CHAPTER 1

From Art to Engineering in Finance

It is often said that investment management is an art, not a science. However since early 1990s the market has witnessed a progressive shift towards a more industrial view of the investment management process. There are several reasons for this change. First, with globalization the universe of investable assets has grown many times over. Asset managers might have to choose from among several thousand possible investments from around the globe. The S&P 500 index is itself chosen from a pool of 8,000 investable U.S. stocks. Second, institutional investors, often together with their investment consultants, have encouraged asset management firms to adopt an increasingly structured process with documented steps and measurable results. Pressure from regulators and the media is another factor. Lastly, the sheer size of the markets makes it imperative to adopt safe and repeatable methodologies. The volumes are staggering. With the recent growth of the world's stock markets, total market capitalization is now in the range of tens of trillions of dollars[1] while derivatives held by U. S. commercial banks topped $65.8 trillion in the second quarter of 2003.[2]

[1] Exact numbers are difficult to come up with as information about many markets is missing and price fluctuations remain large.
[2] Office of the Comptroller of the Currency, *Quarterly Derivatives Report*, Second Quarter 2003.

INVESTMENT MANAGEMENT PROCESS

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 in understanding the activities that the portfolio manager faces and the need for the analytical tools that are described in the chapters that follow in 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 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 Chapter 3 we review the different types of institutional investors. We will also see that in general we can classify institutional investors into two broad categories—those that must 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." Some institutions 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" or "bogey" 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 and are covered in Chapter 2.

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

All of these securities are described in Chapter 2, where what is meant by "investment grade" and "high yield" are also explained. Sometimes, the first three bond asset classes listed above are further divided into "long term" and "short term."

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.

How does one define an asset class? One investment manager, Mark Kritzman, describes how this is done as follows:

> ... some investments take on the status of an asset class simply because the managers of these assets promote them as an asset class. They believe that investors will be more inclined to allocate funds to their products if they are viewed as an asset class rather than merely as an investment strategy.[3]

He then goes on to propose criteria for determining asset class status. We won't review the criteria he proposed here. They involve concepts that are explained in later chapters. After these concepts are explained it will become clear how asset class status is determined. However, it should not come as any surprise that the criteria proposed by Kritzman involve the risk, return, and the correlation of the return of a potential asset class with that of other asset classes.

Along with the designation of an investment as an asset class comes a barometer to be able to quantify performance—the risk, return, and the correlation of the return of the asset class with that of another asset class. The barometer is called a "benchmark index," "market index," or simply "index."

[3] Mark Kritzman, "Toward Defining an Asset Class," *The Journal of Alternative Investments* (Summer 1999), p. 79.

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
- Tax and accounting issues

Client-Imposed Constraints 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.

Regulatory Constraints 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. For depository institutions and insurance companies, the amount of statutory capital required is related to the quality of the assets in which the institution has invested. There are two types of risk-based capital requirements: credit risk-based capital requirements and interest rate-risk based capital requirements. The former relates statutory capital requirements to the credit-risk associated with the assets in the portfolio. The greater the credit risk, the greater the statutory capital required. Interest rate-risk based capital requirements relate the statutory capital to how sensitive the asset or portfolio is to changes in interest rates. The greater the sensitivity, the higher the statutory capital required.

Tax and Accounting Issues Tax considerations are important for several reasons. First, in the United States, certain institutional investors such as pension funds, endowments, and foundations are exempt from federal income taxation. Consequently, the assets in which they invest will not be those that are tax-advantaged investments. Second, there are tax factors that must be incorporated into the investment policy. For example, while a pension fund might be tax-exempt, there may be certain assets or the use of some investment vehicles in which it invests whose earnings may be taxed.

Generally accepted accounting principles (GAAP) and regulatory accounting principles (RAP) are important considerations in developing investment policies. An excellent example is a defined benefit plan for a corporation. GAAP specifies that a corporate pension fund's surplus is equal to the difference between the market value of the assets and the present value of the liabilities. If the surplus is negative, the corporate sponsor must record the negative balance as a liability on its balance sheet. Consequently, in establishing its investment policies, recognition must be given to the volatility of the market value of the fund's portfolio relative to the volatility of the present value of the liabilities.

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 marketplace 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 3.

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. All of the investment tools described in the chapters that follow 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 analytical tools that will be discussed throughout this book to derive the future expected return of an asset. For example, we will see in Chapter 18 that there are various asset pricing models that provide expected return estimates based on factors that historically have been found to systematically affect the return on all assets. Investors can use historical average returns as their estimate of future expected returns. Investors can modify historical average returns with their judgment of the future to obtain a future expected return. Another approach is for investors to simply use their intuition without any formal analysis to come up with the future expected return.

