

CHAPTER 1

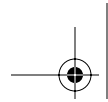
Financial Econometrics: Scope and Methods

Financial econometrics is the econometrics of financial markets. It is a quest for models that describe financial time series such as prices, returns, interest rates, financial ratios, defaults, and so on. The economic equivalent of the laws of physics, econometrics represents the quantitative, mathematical laws of economics. The development of a quantitative, mathematical approach to economics started at the end of the 19th century, in a period of great enthusiasm for the achievements of science and technology.

The World Exhibition held in Paris in 1889 testifies to the faith of that period in science and technology. The key attraction of the exhibition—the Eiffel Tower—was conceived by Gustave Eiffel, an architect and engineer who had already earned a reputation building large metal structures such as the 94-foot-high wrought-iron square skeleton that supports the Statue of Liberty.¹ With its 300-meter-high iron structure, Eiffel's tower was not only the tallest building of its time but also a

¹ Eiffel was a shrewd businessman as well as an accomplished engineer. When he learned that the funding for the 1889 World Exhibition tower would cover only one fourth of the cost, he struck a deal with the French government: He would raise the requisite funds in return for the right to exploit the tower commercially for 20 years. The deal made him wealthy. In the first year alone, revenues covered the entire cost of the project! Despite his sense of business, Eiffel's career was destroyed by the financial scandal surrounding the building of the Panama Canal, for which his firm was a major contractor. Though later cleared of accusations of corruption, Eiffel abandoned his business activities and devoted the last 30 years of his life to research.





monument to applied mathematics. To ensure that the tower would withstand strong winds, Eiffel wrote an integral equation to determine the tower's shape.²

The notion that mathematics is the language of nature dates back 2,000 years to the ancient Greeks and was forcefully expressed by Galileo. In his book *Il saggiaiore (The Assayer)*, published in 1623, Galileo wrote (translation by one of the authors of this book):

[The universe] cannot be read until we have learnt the language and become familiar with the characters in which it is written. It is written in the language of mathematics; the letters are triangles, circles, and other geometrical figures, without which it is humanly impossible to comprehend a single word.

It was only when Newton published his *Principia* some 60 years later (1687) that this idea took its modern form. In introducing the concept of *instantaneous rate of change*³ and formulating mechanics as laws that link variables and their rates of change, Newton made the basic leap forward on which all modern physical sciences are based. Linking variables to their rate of change is the principle of differential equations. Its importance can hardly be overestimated. Since Newton, differential equations have progressively conquered basically all the fields of the physical sciences, including mechanics, thermodynamics, electromagnetism, relativity, and quantum mechanics.

During the 19th century, physics based on differential equations revolutionized technology. It was translated into steam and electrical engines, the production and transmission of electrical power, the transmission of electrical signals, the chemical transformation of substances, and the ability to build ships, trains, and large buildings and bridges. It

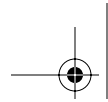
² The design principles employed by Eiffel have been used in virtually every subsequent tall building. Eiffel's equation,

$$\frac{1}{2} \int_x^H f(x)^2 dx - c(H-x) = \int_x^H xw(x)f(x) dx$$

states that the torque from the wind on any part of the Tower from a given height to the top is equal to the torque of the weight of this same part.

³ The instantaneous rate of change, "derivative" in mathematical terminology, is one of the basic concepts of calculus. Calculus was discovered independently by Newton and Leibniz, who were to clash bitterly in claiming priority in the discovery.





changed every aspect of the manufacture of goods and transportation. Faith in the power of science and technology reached a peak.⁴

Enthusiasm for science led to attempts to adopt the principles of the physical sciences to domains as varied as linguistics, the behavioral sciences, and economics. The notion of economic equilibrium had already been introduced by Stanley Jevons⁵ and Carl Menger⁶ when Leon Walras⁷ and Vilfredo Pareto⁸ made the first attempts to write comprehensive mathematical laws of the economy. Engineers by training, Walras and Pareto set themselves the task of explicitly writing down the equation of economic equilibrium. Their objective was well in advance on their time. A reasonable theoretical quantitative description of economic systems had to wait the full development of probability theory and statistics during the first half of the 20th century. And its practical application had to wait the development of fast computers. It was only in the second half of the 20th century that a quantitative description of economics became a mainstream discipline: econometrics (i.e., the quantitative science of economics) was born.

THE DATA GENERATING PROCESS

The basic principles for formulating quantitative laws in financial econometrics are the same as those that have characterized the development of quantitative science over the last four centuries. We write mathematical models, that is, relationships between different variables and/or variables in different moments and different places. The basic tenet of quantitative science is that there are relationships that do not change regardless of the


⁴ The 19th century had a more enthusiastic and naive view of science and the linearity of its progress than we now have. There are two major differences. First, 19th century science believed in unlimited possibilities of future progress; modern science is profoundly influenced by the notion that uncertainty is not eliminable. Second, modern science is not even certain about its object. According to the standard interpretation of quantum mechanics, the laws of physics are considered mere recipes to predict experiments, void of any descriptive power.

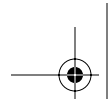
⁵ Stanley Jevons, *Theory of Political Economy* (London: Macmillan, 1871).

