

SFB 649 Discussion Paper 2010-031

Modeling Asset Prices

James E. Gentle*
Wolfgang Karl Härdle**



* George Mason University, USA

** Humboldt-Universität zu Berlin, Germany

This research was supported by the Deutsche
Forschungsgemeinschaft through the SFB 649 "Economic Risk".

<http://sfb649.wiwi.hu-berlin.de>
ISSN 1860-5664

SFB 649, Humboldt-Universität zu Berlin
Spandauer Straße 1, D-10178 Berlin



SFB 649 ECONOMIC RISK BERLIN

Modeling Asset Prices*

James E. Gentle[†] Wolfgang Karl Härdle[‡]

28th May 2010

As an asset is traded, its varying prices trace out an interesting time series. The price, at least in a general way, reflects some underlying value of the asset. For most basic assets, realistic models of value must involve many variables relating not only to the individual asset, but also to the asset class, the industrial sector(s) of the asset, and both the local economy and the general global economic conditions. Rather than attempting to model the value, we will confine our interest to modeling the price. The underlying assumption is that the price at which an asset trades is a "fair market price" that reflects the actual value of the asset. Our initial interest is in models of the price of a basic asset, that is, not the price of a derivative asset. Usually instead of the price itself, we consider the relative change in price, that is, the rate of return, over some interval of time.

The purpose of asset pricing models is not for prediction of future prices; rather the purpose is to provide a description of the stochastic behavior of prices. Models of price changes have a number of uses, including, for investors, optimal construction of portfolios of assets and, for market regulators, maintaining a fair and orderly market. A major motivation for developing models of price changes of given assets is to use those models to develop models of fair value of derivative assets that depend on the given assets.

Keywords: Discrete time series models, continuous time diffusion models, models with jumps, stochastic volatility, GARCH

JEL classification: C15

*The financial support from the Deutsche Forschungsgemeinschaft via SFB 649 "Ökonomisches Risiko", Humboldt-Universität zu Berlin is gratefully acknowledged.

[†]George Mason University, jgentle@gmu.edu

[‡]Humboldt - Universität zu Berlin, CASE - Center for Applied Statistics and Economics, haerdle@wiwi.hu-berlin.de

As an asset is traded, its varying prices trace out an interesting time series. The price, at least in a general way, reflects some underlying value of the asset. For most basic assets, realistic models of value must involve many variables relating not only to the individual asset, but also to the asset class, the industrial sector(s) of the asset, and both the local economy and the general global economic conditions. Rather than attempting to model the value, we will confine our interest to modeling the price. The underlying assumption is that the price at which an asset trades is a "fair market price" that reflects the actual value of the asset.

Our initial interest is in models of the price of a basic asset, that is, not the price of a derivative asset. Usually instead of the price itself, we consider the relative change in price, that is, the rate of return, over some interval of time.

The purpose of asset pricing models is not for prediction of future prices; rather the purpose is to provide a description of the stochastic behavior of prices. Models of price changes have a number of uses, including, for investors, optimal construction of portfolios of assets and, for market regulators, maintaining a fair and orderly market. A major motivation for developing models of price changes of given assets is to use those models to develop models of fair value of derivative assets that depend on the given assets.

The rate of return has a strong stochastic component, and in this chapter, we describe various stochastic models of the rate of return. We also briefly discuss statistical inference in these models, and applications of these models for pricing derivative assets. Our presentation is quite general. We refer to readily-available literature, some in the present volume, for details on the analysis and applications of the models.

The models we consider in this chapter are for the prices of a single asset, although, of course, that asset may be a portfolio of individual assets. Pricing models of more than one asset must take into account the correlations among their prices. Multivariate pricing models are discussed by Hafner and Manner (2010, this volume).

In most models of asset prices such as those we discuss in Sections 2 through 4, the basic observable components are the prices themselves, and the stochastic components of interest are the changes in asset prices. Such models assume rational and independent traders.

Models of asset prices depend on principles of general economic theory such as equilibrium and arbitrage.

Another approach to modeling asset prices is based on modeling the stochastic aspects in terms of behavior of the traders who collectively determine the asset prices. This agent-based approach allows incorporation of human behavior in the model and so instead of relying solely on classical economic theory, the results of behavioral economics can be included in the model. In the agent-based approach, which we briefly discuss in Section 6, the actions of the agents include a random component and their actions determine the prices.

In discussing models, it is always worthwhile to recall the dictum, generally attributed to George Box, “All models are wrong, but some are useful.” The usefulness of models of asset prices is not because of the opportunity for financial gain, but rather for determining fair prices, for better understanding of market dynamics, and possibly for regulatory policy development.

1 Characteristics of Asset Price Data

Asset prices are directly observable and are readily available from the various markets in which trading occurs. Instead of the prices themselves, however, we are often more interested in various derived data and statistical summaries of the derived data. The most common types of derived data are a first-order measure of change in the asset prices in time, and a second-order measure of the variation of the changes.

The scaled change in the asset price is called the rate of return, which in its simplest form is just the price difference between two time points divided by the price at the first time point, but more often is the difference in the logarithm of the price at the first time point and that at the second time point. The length of the time period of course must be noted. Rates of return are often scaled in some simple way to correspond to an annual rate. In the following, when we refer to “rate of return”, we will generally mean the *log-return*, that is, the difference in the logarithms. This derived measure is one of the basic quantities we seek to model.

The log-return depends on the length of the time interval, and so we may speak of “weekly” log-returns, “daily” returns, and so on. As the time interval becomes very short, say of the order of a few minutes, the behavior of the returns changes in a significant way. We will briefly comment on that high-frequency property in Section 2.5 below.

One of the most important quantities in financial studies is some measure of the variability of the log-returns. The standard deviation of the log-return is called the *volatility*.

A standard deviation is not directly observable, so an important issue in financial modeling is what derived measures of observable data can be used in place of the standard deviation. The sample standard deviation of measured log-returns over some number of time intervals, of course, is an obvious choice. This measure is called *statistical volatility* or *realized volatility*.

Before attempting to develop a model of an empirical process, we should examine data from the process. Any reasonable model must correspond at least to the grossest aspects of the process. In the case of asset prices, there may be various types of empirical processes. We will just focus on one particular index of the price of a set of assets, the S&P 500 Index.

