RC24306 (W0707-086) July 11, 2007
Mathematics

IBM Research Report

Applications of OR to Finance

Aliza R. Heching, Alan J. King
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

1. Introduction

Operations Research provides a rich set of tools and techniques that are applied to financial decision making. The first topic that likely comes to mind for most readers is Markowitz's Nobel Prize winning treatment of the problem of portfolio diversification using quadratic programming techniques. This treatment, which first appeared in 1952, underlies almost all of the subsequent research into the pricing of risk in financial markets. Linear programming, of course, has been applied in many financial planning problems, from the management of working capital to formulating a bid for the underwriting of a bond issue. Less well known is the fundamental role that duality theory plays in the theoretical treatment of the pricing of options and contingent claims, both in its discrete state and time formulation using linear programming and in its continuous time counterparts. This duality leads directly to the Monte Carlo simulation method for pricing and evaluating the risk of options portfolios for investment banks; this activity probably comprises the single greatest use of computing resources in any industry.

This chapter does not cover every possible topic in the applications of Operations Research ("OR") to Finance. We have chosen to highlight the main topics in investment theory and to give an elementary, mostly self-contained, exposition of each. A comprehensive perspective of the application of OR techniques to financial markets along with an excellent bibliography of the recent literature in this area can be found in the survey by Board et al. (2003). In this chapter we chose not to cover the more traditional applications of OR to financial management for firms, such as the management of working capital, capital investment, taxation, and financial planning. For these, we direct the reader to consult Ashford et al. (1988). We also excluded financial forecasting models; the reader may refer to Campbell et al. (1997) and Mills (1999) for recent treatments of these topics. Finally, Board et al. (2003) provide a survey of the application of OR techniques for the allocation of investment budgets between a set of projects. Complete and up-to-date coverage of finance and financial engineering topics for readers in Operations Research and Management Science may be found in the handbooks of Jarrow, Maksimovic and Ziemba (1995) and the forthcoming volume of Birge and Linetsky (forthcoming, 2007).

We begin this chapter by introducing some basic concepts in investment theory. In Section 2 we present the formulas for computing the return and variance of return on a portfolio. The formulas for a portfolio's mean and variance presume that these parameters are known for the individual assets in the portfolio. In Section 3 we discuss two methods for estimating these parameters when they are not known.

Section 4 explains how a portfolio's overall risk can be reduced by including a diverse set of assets in the portfolio. In Section 5 we introduce the risk-reward tradeoff efficient frontier and the Markowitz problem. Up to this point, we have assumed that the investor is able to specify a mathematical function describing his attitude toward risk. In Section 6, we consider utility theory which does not require an explicit specification of a risk function. Instead, utility theory assumes that investors specify a utility, or *satisfaction*, with any cash payout. The associated optimal portfolio selection problem will seek to maximize the investor's expected utility.

Section 7 discusses the Black-Litterman model for asset allocation. Black and Litterman use Bayesian updating to combine historical asset returns with individual investor views to determine a posterior distribution on asset returns which is used to make asset allocation decisions. Section 8 considers the challenges of risk management. We introduce the notion of coherent risk measures and conditional value-at-risk (*CVaR*), and show how a portfolio selection problem with a constraint on *CVaR* can be formulated as a stochastic program.

In Sections 9 through 13 we turn to the problem of options valuation. Options valuation combines a mathematical model for the behavior of the underlying uncertain market factors with simulation or dynamic programming (or combinations thereof) to determine options prices. Section 14 considers the problem of asset-liability matching in a multi-period setting. The solution uses stochastic optimization based upon Monte Carlo simulation. Finally, in Section 15 we present some concluding remarks.

2. Return

Suppose that an investor invests in asset *i* at time 0 and sells the asset at time *t*. The rate of return (more simply referred to as the return) on asset *i* over time period *t* is given by:

$$r_i = \frac{amount\ received\ at\ time\ \ t - amount\ invested\ in\ asset\ i\ at\ time\ 0}{amount\ invested\ in\ asset\ i\ at\ time\ 0}. \qquad (1)$$

Now suppose that an investor invests in a portfolio of N assets. Let f_i denote the fraction of the portfolio that is comprised of asset *i*. Assuming that no short selling is allowed, $f_i \geq 0$. Clearly, $\sum_{i=1}^{N} f_i = 1$.

The portfolio return is given by the weighted sum of the returns on the individual assets in the portfolio:

$$r_p = \sum_{i=1}^{N} f_i r_i \ . \tag{2}$$

We have described asset returns as if they are known with certainty. However, there is typically uncertainty surrounding the amount that will be received at the time that an asset is sold. We can use a probability distribution to describe this unknown rate of return. If return is normally distributed, then only two parameters – its expected return and its standard deviation (or variance) - are needed to describe this distribution. The expected return is the return around which the probability distribution is centered; it is the expected value of the probability distribution of possible returns. The standard deviation describes the dispersion of the distribution of possible returns.

2.1. Expected Portfolio Return

Suppose there are N assets with random returns $r_1, \ldots, r_N$. The corresponding expected returns are $E(r_1), \ldots, E(r_N)$. An investor wishes to create a portfolio of these N assets, by investing a fraction f_i of his wealth in asset i.

Using the properties of expectation, we may compute the expected portfolio return using equation (2):

$$E(r_p) = \sum_{i=1}^{N} f_i E(r_i) \, ,$$

i.e., the expected portfolio return is equal to the weighted sum of the expected returns of its individual asset components.

2.2. Portfolio Variance

The volatility of an asset's return can be measured by its variance. Variance is often adopted as a measure of an asset's risk. If σ_i^2 denotes the variance of asset *i*'s return, then the variance of the portfolio's return is given by:

$$\begin{aligned}
\sigma_p^2 &= E\left[\left(\sum_{i=1}^{N} f_i r_i - \sum_{i=1}^{N} f_i \overline{r_i}\right)^2\right] \\
&= E\left[\left(\sum_{i=1}^{N} f_i (r_i - \overline{r_i})\right)\left(\sum_{j=1}^{N} f_j (r_j - \overline{r_j})\right)\right] \\
&= E\left[\sum_{i,j=1}^{N} f_i f_j (r_i - \overline{r_i})(r_j - \overline{r_j})\right] \\
&= \sum_{i,j=1}^{N} f_i f_j \sigma_{ij} \\
&= \sum_{i=1}^{N} f_i^2 \sigma_i^2 + \sum_{i=1}^{N} \sum_{\substack{j=1 \\ j \neq i}}^{N} f_i f_j \sigma_{ij}
\end{aligned} \quad , \qquad (3)$$

where $\overline{r_i} = E(r_i)$. Note that portfolio variance is a combination of the variance of the returns of each individual asset in the portfolio plus their covariance.

3. Estimating an asset's mean and variance

Of course, asset *i*'s rate of return and variance are not known and must be estimated. These values can be estimated based upon historical data using standard statistical methods. Alternatively, one can use a scenario-based approach. We describe the two methods below.

3.1. Estimating statistics using historical data

To estimate statistics using historical data, one must collect several periods of historical returns on the assets. The estimated average return on asset *i* is computed as the sample average of returns on asset *i*, $\overline{X}_i$

$$\overline{X}_i = \frac{\sum_t x_{it}}{T},$$

where x_{it} is the historical return on asset i in period t and there are T periods of historical data.

The variance of return on asset i is estimated by s_i^2, the historical sample variance of returns on investment i:

$$s_i^2 = \frac{\sum_t (x_{it} - \overline{X}_i)^2}{T-1}.$$

For example, **Table 1** contains the monthly closing stock prices and monthly returns for Sun Microsystems and Continental Airlines for the months January 2004 through February 2006. The first column of this table indicates the month, the second and third columns contain the closing stock prices for SUN and Continental Airlines, respectively. The fourth and fifth columns contain the monthly stock returns for SUN ($X_{SUN,t}$) and Continental Airlines ($X_{CAL,t}$), respectively; these columns were populated using Equation (1). In this example, $T=26$.

Month	SUN Stock Price	CAL Stock Price	SUN Return	CAL Return
Jan-04	55.55	15.65	8.12%	-5.04%
Feb-04	61.66	14.97	11.00%	-4.35%
Mar-04	62.38	12.58	1.17%	-15.97%
Apr-04	63.00	10.66	0.99%	-15.26%
May-04	62.04	10.50	-1.52%	-1.50%
Jun-04	63.63	11.37	2.56%	8.29%
Jul-04	68.17	8.75	7.13%	-23.04%
Aug-04	61.60	9.65	-9.64%	10.29%
Sep-04	73.98	8.60	20.10%	-10.88%
Oct-04	74.76	9.24	1.05%	7.44%
Nov-04	82.45	11.16	10.29%	20.78%
Dec-04	81.71	13.75	-0.90%	23.21%

Jan-05	88.04	10.50	7.75%	-23.64%
Feb-05	98.80	11.14	12.22%	6.10%
Mar-05	105.00	12.00	6.28%	7.72%
Apr-05	98.74	12.00	-5.96%	0.00%
May-05	103.00	13.70	4.31%	14.17%
Jun-05	114.14	13.32	10.82%	-2.77%
Jul-05	126.33	15.81	10.68%	18.69%
Aug-05	75.35	13.28	-40.35%	-16.00%
Sep-05	78.20	9.70	3.78%	-26.96%
Oct-05	73.76	13.20	-5.68%	36.08%
Nov-05	78.10	15.59	5.88%	18.11%
Dec-05	79.34	21.27	1.59%	36.43%
Jan-06	95.20	20.50	19.99%	-3.62%
Feb-06	74.35	23.79	-21.90%	16.05%

Table 1: Monthly closing stock prices and returns for Sun Microsystems and Continental Airlines

Table 2 shows the mean and standard deviation of returns for these two stocks, based upon the 26 months of historical data. The average monthly return for SUN, $\overline{X}_{SUN}$, and the average monthly return for Continental, $\overline{X}_{CAL}$, is computed as the arithmetic average of the monthly returns in the fourth and fifth columns, respectively. An estimate of the variance of monthly return on SUN's (Continental's) stock, is computed as the variance of the returns in the fourth (fifth) column of Table 1. If variance is used as a measure of risk, then Continental is a riskier investment since it has a higher volatility (its variance is higher).