⁶ Carl Menger, *Principles of Economics* (available online at <http://www.mises.org/etexts/menger/Mengerprinciples.pdf>). Translated by James Dingwall and Bert F. Hoselitz from *Grundsätze der Volkswirtschaftslehre* published in 1871.

⁷ Léon Walras. *Éléments d'économie politique pure; ou, Théorie de la richesse sociale* (Elements of Pure Economics or The Theory of Social Wealth) (Lausanne: Rouge, 1874).

⁸ Vilfredo Pareto, *Manuel d'économie Politique* (Manual of Political Economy), translated by Ann S. Schwier from the 1906 edition (New York: A.M. Kelley, 1906).






moment or the place under consideration. For example, while sea waves might look like an almost random movement, in every moment and location the basic laws of hydrodynamics hold without change. Similarly, asset price behavior might appear to be random, but econometric laws should hold in every moment and for every set of assets.

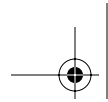
There are similarities between financial econometric models and models of the physical sciences but there are also important differences. The physical sciences aim at finding immutable laws of nature; econometric models model the economy or financial markets—artifacts subject to change. For example, financial markets in the form of stock exchanges have been in operation for two centuries. During this period, they have changed significantly both in the number of stocks listed and the type of trading. And the information available on transactions has also changed. Consider that in the 1950s, we had access only to daily closing prices and this typically the day after; now we have instantaneous information on every single transaction. Because the economy and financial markets are artifacts subject to change, econometric models are not unique representations valid throughout time; they must adapt to the changing environment.

While basic physical laws are expressed as differential equations, financial econometrics uses both continuous time and discrete time models. For example, continuous time models are used in modeling derivatives where both the underlying and the derivative price are represented by stochastic (i.e., random) differential equations. In order to solve stochastic differential equations with computerized numerical methods, derivatives are replaced with finite differences.⁹ This process of discretization of time yields discrete time models. However, discrete time models used in financial econometrics are not necessarily the result of a process of discretization of continuous time models.

Let's focus on models in discrete time, the bread-and-butter of econometric models used in asset management. There are two types of discrete-time models: static and dynamic. Static models involve different variables at the same time. The well-known *capital asset pricing model* (CAPM), for example, is a static model. Dynamic models involve one or more vari-

⁹The stochastic nature of differential equations introduces fundamental mathematical complications. The definition of stochastic differential equations is a delicate mathematical process invented, independently, by the mathematicians Ito and Stratonovich. In the Ito-Stratonovich definition, the path of a stochastic differential equation is not the solution of a corresponding differential equation. However, the numerical solution procedure yields a discrete model that holds pathwise. See Sergio M. Focardi and Frank J. Fabozzi, *The Mathematics of Financial Modeling and Investment Management* (Hoboken, NJ: John Wiley & Sons, 2004) and the references therein for details.





ables at two or more moments.¹⁰ Momentum models, for example, are dynamic models.

In a dynamic model, the mathematical relationship between variables at different times is called the *data generating process* (DGP). This terminology reflects the fact that, if we know the DGP of a process, we can simulate the process recursively, starting from initial conditions. Consider the time series of a stock price p_t , that is, the series formed with the prices of that stock taken at fixed points in time, say daily. Let's now write a simple econometric model of the prices of a stock as follows:¹¹

$$p_{t+1} = \mu + \rho p_t + \varepsilon_{t+1}$$

This model tells us that if we consider any time $t + 1$, the price of that stock at time $t + 1$ is equal to a constant plus the price in the previous moment t multiplied by ρ plus a zero-mean random disturbance independent from the past, which always has the same statistical characteristics.¹² A random disturbance of this type is called a *white noise*.¹³

If we know the initial price p_0 at time $t = 0$, using a computer program to generate random numbers, we can simulate a path of the price process with the following recursive equations:

$$\begin{aligned} p_1 &= \mu + \rho p_0 + \varepsilon_1 \\ p_2 &= \mu + \rho p_1 + \varepsilon_2 \\ &\vdots \end{aligned}$$

That is, we can compute the price at time $t = 1$ from the initial price p_0 and a computer-generated random number ε_1 and then use this new price to compute the price at time $t = 2$, and so on.¹⁴ It is clear that we


¹⁰ This is true in discrete time. In continuous time, a dynamic model might involve variables and their derivatives at the same time.

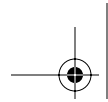
¹¹ In this example, we denote prices with lower case p and assume that they follow a simple linear model. In the following chapters, we will make a distinction between prices, represented with upper case letter P and the logarithms of prices, represented by lower case letters. Due to the geometric compounding of returns, prices are assumed to follow nonlinear processes.

¹² If we want to apply this model to real-world price processes, the constants μ and ρ must be estimated. μ determines the trend and ρ defines the dependence between the prices. Typically ρ is less than but close to 1.

¹³ The concept of white noise will be made precise in the following chapters where different types of white noise will be introduced.