We will examine some empirical data for the S&P 500. First we compute the log-rate for the S&P 500 from January 1, 1990, to December 31, 2005. A histogram for this 15 year period is shown in Figure 1.

With a first glance at the histogram, one may think that the log-returns have a distribution similar to a Gaussian. This belief, however, does not receive affirmation by the q-q plot in Figure 2.

Some may argue, however, that data models based on a normal distribution are often robust, and can accommodate a wide range of distributions that are more-or-less symmetric and unimodal.

One who is somewhat familiar with the performance of the U. S. stock market will recognize that we have been somewhat selective in our choice of time period for examining the log-return of the S&P 500. Let us now look at the period from January 1, 1987, to September 30, 2009. The belief — or hope — that a normal distribution is an adequate model of the stochastic component is quickly dispelled by looking at the q-q plot in Figure 3.

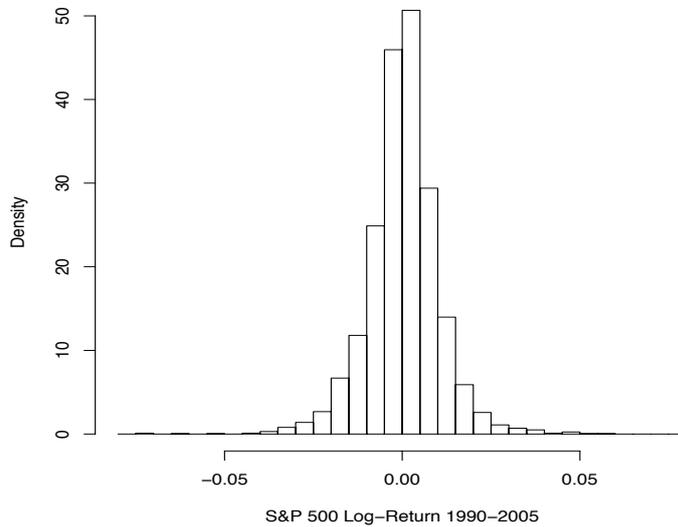


Figure 1: Histogram of Log-Rates of Return 1990 to 2005

Figure 3 indicates that the log-rates of the S&P 500 form a distribution with very heavy tails. We had only seen a milder indication of this in Figures 1 and 2 of the histogram and q-q plots for the 1990 to 2005 period.

The previous graphs have shown only the static properties of the log-return over fixed periods. It is instructive to consider a simple time series plot of the rates of log-returns of the S&P 500 over the same multi-year period, as shown in Figure 4.

Even a cursory glance at the data in Figure 4 indicates the modeling challenges that it presents. We see the few data points with very large absolute values relative to the other data. A visual assessment of the range of the values in the time series gives us a rough measure of the volatility, at least in a relative sense. Figure 4 indicates that the volatility varies over time and that it seems to be relatively high for some periods and relatively low for other periods. The extremely large values of the log-returns seem to occur in close time-proximity to each other.

Of course there are many more ways that we could look at the data in order to develop ideas

for modeling it, but rather than doing that, in the next two sections we will just summarize some of the general characteristics that have been observed. Many of these properties make the data challenging to analyze.

1.1 Stylized Properties of Rates of Return

We have only used a single index of one class of asset prices for illustrations, but the general properties tend to hold to a greater or lesser degree for a wide range of asset classes. From Figures 1 through 4, we can easily observe the following characteristics.

- 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. (Possibly because traders react more strongly to negative information than to positive information.)

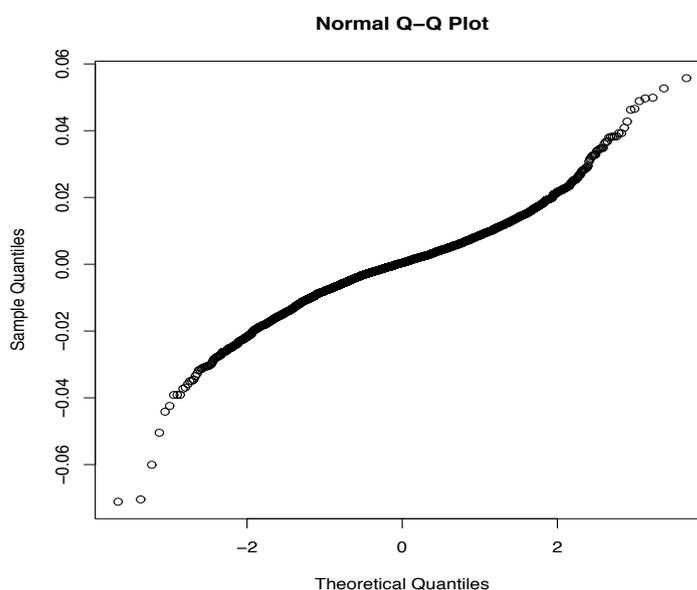


Figure 2: Normal q-q Plot of Log-Rates of Return 1990 to 2005

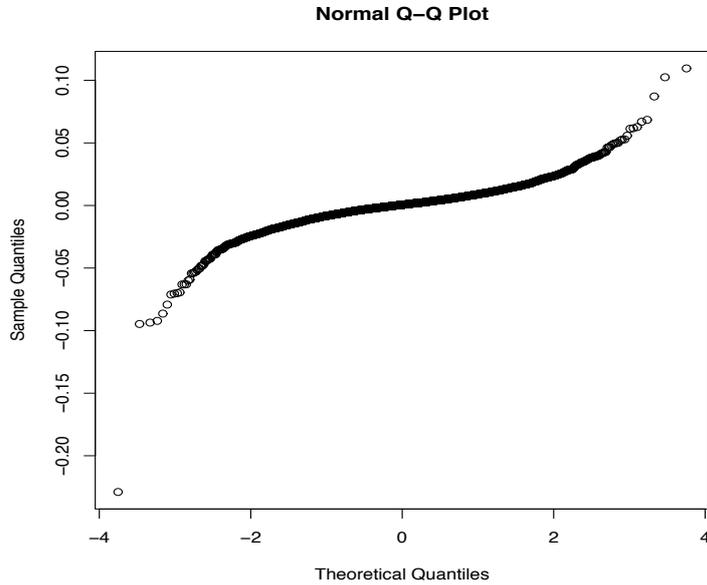


Figure 3: Normal q-q Plot of Log-Rates of Return 1987 to 2009

- Nonconstant volatility. (This is called “stochastic volatility”.)
- Clustering of volatility. (It is serially correlated.)