In Chapter 16, the reason why the variance of asset returns should be used as a measure of an asset's risk will be explained. This input can be obtained for each asset by calculating the historical variance of asset returns. There are sophisticated time series statistical techniques that can be used to improve the estimated variance of asset returns that are discussed in Chapter 18. Some investors calculate the historical variance of asset returns and adjust them based on their intuition.

The covariance (or correlation) of returns is a measure of how the return of two assets vary together. Typically, investors use historical covariances of asset returns as an estimate of future covariances. But why is a covariance of asset returns needed? As will be explained in Chapter 16, the covariance is important because the variance of a portfolio's return depends on it and the key to diversification is the covariance of asset returns.

Approaches to Portfolio Construction

Constructing an efficient portfolio based on the expected return for a portfolio (which depends on the expected return of all the asset returns in the portfolio) and the variance of the portfolio's return (which depends on the variance of the return of all of the assets in the portfolio and the covariance of returns between all pairs of assets in the portfolio) are referred to as "mean-variance" portfolio management. The term "mean" is used because the expected return is equivalent to the "mean" or "average value" of returns. 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. Mean-variance 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 and is explained in Chapter 19 for common stock portfolio management and Chapter 21 for fixed-income portfolio management.

With either the full mean-variance approach or the multifactor risk approach there are two variations. First, the analysis can be performed by investors using individual assets (or securities) or the analysis can be performed on asset classes.

The second variation is one in which the input used to measure risk is the tracking error of a portfolio relative to a benchmark index, rather than the variance of the portfolio return. By a benchmark index it is meant the benchmark that the investor's performance is compared against.

As explained in Chapter 19, tracking error is the variance of the difference in the return on the portfolio and the return on the benchmark index. When this "tracking error multifactor risk approach" to portfolio construction is applied to individual assets, the investor can identify the set of efficient portfolios in terms of a portfolio that matches the risk profile of the benchmark index for each level of tracking error. Selecting assets that intentionally cause the portfolio's risk profile to differ from that of the benchmark index is the way a manager actively manages a portfolio. In contrast, indexing means matching the risk profile. "Enhanced" indexing basically means that the assets selected for the portfolio do not cause the risk profile of the portfolio constructed to depart materially from the risk profile of the benchmark. This tracking error multifactor risk approach to common stock and fixed-income portfolio construction will be explained and illustrated in Chapters 19 and 21, respectively.

At the other extreme of the full mean-variance approach to portfolio management is the assembling of a portfolio in which investors ignore all of the inputs—expected returns, variance of asset returns, and covariance of asset returns—and use their intuition to construct a portfolio. We refer to this approach as the "seat-of-the-pants approach" to portfolio construction. In a rising stock market, for example, this approach is too often confused with investment skill. It is not an approach we recommend.

Step 5: Measuring and Evaluating Performance

The measurement and evaluation of investment performance is the last step in the investment management process. Actually, it is misleading to say that it is the last step since the investment management process is an ongoing process. This step involves measuring the performance of the portfolio and then evaluating that performance relative to some benchmark.

Although a portfolio manager may have performed better than a benchmark, this does not necessarily mean that the portfolio manager satisfied the client's investment objective. For example, suppose that a financial institution established as its investment objective the maximization of portfolio return and allocated 75% of its funds to common stock and the balance to bonds. Suppose further that the manager responsible for the common stock portfolio realized a 1-year return that was 150 basis points greater than the benchmark.[4] Assuming that the risk of the portfolio was similar to that of the benchmark, it would appear that the manager outperformed the benchmark. However, suppose that in spite of this performance, the financial institution cannot

[4] A basis point is equal to 0.0001 or 0.01%. This means that 1% is equal to 100 basis points.

meet its liabilities. Then the failure was in establishing the investment objectives and setting policy, not the failure of the manager.

FINANCIAL ENGINEERING IN HISTORICAL PERSPECTIVE

In its modern sense, financial engineering is the design (or engineering) of contracts and portfolios of contracts that result in predetermined cash flows contingent to different events. Broadly speaking, financial engineering is used to manage investments and risk. The objective is the transfer of risk from one entity to another via appropriate contracts. Though the aggregate risk is a quantity that cannot be altered, risk can be transferred if there is a willing counterparty. Just why and how risk transfer is possible will be discussed in Chapter 23 on risk management.

