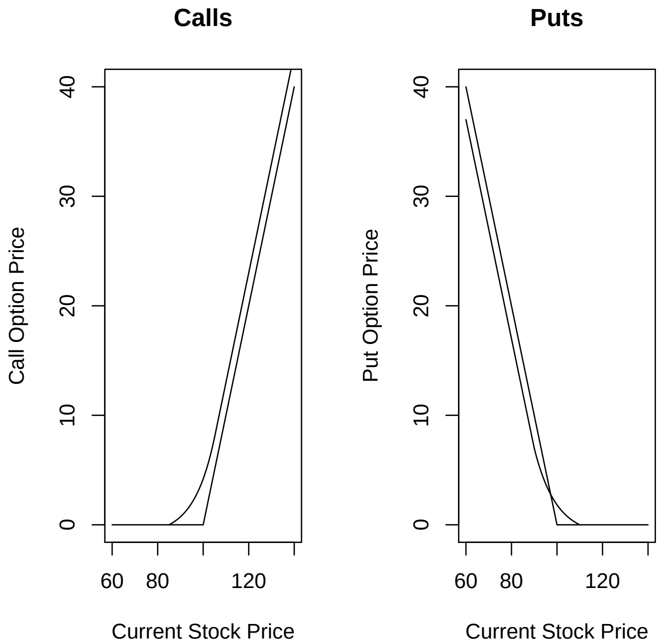

Fig. 8. The Black–Scholes Call Pricing Function.

price, k. (This is called the "implied volatility".) It turns out that for different values of t and $X_t - k$, the implied volatility is different. (Because the implied volatility increases more or less smoothly as $|X_t + d - k|$ increases, where d is some positive number, the relationship is called the "volatility smile".)

One approach is to modify equation (11) to allow for nonconstant σ and augment the model by a second stochastic differential equation for changes in σ. There are various ways this can be done. Fouque, Papanicolaou, and Sircar (2000) describe a simple model in which the volatility is a function of a separate mean-reverting Ornstein-Uhlenbeck process:

$$\begin{aligned} \mathrm{d}X_t &= \mu X_t \mathrm{d}t + \sigma_t X_t \mathrm{d}B_t, \\ \mathrm{d}\sigma_t^2 &= f(Y_t), \\ \mathrm{d}Y_t &= \alpha(\mu_Y - Y_t)\mathrm{d}t + \beta \mathrm{d}\tilde{B}_t, \end{aligned} \tag{24}$$

where α and β are constants and $\tilde{B}_t$ is a linear combination of B_t and an independent Brownian motion. The function f can incorporate various degrees of complexity, including the simple identity function. These coupled equations provide a better match for observed stock and option prices. Economists refer to the condition of nonconstant volatility as "stochastic volatility". Pricing in the presence of stochastic volatility is discussed extensively by Fouque, Papanicolaou, and Sircar (2000). The fact that the implied volatility is not constant for a given stock does not mean that a model with an assumption of constant volatility cannot be useful. It only implies that there are some aspects of the model (11) that do not correspond to observational reality.

Another very questionable assumption in the model given by equation (11) is that the changes in stock prices follow an i.i.d. normal distribution.

There are several other simplifications, such as the restriction to European options, the assumption that the stocks do not pay dividends, the assumption that the derivative of the left-hand side of equation (18) can be done as if a_t were constant, and so on. All of these assumptions allow the derivation of a closed-form solution.

More realistic models can be studied by Monte Carlo methods, and this is currently a fruitful area of research. Paths of prices of the underlying can be simulated using a model similar to equation (13), as we did to produce Figure 7, but with different distributions on Z. A very realistic modification of the model is to assume that Z has a superimposed jump or shock on its $\mathrm{N}(0,1)$ distribution. The simulated paths of the price of the underlying provide a basis for determining a fair price for the options. This price is just the break-even value discounted back in time by the risk-free rate r. Thompson (2000), Barndorff-Nielsen and Shephard (2001), and Jäckel (2002) all discuss use of Monte Carlo simulation in pricing options under various models. Thompson uses a model in which ΔX_t is subjected to a fixed relative decrease as a Poisson event. Such "bear jumps", of course, would decrease the fair price of a call option from the Black–Scholes price and would increase the fair price of a put option. Other modifications to the underlying distribution of ΔX_t result in other differences in the fair price of options. Barndorff-Nielsen and Shephard use a nonnormal process in an Ornstein-Uhlenbeck type of model.