These characteristics are apparent in our graphical illustrations, but the detection of other properties requires computations of various statistics. There are some characteristics that we could observe by using two other kinds of similar plots. In one approach, we compare rates of return at different frequencies, and in the other, we study lagged data. Lagged data is just an additional form of derived measure, much like rate of return itself is a derived measure, and like rate of return it may also depend on the frequency; that is, the length of the lag. We will not display plots illustrating these properties, but merely list them.

- Asymmetry in lagged correlations.
- Aggregational normality.
- Long range dependence.

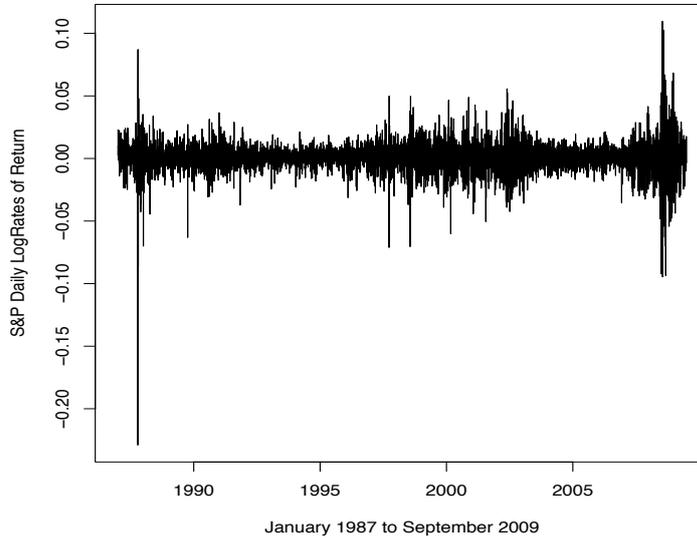


Figure 4: Rates of Return

- Seasonality.
- Dependence of stochastic properties on frequency. Coarse volatility predicts fine volatility better than the other way around.

These stylized properties have been observed through analysis of financial data of various classes over many years. Some of the most interesting of these properties depend on how the volatility changes. We will now note some more properties of the volatility itself.

1.2 Volatility

A standard deviation is defined in terms of a probability model, so defining volatility as the standard deviation of the log-return implies a probability model for the log-return. It is this probability model that is central to more general models of asset prices.

Our preliminary graphical analyses showed that there is a problem with a simple interpreta-

tion of volatility; it is not constant in time. In some cases, it is clear that news events, that is, shocks to financial markets, cause an increase in volatility. In fact, it appears that both “positive” news and “negative” news lead to higher levels of volatility, but negative news tends to increase future volatility more than positive news does. It also appears that there are two distinct components to the effect of news on volatility, one with a rapid decay and one with a slow decay.

Another aspect of volatility, as we mentioned above, is that it is not directly observable, as is the price of an asset or even the change in price of an asset.

The point of this discussion is that the concept of volatility, despite its simple definition, is neither easy to model nor to measure.

Volatility, however, is one of the most important characteristics of financial data, and any useful model of changes in asset prices must include a component representing volatility. Increased volatility, however it is measured, has the practical effect of increasing the risk premium on financial assets.

2 The Basic Models

Asset prices and their rates of change are stochastic processes. We will represent the general form of the stochastic process modeling the asset prices as $\{X_t : t \in \mathcal{I}\}$, for some (well-ordered) index set \mathcal{I} . We assume a general probability space (Ω, \mathcal{F}, P) . The specific form of the stochastic process is determined by the nature of \mathcal{I} and (Ω, \mathcal{F}, P) , and by the stochastic relations between X_t and X_s for $t, s \in \mathcal{I}$ and $s < t$; that is, relations between X_t and the sequence $\{X_s : s \in \mathcal{I}, s < t\}$.

In this section we consider various forms of models of asset prices and of changes in asset prices. We begin with an abstract description. The purpose of this approach is to emphasize that the models used in conventional financial analyses are just particular choices that are made to simplify the analysis.

As we discuss pricing models from simple to more complex, we should bear in mind the

empirical properties discussed in Section 1.1 of the processes we are attempting to model. We will consider various formulations of models to capture various properties, but in the end we see that the models do not fully capture all of those stylized properties.

Systematic Factors and Random Perturbations

Many mathematical models of interesting processes take the form of a systematic component that involves various measurable factors, plus a random component that represents unobservable or non-quantifiable factors and/or truly “random” factors:

$$Y = f(y_s) + E. \tag{1}$$

(Here we are using different notation so as to focus on the abstract model.) The function f may take on a variety of forms. In preliminary models, it is almost always linear. As a model is refined, it may assume more complicated forms. The input y_s may represent directly observable variables or it may represent derived variables such as rates. As models are built or evolve, in addition to changes in the function form of f , the factors included in the input y_s may change. In preliminary models, y_s may include a large number of factors that are of potential interest, and as part of the model-building process, some of these factors are removed from the model. Alternatively, in preliminary models, y_s may include only one or two factors that are believed to be important, and as part of the model-building process, other factors are added to the model.

In many models, the random component E is the most important term in the model. A mathematical model may be very precise in the description of E , for example, the model may state that $E \sim N(0, \sigma^2)$, or the model may be looser, stating only, for example, that the expectation of E is 0, and that in set of E 's, they are exchangeable.

Before we can build models of stochastic processes in time $\{X_t : t \in \mathcal{I}\}$, we must address the nature of the index set \mathcal{I} .

Indexing Time

There are essentially two types of index sets. A “discrete time” index set is countable, and, hence, can be taken as the set of integers. A “continuous time” index can be taken as an interval in the reals. These two ways of treating time lead to two general classes of models.

For discrete time, the models evolve from moving average and autoregressive models. The continuous time models are diffusion processes, possibly in combination with a Poisson process. Although discrete time and continuous time may appear to yield completely different kinds of models, there are various limiting equivalent forms.

For either discrete or continuous time, there are various possibilities for choice of the probability space. A standard approach, of course, is to use a normal distribution, at least as a first approximation, as a model for the stochastic component. The important consideration is the nature of the conditional distribution of X_t given $\{X_s : s \in \mathcal{I}, s < t\}$.

In this chapter we will review the types of models that have been used for changes in asset prices over time. We first describe these briefly, and then indicate some of the ways in which the models are inadequate. Several other papers in this Handbook are concerned with various modifications of these models.