Financial engineering came to the forefront of finance in the 1980s, with the broad diffusion of derivative instruments. However the concept and practice of financial engineering are quite old. Evidence of the use of sophisticated cross-border instruments of credit and payment dating from the time of the First Crusade (1095–1099) has come down to us from the letters of Jewish merchants in Cairo. The notion of the diversification of risk (central to modern risk management) and the quantification of insurance risk (a requisite for pricing insurance policies) were already understood, at least in practical terms, in the 14th century. The rich epistolary of Francesco Datini, a 14th century merchant, banker and insurer from Prato (Tuscany, Italy), contains detailed instructions to his agents on how to diversify risk and insure cargo.[5] It also gives us an idea of insurance costs: Datini charged 3.5% to insure a cargo of wool from Malaga to Pisa and 8% to insure a cargo of malmsey (sweet wine) from Genoa to Southampton, England. These, according to one of Datini's agents, were low rates: He considered 12–15% a fair insurance premium for similar cargo.

What is specific to modern financial engineering is the quantitative management of uncertainty. Both the pricing of contracts and the optimization of investments require some basic capabilities of statistical modeling of financial contingencies. It is the size, diversity, and efficiency of modern competitive markets that makes the use of modeling imperative.

[5] Datini wrote the richest medieval epistolary that has come down to us. It includes 500 ledgers and account books, 300 deeds of partnership, 400 insurance policies, and 120,000 letters. For a fascinating portrait of the business and private life of a medieval Italian merchant, see Iris Onigo, *The Merchant of Prato* (London: Penguin Books, 1963).

THE ROLE OF INFORMATION TECHNOLOGY

Advances in information technology are behind the widespread adoption of modeling in finance. The most important advance has been the enormous increase in the amount of computing power, concurrent with a steep fall in prices. Government agencies have long been using computers for economic modeling, but private firms found it economically justifiable only as of the 1980s. Back then, economic modeling was considered one of the "Grand Challenges" of computational science.[6]

In the late 1980s, firms such as Merrill Lynch began to acquire supercomputers to perform derivative pricing computations. The overall cost of these supercomputing facilities, in the range of several million dollars, limited their diffusion to the largest firms. Today, computational facilities ten times more powerful cost only of a few thousand dollars.

To place today's computing power in perspective, consider that a 1990 run-of-the-mill Cray supercomputer cost several million U.S. dollars and had a clock cycle of 4 nanoseconds (i.e., 4 billionths of a second or 250 million cycles per second, notated as 250 MHz). Today's fast laptop computers are 10 times faster with a clock cycle of 2.5 GHz and, at a few thousand dollars, cost only a fraction of the price. Supercomputer performance has itself improved significantly, with top computing speed in the range of several teraflops[7] compared to the several megaflops of a Cray supercomputer in the 1990s. In the space of 15 years, sheer performance has increased 1,000 times while the price-performance ratio has decreased by a factor of 10,000. Storage capacity has followed similar dynamics.

The diffusion of low-cost high-performance computers has allowed the broad use of numerical methods. Computations that were once performed by supercomputers in air-conditioned rooms are now routinely

[6] Kenneth Wilson, "Grand Challenges to Computational Science," *Future Generation Computer Systems 5* (1989), p. 171. The term "Grand Challenges" was coined by Kenneth Wilson, recipient of the 1982 Nobel Prize in Physics, and later adopted by the U.S. Department Of Energy (DOE) in its High Performance Communications and Computing Program which included economic modeling among the grand challenges. Wilson was awarded the Nobel Prize in Physics for discoveries he made in understanding how bulk matter undergoes "phase transition," i.e., sudden and profound structural changes. The mathematical techniques he introduced—the renormalization group theory—is one of the tools used to understand economic phase transitions. Wilson is an advocate of computational science as the "third way" of doing science, after theory and experiment.

[7] A flops (Floating Point Operations Per Second) is a measure of computational speed. A Teraflop computer is a computer able to perform a trillion floating point operations per second.

performed on desk-top machines. This has changed the landscape of financial modeling. The importance of finding closed-form solutions and the consequent search for simple models has been dramatically reduced. Computationally-intensive methods such as Monte Carlo simulations and the numerical solution of differential equations are now widely used. As a consequence, it has become feasible to represent prices and returns with relatively complex models. Nonnormal probability distributions have become commonplace in many sectors of financial modeling. It is fair to say that the key limitation of financial econometrics is now the size of available data samples or training sets, not the computations; it is the data that limits the complexity of estimates.