If the innovations in the time series model are not i.i.d., various other models may be more appropriate. One approach that seems to work well in some situations is the generalized autoregressive conditional heteroscedastic (GARCH) model,

$$X_t = E_t \sigma_t, \tag{25}$$

where the E_t are i.i.d. $\mathrm{N}(0,1)$,

$$\sigma_t^2 = \sigma^2 + \beta_1 \sigma_{t-1}^2 + \cdots + \beta_p \sigma_{t-p}^2 + \alpha_1 X_{t-1}^2 + \cdots + \alpha_p X_{t-q}^2,$$

$$\sigma > 0,\ \beta_i \geq 0,\ \alpha_i \geq 0,$$

and

$$\sum_{i=1}^{p} \beta_i + \sum_{j=1}^{q} \alpha_j < 1.$$

A GARCH(1,1) model is often adequate. This model is widely used in modeling stock price movements. The assumption of normality for the E_t can easily be changed to some other distribution. Because normality is usually assumed in a GARCH model, if some other distribution is assumed, the name is often changed. For example, if a Student's t distribution is used, the model is sometimes called TGARCH, and if a stable distribution (other than the normal) is used, the model is called SGARCH.

Another type of approach to the problem of a nonconstant σ in the pricing model (11) is to use a GARCH model (equation (25)). This model is easy to simulate with various distributions.

4 Mining Financial Data

Although financial data has been the subject of intense analysis over the years, there still are many remaining challenges. Some of these challenges are in developing better models. Eraker (2004) and Maheu and McCurdy (2004) discuss some of the issues in developing models that accommodate stochastic volatility better.

4.1 Identifying Patterns in Financial Data

Many traders look for specific patterns in data as indicators of near-term market direction. While most research do not believe that these patterns carry any meaning, in the studies of Lo and MacKinlay (1999) indicate that some patterns do not occur completely at random with respect to market moves. The identification of patterns is an interesting problem in computational statistics.

4.2 Identifying Volatility Clusters

The identification of change points in a time series has received much attention in the literature. Many of the methods, however, have been based on the assumption that the random component in the data has a normal distribution. Such methods of inference have very low power for distributions with heavy tails, such as the rates of return, as we have seen very clearly in the qq plot in Figure 3.

The test of Talwar and Gentle (1981) is less sensitive to the presence of outliers. It clearly identifies changes in volatility in 1987, 1996, 2001, and 2007.

Identification of distributional changes in a random sequence, however, is only part of the picture. The real challenge in mining the data is to determine patterns as they begin, not after the fact. More detailed study of the VIX type of implied volatility and other computed measures of volatility may yield some useful insights.

4.3 Mining Text Data

Perhaps the largest challenge in mining financial data is to make sense of the text data from various sources. The effects of arrival of news on the market in general is discussed by Maheu and McCurdy (2004).

Whether or not internet chat data is based on an objective analysis, it influences the market in ways that are not understood. Antweiler and Frank (2004) attempt to measure the effects.

In order to analyze the effects, and then possibly to predict effects, the first step is to identify sources of relevant data. Many are clear; the recognized stock analysts and the standard financial news services. The weight of individual pieces of data must be estimated and then incorporated into analyses of the effect. A major technical task is to integrate the text data with asset price data, and to determine feedback mechanisms.

Financial data of all kinds are streaming data. Mining of streaming data is particularly difficult. Many problems in analysis arise because of the nonstationarity of the data. (See Marchette and Wegman, 2004, for discussions of some of the problems in a different setting.)

The main thrust of this article has been the need to measure and model volatility of returns. The variability of volatility of stock returns makes planning difficult. This has been one of the most interesting features of recent market activity in which the volatility exhibited extreme swings. Exploratory analyses, both of price data and other types of financial data, including chat data, may identify patterns that presage large volatility. Volatility has implications for portfolio hedging, especially in funds intended for retirement, or actually used for retirement income.

References

ANTWEILER, WERNER, and FRANK, MURRAY Z. (2004): Is all that talk just noise? The information content of internet stock message boards, *The Journal of Finance 59, 1259–1294.*

BARNDORFF-NIELSEN, OLE E., and SHEPHARD, NEIL (2001): Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics (with discussion): *Journal of the Royal Statistical Society, Series B 63, 167–241.*

BJURSELL, JOHAN, and GENTLE, JAMES E. (2008): The contribution of jumps to price volatility in the energy futures market, *2008 JSM Proceedings*, American Statistical Association, Alexandria, VA.