2.1 Discrete Time Series Models

Discrete time series models describe the behavior of a stochastic process in terms of a functional relationship of the general form

$$X_t = f(X_{t-1}, \dots, X_{t-p}, \epsilon_t, \epsilon_{t-1}, \dots, \epsilon_{t-q}). \quad (2)$$

In models of this form, the ϵ_i are generally assumed to be random variables, and so if their effects are additive, this is of the same form as model (1). More specific assumptions about their distribution allow various methods of statistical inference to be used in making refinements to the model. In most models of this form, the function f is linear. We will briefly describe various forms of the model (2). These models are the subject of the well-established field of *time series analysis in the time domain*. We begin with a few definitions.

A *white noise* process $\{\epsilon_t\}$ is one in which for each t , $\epsilon_t \sim N(0, 1)$, that is, it has a Gaussian or normal distribution with mean 0 and variance 1, and for $s \neq t$, $\text{Cov}(\epsilon_s, \epsilon_t) = 0$; that is, ϵ_s and ϵ_t are independent (because of normality).

The most useful forms of the function f in equation (2) are linear. A particularly simple form yields a *linear process*. We say $\{X_t\}$ is a linear process if it has the form

$$X_t = \mu + \sum_{i=-\infty}^{\infty} a_i \epsilon_{t-i}, \quad (3)$$

where $\sum_{i=-\infty}^{\infty} a_i < \infty$ and $\{\epsilon_t\}$ is a white noise process.

One of the most important properties of a stochastic process is *stationarity*, which refers to a distributional measure remaining constant in time. The mean of the linear process is stationary: $E(X_t) = \mu$. The linear process is also *covariance stationary* since

$$\text{Cov}(X_t, X_{t+h}) = \sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} a_i a_j I_{\{|i,j|, |i+j|=h\}}(i, j) = \sum_{i=-\infty}^{\infty} a_i a_{i-h}$$

and $V(\epsilon_t) = 1$. Note that covariance stationary means that the covariance between X_s and X_t depends only on $|t - s|$.

In general, we say that a process is *weakly stationary* (or just *stationary*) if it is mean and covariance stationary.

If the linear model involves only the ϵ_i , that is,

$$X_t = \beta_1 \epsilon_{t-1} + \cdots + \beta_q \epsilon_{t-q} + \epsilon_t, \quad (4)$$

it is called a moving average model with q terms. We refer to this model as $\text{MA}(q)$. Assuming $\{\epsilon_t\}$ is a white noise process, the $\text{MA}(q)$ model is a linear process, and the normality of the stochastic components allows use of relatively simple statistical analyses. For example, we can use maximum likelihood methods, which require specification of probability density functions, and these are particularly straightforward when the stochastic components are normal.

If the linear model involves only the X_{t-j} and ϵ_t , that is,

$$X_t = \alpha_0 + \alpha_1 X_{t-1} + \cdots + \alpha_p X_{t-p} + \epsilon_t, \quad (5)$$

it is called an autoregressive model with p terms. We refer to this model as $\text{AR}(p)$. Again, specific assumptions about the distributions of

$$\dots, \epsilon_{t-2}, \epsilon_{t-1}, \epsilon_t, \epsilon_{t+1}, \epsilon_{t+2}, \dots$$

allow various methods for statistical inference about their distribution and about the parameters α_j .

Combining the $\text{MA}(q)$ model of equation (4) with the $\text{AR}(p)$ model of equation (5), we have the autoregressive moving average model of order p and q , that is, $\text{ARMA}(p, q)$,

$$X_t = \alpha_0 + \alpha_1 X_{t-1} + \dots + \alpha_p X_{t-p} + \beta_1 \epsilon_{t-1} + \dots + \beta_q \epsilon_{t-q} + \epsilon_t. \quad (6)$$

Assumptions about the relative values of the β_j and α_k imply certain interesting properties of the time series.

The usefulness of ARMA models can be greatly extended by applying it to differences of the time series. If the X 's in equation (6) are replaced by d^{th} -order differences, the "integrated" model in the same form as equation (6) is called an $\text{ARIMA}(p, d, q)$ model. The differences allow the model to accommodate seasonal effects.

The simple AR, MA, ARMA, and ARIMA models we have just described can be applied to a time series of prices or to a series of returns. The nature of the series and the assumptions about the stochastic component determine the kind of analyses. For example, given the price process $\{X_y\}$, an $\text{AR}(1)$ model of returns $Y_t = (X_t - X_{t-1})/X_{t-1}$ from equation (5) would have the form of a pure noise,

$$Y_t = \delta_t. \quad (7)$$

The random variable δ_t does not have the same distribution as that of ϵ_t . In fact, if $\{\epsilon_t\}$ is a white noise, then δ_t is a Cauchy process, which has infinite moments of all orders. Clearly, the specific assumptions about the distributions of $\{\epsilon_t\}$ determine the methods for statistical inference about their distribution and about the parameters in the model.

2.2 Continuous Time Diffusion Models

Differential equations are effective models of continuous change of quantities over time. Such models are widely used for expressing diffusion of a substance or of energy over a

physical space. At a macro level the laws governing diffusion are deterministic. Furthermore, substances and energy can be treated as ensembles over a physical space, and so the diffusion model represents an average density. Thus, such a model contains no stochastic component.

Empirical evidence clearly indicates that a deterministic differential equation could not effectively model price movements of an asset such as a stock.

The first step must be to introduce a stochastic component into the differential equation, and the simplest way to do this is for the differential to be from a Brownian motion. This is what Bachelier proposed in 1900 (see, for example, Steele, 2001). In Bachelier's stochastic differential equation, the Brownian motion represented the change in price. This model is

$$dX_t = \mu X_t dt + \sigma X_t dW_t, \quad (8)$$

where W_t is a Brownian motion. Clearly, dW_t could represent some other type of stochastic differential, although the existence of a stochastic differential with appropriate properties would need to be established. (Wiener established this for Brownian motion. See, again for example, Steele, 2001.)

Samuelson (1965) modified the model (8) to one he called geometric Brownian motion:

$$\frac{dX_t}{X_t} = \mu dt + \sigma dW_t. \quad (9)$$

This is a model for the *rate of change* of asset prices. Note that this is similar to forming equation (7) from (5), and then changing the assumptions about the distribution of the random component so that the random variable in the derived equation has a simple distribution.