¹⁴ The ε_i are independent and identically distributed random variables with zero mean. Typical choices for the distribution of ε are normal distribution, t -distribution, and stable non-Gaussian distribution. The distribution parameters are estimated from the sample (see Chapter 3).






have a DGP as we can generate any path. An econometric model that involves two or more different times can be regarded as a DGP.

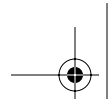
However, there is a more general way of looking at econometric models that encompasses both static and dynamic models. That is, we can look at econometric models from a perspective other than that of the recursive generation of stochastic paths. In fact, we can rewrite our previous model as follows:

$$p_{t+1} - \mu - \rho p_t = \varepsilon_{t+1}$$

This formulation shows that, if we consider any two consecutive instants of time, there is a combination of prices that behave as random noise. More in general, an econometric model can be regarded as a mathematical device that reconstructs a noise sequence from empirical data. This concept is visualized in Exhibit 1.1, which shows a time series of numbers p_t generated by a computer program according to the previous rule with $\rho = 0.9$ and $\mu = 1$ and the corresponding time series ε_t . If we choose any pair of consecutive points in time, say $t + 1, t$, the differ-

EXHIBIT 1.1 DGP and Noise Terms





ence $p_{t+1} - \mu - \rho p_t$ is always equal to the series ε_{t+1} . For example, consider the points $p_{13} = 10.2918$, $p_{14} = 12.4065$. The difference $p_{14} - 0.9p_{13} - 1 = 2.1439$ has the same value as ε_{14} . If we move to a different pair we obtain the same result, that is, if we compute $p_{t+1} - 1 - 0.9p_t$, the result will always be the noise sequence ε_{t+1} .

To help intuition, imagine that our model is a test instrument: probing our time series with our test instrument, we always obtain the same reading. Actually, what we obtain is not a constant reading but a random reading with mean zero and fixed statistical characteristics. The objective of financial econometrics is to find possibly simple expressions of different financial variables such as prices, returns, or financial ratios in different moments that always yield, as a result, a zero-mean random disturbance.

Static models (i.e., models that involve only one instant) are used to express relationships between different variables at any given time. Static models are used, for example, to determine exposure to different risk factors. However, because they involve only one instant, static models cannot be used to make forecasts; forecasting requires models that link variables in two or more instants in time.

FINANCIAL ECONOMETRICS AT WORK

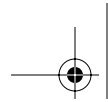
Applying financial econometrics involves three key steps:

1. Model selection
2. Model estimation
3. Model testing

In the first step, model selection, the modeler chooses (or might write *ex novo*) a family of models with given statistical properties. This entails the mathematical analysis of the model properties as well as economic theory to justify the model choice. It is in this step that the modeler decides to use, for example, regression on financial ratios or other variables to model returns.

In general, models include a number of free parameters that have to be estimated from sample data, the second step in applying financial econometrics. Suppose that we have decided to model returns with a regression model, a technique that we discuss in later chapters. This requires the estimation of the regression coefficients, performed using historical data. Estimation provides the link between reality and models. As econometric models are probabilistic models, any model can in principle describe our





empirical data. We choose a family of models in the model selection phase and then determine the optimal model in the estimation phase.

As mentioned, model selection and estimation are performed on historical data. As models are adapted (or fitted) to historical data there is always the risk that the fitting process captures ephemeral features of the data. Thus there is the need to test the models on data different from the data on which the models were estimated. This is the third step in applying financial econometrics, model testing. We assess the performance of models on fresh data.

We can take a different approach to model selection and estimation, namely statistical learning. Statistical learning combines the two steps—model selection and model estimation—insofar as it makes use of a class of universal models that can fit any data. Neural networks are an example of universal models. The critical step in the statistical learning approach is estimation. This calls for methods to restrict model complexity (i.e., the number of parameters used in a model).

Within this basic scheme for applying financial econometrics, we can now identify a number of modeling issues, such as:

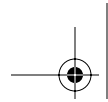
- How do we apply statistics given that there is only one realization of financial series?
- Given a sample of historical data, how do we choose between linear and nonlinear models, or the different distributional assumptions or different levels of model complexity?
- Can we exploit more data using, for example, high-frequency data?
- How can we make our models more robust, reducing model risk?
- How do we measure not only model performance but also the ability to realize profits?

Implications of Empirical Series with Only One Realization

As mentioned, econometric models are *probabilistic* models: Variables are random variables characterized by a probability distribution. Generally speaking, probability concepts cannot be applied to single “individuals.”¹⁵ Probabilistic models describe “populations” formed by many individuals. However, empirical financial time series have only one realization. For example, there is only one historical series of prices for each stock—and we have only one price at each instant of time. This makes problematic the application of probability concepts. How, for example, can we meaningfully discuss the distribution of prices at a specific time given that there is only one price observation? Applying probability concepts to perform estimation and testing would require populations made up of multi-

¹⁵ At least, not if we use a frequentist concept of probability. See Chapter 2.





ple time series and samples made up of different time series that can be considered a random draw from some distribution.