	Expected monthly return	Variance of return	Standard Deviation
SUN	2.30%	1.54	12.40%
Continental Airlines	2.86%	3.04	17.43%

Table 2: Expected historical monthly returns, variances, and standard deviations of returns

3.2. The scenario approach to estimating statistics

Sometimes, historical market conditions are not considered a good predictor of future market conditions. In this case, historical data may not be a good source for estimating expected returns or risk. When historical estimates are determined to be poor predictors of the future, one can consider a scenario approach.

The scenario approach proceeds as follows:

Define a set of *S* future economic scenarios and assign likelihood $p(s)$ that scenario *s* will occur. $\sum_{s \in S} p(s) = 1$, since in the future the economy must be in exactly one of these economic conditions. Next, define each asset's behavior (its return) under each of the defined economic scenarios. Asset *i*'s expected return is computed as:

$$r_i = \sum_s p(s) r_i(s), \tag{4}$$

where $r_i(s)$ is asset *i*'s return under scenario *s*.

Similarly, we compute the variance of return on asset *i* as:

$$v_i = \sum_s p(s)(r_i(s) - r_i)^2 . \tag{5}$$

For example, suppose we use the scenario approach to predict expected monthly return on SUN stock. We have determined that the economy may be in one of three states: weak, stable, or strong, with a likelihood of 0.3, 0.45, and 0.25, respectively. Table 3 indicates the forecasted monthly stock returns under each of these future economic conditions:

Scenario (s)	Likelihood ($p(s)$)	Return($r_i(s)$)	$p(s) * r_i(s)$	$p(s) * (r_i(s) - r_i)^2$
Weak economy	0.30	-15.00%	-4.50%	1.03
Stable economy	0.45	9.00%	4.05%	0.13
Strong economy	0.25	16.00%	4.00%	0.39
			3.55%	1.55

Table 3: Definition of future possible scenarios for states of the economy

The first and second columns in **Table 3** indicate the economic scenario and likelihood that the scenario will occur, respectively. The third column contains the expected return under each of the defined future scenarios. The fourth and fifth columns contain intermediate computations needed to calculate the expected return and standard deviation of returns on SUN stock, based upon Equations (3) and (4). Using these equations we find that the expected monthly return on SUN stock is 3.55%, variance of monthly return is 1.55, with corresponding standard deviation of 12.46%.

Comparing the estimates of mean and standard deviation of SUN's monthly return using the historical data approach versus the scenario-based approach we find that while the estimates of volatility of return are close in value, the estimates of monthly return differ significantly. (2.3% versus 3.55%.) In Section 7 we discuss the negative impact that can result from portfolio allocation based upon incorrect parameter estimation. Thus, care must be taken to determine the correct method and assumptions when estimating these values.

4. Diversification

We now explore how a portfolio's risk, as measured by the variance of the portfolio return, can be reduced when stocks are added to the portfolio. This phenomenon, whereby a portfolio's risk is reduced when assets are added to the portfolio, is known as diversification.

Portfolio return is the weighted average of the returns of the assets in the portfolio, weighted by their appearance in the portfolio. However, portfolio variance (as derived in equation (3)) is given by:

$$\sum_{i=1}^{N} f_i^2 \sigma_i^2 + \sum_{i=1}^{N} \sum_{\substack{j=1 \\ j \neq i}}^{N} f_i f_j \sigma_{ij} .$$

Namely, portfolio variance is comprised of two components. One component is the variances of the individual assets in the portfolio and the other component is the covariance between the returns on the different assets in the portfolio. Covariance of return between asset *i* and *j* is the expected value of the product of the deviations of each of the assets from their respective means. If the two assets deviate from their respective means in identical fashions (i.e., both are above their means or below their means at the same time) then the covariance is highly positive, and its contribution to portfolio variance is highly positive. If the return on one asset deviates below its mean at the time that the return on the other asset deviates above its mean, then the covariance is highly negative which reduces the overall portfolio variance.

Let us explore the impact of the covariance term on overall portfolio variance. If all of the assets are independent, then the covariance terms equal zero and only the variance of the individual assets in the portfolio contribute to overall portfolio variance. In this case, the variance formula is:

$$\sum_{i=1}^{N} f_i^2 \sigma_i^2 .$$

If the investor invests an equal amount in each of the N independent assets, then the portfolio variance is:

$$\sum_{i=1}^{N} (\frac{1}{N})^2 \sigma_i^2 = \frac{1}{N} \sum_{i=1}^{N} \frac{1}{N} \sigma_i^2 = \frac{1}{N} \overline{\sigma}_i^2 ,$$

where $\overline{\sigma}_i^2$ is the average variance of the assets in the portfolio. As N gets larger, the portfolio variance goes to zero. Thus, for a portfolio comprised only of independent assets, when the number of assets in the portfolio is large enough the variance of the portfolio return is zero.

Now consider N assets that are not independent. Without loss of generality, assume the assets appear in the portfolio with equal weight. Then, variance of portfolio return is:

$$
\begin{aligned}
&\frac{1}{N}\overline{\sigma}_i^2 + \sum_{i=1}^{N}\sum_{\substack{j=1\\ j\neq i}}^{N}\left(\frac{1}{N}\right)^2 \sigma_{ij} \\
&= \frac{1}{N}\overline{\sigma}_i^2 + \frac{N-1}{N}\sum_{i=1}^{N}\sum_{\substack{j=1\\ j\neq i}}^{N}\frac{\sigma_{ij}}{N(N-1)} \\
&= \frac{1}{N}\overline{\sigma}_i^2 + \frac{N-1}{N}\overline{\sigma}_{ij}
\end{aligned}
\qquad (6)
$$

The final equality in (6) reveals that as the number of assets in the portfolio increases, the contribution of the first component (variance) becomes negligible – it is diversified away – and the second term (covariance) approaches average covariance. Thus, while increasing the number of assets in a portfolio will diversify away the individual risk of the assets, the risk attributed to the covariance terms cannot be diversified away.

As a numerical example, consider the SUN and Continental Airlines monthly stock returns from Section 3.1. Average monthly return on Sun Microsystems (Continental Airlines) stock from January 2004 through February 2006 was 2.3% (2.86%). A portfolio comprised of equal investments in Sun Microsystems and Continental Airlines yields an average monthly return of 2.58%. The variance of the monthly returns on Sun Microsystems stock over the 26 months considered is 1.54. The variance of the monthly returns on Continental Airlines stock over that same time period is 3.04. However, for the portfolio consisting of equal investments in Sun Microsystems and Continental Airlines, the variance of portfolio returns is 1.09. This represents a significant reduction in risk from the risk associated with either of the stocks alone. The explanation lies in the value of the covariance between the monthly returns on these two stocks; the value of the covariance is -0.12. Because the returns on the stocks have a negative covariance, diversification reduces the portfolio risk.

5. Efficient frontier

The discussion in Section 4 illustrates the potential benefits of combining assets in a portfolio. For a *risk averse* investor, diversification provides the opportunity to reduce portfolio risk while maintaining a minimum level of return. An investor can consider different combinations of assets, each of which has an associated risk and return.

This naturally leads one to question whether there is an optimal way to combine assets. We address this question within the context of assuming that: (i) for a fixed return investors prefer the lowest possible risk, that is, investors are risk averse, (ii) for a given level of risk, investors always prefer the highest possible return. This property is referred to as nonsatiation, and (iii) the first two moments of the distribution of an asset's return are sufficient to describe the asset's character; there is no need to gather information about higher moments such as skew.

Given these assumptions, the risk-return trade-off of portfolios of assets can be graphically displayed by constructing a plot with risk (as measured by standard deviation) on the horizontal axis and return on the vertical axis. We can plot every possible portfolio on this risk-return space. The set of all portfolios plotted form a *feasible region*. By the risk averse assumption, for a given level of return investors prefer portfolios that lie as far to the left as possible, since these have the lowest risk. Similarly, by the nonsatiation property for a given level of risk investors prefer portfolios that lie higher on the graph since these yield a greater return. The upper left perimeter of the feasible region is called the *efficient frontier*. It represents the least risk combination for a given level of return. The efficient frontier is concave.