The geometric Brownian motion model (9) has been widely used in financial analysis. In the context of a riskless portfolio of an asset and an option on the asset, the geometric Brownian motion model leads to the Black-Scholes-Merton differential equation for the fair price P of an option:

$$\frac{\partial P_t}{\partial t} + r X_t \frac{\partial P_t}{\partial X_t} + \frac{1}{2} \sigma^2 X_t^2 \frac{\partial^2 P_t}{\partial X_t^2} = r P, \quad (10)$$

where r is a risk-free interest rate.

Detemple and Rindisbacher (2010, this volume) provide a more extensive discussion of diffusion models. We will briefly consider some modifications of the basic diffusion models in Section 4.

2.3 Accounting for Jumps

Looking at the data in Figure 4, we notice a few extremely large returns, both positive and negative. These outliers are called “jumps”. Figure 2 and 3 indicate that the presence of these outliers is inconsistent with the assumption that the underlying random variables in either model (6) or model (9) have Gaussian distributions. (In model (6) the random variable is ϵ , and in model (9) it is dW_t .)

In standard statistical analyses, there are two simple ways of accounting for outliers. One way is to use an “outlier-generating distribution” or “jump process”, that is, a heavy-tailed distribution, such as stable distribution other than the Gaussian. Figueroa-López (2010, this volume) describes the use of Lévy processes in diffusion models. Other discussions of models with non-Gaussian random components are in Jondeau, Poon, and Rockinger (2007) and Rachev, Menn, and Fabozzi (2005).

Another method of accounting for jumps is to use a mixture of distributions. Even mixtures of Gaussian distributions result in outlier-generating distributions. Instead of using simple mixtures of Gaussians, however, a more common approach is to use a mixture of a continuous distribution, such as a Gaussian, and a discrete Poisson process, possibly associated with an effect with a random magnitude. Bjursell and Gentle (2010, this volume) and Cont and Tankov (2004) describe the use of mixtures that include Poisson processes. We will briefly consider jump-diffusion models in Section 4.2.

Either of these modifications to the models results in more difficult data analyses.

2.4 Accounting for Stochastic Volatility

The ARMA model of equation (6) incorporates the volatility of the stochastic process in the standard deviation of the random variables ϵ , and the diffusion model of equation (9) incorporates the volatility in the standard deviation of the random variables σdW_t . An assumption of either model is that this standard deviation is constant; hence, a serious deficiency of either of the two basic models (6) and (9) is that the model does not account for the stochastic volatility that is apparent in Figure 4.

To be realistic, either type of model must be modified to allow for the volatility to be nonconstant. Further, as we note from Figure 4, the modification must include a serial correlation of the volatility.

2.5 Market MicroStructure

Pricing data represent the value exchanged in a specific trade. The price at which a specific transaction occurs should be exactly the same as the price (within the minimum unit of money) of the same transaction at the same time. It turns out, for a variety of reasons, that this is not the case. *Tick* data, that is, data on each transaction (also called “high-frequency data”) exhibit characteristics that are different from price data collected less frequently, say at the close of each trading day.

Some stylized properties of tick data include intraday periodicity; nonsynchronicity, that is, a sequence of prices over a short time interval do not form an equally-spaced time series; price clustering; and negative lag-1 autocorrelations. These properties constitute what is called “market microstructure”. See Lai and Xing (2008) for more discussion of microstructure.

Bjursell and Gentle (2010, this volume) discuss the use of microstructure noise to test for jumps superimposed on a diffusion model.

3 GARCH-Type Models

The AR, MA, ARMA, and ARIMA models described in Section 2.1 assume a constant variance. There are various ways of modifying the model to make the variance change over time.

For a model of the form (7), we first introduce a scale on the random component:

$$Y_t = \sigma_t \delta_t. \quad (11)$$

Then, following the empirical observation that the standard deviation of a process is proportional to the magnitude (that is, the coefficient of variation is relatively constant), we may assume a model for the variance of the form

$$\sigma_t^2 = \alpha_0 + \alpha_1 Y_{t-1}^2. \quad (12)$$

The variance is conditional on the value of Y_{t-1}^2 , and so this kind of model is called an ARCH (autoregressive conditionally heteroscedastic) model; specifically the model of equations (11) and (12) is called an ARCH(1) model (recall that it originated as an AR(1) model).

The ARCH models can be generalized further by modeling the variance as an AR process; that is, equation (12) may become, for example,

$$\sigma_t^2 = \alpha_0 + \alpha_1 Y_{t-1}^2 + \beta_1 \sigma_{t-1}^2. \quad (13)$$

Such models are called GARCH (generalized autoregressive conditionally heteroscedastic) models; specifically, the model of equations (11) and (13) is a GARCH(1,1) model, because both components are lag 1 processes.

Notice that the simple ARCH(1) model of equations (11) and (12) could be reformulated by squaring both sides of equation (11), then subtracting equation (12) and then rearrange terms to obtain

$$Y_t^2 = \alpha_0 + \alpha_1 Y_{t-1}^2 + \gamma_t, \quad (14)$$

in which, if δ_t is a $N(0, 1)$ random variable, then γ_t is a scaled and shifted chi-squared random variable with one degree of freedom.

The purpose of this re-expression is only to show that the ARCH(1) model is related to an AR(1) model with a change of distribution of the random component. The ARCH and GARCH models, while they do incorporate stochastic volatility, if the underlying distribution of the stochastic component is normal, the models will not display the heavy-tailed and asymmetric returns that are observed empirically.

Many variations of GARCH models have been studied; see, for example, Christoffersen, Jacobs, and Ornathanalai (2010, this volume) and Gouriéroux (1997). Most of these variations are still based on an underlying normal distribution, however.

3.1 GARCH with Jumps

As we mentioned previously, jumps can be modeled either through an outlier-generating distribution or by superimposition of a jump process. The most common way of incorporating jumps in a discrete time series model is by use of a heavy-tailed distribution, such as stable distribution other than the Gaussian. This, of course, presents problems in the statistical analysis of data using such models.

3.2 Inference on the Parameters

Statistical inference on autoregressive moving average models is usually based on the likelihood. Given a distribution for the random components in any such model, it is usually rather simple to formulate the associated likelihood. The likelihood rarely can be maximized analytically, but there are efficient numerical methods. These methods are usually two-stage optimizations, and are similar to methods originally used in the ARIMA models of Box and Jenkins. Gouriéroux (1997) describes maximum likelihood methods for various GARCH models.