CATALDO, ANTHONY J., and SAVAGE, ARLINE A. (2000): *The January Effect and Other Seasonal Anomalies: A Common Theoretical Framework*, Jai Press, Inc., Stamford, CT.

CHRISS, NEIL A. (1997): *Black-Scholes and Beyond*, Irwin Professional Publishing. Chicago.

CONT, RAMA, and TANKOV, PETER (2004): *Financial Modelling with Jump Processes*, Chapman & Hall/CRC, Boca Raton.

ERAKER, BJORN (2004): Do stock prices and volatility jump? Reconciling evidence from spot and option prices, *The Journal of Finance 59, 1367–1403.*

FOUQUE, JEAN-PIERRE, PAPANICOLAOU, GEORGE, and sIRCAR, k. RONNIE (2000): *Derivatives in Financial Markets with Stochastic Volatility*, Cambridge University Press, Cambridge, UK.

GATHERAL, JIM (2006): *The Volatility Surface*, John Wiley & Sons, New York.

GENTLE, JAMES E. (2002): *Elements of Computational Statistics*, Springer, New York.

GENTLE, JAMES E. (2003): *Random Number Generation and Monte Carlo Methods*, Springer, New York.

GENTLE, JAMES E. (2008): Challenges in Mining Financial Data, In: Hillol Kargupta, Jiawei Han, Philip S. Yu, Rajeev Motwani, and Vipin Kumar (Eds.): *Next Generation of Data Mining*, Chapman & Hall, New York.

GLASSERMAN, PAUL (2004): *Monte Carlo Methods in Financial Engineering*, Springer, New York.

HENTSCHEL, LUDGER (2003): Errors in implied volatility estimation, *The Journal of Financial and Quantitative Analysis 38, 779–810.*

HULL, JOHN C. (2005): *Options, Futures and Other Derivatives* (sixth edition), Prentice-Hall, Englewood Cliffs, New Jersey.

HUNT, P. J., and KENNEDY, J. E. (200): *Financial Derivatives in Theory and Practice*, John Wiley & Sons, New York.

JÄCKEL, PETER (2002): *Monte Carlo Methods in Finance*, John Wiley & Sons Ltd., Chichester.

JONDEAU, ERIC, POON, SER-HUANG, and ROCKINGER, MICHAEL (2007): *Financial Modeling under Non-Gaussian Distributions*, Springer, New York.

KLOEDEN, PETER E., PLATEN, ECKHARD, and SCHURZ, HENRI (1994): *Numerical Solution of SDE through Computer Experiments*, Springer, Heidelberg.

LO, ANDREW W., and MACKINLAY, A. CRAIG (1999): *A Non-Random Walk Down Wall Street*, Princeton University Press, Princeton, NJ.

MARCHETTE, DAVID J., and WEGMAN, EDWARD J. (2004): Statistical analysis of network data for cybersecurity, *Chance 17, 8–18.*

MAHEU, JOHN M., and MCCURDY, THOMAS H. (2004): News arrival, jump dynamics, and volatility components for individual stock returns, *The Journal of Finance 59, 755–793.*

MCQUEEN, GRANT, and VORKINK, KEITH (2004): Whence GARCH? A preference-based explanation for conditional volatility, *The Review of Financial Studies 17, 915–949.*

ØKSENDAL, BERNT (1998): *Stochastic Differential Equations. An Introduction with Applications*, fifth edition, Springer, Heidelberg.

RACHEV, SVETLOZAR T., MENN, CHRISTIAN, and FABOZZI, FRANK J. (2005): *Fat-Tailed and Skewed Asset Return Distributions*, John Wiley & Sons, New York.

SEYDEL, RÜDIGER (2006): *Tools for Computational Finance*, third edition, Springer, Heidelberg.

STENTOFT, LARS (2005): Pricing American options when the underlying asset follows GARCH processes, *Journal of Empirical Finance 12, 576–611.*

TALWAR, PREM P., and GENTLE, JAMES E. (1981): Detecting a scale shift in a random sequence at an unknown time point, *Applied Statistics 30, 301–304.*

TAVELLA, DOMINGO. (2002): *Quantitative Methods in Derivatives Pricing*, John Wiley & Sons, New York.
THOMPSON, JAMES R. (2000): *Simulation: A Modeler's Approach*, John Wiley & Sons, New York.