Figure 1 shows a mean-standard deviation plot for 10,000 random portfolios created from the thirty stocks in the Dow from 1986 through 1991. It represents the feasible region of portfolios. **Figure 2** is the efficient frontier for the stocks from the Dow Industrials from 1985 through 1991. (Both figures were taken from the NEOS Server for Optimization website.)

Figure 1: Return vs. standard deviation for 10,000 random portfolios from the Dow Industrials

Figure 2: Efficient frontier for Dow Industrials from 1985-1991

5.1. Risk-reward portfolio optimization

Markowitz (1952) developed a single period portfolio choice problem where the objective is to minimize portfolio risk (variance) for a specified return on the portfolio. For this model, it is assumed that all relevant information required by investors to make portfolio decisions is captured in the mean, standard deviation, and correlation of assets. This method for portfolio selection is often referred to as mean-variance optimization since it trades off an investor's desire for higher mean return against his aversion to greater risk as measured by portfolio variance.

The Markowitz model for portfolio optimization is given by:

Minimize $\sum_{i,j} f_i f_j \sigma_{ij}$

Subject to: $\sum_{i=1}^{N} f_i r_i = \underline{R}$ (7)

$$\sum_{i=1}^{N} f_i = 1$$

The objective is to minimize variance subject to two constraints: (i) portfolio return must equal the targeted return R̲ and (ii) total allocation must equal 1. Negative values for f_i correspond to short selling.

One method that can be employed to solve this constrained optimization problem is to form an auxiliary function L called the Lagrangian by (i) rearranging each constraint so that the right hand side equals to zero and (ii) introducing a single Lagrange multiplier for each constraint in the problem as follows:

$$L = \sum_{i,j} f_i f_j \sigma_{ij} \rho_{ij} - \lambda\left(\sum_{i=1}^{N} f_i r_i - \underline{R}\right) - \mu\left(\sum_{i=1}^{N} f_i - 1\right). \tag{8}$$

We now treat the Lagrangian (8) as an unconstrained minimization problem. A necessary condition for a point to be optimal is that the partial derivates of the Lagrangian with respect to each of the variables must equal zero. Thus, we take the partial derivative of L with respect to each of the i asset weights f_i, μ, and λ and set each partial derivative equal to zero. (Notice that the partial derivative of L with respect to λ yields the portfolio return constraint, and the partial derivative of L with respect to μ yields the constraint on the asset weights.) The result is a system of $i+2$ constraints. We use this set of constraints to solve for the $i+2$ unknowns: f_i, μ, and λ.

The Lagrangian method can often only be successfully implemented for small problems. For larger problems with many variables it may be virtually impossible to solve the set of equations for the $i+2$ unknowns. Instead, one can find the solution to Problem (7) using methods developed for optimizing quadratic programs. A *quadratic* *program* is a mathematical optimization model where the objective function is quadratic and all constraints are linear equalities or inequalities. The constraints of the optimization problem define the *feasible region* within which the optimal solution must lie. Quadratic programs are, in general, difficult to solve. Quadratic programming solution methods work in two phases. In the first phase a feasible solution is found. In the second phase, the method searches along the edges and surfaces of the feasible region for another solution that improves upon the current feasible solution. Unless the objective function is convex, the method will often identify a local optimal solution.

The solution to the Markowitz problem yields a point that lies on the efficient frontier. By varying the value of $\underline{R}$ in Problem (7), one can map out the entire efficient frontier.

6. Utility analysis

The solution to the Markowitz problem provides one means for making investment decisions in the mean-variance space. It requires the investor to define a measure of risk and a measure of value and then utilizes an explicitly defined trade-off between these two measures to determine the investor's preference decisions. Utility theory provides an alternative way to establish an investor's preferences without explicitly defining risk functions.

Utility describes an investor's attitude toward risk by translating the investor's satisfaction associated with different cash payouts into a utility value. The application of utility to uncertain financial situations was first introduced by von Neumann and Morgenstern (1944). Utility functions can be used to explain how investors make choices between different portfolios.

Utility theory is often introduced by way of the concept of *certainty equivalent*. Certainty equivalent is the amount of wealth that is equally preferred to an uncertain alternative. Or, this certain amount has the same *utility* as the uncertain alternative. Risk averse investors will prefer a lower certain cash payout to a higher risky cash payout. That is, their certainty equivalent is lower than the expected value of uncertain alternatives. Risk seeking individuals have a certainty equivalent that is higher than the expected value of the uncertain alternatives. A utility function captures these attitudes toward risk. It assigns different weights to different outcomes according to the risk profile of the individual investor. The shape of a utility function is defined by the risk profile of the investor.

Each individual investor may have a different utility function, since each investor may have a different attitude toward risk. However, all utility functions U satisfy the following properties:

- *Nonsatiation*: Utility functions must be consistent with more being preferred to less. If x and y are two cash payouts and $x>y$, then $U(x)>U(y)$. This property is equivalent to stating that the first derivative of the utility function, with respect to cash payout, is positive.
- *Risk preference*: Economic theory presumes that an investor will seek to maximize the utility of his investment. However, all investors will not make identical investment decisions because they will not all share the same attitudes toward risk. Investors can be classified into three classes, according to their willingness to accept risk: risk averse, risk neutral, or risk taking. Risk averse investors invest in investments where the utility of the return exceeds the risk-free

rate. If no such investment exists, the investor invests in risk-free investments. The utility function of a risk averse investor is increasing and concave in the cash payout ($U''(x) < 0$); the value assigned to each additional dollar received decreases due to the risk averse nature of the individual. Risk-neutral investors ignore risk when making investment decisions. They seek investments with a maximum return, irrespective of the risk involved. The utility function of a risk neutral individual is linear increasing in the cash payout; the same value is assigned to each additional dollar. Risk-taking investors are more likely to invest in investments with a higher risk involved. The utility function of a risk-taking individual is convex ($U''(x) > 0$).

- *Changes in wealth*: Utility functions also define how an investor's investment decisions will be affected by changes in his wealth. Specifically, if an investor has a larger amount of capital will this change his willingness to accept risk? The Arrow-Pratt absolute risk aversion coefficient, given by

$$A(x) = -\frac{U''(x)}{U'(x)}$$

is a measure of an investor's absolute risk aversion. $A'(x)$ measures changes in an investor's absolute risk aversion as a function of changes in his wealth. If $A'(x) > 0$ ($A'(x) < 0$), the investor has increasing (decreasing) absolute risk aversion and he will hold fewer (more) dollars in risky assets as his wealth increases. If $A'(x) = 0$, the investor has constant absolute risk aversion and he will hold the same dollar amount of risky assets as his wealth increases.

Unlike mean-variance optimization, where the cardinal value of the objective function is meaningful, a utility function only *ranks* alternatives according to risk preferences; its numerical value has no real meaning. Thus, utility functions are unique up to a positive affine transformation. More specifically, if $U(x)$ is a utility function then $V(x)=a+bU(x)$ $(b>0)$ will yield the same rankings (and hence the same investment decisions) as $U(x)$.

6.1. Utility functions

Although each investor can define his own utility function, there are a number of predefined utility functions that are commonly used in the finance and economics literature. In this section we describe the exponential utility function. We also reconcile between mean-variance optimization and utility optimization.

The exponential utility function is often adopted for financial analysis. The exponential utility function is defined as:

$$U(x) = 1 - e^{-\frac{x}{R}}, \text{ for all } x \tag{9}$$

where x is the wealth and $R>0$ represents the investor's risk tolerance. Greater values of R mean that the investor is less risk averse. (In fact, as $R\to\infty$ the investor becomes risk neutral.) $U'(x) = \frac{1}{R}e^{-\frac{x}{R}}$ and $U''(x) = -(\frac{1}{R})^2 e^{-\frac{x}{R}}$. The second derivative of (9) is strictly negative so the exponential utility function is concave; the exponential utility function describes the behavior of a risk averse investor. The absolute risk aversion is $(\frac{1}{R})$, which is constant with wealth; the investor invests constant dollar amounts in risky investments as his wealth increases.

We now construct the relationship between utility and mean-variance for concave (i.e., risk-averse) investors. The relationship can be constructed following King (1993). Suppose *U(X)* is a strictly concave utility function. The Taylor expansion of *U(X)* is an approximation of the function at a particular point, say *M*, using its derivatives. The second order Taylor expansion of *U(X)* about point *M* is given by:

$$U(X) = U(M) + U'(M)(X - M) + \frac{1}{2}U''(M)(X - M)^2 . \qquad (10)$$

Now suppose that the point *M* is equal to expected wealth, i.e., *E(X)=M*. Then the expected value of the second-order Taylor expansion expression (10) is equal to

$$\begin{aligned} E[U(X)] &= U(M) + \frac{1}{2}U''(M)E(X - M)^2 \\ &= U(M) + \frac{1}{2}U''(M)\sigma^2 \end{aligned}$$

(The middle term drops out since *E(X)=M*.) The second derivative is negative for a strictly concave utility function. This implies that maximizing the second order Taylor expansion is equivalent to

$$\min E(X - M)^2 = \sigma^2$$

for all *X* with *E(X)=M*.