As each financial time series is unique, the solution is to look at the single elements of the time series as the individuals of our population. For example, because there is only one realization of each stock's price time series, we have to look at the price of each stock at different moments. However, the price of a stock (or of any other asset) at different moments is not a random independent sample. For example, it makes little sense to consider the distribution of the prices of a single stock in different moments because the level of prices typically changes over time. Our initial time series of financial quantities must be transformed; that is, a unique time series must be transformed into populations of individuals to which statistical methods can be applied. This holds not only for prices but for any other financial variable.


Econometrics includes transformations of the above type as well as tests to verify that the transformation has obtained the desired result. The DGP is the most important of these transformations. Recall that we can interpret a DGP as a method for transforming a time series into a sequence of noise terms. The DGP, as we have seen, constructs a sequence of random disturbances starting from the original series; it allows one to go backwards and infer the statistical properties of the series from the noise terms and the DGP. However, these properties cannot be tested independently.

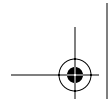
The DGP is not the only transformation that allows statistical estimates. Differencing time series, for example, is a process that, as we will see in Chapter 6, may transform nonstationary time series into stationary time series. A stationary time series has a constant mean that, under specific assumptions, can be estimated as an empirical average.

Determining the Model

As we have seen, econometric models are mathematical relationships between different variables at different times. An important question is whether these relationships are linear or nonlinear. Consider that every econometric model is an approximation. Thus the question is: Which approximation—linear or nonlinear—is better?

To answer this, it is generally necessary to consider jointly the linearity of models, the distributional assumptions, and the number of time lags to introduce. The simplest models are linear models with a small number of lags under the assumption that variables are normal variables. A widely used example of normal linear models are regression models where returns are linearly regressed on lagged factors under the assumption that noise terms are normally distributed. A model of this type can be written as:





$$r_{t+1} = \beta f_t + \varepsilon_{t+1}$$

where r_t are the returns at time t and f_t are factors, that is economic or financial variables. Given the linearity of the model, if factors and noise are jointly normally distributed, returns are also normally distributed.

However, the distribution of returns, at least at some time horizons, is not normal. If we postulate a nonlinear relationship between factors and returns, normally distributed factors yield a nonnormal return distribution. However, we can maintain the linearity of the regression relationship but assume a nonnormal distribution of noise terms and factors. Thus a nonlinear models transforms normally distributed noise into non-normal variables but it is not true that nonnormal distributions of variables implies nonlinear models.


If we add lags (i.e., a time space backwards), the above model becomes sensitive to the shape of the factor paths. For example, a regression model with two lags will behave differently if the factor is going up or down. Adding lags makes models more flexible but more brittle. In general, the optimal number of lags is dictated not only by the complexity of the patterns that we want to model but also by the number of points in our sample. If sample data are abundant, we can estimate a rich model.

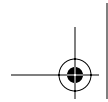
Typically there is a trade-off between model flexibility and the size of the data sample. By adding time lags and nonlinearities, we make our models more flexible, but the demands in terms of estimation data are greater. An optimal compromise has to be made between the flexibility given by nonlinear models and/or multiple lags and the limitations due to the size of the data sample.

TIME HORIZON OF MODELS

There are trade-offs between model flexibility and precision that depend on the size of sample data. To expand our sample data, we would like to use data with small time spacing in order to multiply the number of available samples. High-frequency data or HFD (i.e., data on individual transactions) have the highest possible frequency (i.e., each individual transaction) and are irregularly spaced. To give an idea of the ratio in terms of numbers, consider that there are approximately 2,100 ticks per day for the median stock in the Russell 3000.¹⁶ Thus the size of the HDF data set of one day for a typical stock in the Russell 3000 is 2,100 times larger than the size of closing data for the same day!

¹⁶Thomas Neal Falkenberry, "High Frequency Data Filtering," Tick Data Inc., 2002.





In order to exploit all available data, we would like to adopt models that work over time intervals of the order of minutes and, from these models, compute the behavior of financial quantities over longer periods. Given the number of available sample data at high frequency, we could write much more precise laws than those established using longer time intervals. Note that the need to compute solutions over forecasting horizons much longer than the time spacing is a general problem which applies at any time interval. For example, as will be discussed in Chapter 5, in asset allocation we need to understand the behavior of financial quantities over long time horizons. The question we need to ask is if models estimated using daily intervals can correctly capture the process dynamics over longer periods, such as years.

It is not necessarily true that models estimated on short time intervals, say minutes, offer better forecasts at longer time horizons than models estimated on longer time intervals, say days. This is because financial variables might have a complex short-term dynamics superimposed on a long-term dynamics. It might be that using high-frequency data one captures the short-term dynamics without any improvement in the estimation of the long-term dynamics. That is, with high-frequency data it might be that models get more complex (and thus more data-hungry) because they describe short-term behavior superimposed on long-term behavior. This possibility must be resolved for each class of models.


Another question is if it is possible to use the *same* model at different time horizons. To do so is to imply that the behavior of financial quantities is similar at different time horizons. This conjecture was first made by Benoit Mandelbrot who observed that long series of cotton prices were very similar at different time aggregations.¹⁷ This issue will be discussed in Chapter 14 where we review families of variables and processes that exhibit self-similarity.