Just fitting the parameters, of course, is only one part of the problem of statistical inference. Various assumptions about the distributions of the stochastic components require different methods for statistical inference such as tests and confidence regions. Even if the underlying distribution is not assumed to be normal, most inference methods end up using approximate

normal distributions.

4 Diffusion Models

The basic geometric Brownian motion diffusion model (9),

$$\frac{dX_t}{X_t} = \mu dt + \sigma dW_t,$$

misses most of the salient empirical properties of Section 1.1.

Brownian motion is a rather complex process, and given our understanding of it — and our lack of understanding of a similar process not based on Gaussianity — we would seek to build modifications onto the Brownian motion, rather than to replace the Gaussian distribution with some other distribution that is either heavy-tailed or asymmetric. (Recall our elliptical reference above to the *existence* of Brownian motion.)

There are several possible modifications of the Brownian motion. We will formulate two modifications below that address stochastic volatility and jumps. Before doing so, however, we mention a simple modification that allows for long range dependencies in a model of the form (9). In this modification, instead of the Brownian motion W_t , we use a fractional Brownian motion, W_t^H , where $0 < H < 1$ is the Hurst index. (An index of 0.5 is ordinary Brownian motion.) The essential characteristic of a fractional Brownian motion,

$$\text{Cov}(W_t^H, W_s^H) = \frac{1}{2} (|t|^{2H} + |s|^{2H} - |s - t|^{2H}),$$

allows for the modified model (9) to exhibit long range dependencies, which, as we remarked without elaboration in Section 1.1, is an empirical property of rates of return. Fractional Brownian motion is in spirit related to the reformulation of the ARCH(1) model of equations (11) and (12) as the AR(1) model (14).

4.1 Coupled Diffusion Models

The modification of an AR model that yields a GARCH model is merely to apply to a function of the volatility the same basic time series model that is used for returns. This

way of handling stochastic volatility in the case of diffusion models would result in coupled diffusion models in which a secondary diffusion model is applied to a function of the volatility:

$$\frac{dX_t}{X_t} = \mu dt + \sigma_t d(W_1)_t \quad (15)$$

$$d\sigma_t^2 = \alpha(\mu_{\sigma_t^2} - \sigma_t^2)dt + \beta(\sigma_t^2)^\gamma d(W_2)_t, \quad (16)$$

where α , $\mu_{\sigma_t^2}$, β , and γ are constants and $(W_1)_t$ and $(W_2)_t$ are Brownian motions.

Equations (15) and (16) are sometimes called the Hull and White model (although that term is usually used for a different model used for interest rate derivatives). For the special case of $\gamma = 0.5$, it is also called the Heston model.

There are many variations on models of this form. Notice that this model does not tie the magnitude of the volatility to the magnitude of the return, as the simple ARCH model did. This could be remedied by an incorporation of X into Equation (16). An important consideration is the relationship between the two Brownian motions $(W_1)_t$ and $(W_2)_t$. The simplest assumption is that they are independent. An alternative, but still very simple assumption, is that $(W_2)_t$ is a linear combination of $(W_1)_t$ and an independent Brownian motion.

While the coupled diffusion model do incorporate stochastic volatility, just as with the ARCH and GARCH models, because the underlying distribution of the stochastic component is normal, the models will not display the heavy-tailed and asymmetric returns that are observed empirically.

4.2 Diffusion with Jumps

A modification of any of the models that we have discussed above that can display both heavy-tailed and asymmetric returns is to superimpose a Poisson process onto the model. Starting with the simple geometric Brownian motion diffusion model (9), we write

$$\frac{dX_t}{X_t} = \mu dt + \sigma dW_t + \kappa_t dq_t, \quad (17)$$

where W_t is the standard Wiener process; q_t is a counting process with intensity λ_t , that is, $P(dq_t = 1) = \lambda_t dt$; and κ_t is the size of the price jump at time t if a jump occurred. If X_{t-}

denotes the price immediately prior to the jump at time t , then $\kappa_t = X_t - X_{t-}$.

4.3 Inference on the Parameters

If restrictive assumptions are made about the constancy of parameters and independence of the events in the process, there are fairly simple statistical estimators for most of the parameters in the single-equation models. Parameters in coupled equations can often be estimated using two-stage likelihood methods. The parameters in a model such as equation (17) are difficult to estimate because we do not know which of the two processes is operating. One approach to the fitting the parameters in a model with a superimposed process is to set an arbitrary threshold for the return, and to assume the Poisson process generates any realization greater than that threshold.

For models with time-varying parameters, analysis generally depends on the use of Monte Carlo methods.

5 How Simple Can a Realistic Model Be?

At this point, we must ask how simple can a pricing model be and still capture all of the empirical properties that we have observed. Clearly, the basic models of Section 2 fail drastically.

The first modification to the simple ARMA or geometric Brownian motion model is usually to address the stochastic volatility. An approach in either case is to couple the basic process with a similar process for the volatility. So long as the underlying stochastic components are Gaussian, two-stage maximum likelihood methods can be used in the analysis.

The issue of heavy tails and asymmetric distributions could perhaps be addressed by replacing the Gaussian processes with some asymmetric heavy-tailed process, perhaps a stable process. The loss of the simplicity of the normal distribution, however, is a very steep price to pay. An alternative approach is to superimpose a Poisson jump process, as in model (17). Such a model has a stochastic volatility (due to the firing of the Poisson process), but it

is not the slowly-varying volatility that we observe. Hence, the jump process needs to be superimposed on a model that already accounts for stochastic volatility, such as a GARCH model or a coupled diffusion model.

It is clear that the amount of a jump, κ_t , is not constant. A simple modification would be to take κ_t as an independent random variable. Its distribution would seem to be heavy-tailed and to have a negative mean. Empirically (see Figure 4) a negative (positive) jump tends to be followed immediately by a positive (negative) jump. This may suggest that jumps be modeled as paired events instead of trying to accommodate these positive and negative values in the distribution of κ_t .