However, this is equivalent to the mean-variance problem with a given mean return *M*. It follows that mean-variance is a second-order approximation to utility maximization. An investor who makes investment decisions by maximizing his expected utility is equivalently performing a mean-variance optimization. Of course, due to the two-sided nature of the variance, eventually this approximation will become negatively-sloped – and hence not really valid as a utility function – as *X* increases in value. However, the range over which the approximation is valid can be pretty wide. The upper bound of the range is the point where $U'(X)+U''(X)(X-M)=0$, in other words this is the maximum point of the approximating quadratic. This range can be quite wide enough in practice. For example, for the logarithmic utility $U(X)=\log(X)$, the upper bound of the range where the mean-variance approximation remains quadratic is $X=2M$. In other words, the mean-variance approximation for the logarithm is valid for a range that includes twice the mean value of the return!

6.2. Utility in Practice

In practice, utility is not often used as an objective criterion for investment decision making because utility curves are difficult to estimate. However, Holmer (1998) reports that Fannie Mae uses expected utility to optimize its portfolio of assets and liabilities. Fannie Mae faces a somewhat unique set of risks due to the specific nature of its business. Fannie Mae buys mortgages on the secondary market, pools them, and then sells them on the open market to investors as mortgaged backed securities. Fannie Mae faces many risks such as: prepayment risk, risks due to potential gaps between interest due and interest owed, and long term asset and liability risks due to interest rate movements. Utility maximization allows Fannie Mae to explicitly consider degrees of risk aversion against expected return to determine its risk adjusted optimal investment portfolio.

7. Black - Litterman Asset Allocation Model

Markowitz mean-variance portfolio optimization requires mean and covariance as input and outputs optimal portfolio weights. The method has been criticized because:

(i) The optimal portfolio weights are highly dependent upon the input values. However, it is difficult to accurately estimate these input values. Chopra and Ziemba (1993), Kallberg and Ziemba (1981, 1984), and Michaud (1989) use simulation to demonstrate the significant cash-equivalent losses due to incorrect estimates of the mean. Bengtsson (2004) showed that incorrect estimates of variance and covariance also have a significant negative impact on cash returns.

(ii) Markowitz mean-variance optimization requires the investor to specify the universe of return values. It is unreasonable to expect an investor to know the universe of returns. On the other hand, mean-variance optimization is sensitive to the input values so incorrect estimation can significantly skew the results.

(iii) Black and Litterman (1992) and He and Litterman (1999) have studied the optimal Markowitz model portfolio weights assuming different methods for estimating the assets means' and found that the resulting portfolio weights were unnatural. Unconstrained mean-variance optimization typically yields an optimal portfolio that takes many large long and short positions. Constrained mean-variance optimization often results in an extreme portfolio that is highly concentrated in a small number of assets. Neither of these portfolio profiles is typically considered acceptable to investors.

(iv) Due to the intricate interaction between mean and variance, the optimal weights determined by Markowitz's mean-variance estimation are often non-intuitive. A small change in an estimated mean of a single asset can drastically change the weights of many assets in the optimal portfolio.

Black and Litterman observed the potential benefit of using mathematical optimization for portfolio decision-making, yet understood investment manager's hesitations in implementing Markowitz's mean-variance optimization model. Black and Litterman (1992) developed a Bayesian method for combining individual investor subjective views on asset performance with market equilibrium returns to create a mixed estimate of expected returns. The Bayes approach works by combining a prior belief with additional information to create an updated "posterior" distribution of expected asset returns. In the Black-Litterman framework the equilibrium returns are the prior and investor subjective views are the additional information. Together, these form a posterior distribution on expected asset returns. These expected returns can then be used to make asset allocation decisions. If the investor has no subjective views on asset performance, then the optimal allocation decision is determined solely according to the market equilibrium returns. Only if the investor expresses opinions on specific assets will the weights for those assets shift away from the market equilibrium weights in the direction of the investor's beliefs.

The Black-Litterman model is based on taking a market equilibrium perspective on asset returns. Asset "prior" returns are derived from the market capitalization weights of the optimal holdings of a mean-variance investor, given historical variance. Then, if the investor has specific views on the performance of any assets, the model combines the equilibrium returns with these views, taking into consideration the level of confidence that the investor associates with each of the views. The model then yields a set of updated expected asset returns as well as updated optimal portfolio weights, updated according to the views expressed by the investor.

The key inputs to the Black-Litterman model are market equilibrium returns and the investor views. We now consider these inputs in more detail.

7.1. Market equilibrium returns

Black and Litterman use the market equilibrium expected returns, or CAPM returns, as a neutral starting point in their model. (See, e.g., Sharpe (1964).) The basic assumptions are that (i) security markets are frictionless, (ii) investors have full information relevant to security prices, and (iii) all investors process the information as if they were mean-variance investors. The starting point for the development of the CAPM is to form the efficient frontier for the market portfolios and to draw the Capital Market Line (CML). The CML begins at the risk free rate on the vertical axis (which has risk 0) and is exactly tangent to the efficient frontier. The point where the CML touches the efficient frontier is the pair (σ_m, r_m), which is defined to be the "market" standard deviation and "market" expected return. By changing the relative proportions of riskless asset and market portfolio, an investor can obtain any combination of risk and return that lies on the CML. Because the market point is tangent to the line, there are no other combinations of risky and riskless assets that can provide better expected returns for a given level of risk. Now consider a particular investor portfolio *i* with expected return $E(r_i)$ and standard deviation σ_i. For an investor to choose to hold this portfolio it must have returns comparable to the returns that lie on the Capital Market Line. Thus, the following must hold:

$$E(r_i) = r_f + \frac{\sigma_i}{\sigma_m}(r_m - r_f)$$

where r_f is the risk free rate. This equation is called the Capital Asset Pricing Model, or CAPM.

The interpretation of the CAPM is that investor's portfolios have an expected return that includes a reward for taking on risk. This reward, by the CAPM hypothesis, must be equal to the return that would be obtained from holding a portfolio on the Capital Market Line that has an equivalent risk. Any remaining risk in the portfolio can be diversified away (by, for example, holding the market portfolio) so the investor does not gain any reward for the non-systematic, or diversifiable, risk of the portfolio. The CAPM argument applied to individual securities implies that the holders of individual securities will be compensated only for that part of the risk that is correlated with the market, or the so-called systematic risk. For an individual security j the CAPM relationship is

$$E(r_j) = r_f + \beta_j (r_m - r_f)$$

where $\beta_j = \frac{\sigma_j}{\sigma_m} \rho_{mj}$ and ρ_{mj} is the correlation between asset j returns and the market returns.

The Black-Litterman approach uses the CAPM in reverse, by assuming that in equilibrium the market portfolio is held by mean-variance investors and by using optimization to back out the expected returns that such investors would require given their observed holdings of risky assets. Let N denote the number of assets and let the excess equilibrium market returns (above the risk free rate) be defined by:

$$\Pi = \lambda \Sigma w, \tag{11}$$

where

$\Pi = Nx1$ vector of implied excess returns

Σ = NxN covariance matrix of returns

w = $Nx1$ vector of market capitalization weights of the assets

λ = risk aversion coefficient that characterizes the expected risk-reward tradeoff.

λ is the price of risk as it measures how risk and reward can be traded off when making portfolio choices. It measures the rate at which an investor will forego expected return for less variance. λ is calculated as $\lambda=(r_m - r_f)/\sigma^2_m$, where σ^2_m is the variance of the market return. The elements of the covariance matrix are computed using historical correlations and standard deviations. Market CAP weights are determined by measuring the dollar value of the global holdings of all equity investors in the large public stock exchanges. The CAP weight of a single equity name is the dollar-weighted market-closing value of its equity share times the outstanding shares issued. Later we will show that Equation (11) is used to determine the optimal portfolio weights in the Black-Litterman model.

7.2. Investor views

The second key input to the Black-Litterman model is individual investor views.

Assume that an investor has K views, denoted by a $Kx1$ vector Q. Uncertainty regarding these views is denoted by an error term ε, where ε is normally distributed with mean zero and KxK covariance matrix Ω. Thus, a view has the form:

$$Q + \varepsilon = \begin{bmatrix} Q_1 \\ \vdots \\ \\ Q_K \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \vdots \\ \\ \varepsilon_K \end{bmatrix}.$$

$\varepsilon=0$ means that the investor is 100% confident about his view; in the more likely case that the investor is uncertain about his view, ε takes on some positive or negative value.

ω denotes the variance of each error term. We assume that the error terms are independent of each other. (This assumption can be relaxed.) Thus, the covariance matrix Ω is a diagonal matrix where the elements on the diagonal are ω, the variances of each error term. A higher variance indicates greater investor uncertainty with the associated view. The error terms ε do not enter directly into the Black-Litterman formula; only their variances enter via the covariance matrix Ω.

$$\Omega = \begin{bmatrix} \omega_1 & 0 & \cdots & & 0 \\ 0 & \ddots & & & \vdots \\ \vdots & & & & \\ & & & & 0 \\ 0 & \cdots & & 0 & \omega_K \end{bmatrix}.$$

The Black-Litterman model allows investors to express views such as:

View 1: Asset A will have an excess return of 5.5% with 40% confidence.