Mathematical modeling has also undergone major changes. Techniques such as equivalent martingale methods are being used in derivative pricing (Chapter 15) and cointegration (Chapter 11), the theory of fat-tailed processes (Chapter 13), and state-space modeling (including ARCH/GARCH and stochastic volatility models) are being used in econometrics (Chapter 11).

Powerful specialized mathematical languages and vast statistical software libraries have been developed. The ability to program sequences of statistical operations within a single programming language has been a big step forward. Software firms such as Mathematica and Mathworks, and major suppliers of statistical tools such as SAS, have created simple computer languages for the programming of complex sequences of statistical operations. This ability is key to financial econometrics which entails the analysis of large portfolios.[8]

Presently only large or specialized firms write complex applications from scratch; this is typically done to solve specific problems, often in the derivatives area. The majority of financial modelers make use of high-level software programming tools and statistical libraries. It is difficult to overestimate the advantage brought by these software tools; they cut development time and costs by orders of magnitude.

In addition, there is a wide range of off-the-shelf financial applications that can be used directly by operators who have a general understanding of the problem but no advanced statistical or mathematical training. For example, powerful complete applications from firms such as Barra and component applications from firms such as FEA make sophisticated analytical methods available to a large number of professionals.

Data have, however, remained a significant expense. The diffusion of electronic transactions has made available large amounts of data,

[8] A number of highly sophisticated statistical packages are available to economists. These packages, however, do not serve the needs of the financial econometrician who has to analyze a large number of time series.

including high-frequency data (HFD) which gives us information at the transaction level. As a result, in budgeting for financial modeling, data have become an important factor in deciding whether or not to undertake a new modeling effort.

A lot of data are now available free on the Internet. If the required granularity of data is not high, these data allow one to study the viability of models and to perform rough tuning. However, real-life applications, especially applications based on finely grained data, require data streams of a higher quality than those typically available free on the Internet.

INDUSTRY'S EVALUATION OF MODELING TOOLS

A recent study by The Intertek Group[9] tried to assess how the use of financial modeling in asset management had changed over the highly volatile period from 2000 to 2002. Participants in the study included 44 heads of asset management firms in Europe and North America; more than half were from the biggest firms in their home markets.

The study found that the role of quantitative methods in the investment decision-making process had increased at almost 75% of the firms while it had remained stable at about 15% of the firms; five reported that their process was already essentially quantitative. Demand pull and management push were among the reasons cited for the growing role of models. The head of risk management and product control at an international firm said, "There is genuinely a portfolio manager demand pull plus a top-down management push for a more systematic, robust process." Many reported that fund managers have become more eager consumers of modeling. "Fund managers now perceive that they gain increased insights from the models," the head of quantitative research at a large northern European firm commented.

In another finding, over one half of the participants evaluated that models had performed better in 2002 than two years ago; some 20% evaluated 2002 model performance to be stable with respect to the previous two years while another 20% considered that performance worsened. Performance was widely considered to be model-dependent. Among those that believed that model performance had improved, many attributed better performance to a better understanding of models and the modeling process at asset management firms. Some firms reported hav-

[9] Caroline Jonas and Sergio Focardi, *Trends in Quantitative Methods in Asset Management, 2003*, The Intertek Group, Paris, 2003.

ing in place a formal process in which management was systematically trained in modeling and mathematical methods.

The search for a silver bullet typical of the early days of "rocket science" in finance has passed; modeling is now widely perceived as an approximation, with the various models shedding different light on the same phenomena. Just under 60% of the participants in the 2002 study indicated having made significant changes to their modeling approach from 2000 to 2002; for many others, it was a question of continuously recalibrating and adapting the models to the changing environment.[10]

Much of the recent attention on quantitative methods has been focused on risk management—a relatively new function at asset management firms. More than 80% of the firms participating in the Intertek study reported a significant evolution of the role of risk management from 2000 to 2002. Some of the trends revealed by the study included daily or real-time risk measurement and the splitting of the role of risk management into two separate functions, one a support function to the fund managers, the other a central control function reporting to top management. These issues will be discussed in Chapter 23.