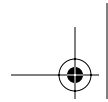
Model Risk and Model Robustness

Not only are econometric models probabilistic models, as we have already noted; they are only *approximate* models. That is, the probability distributions themselves are only approximate and uncertain. The theory of model risk and model robustness assumes that all parameters of a model are subject to uncertainty, and attempts to determine the consequence of model uncertainty and strategies for mitigating errors.

The growing use of models in finance over the last decade has heightened the attention to model risk and model-risk mitigation techniques. Asset management firms are beginning to address the need to

¹⁷ Benoit Mandelbrot, "The Variation of Certain Speculative Prices," *Journal of Business* 36 (1963), pp. 394–419.





implement methodologies that allow both robust estimation and robust optimization in the portfolio management process.

Performance Measurement of Models

It is not always easy to understand *ex ante* just how well (or how poorly) a forecasting model will perform. Because performance evaluations made on training data are not reliable, the evaluation of model performance requires separate data sets for training and for testing. Models are estimated on training data and tested on the test data. Poor performance might be due to model misspecification, that is, models might not reflect the true DGP of the data (assuming one exists), or there might simply be no DGP.

Various measures of model performance have been proposed. For example, one can compute the correlation coefficient between the forecasted variables and their actual realizations. Each performance measure is a single number and therefore conveys only one aspect of the forecasting performance. Often it is crucial to understand if errors can become individually very large or if they might be correlated. Note that a simple measure of model performance does not ensure the profitability of strategies. This can be due to a number of reasons, including, for example, the risk inherent in apparently profitable forecasts, market impact, and transaction costs.



APPLICATIONS


There has been a greater use of econometric models in investment management since the turn of the century. Application areas include:

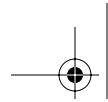
- Portfolio construction and optimization
- Risk management
- Asset and liability management

Each type of application requires different modeling approaches. In the appendix to this chapter, we provide a more detailed description of the investment management process and some investment concepts that will be used in this book.

Portfolio Construction and Optimization

Portfolio construction and optimization require models to forecast returns: There is no way to escape the need to predict future returns. Passive strategies apparently eschew the need to forecast future returns of





individual stocks by investing in broad indexes. They effectively shift the need to forecast to a higher level of analysis and to longer time horizons.

Until recently, the mainstream view was that financial econometric models could perform dynamic forecasts of volatility but not of expected returns. However, volatility forecasts are rarely used in portfolio management. With the exception of some proprietary applications, the most sophisticated models used in portfolio construction until recently were factor models where forecasts are not dynamic but consist in estimating a drift (i.e., a constant trend) plus a variance-covariance matrix.

Since the late 1990s, the possibility of making dynamic forecasts of both volatility and expected returns has gained broad acceptance. During the same period, it became more widely recognized that returns are not normally distributed, evidence that had been reported by Mandelbrot in the 1960s. Higher moments of distributions are therefore important in portfolio management. We discuss the representation and estimation of nonnormal distributions in Chapter 14.




As observed above, the ability to correctly forecast expected returns does not imply, per se, that there are profit opportunities. In fact, we have to take into consideration the interplay between expected returns, higher moments, and transaction costs. As dynamic forecasts typically involve higher portfolio turnover, transaction costs might wipe out profits. As a general comment, portfolio management based on dynamic forecasts calls for a more sophisticated framework for optimization and risk management with respect to portfolio management based on static forecasts.

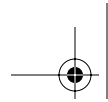
At the writing of this book, regression models form the core of the modeling efforts to predict future returns at many asset management firms. Regression models regress returns on a number of predictors. Stated otherwise, future returns are a function of the value of present and past predictors. Predictors include financial ratios such as earning-to-price ratio or book-to-price ratio and other fundamental quantities; predictors might also include behavioral variables such as market sentiment. A typical formula of a regressive model is the following:

$$r_{i,t+1} = \alpha_i + \sum_{j=1}^s \beta_{ij} f_{j,t} + \varepsilon_{i,t+1}$$

where

$$r_{i,t+1} = \frac{P_{i,t+1} - P_{i,t}}{P_{i,t}}$$





is the return at time $t + 1$ of the i -th asset and the $f_{j,t}$ are factors observed at time t . While regressions are generally linear, nonlinear models are also used.


In general, the forecasting horizon in asset management varies from a few days for actively managed or hedge funds to several weeks for more traditionally managed funds. Dynamic models typically have a short forecasting horizon as they capture a short-term dynamics. This contrasts with static models, such as the widely used multifactor models, which tend to capture long-term trends and ignore short-term dynamics.


The evolution of forecasting models over the last two decades has also changed the way forecasts are used. A basic utilization of forecasts is in stock picking/ranking systems, which have been widely implemented at asset management firms. The portfolio manager builds his or her portfolio combining the model ranking with his or her personal views and within the constraints established by the firm. A drawback in using such an approach is the difficulty in properly considering the structure of correlations and the role of higher moments.

Alternatively, forecasts can be fed to an optimizer that automatically computes the portfolio weights. But because an optimizer implements an optimal trade-off between returns and some measure of risk, the forecasting model must produce not only returns forecasts but also measures of risk. If risk is measured by portfolio variance or standard deviation, the forecasting model must be able to provide an estimated variance-covariance matrix.