View 2: Asset B will outperform asset C by 3% with 15% confidence.

View 3: Asset D will outperform assets E and F by 1% with 20% confidence.

The first view is called an *absolute* view, while the second and third views are called *relative* views. Notice that the investor assigns a level of confidence to each view.

Each view can be seen as a portfolio of long and short positions. If the view is an absolute view then the portfolio position will be long. If the view is a relative view, then the portfolio will take a long position in the asset that is expected to “overperform” and a short position in the asset that is expected to “underperform.” In general, the impact on the optimal portfolio weights is determined by comparing the equilibrium difference in performance of these assets to the performance expressed by the investor view. If the performance expressed in the view is better than the equilibrium performance, the model will tilt the portfolio toward the outperforming asset. More specifically, consider View 2 which states that asset B will outperform asset C by 3%. If the equilibrium returns indicate that asset B will outperform asset C by more than 3% then the view represents a weakening view in performance of asset B and the model will tilt the portfolio away from asset B.

One of the most challenging questions in applying the Black-Litterman model is how to populate the covariance matrix *Ω* and how to translate the user specified expressions of confidence into uncertainty in the views. We will discuss this further below.

7.3. An example of investor view

We now illustrate how one builds the inputs for the Black-Litterman model, given the three views expressed. Suppose that there are *N=7* assets: Assets A-H. The *Q* matrix is given by:

$$Q = \begin{bmatrix} 5.5 \\ 3 \\ 1 \end{bmatrix}.$$

Note that the investor only has views on six of the seven assets. We use the matrix *P* to match the views to the individual assets. Each view results in a *1xN* vector so that *P* is a *KxN* matrix. In our case, where there are seven assets and three views, *P* is a *3x7* matrix. Each column corresponds to one of the assets; column 1 corresponds to

Asset A, column 2 corresponds to Asset B, etc. In the case of absolute views, the sum of the elements in the row equals 1. In our case, View 1 yields the vector:

$P_1 = [1 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0]$.

In the case of relative views, the sum of the elements equals zero. Elements corresponding to relatively outperforming assets have positive values; elements corresponding to relatively underperforming assets take negative values. We determine the values of the individual elements by dividing 1 by the number of outperforming and underperforming assets, respectively. For view 2, we have one outperforming asset and one underperforming asset. Thus, View 2 yields the vector:

$P_2 = [0 \quad 1 \quad -1 \quad 0 \quad 0 \quad 0 \quad 0]$.

View 3 has one outperforming asset (Asset D) and two relatively underperforming assets (Assets E and F). Thus, Asset D is assigned a value of +1 and Assets E and F are assigned values of -0.5 each. View 3 yields the vector:

$P_3 = [0 \quad 0 \quad 0 \quad 1 \quad -0.5 \quad -0.5 \quad 0]$.

Matrix P is given by:

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -0.5 & -0.5 & 0 \end{bmatrix}.$$

The variance of the k^{th} view portfolio can be calculated according to the formula $p_k \Sigma p'_k$, where p_k is the k^{th} row of the P matrix and Σ is the covariance matrix of the excess equilibrium market returns. (Recall, these form the neutral starting point of the Black-Litterman model). The variance of each view portfolio is an important source of information regarding the confidence that should be placed in the investor's view k.

7.4. Combining equilibrium returns with investor views

We now state the Black-Litterman equation for combining equilibrium returns with investor views to determine a vector of expected asset returns that will be used to determine optimal portfolio weights. The vector of combined asset returns is given by:

$$E[R] = [(\tau\Sigma)^{-1} + P'\Omega^{-1}P]^{-1}[(\tau\Sigma)^{-1}\Pi + P'\Omega^{-1}Q] \tag{12}$$

where:

E[R] = *Nx1* vector of combined returns

τ = scalar, indicating uncertainty of the CAPM prior

Σ = *NxN* covariance matrix of equilibrium excess returns

P = *KxN* matrix of investor views

Ω = *KxK* diagonal covariance matrix of view error terms (uncorrelated view uncertainty)

Π = *Nx1* vector of equilibrium excess returns

Q = *Kx1* vector of investor views

Examining this formula, we have yet to describe how the value of τ should be set and how the matrix Ω should be populated. Recall that if an investor has no views, the Black-Litterman model suggests that the investor does not deviate from the market equilibrium portfolio. Only weights on assets for which the investor has views should change from their market equilibrium weights. The amount of change depends upon τ, the investor's confidence in the CAPM prior, and ω, the uncertainty in the views expressed.

The literature does not have a single view regarding how the value of τ should be set. Black and Litterman (1992) suggest a value close to zero. He and Litterman (1999) set τ equal to 0.025 and populate the covariance matrix Ω so that

$$\Omega = \begin{bmatrix} (p_1 \Sigma p_1')\tau & 0 & \dots & & 0 \\ 0 & \ddots & & & \vdots \\ \vdots & & & & \\ & & & & 0 \\ 0 & & \dots & 0 & (p_K \Sigma p_K')\tau \end{bmatrix}.$$

We note that the implied assumption is that the variance of the view portfolio is the information that determines an investor's confidence in his view. There may be other information that contributes to the level of an investor's confidence but it is not accounted for in this method for populating Ω.

Formula (12) uses Bayes approach to yields posterior estimates of asset returns that reflect a combination of the market equilibrium returns and the investor views. These updated returns are now used to compute updated optimal portfolio weights.

In the case that the investor is unconstrained, we use formula (11). Using formula (11) w^*, the optimal portfolio weights, are given by:

$$w^* = (\lambda\Sigma)^{-1}\mu, \tag{13}$$

where μ is the vector of combined returns. Equation (13) is the solution to the unconstrained maximization problem $\max_w w'\mu - \lambda w'\Sigma w/2$.

In the presence of constraints (e.g., risk, short selling, etc.) Black and Litterman suggest that the vector of combined returns be input into a mean-variance.

We note two additional comments on the updated weights w^*:

(i) Not all view portfolios necessarily have equal impact on the optimal portfolio weights derived using the Black-Litterman model. A view with a higher level of uncertainty is given less weight. Similarly, a view portfolio that has a covariance with the market equilibrium portfolio is given less weight. This is because such a view represents less new information and hence should have a smaller impact in moving the optimal portfolio weights away from the market equilibrium weights. Finally, following the same reasoning, a view portfolio that has a covariance with another view portfolio has less weight.

(ii) A single view causes all returns to change, because all returns are linked via the covariance matrix Σ. However, only the weights for assets for which views were expressed change from their original market capitalization weights. Thus, the Black-Litterman model yields a portfolio that is intuitively understandable to the investor. The optimal portfolio represents a combination of the market portfolio and a weighted sum of the view portfolios expressed by the investor.

7.5. Application of Black-Litterman Model

The Black-Litterman model was developed at Goldman Sachs in the early 1990s and is used by the Quantitative Strategies group at Goldman Sachs Asset Management. This group develops quantitative models to manage portfolios. The group creates views and then uses the Black-Litterman approach to transform these views into expected asset returns. These expected returns are used to make optimal asset allocation decisions for all of the different portfolios managed by the group. Different objectives or requirements (such as liquidity requirements, risk aversion, etc.) are incorporated via constraints on the portfolio. The Black-Litterman model has gained widespread use in other financial institutions.

8. Risk Management

Risk exists when more than one outcome is possible from the investment. Sources of risk may include business risk, market risk, liquidity risk, and the like. Variance or standard deviation of return is often used as a measure of the risk associated with an asset's return. If variance is small, there is little chance that the asset return will differ from what is expected; if variance is large then the asset returns will be highly variable.

Financial institutions manage their risk on a regular basis both to meet regulatory requirements as well as for internal performance measurement purposes. However, while variance is a traditional measure of risk in economics and finance, in practice it is typically not the risk measure of choice. Variance assumes symmetric deviations above and below expected return. In practice, one does not observe deviations below expected return as often as deviations above expected return due to positions in options and options-like instruments in portfolios. Moreover, variance assigns equal penalty to deviations above and below the mean return. However, investors typically are not averse to receiving higher than anticipated returns. Investors are more interested in shortfall risk measures. These are risk measures that measure either the distance of return below a target or measure the likelihood that return will fall below a threshold.

One measure of shortfall risk is downside risk. Downside risk measures the expected amount by which the return falls short of a target. Specifically, if z is the realized return and X is the target then downside risk is given by $E[\max(X - z, 0)]$.

Semivariance is another measure of shortfall risk. Semivariance measures the variance of the returns that fall below a target value. Semivariance is given by $E\left[\max(X - z, 0)^2\right]$.