In another area which is a measure of an increasingly systematic process, more than 60% of the firms in the 2002 study reported having formalized procedures for integrating quantitative and qualitative input, though half mentioned that the process had not gone very far and 30% reported no formalization at all. One way the integration is being handled is through management structures for decision-making. A source at a large player in the bond market said, "We have regularly scheduled meetings where views are expressed. There is a good combination of views and numbers crunched. The mix between quantitative and qualitative input will depend on the particular situation. For example, if models are showing a 4 or 5 standard deviation event, fundamental analysis would have to be very strong before overriding the models."

Many firms have cast integration in a quantitative framework. The head of research at a large European firm said, "One year ago, the integration was totally fuzzy, but during the past year we have made the integration extremely rigorous. All managers now need to justify their statements and methods in a quantitative sense." Some firms are prioritizing the inputs from various sources. A business manager at a Swiss firm said, "We have recently put in place a scoring framework which pulls together the gut feeling of the fund manager and the quantitative

[10] Financial models are typically statistical models that have to be estimated and calibrated. The estimation and calibration of models will be discussed in Chapter 23. The above remarks reflect the fact that financial models are not "laws of nature" but relationships valid only for a limited span of time.

models. We will be taking this further. The objective is to more tightly link the various inputs, be they judgmental or model results."

Some firms see the problem as one of model performance evaluation. "The integration process is becoming more and more institutionalized," said the head of quantitative research at a big northern European firm. "Models are weighted in terms of their performance: if a model has not performed so well, its output is less influential than that of models which have performed better."

In some cases, it is the portfolio manager himself who assigns weights to the various inputs. A source at a large firm active in the bond markets said, "Portfolio managers weight the relative importance of quantitative and qualitative input in function of the security. The more complex the security, the greater the quantitative weighting; the more macro, long-term, the less the quantitative input counts: Models don't really help here." Other firms have a fixed percentage, such as 50/50, as corporate policy. Outside of quantitatively run funds, the feeling is that there is a weight limit in the range of 60–80% for quantitative input. "There will always be a technical and a tactical element," said one source.

Virtually all firms reported a partial automation in the handling of qualitative information, with some 30% planning to add functionality over and above the filtering and search functionality now typically provided by the suppliers of analyst research, consensus data and news. About 25% of the participants said that they would further automate the handling of information in 2003. The automatic summarization and analysis of news and other information available electronically was the next step for several firms that had already largely automated the investment process.

INTEGRATING QUALITATIVE AND QUANTITATIVE INFORMATION

Textual information has remained largely outside the domain of quantitative modeling, having long been considered the domain of judgment. This is now changing as financial firms begin to tackle the problem of what is commonly called *information overload*; advances in computer technology are again behind the change.[11]

Reuters publishes the equivalent of three bibles of (mostly financial) news daily; it is estimated that five new research documents come out of Wall Street every minute; asset managers at medium-sized firms report receiving up to 1,000 e-mails daily and work with as many as five

[11] Caroline Jonas and Sergio Focardi, *Leveraging Unstructured Data in Investment Management*, The Intertek Group, Paris, 2002.

screens on their desk. Conversely, there is also a lack of "digested" information. It has been estimated that only one third of the roughly 10,000 U.S. public companies are covered by meaningful Wall Street research; there are thousands of companies quoted on the U.S. exchanges with no Wall Street research at all. It is unlikely the situation is better relative to the tens of thousands of firms quoted on other exchanges throughout the world. Yet increasingly companies are providing information, including press releases and financial results, on their Web sites, adding to the more than 3.3 billion pages on the World Wide Web as of mid-2003.

Such unstructured (textual) information is progressively being transformed into self-describing, semistructured information that can be automatically categorized and searched by computers. A number of developments are making this possible. These include:

- The development of XML (eXtensible Markup Language) standards for tagging textual data. This is taking us from free text search to queries on semi-structured data.
- The development of RDF (Resource Description Framework) standards for appending metadata. This provides a description of the content of documents.
- The development of algorithms and software that generate taxonomies and perform automatic categorization and indexation.
- The development of database query functions with a high level of expressive power.
- The development of high-level text mining functionality that allows "discovery."

The emergence of standards for the handling of "meaning" is a major development. It implies that unstructured textual information, which some estimates put at 80% of all content stored in computers, will be largely replaced by semistructured information ready for machine handling at a semantic level. Today's standard structured databases store data in a prespecified format so that the position of all elementary information is known. For example, in a trading transaction, the date, the amount exchanged, the names of the stocks traded and so on are all stored in predefined fields. However, textual data such as news or research reports, do not allow such a strict structuring. To enable the computer to handle such information, a descriptive metafile is appended to each unstructured file. The descriptive metafile is a structured file that contains the description of the key information stored in the unstructured data. The result is a semistructured database made up of unstructured data plus descriptive metafiles.