Estimating the variance-covariance matrix is the most delicate of the estimation tasks. Here is why. The number of entries of a variance-covariance matrix grows with the square of the number of stocks. As a consequence, the number of entries in a variance-covariance matrix rapidly becomes very large. For example, the variance-covariance matrix of the stocks in the S&P 500 is a symmetric matrix that includes some 125,000 entries. If our universe were the Russell 5000, the variance-covariance matrix would include more than 12,000,000 entries. The problem with estimating matrices of this size is that estimates are very noisy because the number of sample data is close to the number of parameters to estimate. For example, if we use three years of data for estimation, we have, on average, less than three data points per estimated entry in the case of the S&P 500; in the case of the Russell 5000, the number of data points would be one fourth of the number of entries to estimate! Robust estimation methods are called for.

Note that if we use forecasting models we typically have (1) an equilibrium variance-covariance matrix that represents the covariances of the long-run relationships between variables plus (2) a short-term, time-dependent, variance-covariance matrix. If returns are not normally distributed, optimizers might require the matrix of higher moments.





A third utilization of forecasting models and optimizers is to construct model portfolios. In other words, the output of the optimizer is used to construct not an actual but a model portfolio. This model portfolio is used as input by portfolio managers.

Risk Management

Risk management has different meanings in different contexts. In particular, when optimization is used, risk management is intrinsic to the optimization process, itself a risk-return trade-off optimization. In this case, risk management is an integral part of the portfolio construction process.

However, in most cases, the process of constructing portfolios is entrusted to human portfolio managers who might use various inputs including, as noted above, ranking systems or model portfolios. In these cases, portfolios might not be optimal from the point of view of risk management and it is therefore necessary to ensure independent risk oversight. This oversight might take various forms. One form is similar to the type of risk oversight adopted by banks. The objective is to assess potential deviations from expectations. In order to perform this task, the risk manager receives as input the composition of portfolios and makes return projections using static forecasting models.





Another form of risk oversight, perhaps the most diffused in portfolio management, assesses portfolio exposures to specific risk factors. As portfolio management is often performed relative to a benchmark and risk is defined as underperformance relative to the benchmark, it is important to understand the sensitivity of portfolios to different risk factors. This type of risk oversight does not entail the forecasting of returns. The risk manager uses various statistical techniques to estimate how portfolios move in function of different risk factors. In most cases, linear regressions are used. A typical model will have the following form:


$$r_{i,t} = \alpha_i + \sum_{j=1}^s \beta_{ij} f_{j,t} + \varepsilon_{i,t}$$

where

$$r_{i,t} = \frac{P_{i,t} - P_{i,t-1}}{P_{i,t-1}}$$

is the return observed at time t of the i -th asset and the $f_{j,t}$ are factors observed at time t . Note that this model is fundamentally different from a regressive model with time lags as written in the previous section.





Asset-Liability Management

Asset-liability management (ALM) is typical of those asset management applications that require the optimization of portfolio returns at some fixed time horizon plus a stream of consumption throughout the entire life of the portfolio. ALM is important for managing portfolios of institutional investors such as pension funds or foundations. It is also important for wealth management, where the objective is to cover the investor's financial needs over an extended period.

ALM requires forecasting models able to capture the asset behavior at short-, medium-, and long-term time horizons. Models of the long-term behavior of assets exist but are clearly difficult to test. Important questions related to these long-term forecasting models include:

- Do asset prices periodically revert to one or many common trends in the long run?
- Can we assume that the common trends (if they exist) are deterministic trends such as exponentials or are common trends stochastic (i.e., random) processes?
- Can we recognize regime shifts over long periods of time?



APPENDIX: INVESTMENT MANAGEMENT PROCESS

Finance is classified into two broad areas: investment management (or portfolio management) and corporate finance. While financial econometrics has been used in corporate finance primarily to test various theories having to do with the corporate policy, the major use has been in investment management. Accordingly, our primary focus in this book is on applications to investment management.

The investment management process involves the following five steps:

Step 1: Setting investment objectives


Step 2: Establishing an investment policy


Step 3: Selecting an investment strategy

Step 4: Selecting the specific assets

Step 5: Measuring and evaluating investment performance

The overview of the investment management process described below should help understand how the econometric tools presented in this book are employed by portfolio managers, analysts, plan sponsors, and researchers. In addition, we introduce concepts and investment terms that are used in the investment management area throughout this book.





Step 1: Setting Investment Objectives

The first step in the investment management process, setting investment objectives, begins with a thorough analysis of the investment objectives of the entity whose funds are being managed. These entities can be classified as *individual investors* and *institutional investors*. Within each of these broad classifications, there is a wide range of investment objectives.

The objectives of an individual investor may be to accumulate funds to purchase a home or other major acquisitions, to have sufficient funds to be able to retire at a specified age, or to accumulate funds to pay for college tuition for children. An individual investor may engage the services of a financial advisor/consultant in establishing investment objectives.