8.1. Value at Risk

Value at Risk (*VaR*) is a risk measure that is used for regulatory reporting. Rather than measuring risk as deviation from a target return, *VaR* measures risk as required capital. *VaR* is a risk measurement on the loss distribution of a portfolio. Let $L(f, \tilde{r})$ be the random loss on a portfolio with allocation vector *f* and random return vector $\tilde{r}$. Let *F* be its distribution function so that $F(f, u) = \Pr\{L(f, \tilde{r}) \leq u\}$. Value at risk is the α-quantile of the loss distribution and is defined by:

$$VaR_{\alpha}(f, \varsigma) = \min\{\varsigma : F(f, \varsigma) \geq \alpha\} = \min\{\varsigma : P\{L(f, \tilde{r}) \geq \varsigma\} \geq 1 - \alpha\}, \qquad (14)$$

where ζ is the value at risk and $\tilde{r}$ are the random asset returns. Thus, value-at-risk is the smallest amount *u* such that with probability *α* the loss will not exceed *u*.

The first step in computing *VaR* is to determine the underlying market factors that contribute to the risk (uncertainty) in the portfolio. The next step is to simulate these sources of uncertainty and the resulting portfolio loss. Monte Carlo simulation is used largely because many of the portfolios under management by the more sophisticated banks include a preponderance of instruments that have optional features. As we shall see in Section 10, the price changes of these instruments can best be approximated by simulation. *VaR* can then be calculated by determining the distribution of the portfolio losses. The key question is what assumptions to make about the distributions of these uncertain market factors. Similar to the two methods that we discuss for estimating asset returns and variances, one can use historical data or a scenario approach to build the distributions.

Using historical data, one assumes that past market behavior is a good indicator of future market behavior. Take T periods of historical data. For each period, simulate the change in the portfolio value using the actual historical data. Use these T data points of portfolio profit/loss to compute the loss distribution and hence *VaR*. The benefit of this method is that there is no need for artificial assumptions about the distribution of the uncertainty of the underlying factors that impact the value of the portfolio. On the other hand, this method assumes that future behavior will be identical to historical behavior.

An alternative approach is to specify a probability distribution for each of the sources of market uncertainty and to then randomly generate events from those distributions. The events translate into behavior of the uncertain factors, which result in a change in the portfolio value. One would then simulate the portfolio profit/loss assuming that these randomly generated events occur and construct the loss distribution. This distribution is used to compute *VaR*.

8.2. Coherent Risk Measures

VaR is a popular risk measure. However, *VaR* does not satisfy one of the basic requirements of a good risk measure: *VaR* is not subadditive for all distributions (i.e., it is not always the case that $VaR(A+B)<VaR(A)+VaR(B)$), a property one would hope to hold true if risk is reduced by adding assets to a portfolio. This means that the *VaR* of a diversified portfolio may exceed the sum of the *VaR* of its component assets.

Artzner, Delbaen, Eber, and Heath (1999) specify a set of axioms satisfied by all *coherent* risk measures. These are:

- Subadditivity: $\rho(A+B) \leq \rho(A)+\rho(B)$
- Positive Homogeneity: $\rho(\lambda A)= \lambda\rho(A)$ *for* $\lambda\geq 0$.
- Translation invariance: $\rho(A+c) =\rho(A)-c$ *for all* c.
- Monotonicity: $A\leq B$ *then* $\rho(B) \leq \rho(A)$

Subadditivity implies that the risk of a combined position of assets does not exceed the combined risk of the individual assets. This allows for risk reduction via diversification, as we discuss in Section 4.

Conditional Value-at-risk (*CVaR*), also known as Expected Tail Loss, is a coherent risk measure. *CVaR* measures the expected losses conditioned on the fact that the losses exceed *VaR*. Following the definition of *VaR* in equation (14), if *F* is continuous then *CVaR* is defined as:

$$CVaR(f,\alpha) = E\{L(f,\tilde{r}) \mid L(f,\tilde{r}) \geq VaR(f,\alpha)\}. \tag{15}$$

An investment strategy that minimizes *CVaR* will minimize *VaR* as well.

An investor wishing to maximize portfolio return subject to a constraint on maximum *CVaR* would solve the following mathematical program:

$$\begin{aligned}
&\max E(\sum_{i=1}^{N} f_i r_i) \\
&\textit{subject to}: \\
&CVaR_\alpha(f_1,\ldots,f_N) \leq C \\
&\textstyle\sum_{i=1}^{N} f_i = 1 \\
&0 \leq f_i \leq 1
\end{aligned} \tag{16}$$

where f_i is the fraction of wealth allocated to asset i, r_i is the return on asset i, and C is the maximum acceptable risk. Formulation (16) is a nonlinear formulation due to the constraint on *CVaR,* and is a hard problem to solve.

Suppose, instead of assuming that the loss distribution *F* is continuous, we discretize the asset returns by considering a finite set of $s=1,\ldots,S$ scenarios of the portfolio performance. Let $p(s)$ denote the likelihood that scenario *s* occurs; $0 \le p(s) \le 1; \sum_{s\in S} p(s) = 1$. $\varphi(s)$ denotes the vector of asset return under scenario *s* and $\rho(f,\varphi(s))$ denotes the portfolio return under scenario *s* assuming asset allocation vector *f*. Then, using this discrete set of asset returns, the expected portfolio return is given by

$$\sum_{s\in S} p(s)\rho(f,\varphi(s)) . \tag{17}$$

The investor will wish to maximize (17). Definition (15) applies to continuous loss distributions. Rockafellar and Uryasev (2002) have proposed an alternative definition of *CVaR* for any general loss distribution:

$$CVAR(f,\alpha) = (1-\frac{1}{1-\alpha}\sum_{\{s\in S|\rho(f,\varphi(s))\le\varsigma\}} p(s))\varsigma + \frac{1}{1-\alpha}\sum_{\{s\in S|\rho(f,\varphi(s))\le\varsigma\}} p(s)\rho(f,\varphi(s)) \tag{18}$$

Using definition (18) we can solve problem (16) using a stochastic programming approach. First, we define an auxiliary variable z_s for each scenario *s*, which denotes the shortfall in portfolio return from the value at risk ζ:

$$z(s) = \max(0, \varsigma - \rho(f,\varphi(s))) \tag{19}$$

Following Rockafellar and Uryasev (2002), CVaR can be expressed using the shortfall variables (19) as:

$$CVaR(f,\alpha) = \varsigma - \frac{1}{1-\alpha}\sum_{s=1}^{S} p(s)z(s) .$$

The linear program is given by:

$$
\begin{aligned}
&\max \sum_{s \in S} p(s)\rho(f,\varphi(s)) \\
&\textit{subject to}: \\
&\textstyle\sum_{i=1}^{N} f_i = 1 \\
&0 \le f_i \le 1 \qquad (20)\\
&\varsigma - \frac{1}{1-\alpha}\textstyle\sum_{s=1}^{S} p(s)z(s) \ge C \\
&z(s) \ge \varsigma - \rho(f,\varphi(s)) \quad s = 1,\dots,S \\
&z(s) \ge 0 \quad s = 1,\dots,S
\end{aligned}
$$

where the last two inequalities follow from the definition of the shortfall variables (19).

Formulation (20) is a linear stochastic programming formulation of the *CVaR* problem. To solve this problem requires an estimate of the distribution of the asset returns, which will be used to build the scenarios. If historical data is used to develop the scenarios then it is recommended that as time passes and more information is available, the investor reoptimize problem (20) (rebalance his portfolio) using these additional (in-sample) scenarios. We direct the reader to Rockafellar and Uryasev (2000,2002) for additional information on this subject.

8.3. Risk Management in Practice

VaR and *CVaR* are popular risk measures. *VaR* is used in regulatory reporting and to determine the minimum capital requirements to hedge against market, operational, and credit risk. Financial institutions may be required to report portfolio risks such as the 30-day 95% VaR or the 5% quantile of 30-day portfolio returns and to hold reserve accounts in proportion to these calculated amounts. *VaR* is used in these contexts for historical reasons. But even though as we saw above it is not a coherent risk measure, there is possibly some justification in continuing to use it. *VaR* is a frequency measure, so regulators can easily track whether the bank is reporting *VaR* accurately; *CVaR* is an integral over the tail probabilities and is likely not as easy for regulators to track.

In addition to meeting regulatory requirements, financial institutions may use *VaR* or *CVaR* to measure the performance of its business units that control financial portfolios by comparing the profit generated by the portfolio actions versus the risk of the portfolio itself.

For purposes of generating risk management statistics, banks will simulate from distributions that reflect views of the market and the economy. Banks will also incorporate some probability of extreme events such as the wild swings in correlations and liquidity that occur in market crashes.

9. Options

In this section we discuss options. An option is a derivative security, which means that its value is derived from the value of an underlying variable. The underlying variable may or may not be a traded security. Stocks or bonds are examples of traded securities; interest rates or the weather conditions are examples of variables upon which the value of an option may be contingent but that are not traded securities. Derivative securities are sometimes referred to as contingent claims, since the derivative represents a claim whose payoff is contingent on the history of the underlying security.

The two least complex types of option contracts for individual stocks are calls and puts. Call options give the holder the right to buy a specified number of shares (typically, 100 shares) of the stock at the specified price (known as the exercise or strike price) by the expiration date (known as the exercise date or maturity). A put option gives the holder the right to sell a specified number of shares of the stock at the strike price by maturity. American options allow the holder to exercise the option at any time until maturity; European options can only be exercised at maturity. The holder of an option contract may choose whether or not he wishes to exercise his option contract. However, if the holder chooses to exercise, the seller is obligated to deliver (for call options) or purchase (for put options) the underlying securities.