Industry-specific and application-specific standards are being developed around the general-purpose XML. At the time of this writing, there are numerous initiatives established with the objective of defining XML standards for applications in finance, from time series to analyst and corporate reports and news. While it is not yet clear which of the competing efforts will emerge as the de facto standards, attempts are now being made to coordinate standardization efforts, eventually adopting the ISO 15022 central data repository as an integration point.

Technology for handling unstructured data has already made its way into the industry. Factiva, a Dow Jones-Reuters company, uses commercially available text mining software to automatically code and categorize more than 400,000 news items daily, in real time (prior to adopting the software, they manually coded and categorized some 50,000 news articles daily). Users can search the Factiva database which covers 118 countries and includes some 8,000 publications, and more than 30,000 company reports with simple intuitive queries expressed in a language close to the natural language. Suppliers such as Multex use text mining technology in their Web-based research portals for clients on the buy and sell sides. Such services typically offer classification, indexation, tagging, filtering, navigation, and search.

These technologies are helping to organize research flows. They allow to automatically aggregate, sort, and simplify information and provide the tools to compare and analyze the information. In serving to pull together material from myriad sources, these technologies will not only form the basis of an internal knowledge management system but allow to better structure the whole investment management process. Ultimately, the goal is to integrate data and text mining in applications such as fundamental research and event analysis, linking news, and financial time series.

PRINCIPLES FOR ENGINEERING A SUITE OF MODELS

Creating a suite of models to satisfy the needs of a financial firm is engineering in full earnest. It begins with a clear statement of the objectives. In the case of financial modeling, the objective is identified by the type of decision-making process that a firm wants to implement. The engineering of a suite of financial models requires that the process on which decisions are made is fully specified and that the appropriate information is supplied at every step. This statement is not as banal as it might seem.

We have now reached the stage where, in some markets, financial decision–making can be completely automated through optimizers. As we will see in the following chapters, one can define models able to construct a conditional probability distribution of returns. An optimizer will then translate the forecast into a tradable portfolio. The manager becomes a kind of high-level supervisor of an otherwise automated process.

However, not all financial decision-making applications are, or can be, fully automated. In many cases, it is the human operator who makes the decision, with models supplying the information needed to arrive at the decision. Building an effective suite of financial models requires explicit decisions as to (1) what level of automation is feasible and desirable and (2) what information or knowledge is required.

The integration of different models and of qualitative and quantitative information is a fundamental need. This calls for integration of different statistical measures and points of view. For example, an asset management firm might want to complement a portfolio optimization methodology based on Gaussian forecasting with a risk management process based on Extreme Value Theory (see Chapter 13). The two processes offer complementary views. In many cases, however, different methodologies give different results though they work on similar principles and use the same data. In these cases, integration is delicate and might run against statistical principles.

In deciding which modeling efforts to invest in, many firms have in place a sophisticated evaluation system. "We look at the return on investment [ROI] of a model: How much will it cost to buy the data necessary to run the model? Then we ask ourselves: What are the factors that are remunerated? Our decision on what data to buy and where to spend on models is made in function of what indicators are the most 'remunerated,'" commented the head of quantitative management at a major European asset management firm.

SUMMARY

- The investment management process is becoming increasingly structured; the objective is a well-defined, repeatable investment process.
- This requires measurable objectives and measurable results, financial engineering, risk control, feedback processes and, increasingly, knowledge management.
- In general, the five steps in the investment management process are setting investment objectives, establishing an investment policy, selecting an investment strategy, selecting the specific assets, and measuring and evaluating investment performance.

- Changes in the investment management business are being driven by the explosion in the universe of investable assets brought about by globalization, investors, and especially institutional investors and their consultants, pressure from regulators and the media, and the sheer size of the markets.
- Given the size, diversity, and efficiency of modern markets, a more disciplined process can be achieved only in a quantitative framework.
- Key to a quantitative framework is the measurement and management of uncertainty (i.e., risk) and financial engineering.
- Modeling is the tool to achieve these objectives; advances in information technology are the enabler.
- Unstructured textual information is progressively being transformed into self-describing, semistructured information, allowing a better structuring of the research process.
- After nearly two decades of experience with quantitative methods, market participants now more clearly perceive the benefits and the limits of modeling; given today's technology and markets, the need to better integrate qualitative and quantitative information is clearly felt.