In general, we can classify institutional investors into two broad categories—those that have to meet contractually specified liabilities and those that do not. We can classify those in the first category as institutions with “liability-driven objectives” and those in the second category as institutions with “nonliability-driven objectives.” Many firms have a wide range of investment products that they offer investors, some of which are liability-driven and others that are nonliability-driven. Once the investment objective is understood, it will then be possible to (1) establish a benchmark by which to evaluate the performance of the investment manager and (2) evaluate alternative investment strategies to assess the potential for realizing the specified investment objective.


Step 2: Establishing an Investment Policy


The second step in the investment management process is establishing policy guidelines to satisfy the investment objectives. Setting policy begins with the asset allocation decision. That is, a decision must be made as to how the funds to be invested should be distributed among the major classes of assets.

Asset Classes

Throughout this book we refer to certain categories of investment products as an “asset class.” From the perspective of a U.S. investor, the convention is to refer the following as *traditional asset classes*:

- U.S. common stocks
- Non-U.S. (or foreign) common stocks
- U.S. bonds
- Non-U.S. (or foreign) bonds
- Cash equivalents
- Real estate





Cash equivalents are defined as short-term debt obligations that have little price volatility. Common stocks and bonds are further divided into asset classes. For U.S. common stocks (also referred to as U.S. equities), the following are classified as asset classes:



- Large capitalization stocks
- Mid-capitalization stocks
- Small capitalization stocks
- Growth stocks
- Value stocks

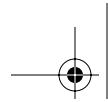
By “capitalization,” it is meant the market capitalization of the company’s common stock. This is equal to the total market value of all of the common stock outstanding for that company. For example, suppose that a company has 100 million shares of common stock outstanding and each share has a market value of \$10. Then the capitalization of this company is \$1 billion (100 million shares times \$10 per share). The market capitalization of a company is commonly referred to as the “market cap” or simply “cap.”

For U.S. bonds, also referred to as fixed-income securities, the following are classified as asset classes:

- U.S. government bonds
- Investment-grade corporate bonds
- High-yield corporate bonds
- U.S. municipal bonds (i.e., state and local bonds)
- Mortgage-backed securities
- Asset-backed securities

Corporate bonds are classified by the type of issuer. The four general classifications are (1) public utilities, (2) transportations, (3) banks/finance, and (4) industrials. Finer breakdowns are often made to create more homogeneous groupings. For example, public utilities are subdivided into electric power companies, gas distribution companies, water companies, and communication companies. Transportations are divided further into airlines, railroads, and trucking companies. Banks/finance include both money center banks and regional banks, savings and loans, brokerage firms, insurance companies, and finance companies. Industrials are the catchall class and the most heterogeneous of the groupings with respect to investment characteristics. Industrials include manufacturers, mining companies, merchandising, retailers, energy companies, and service-related industries.





Corporate bonds expose investors to credit risk. There are private companies that rate bonds with respect to their likelihood to default. They are Moody's, Standard & Poor's, and Fitch. These firms perform credit analysis and issue their conclusions about the credit risk of a company in the form of a rating. The rating systems use similar symbols. In all three systems, the term "high grade" means low credit risk, or conversely, high probability of future payments. The highest-grade bonds are designated by Moody's by the letters Aaa, and by the other two rating agencies by AAA. The next highest grade is Aa or AA; for the third grade all rating agencies use A. The next three grades are Baa or BBB, Ba or BB, and B, respectively. There are also C grades. Standard & Poor's and Fitch uses plus or minus signs to provide a narrower credit quality breakdown within each class, and Moody's uses 1, 2, or 3 for the same purpose. Bonds rated triple A (AAA or Aaa) are said to be prime; double A (AA or Aa) are of high quality; single A issues are called upper medium grade, and triple B are medium grade. Lower-rated bonds are said to have speculative elements or to be distinctly speculative.


Bond issues that are assigned a rating in the top four categories are referred to as investment-grade bonds. Issues that carry a rating below the top four categories are referred to as noninvestment-grade bonds, or more popularly as high-yield bonds or junk bonds. Thus, the corporate bond market can be divided into two sectors: the investment-grade and noninvestment-grade markets.

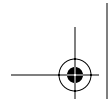
Mortgage-backed and asset-backed securities are referred to as securitized products. Agency mortgage-backed securities carry little credit risk and represent the largest spread sector in the bond market. By spread sector it is meant sectors of the bond market that offer a spread to U.S. Treasuries. The key use of econometric tools in analyzing mortgage-backed securities is to forecast prepayments. In the case of nonagency and asset-backed securities, econometric tools are used to forecast defaults and recoveries in addition to prepayments.

For non-U.S. stocks and bonds, the following are classified as asset classes:

- Developed market foreign stocks
- Emerging market foreign stocks
- Developed market foreign bonds
- Emerging market foreign bonds

In addition to the traditional asset classes, there are asset classes commonly referred to as *alternative investments*. Two of the more popular ones are hedge funds and private equity.





Constraints

There are some institutional investors that make the asset allocation decision based purely on their understanding of the risk-return characteristics of the various asset classes and expected returns. The asset allocation will take into consideration any investment constraints or restrictions. Asset allocation models are commercially available for assisting those individuals responsible for making this decision.