When two investors enter into an options contract, the buyer pays the seller the option price and takes a *long* position; the seller takes a *short* position. The buyer has large potential upside from the option, but his downside loss is limited by the price that he paid for the option. The seller's profit or loss is the reverse of that of the seller's. The seller receives cash upfront (the price of the option) but has a potential future liability in the case that the buyer exercises the option.

9.1. Call option payoffs

We first consider European options. We will define European option payoffs at their expiration date.

Let C represent the option cost, S_T denote the stock price at maturity, and K denote the strike price. An investor will only exercise the option if the stock price exceeds the strike price, i.e., $S_T > K$. If $S_T > K$, the investor will exercise his option to buy the stock at price K and gain $(S_T - K)$ for each share of stock that he purchases. If $S_T < K$, then the investor will not exercise the option to purchase the stock at price K. Thus, the option payoff for each share of stock is

$$max(S_T - K, 0). \qquad (21)$$

The payoff for the option writer (who has a short position in a call option) is the opposite of the payoff for the long position. If, at expiration, the stock price is below the strike price the holder will not exercise the option. However, if the stock price is above the strike price the owner will exercise his option. The writer must sell the stock to the owner at the strike price. For each share of stock that he sells, the writer must purchase the stock in the open market at a price per share of S_T and then sell it to the owner at a price of K. Thus, the writer loses $(S_T - K)$ for each share that he is obligated to sell to the owner. The writer thus has an unlimited potential loss depending upon the final stock price of the underlying asset; the writer's payoff is

$$-max(S_T - K, 0) = min(K - S_T, 0).$$

Thus, an investor with a short position in a European call option has potentially unlimited loss depending upon the final stock price of the underlying stock. This risk must be compensated by the price of the option C.

The graph on the left of **Figure 3** depicts the profit at maturity for the owner of a call option. The point of inflection occurs when the ending stock price equals the strike price. The negative payoff is the price the investor paid for the option. As mentioned, the option holder's downside loss is limited by the price paid for the option. The payoff for the call option writer is depicted in the graph on the right of **Figure 3** and is the reverse of the payoff to the call option buyer.

Figure 3 about here

Figure 3: Profit from European Call Options

When the price of the underlying stock is above the strike price, we say that the option is "in the money." If the stock price is below the strike price we say that the option is "out of the money." If the stock price is exactly equal to the strike price we say that the call option is "at the money."

American options can be exercised at any time prior to maturity. The decision of whether or not to exercise hinges upon a comparison of the value of exercising immediately (the *intrinsic value* of the option) against the expected future value of the option if the investor continues to hold the options. We will discuss this in further detail in Section 12 when we discuss pricing American options.

9.2. Put option payoffs

An investor with a long position in a put option profits if the price of the underlying stock drops below the option's strike price. Similar to the definitions for a call option, we say that a put option is "in the money" if the stock price at maturity is lower than the strike price. The put option is "out of the money" if the stock price exceeds the strike price. The option is "at the money" if the strike price equals the stock price at maturity.

Let P denote the cost of the put option, K is its strike price, and S_T the stock price at expiration. The holder of a put option will only sell if the stock price at expiration is lower than the strike price. In this case, the owner can sell the shares to the writer at the strike price and will gain $(K- S_T)$ per share. Thus, payoff on a put option is $max\ (K- S_{T},0)$; we note that positive payoff on a put is limited at K. If the stock price is higher than the strike price at maturity then the holder will not exercise the put option since he can sell his shares on the open market at a higher price. In this case, the option will expire worthless.

A put writer has opposite payoffs. If the stock price exceeds the strike price, the put holder will never exercise his option, however, if the stock price declines, the writer will lose money. The price P paid by the owner must compensate the writer for this risk.

Figure 4 depicts the option payoff for a put holder and writer.

Figure 4 about here

Figure 4: Profit from European put options

10. Valuing options

The exact value of a stock option is easy to define at maturity. Valuing options prior to expiration is more difficult and depends upon the distribution of the underlying stock price, amongst other factors. Black and Scholes (1973) derived a differential equation that can be used to price options on non-dividend paying stocks. We discuss the Black-Scholes formula in Section 10.1. However, in general exact formulas are not available for valuing options. In most cases, we rely on numerical methods and approximations for options valuation. In Sections 1010.3 and 1112 we discuss two numerical methods for pricing options: Monte Carlo simulation and dynamic programming. Dynamic programming is useful for pricing American options, where the holder has the ability to exercise prior to the expiration date. Monte Carlo simulation is useful for pricing a European option, where the option holder cannot exercise the option prior to the maturity date. In the following sections we describe Monte Carlo simulation and dynamic programming and show how they can be used to price options. We first will lay out some background and basic assumptions that are required.

10.1. Risk Neutral Valuation in efficient and complete markets

We base our discussion of options pricing on the assumption that markets are efficient and complete. Efficient markets are arbitrage-free. Arbitrage provides an investor with a riskless investment opportunity with unlimited return, without having to put up any money. We assume that if any such opportunities exist there would be infinite demand for such assets. This would immediately raise the price of the investments and the arbitrage opportunity would disappear.

Black and Scholes derived a differential equation that describes the value of a trader's portfolio who has a short position in the option and who is trading in the underlying asset and a cash-like instrument. Efficiency of markets is one of the key assumptions required in their derivation. In addition, they assume that instantaneous and continuous trading opportunities exist, no dividends, transaction costs, or taxes are paid, and that short selling is permitted. Finally, they assume that the price of the underlying stock follows a specific stochastic process called Brownian motion. (See Section 10.2 for discussion of Brownian motion.) In this Black-Scholes framework it turns out that there is a trading strategy (called "delta-hedging") that makes the portfolio's return completely riskless. In an efficient market, a riskless portfolio will return the risk-free rate. This arbitrage reasoning combined with the delta-hedging strategy leads to a partial differential equation that resembles the heat equation of classical physics. Its solution provides the option's value at any point in time. In the case of European-style options (those that have a fixed exercise date) the solution can be achieved in closed form – this is the famous Black-Scholes formula. The Black-Scholes formula for the values of a European call C or put P are:

$$C = S\Phi(d_1) - Ke^{-rT}\Phi(d_2)$$
$$P = Ke^{-rT}\Phi(-d_2) - S\Phi(-d_1)$$

where:

$$d_1 = [\log(S/K) + T(r + \sigma^2)/2]/\sigma\sqrt{T}$$
$$d_2 = d_1 - \sigma\sqrt{T}$$

Here, r is the risk-free rate (the rate of return of a riskless security such as a US Treasury security over time T), log denotes the natural logarithm, and $\Phi()$ is the cumulative distribution function for the standard normal distribution $N(0,1)$.

Exotic option contracts, especially those with exercise rules that give the owner the discretion of when to exercise, or options with underlying assets that are more complicated than equity stocks with no dividends, or options that depend on multiple assets, turn out to be very difficult to solve using the Black-Scholes partial differential equation.

Harrison and Pliska (1981) developed a more general perspective on options pricing that leads to a useful approach for these more general categories of options. It turns out that the existence of the riskless trading strategy in the Black-Scholes framework can be viewed as mathematically equivalent to the existence of a dual object called a "risk-neutral measure", and also that the options price is the integral of the option payouts with this risk-neutral measure. When the risk-neutral measure is unique, the market is called "complete". This assumption means that there is a single risk-neutral measure that can be used to price all the options.

This perspective leads to the following methodology for options pricing. Observe the prices of the traded options. Usually these are of fairly simple type (European or American calls and puts), for which closed-form expressions like the Black-Scholes formula can be used. Then invert the formula to find the parameters of the risk-neutral distribution. This distribution can then be used to simulate values of any option – under the assumption that the market is efficient (arbitrage-free) and complete.

10.2. Brownian Motion

A key component of valuing stock options is a model of the price process of the underlying stock. In this section, we describe the Brownian motion model for the stock prices.