A further glance at Figure 4 indicates two additional considerations (assuming a somewhat arbitrary visual identification of jumps): jumps do not follow a time-homogeneous Poisson process, and jumps and (ordinary) volatility are not independent. This means that λ_t (the Poisson intensity) must be stochastic and it must depend on q_s , for $s < t$. Also, σ_t must depend on q_s , for $s < t$. Furthermore, σ_t and λ_t must be correlated.

Rather than suggesting a comprehensive and realistic model, in this section, we have just discussed some of the relevant considerations. We seek a realistic model that accounts for the peculiar properties of the rate-of-return process, but we must realistically limit the degrees of freedom in the model.

6 Agent-Based Models

The pricing models discussed in Sections 2 through 5 are developed from a macro perspective on the prices themselves. This perspective excludes aspects of the market that results from irrational human behavior, where “irrational” is defined subjectively and usually means that the market participants are attempting to optimize a simple objective function. In a rational approach to modeling market behavior, what individual traders are doing has no affect on the decision of a trader to buy or sell; that is the market does not have “momentum”. There is an instantaneous adjustment of prices to some “fair market value”. No matter how attractive a rational approach to financial modeling is, its attractive simplicity cannot make

it so. Market participants do not act independently of each other. Traders do not have share the same processed data. Traders do not identify the same objective function. Traders do not all share a similar model of the market. The proportion of traders who behave in a certain way, that is, who do share a similar model varies in time.

The ultimate dependence of prices on the beliefs and actions of individual traders suggests another approach to financial modeling. This approach begins with models of behavior of the market participants. In this kind of approach to scientific modeling, called “agent-based”, the actions of a set of individual “agents” are governed by control parameters that can depend on the actions of other agents.

We will not pursue this approach here. LeBaron (2006) provides a survey of the micro perspective modeling incorporated in an agent-based approach.

7 Applications of Pricing Models

We must emphasize again that the role of pricing models is not to predict prices. Pricing models provide a description of stochastic behavior, and for that reason they have important applications in a number of areas, such as in the regulation of financial markets, in management of risk, and in pricing of derivative assets.

Options pricing is probably the highest profile application of asset pricing models. This application soared to prominence in the early 1970’s when Black and Scholes used the differential equation (10) derived from the geometric Brownian motion model (9) to develop exact formulas for fair prices of European puts and calls.

As we have pointed out, the simple geometric Brownian motion model does not correspond very well with empirical data. Although prices yielded by the Black-Scholes options pricing formulas were useful for traders, they quickly noticed that the prices set by the market differed from the Black-Scholes prices in systematic ways. If the market price is inserted as the price in a Black-Scholes formula, any other single variable in the formula can be solved for. The time to expiry, the current market price of the asset, and the strike price are all directly observable, so the only variable in the model that might be considered questionable

is the volatility. An interesting fact emerged; if the formula is applied to options on the same underlying asset and at the same time to expiry but at different strike prices, the value of the volatility that satisfies the formula is not constant, but rather a convex function of the strike price. This was called the “volatility smile”. Likewise, if the same strike price but different times to expiry are entered into the formula, the volatility exhibits systematic curvature. Fengler (2010, this volume) provides more details on this kind of result from the Black-Scholes formula.

Although we have taken the definition of “volatility” simply to be “standard deviation of rates of returns”, we have already indicated in Section 1.2 the difficulties in assigning a value to volatility. The value of volatility implied by the inverted use of the Black-Scholes formula with observed prices of derivatives therefore has intrinsic interest. Volatility defined by inversion of a pricing formula is called “implied volatility”, and so volatility defined as originally in terms of a standard deviation is now often called “statistical volatility”. The inverted use of pricing models together with observed prices of derivatives to define a type of asset price volatility is probably more common now than use of the pricing models for their earlier purpose of determining fair prices for derivatives.

There are now markets in implied volatility of various market indexes, and this kind of market provides another tool for hedging investment risks. The most widely traded such implied volatility index is the VIX, which follows the implied volatility of the S&P 500. Traded implied volatility indexes use rather complicated asset pricing models; none currently use the simple Black-Scholes formula.

The simpler models such as ARMA/ARIMA or geometric Brownian motion can often be analyzed by well-established statistical methods. The most impressive result of such an analysis is probably the Black-Scholes formulas. For more realistic models, the analysis is often by Monte-Carlo methods. In the case of stochastic models, the Monte Carlo methods are often coupled with numerical solutions to the stochastic differential equations; see, for example, Sauer (2010, this volume).

Realistic asset pricing models generally present analysis problems that can feasibly be addressed only by Monte Carlo methods. See Yu (2010, this volume) or Glasserman (2004) for more detailed discussion of Monte Carlo methods in the application of asset pricing models.

References

Bjursell, Johan, and James E. Gentle (2010), Identifying jumps in asset prices, *Handbook of Computational Finance* (edited by Jin-Chuan Duan, James E. Gentle, and Wolfgang Härdle), Springer, Berlin, xxx–xxx.

Christoffersen, Peter; Kris Jacobs; and Chayawat Ornthanalai (2010), GARCH Option Pricing: Theory and Evidence *Handbook of Computational Finance* (edited by Jin-Chuan Duan, James E. Gentle, and Wolfgang Härdle), Springer, Berlin, xxx–xxx.

Cont, Rama, and Peter Tankov (2004), *Financial Modelling with Jump Processes*, Chapman & Hall/CRC, Boca Raton.

Detemple, Jérôme, and Marcel Rindisbacher (2010), Diffusion models of asset prices, *Handbook of Computational Finance* (edited by Jin-Chuan Duan, James E. Gentle, and Wolfgang Härdle), Springer, Berlin, xxx–xxx.

Fengler, Matthias (2010), Option data and modelling BSM implied volatility, *Handbook of Computational Finance* (edited by Jin-Chuan Duan, James E. Gentle, and Wolfgang Härdle), Springer, Berlin, xxx–xxx.

Figuroa-López, José E. (2010), Jump-diffusion models driven by Lévy processes, *Handbook of Computational Finance* (edited by Jin-Chuan Duan, James E. Gentle, and Wolfgang Härdle), Springer, Berlin, xxx–xxx.

Glasserman, Paul (2004), *Monte Carlo Methods in Financial Engineering*, Springer-Verlag, New York.

Gouriéroux, Christian (1997), *ARCH Models and Financial Applications*, Springer-Verlag, New York.

Hafner, Christian M., and Hans Manner (2010), Multivariate time series models for asset prices, *Handbook of Computational Finance* (edited by Jin-Chuan Duan, James E. Gentle, and Wolfgang Härdle), Springer, Berlin, xxx–xxx.