In the development of an investment policy, the following factors must be considered: client constraints, regulatory constraints, and tax and accounting issues.

Examples of client-imposed constraints would be restrictions that specify the types of securities in which a manager may invest and concentration limits on how much or little may be invested in a particular asset class or in a particular issuer. Where the objective is to meet the performance of a particular market or customized benchmark, there may be a restriction as to the degree to which the manager may deviate from some key characteristics of the benchmark.


There are many types of regulatory constraints. These involve constraints on the asset classes that are permissible and concentration limits on investments. Moreover, in making the asset allocation decision, consideration must be given to any risk-based capital requirements.

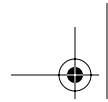
Step 3: Selecting a Portfolio Strategy

Selecting a portfolio strategy that is consistent with the investment objectives and investment policy guidelines of the client or institution is the third step in the investment management process. Portfolio strategies can be classified as either active or passive.

An *active portfolio strategy* uses available information and forecasting techniques to seek a better performance than a portfolio that is simply diversified broadly. Essential to all active strategies are expectations about the factors that have been found to influence the performance of an asset class. For example, with active common stock strategies this may include forecasts of future earnings, dividends, or price-earnings ratios. With bond portfolios that are actively managed, expectations may involve forecasts of future interest rates and sector spreads. Active portfolio strategies involving foreign securities may require forecasts of local interest rates and exchange rates.

A *passive portfolio strategy* involves minimal expectational input, and instead relies on diversification to match the performance of some market index. In effect, a passive strategy assumes that the marketplace will reflect all available information in the price paid for securities. Between these extremes of active and passive strategies, several strategies





have sprung up that have elements of both. For example, the core of a portfolio may be passively managed with the balance actively managed.

In the bond area, several strategies classified as *structured portfolio strategies* have been commonly used. A structured portfolio strategy is one in which a portfolio is designed to achieve the performance of some predetermined liabilities that must be paid out. These strategies are frequently used when trying to match the funds received from an investment portfolio to the future liabilities that must be paid.

Given the choice among active and passive management, which should be selected? The answer depends on (1) the client's or money manager's view of how "price-efficient" the market is; (2) the client's risk tolerance; and (3) the nature of the client's liabilities. By market-price efficiency we mean how difficult it would be to earn a greater return than passive management after adjusting for the risk associated with a strategy and the transaction costs associated with implementing that strategy. Market efficiency is explained in Chapter 5. Econometric tools are used to test theories about market efficiency.

Step 4: Selecting the Specific Assets

Once a portfolio strategy is selected, the next step is to select the specific assets to be included in the portfolio. It is in this phase of the investment management process that the investor attempts to construct an *efficient portfolio*. An efficient portfolio is one that provides the greatest expected return for a given level of risk or, equivalently, the lowest risk for a given expected return.


Inputs Required


To construct an efficient portfolio, the investor must be able to quantify risk and provide the necessary inputs. As will be explained in the next chapter, there are three key inputs that are needed: future expected return (or simply expected return), variance of asset returns, and correlation (or covariance) of asset returns. Many of the financial econometric tools described in this book are intended to provide the investor with information with which to estimate these three inputs.

There are a wide range of approaches to obtain the expected return of assets. Investors can employ various econometric tools discussed in this book to derive the future expected return of an asset.

Approaches to Portfolio Construction

Based on the expected return for a portfolio (which depends on the expected returns of all the asset returns in the portfolio) and some risk measure of the portfolio's return (which depends on the covariance of





returns between all pairs of assets in the portfolio) an efficient portfolio can be constructed. This approach also allows for the inclusion of constraints such as lower and upper bounds on particular assets or assets in particular industries or sectors. The end result of the analysis is a set of efficient portfolios—alternative portfolios from which the investor can select—that offer the maximum expected portfolio return for a given level of portfolio risk.

There are variations on this approach to portfolio construction. The analysis can be employed by estimating risk factors that historically have explained the variance of asset returns. The basic principle is that the value of an asset is driven by a number of systematic factors (or, equivalently, risk exposures) plus a component unique to a particular company or industry. A set of efficient portfolios can be identified based on the risk factors and the sensitivity of assets to these risk factors. This approach is referred to the “multifactor risk approach” to portfolio construction.


Step 5: Measuring and Evaluating Performance


The measurement and evaluation of investment performance is the last step in the investment management process. This step involves measuring the performance of the portfolio and then evaluating that performance relative to some benchmark. Econometric tools are used to construct models that can be employed to evaluate the performance of managers. We discuss this in Chapter 5.



CONCEPTS EXPLAINED IN THIS CHAPTER (IN ORDER OF PRESENTATION)

- Data generating process
- White noise
- High frequency data
- Models
 - Mathematical model
 - Model estimation
 - Model selection
 - Model testing
 - Static models
 - Dynamic models
 - Models as probes that recover IID sequences
 - Linear and nonlinear models
 - Model risk





Model robustness
Exact and approximate models
Model performance
Statistical learning
Stationary time series
Nonstationary time series
Differencing
Portfolio construction and optimization
Risk management
Asset-liability management