Jondeau, Eric; Ser-Huang Poon; and Michael Rockinger (2007), *Financial Modeling under*

Non-Gaussian Distributions, Springer, London.

Lai, Tze-Leung, and Haipeng Xing (2008), *Statistical Models and Methods for Financial Markets*, Springer Science+Business Media LLC, New York.

LeBaron, Blake (2006), Agent-based computational finance, *Handbook of Computational Economics*, Volume 2, (edited by Leigh Tesfatsion and Kenneth L. Judd), North-Holland, 1187-1232.

Rachev, Svetlozar T.; Christian Menn; and Frank J. Fabozzi (2005), *Fat-Tailed and Skewed Asset Return Distributions*, John Wiley & Sons, Inc., Hoboken.

Samuelson, P. A. (1965), Proof that properly anticipated prices fluctuate randomly, *Industrial Management Review* **6**, 41–49.

Sauer, Timothy (2010), Numerical solution of stochastic differential equations in finance, *Handbook of Computational Finance* (edited by Jin-Chuan Duan, James E. Gentle, and Wolfgang Härdle), Springer, Berlin, xxx–xxx.

Steele, J. Michael (2001), *Stochastic Calculus and Financial Applications*, Springer-Verlag, New York.

Yu, Jun (2010), Simulation-based estimation methods for financial time series models, *Handbook of Computational Finance* (edited by Jin-Chuan Duan, James E. Gentle, and Wolfgang Härdle), Springer, Berlin, xxx–xxx.

SFB 649 Discussion Paper Series 2010

For a complete list of Discussion Papers published by the SFB 649, please visit <http://sfb649.wiwi.hu-berlin.de>.

- 001 "Volatility Investing with Variance Swaps" by Wolfgang Karl Härdle and Elena Silyakova, January 2010.
- 002 "Partial Linear Quantile Regression and Bootstrap Confidence Bands" by Wolfgang Karl Härdle, Ya'acov Ritov and Song Song, January 2010.
- 003 "Uniform confidence bands for pricing kernels" by Wolfgang Karl Härdle, Yarema Okhrin and Weining Wang, January 2010.
- 004 "Bayesian Inference in a Stochastic Volatility Nelson-Siegel Model" by Nikolaus Hautsch and Fuyu Yang, January 2010.
- 005 "The Impact of Macroeconomic News on Quote Adjustments, Noise, and Informational Volatility" by Nikolaus Hautsch, Dieter Hess and David Veredas, January 2010.
- 006 "Bayesian Estimation and Model Selection in the Generalised Stochastic Unit Root Model" by Fuyu Yang and Roberto Leon-Gonzalez, January 2010.
- 007 "Two-sided Certification: The market for Rating Agencies" by Erik R. Fasten and Dirk Hofmann, January 2010.
- 008 "Characterising Equilibrium Selection in Global Games with Strategic Complementarities" by Christian Basteck, Tijmen R. Daniels and Frank Heinemann, January 2010.
- 009 "Predicting extreme VaR: Nonparametric quantile regression with refinements from extreme value theory" by Julia Schaumburg, February 2010.
- 010 "On Securitization, Market Completion and Equilibrium Risk Transfer" by Ulrich Horst, Traian A. Pirvu and Gonçalo Dos Reis, February 2010.
- 011 "Illiquidity and Derivative Valuation" by Ulrich Horst and Felix Naujokat, February 2010.
- 012 "Dynamic Systems of Social Interactions" by Ulrich Horst, February 2010.
- 013 "The dynamics of hourly electricity prices" by Wolfgang Karl Härdle and Stefan Trück, February 2010.
- 014 "Crisis? What Crisis? Currency vs. Banking in the Financial Crisis of 1931" by Albrecht Ritschl and Samad Sarferaz, February 2010.
- 015 "Estimation of the characteristics of a Lévy process observed at arbitrary frequency" by Johanna Kappusl and Markus Reiß, February 2010.
- 016 "Honey, I'll Be Working Late Tonight. The Effect of Individual Work Routines on Leisure Time Synchronization of Couples" by Juliane Scheffel, February 2010.
- 017 "The Impact of ICT Investments on the Relative Demand for High-Medium-, and Low-Skilled Workers: Industry versus Country Analysis" by Dorothee Schneider, February 2010.
- 018 "Time varying Hierarchical Archimedean Copulae" by Wolfgang Karl Härdle, Ostap Okhrin and Yarema Okhrin, February 2010.
- 019 "Monetary Transmission Right from the Start: The (Dis)Connection Between the Money Market and the ECB's Main Refinancing Rates" by Puriya Abbassi and Dieter Nautz, March 2010.
- 020 "Aggregate Hazard Function in Price-Setting: A Bayesian Analysis Using Macro Data" by Fang Yao, March 2010.
- 021 "Nonparametric Estimation of Risk-Neutral Densities" by Maria Grith, Wolfgang Karl Härdle and Melanie Schienle, March 2010.

SFB 649 Discussion Paper Series 2010

For a complete list of Discussion Papers published by the SFB 649, please visit <http://sfb649.wiwi.hu-berlin.de>.

- 022 "Fitting high-dimensional Copulae to Data" by Ostap Okhrin, April 2010.
- 023 "The (In)stability of Money Demand in the Euro Area: Lessons from a Cross-Country Analysis" by Dieter Nautz and Ulrike Rondorf, April 2010.
- 024 "The optimal industry structure in a vertically related market" by Raffaele Fiocco, April 2010.
- 025 "Herding of Institutional Traders" by Stephanie Kremer, April 2010.
- 026 "Non-Gaussian Component Analysis: New Ideas, New Proofs, New Applications" by Vladimir Panov, May 2010.
- 027 "Liquidity and Capital Requirements and the Probability of Bank Failure" by Philipp Johann König, May 2010.
- 028 "Social Relationships and Trust" by Christine Binzel and Dietmar Fehr, May 2010.
- 029 "Adaptive Interest Rate Modelling" by Mengmeng Guo and Wolfgang Karl Härdle, May 2010.
- 030 "Can the New Keynesian Phillips Curve Explain Inflation Gap Persistence?" by Fang Yao, June 2010.
- 031 "Modeling Asset Prices" by James E. Gentle and Wolfgang Karl Härdle, June 2010.