The efficient market hypothesis, which states that stock prices reflect all history and that any new information is immediately reflected in the stock prices, ensures that stock prices follow a Markov process so the next stock price depends only upon the current stock price and does not depend upon the historical stock price process. A Markov process is a stochastic process with the property that only the current value of the random variable is relevant for the purposes of determining the next value of the variable. A Wiener process is a type of Markov process. A Wiener process $Z(t)$ has normal and independent increments with variance proportional to the square root of time, i.e., $Z(t)-Z(s)$ has a normal distribution with mean zero and variance $\sqrt{t-s}$. It turns out that $Z(t)$, $t>0$ will be a continuous function of t. If Δt represents an increment in time and ΔZ represents the change in Z over that increment in time then the relationship between ΔZ and Δt can be expressed by:

$$\Delta Z(t) = \varepsilon\sqrt{\Delta t} \tag{22}$$

where ε is drawn from a standard normal distribution. A Wiener process is the limit of the above stochastic process as the time increments get infinitesimally small, i.e., as $\Delta t \rightarrow 0$. Equation (22) is expressed as

$$dZ(t) = \varepsilon\sqrt{dt} \ . \tag{23}$$

If *x(t)* is a random variable and *Z* is a Wiener process, then a generalized Wiener process is defined as

$$dx(t) = adt + bdZ$$

where *a* and *b* are constants. An Ito process is a further generalization of a generalized Wiener process. In an Ito process, *a* and *b* are not constants rather, they can be functions of *x* and *t*. An Ito process is defined as

$$dx(t) = a(x,t)dt + b(x,t)dZ$$

Investors are typically interested in the rate of return on a stock, rather than the absolute change in stock price. Let *S* be the stock price and consider the change in stock price *dS* over a small period of time *dt*. The rate of return on a stock, *dS/S*, is often modeled as being comprised of a deterministic and stochastic component. The deterministic component, μdt, represents the contribution of the average growth rate of the stock. The stochastic component captures random changes in stock price due to unanticipated news. This component is often modeled by σdZ, where *Z* is a Brownian motion. Combining the deterministic growth rate (also known as drift) with the stochastic contribution to rate of change in stock price yields the equation:

$$dS / S = \mu dt + \sigma dZ \ , \tag{24}$$

an Ito process. μ and σ can be estimated using the methods described in Section 3.

The risk-neutral measure Q of Harrison and Pliska as applied to the Black-Scholes framework is induced by the risk-neutral process *X* that satisfies the modified Brownian motion

$$dX / X = (r - \sigma^2 / 2)dt + \sigma dZ \tag{25}$$

It is important to note that this process is not the same as the as the original process followed by the stock – the drift has been adjusted. This adjustment is what is required to generate the probability measure that makes the delta-hedging portfolio process into a martingale. According to the theory discussed above in Section 10.1, we price options in the Black-Scholes framework by integrating their cash flows under the risk-neutral measure generated by equation (25). In the following section we discuss how efficient markets and risk neutral valuation are used to compute an option's value using Monte Carlo simulation.

10.3. Simulating risk-neutral paths for options pricing

In this section we discuss how simulation can be used to price options on a stock by simulating the stock price under the risk-neutral measure over T periods, each of length $\Delta t=1/52$. At each time interval, we simulate the current stock price and then step the process forward so that there are a total of T steps in the simulation. To simulate the path of the stock price over the T week period, we consider the discrete time version of equation (25): $\Delta S/S= (r - \sigma^2/2)\Delta t+ \sigma dZ= (r - \sigma^2/2)\Delta t+ \sigma \varepsilon\sqrt{\Delta t}$. Since ε is distributed like a standard normal random variable, $\Delta S/S \sim N((r - \sigma^2/2)\Delta t, \sigma \sqrt{\Delta t})$.

Each week, use the following steps to determine the simulated stock price:

(i) Set $i=0$.

(ii) Generate a random value v_1 from a standard normal distribution. (Standard spreadsheet tools include this capability.)

(iii) Convert v_1 to a sample v_2 from a $N((r - \sigma^2/2)\Delta t, \sigma \sqrt{\Delta t})$ by setting $v_2 = (r - \sigma^2/2)\Delta t + \sigma \sqrt{\Delta t}\, v_1$.

(iv) Set $\Delta S=v_2 S$. ΔS represents the incremental change in stock price from the prior period to the current period.

(v) $S' = S + \Delta S$, where S' is the simulated updated value of the stock price after one period.

(vi) Set $S = S'$, $i=i+1$.

(vii) If $i=T$ then stop. S is the simulated stock price at the end of six months. If $i<T$, return to step (i).

Note that randomness only enters in step (ii) when generating a random value v_1. All other steps are mere transformations or calculations and are deterministic. The payoff of a call option at expiration is given by equation (21). Further, in the absence of arbitrage opportunities (i.e., assuming efficient markets) and by applying the theory of risk neutral valuation we know that the value of the option is equal to its expected payoff discounted by the risk free rate. Using these facts, we apply the Monte Carlo simulation method to price the option. The overall methodology is as follows:

(i) Define the number of time periods until maturity, T.

(ii) Use Monte Carlo simulation to simulate a path of length T describing the evolution of the underlying stock price, as described above. Denote the final stock price at the end of this simulation by S^F.

(iii) Determine the option payoff according to equation (21), assuming S^F, the final period T stock price determined in step (ii).

(iv) Discount the option payoff from step (iii) assuming the risk-free rate. The resulting value is the current value of the option.

(v) Repeat steps (ii)-(iv) until the confidence bound on the estimated value of the option is within an acceptable range.

10.4. A numerical example

A stock has expected annual return of $\mu=15\%$ per year and standard deviation of $\sigma=20\%$ per year. The current stock price is $S=\$42$. An investor wishes to determine the value of a European call option with a strike price of $40 that matures in six months. The risk free rate is 8%.

We will use Monte Carlo simulation to simulate the path followed by the stock price and hence the stock price at expiration which determines the option payoff. We consider weekly time intervals, i.e., $\Delta t=1/52$. Thus $T=24$ assuming, for the sake of simplicity, that there are four weeks in each month.

To simulate the path of the stock price over the 24 week period, we follow the algorithm described in Section 10.3. We first compute the risk-neutral drift $(r-\sigma^2/2)\Delta t$, which with these parameter settings works out to be 0.0012. The random quantity ε is distributed like a standard normal random variable, so $\Delta S/S \sim N(.0012,.0277)$.

The starting price of the stock is $42. Each week, use the following steps to determine the simulated updated stock price:

(i) Generate a random value v_1 from a standard normal distribution. (Standard methods can be used.)

(ii) Convert v_1 to a sample v_2 from a $N(.0012,.0277)$ by setting $v_2 = .0012+.0277v_1$.

(iii) Set $\Delta S=v_2 S$.

(iv) Set $S = S + \Delta S$

Steps (i)-(iv) yield the simulated updated stock price after one week. Repeat this process *T=24* times to determine S^F, the stock price at the end of six months. Then, the option payoff equals $P=max(S^F-40,0)$. *P* is the option payoff based upon a single simulation of the ending stock price after six months, i.e., based upon a single simulation run. Perform many simulation runs and after each run compute the arithmetic average and confidence bounds of the simulated values of *P*. When simulation runs have been performed so that the confidence bounds are acceptable, the value of the option can be computed based upon the expected value of *P*: $V = e^{-.08(.5)} E(P)$.

11. Dynamic programming

Dynamic programming is formal method for performing optimization over time. The algorithm involves breaking a problem into a number of subproblems, solving the smaller subproblems, and then using those solutions to help solve the larger problem. Similar to stochastic programming with recourse, dynamic programming involves sequential decision making where decisions are made, information is revealed, and then new decisions are made. More formally, the dynamic programming approach solves a problem in *stages*. Each stage is comprised of a number of possible *states*. The optimal solution is given in the form of a policy which defines the optimal action for each stage. Taking action causes the system to transition from one stage to a new state in the next stage.

There are two types of dynamic programming settings: deterministic and stochastic. In a deterministic setting, there is no system uncertainty. Given the current state and the action taken, the future state is known with certainty. In a stochastic setting, taking action will select the probability distribution for the next state. For the remainder we restrict our discussion to a stochastic dynamic programming setting, since finance problems are generally not deterministic. If the current state is the value of a portfolio, and possible actions are allocations to different assets, the value of the portfolio in the next stage is not known with certainty (assuming that some of the assets under consideration contain risk).

A dynamic program typically has the following features:

(i) The problem is divided into $t=1,\dots,T$ stages. x_t denotes the state at the beginning of stage t and $a_t(x_t)$ denotes the action taken during stage t given state x_t. Taking action transitions the system to a new state in the next stage so that $x_{t+1}=f(x_t, a_t(x_t), \varepsilon_t)$, where ε_t is a random noise term. The initial state x_0 is known.

(ii) The cost (or profit) function in period t is given by $g_t(x_t, a_t(x_t), \varepsilon_t)$. This cost function is additive in the sense that the total cost (or profit) over the entire T stages is given by:

$$g_T(x_T, a_T(x_T), \varepsilon_T) + \sum_{t=1}^{T-1} g_t(x_t, a_t(x_t), \varepsilon_t). \qquad (26)$$

The objective is to optimize the expected value of equation (26).

(iii) Given the current state, the optimal solution for the remaining states is independent of any previous decisions or states. The optimal solution can be found by backward recursion. Namely, the optimal solution is found for the period T subproblem, then for the periods $T-1$ and T subproblem, etc. The final period T subproblem must be solvable.

The features dynamic program that defines the options pricing problem differ somewhat from the features described here. In Section 12.1 we note these differences.

12. Pricing American options using dynamic programming

Monte Carlo simulation is a powerful tool for options pricing. It performs well even in the presence of a large number of underlying stochastic factors. However, at each step simulation progresses forward in time. On the other hand, options that allow for early exercise must be evaluated backward in time where in each period the option holder must compare the intrinsic value of the option against the expected future value of the option.

The traditional approach to pricing American options has been to use binomial trees and dynamic programming. Dealing with early exercise requires one to go backwards in time, as at each decision point the option holder must compare the value of exercising immediately against the value of holding the option. The value of holding the option is simply the price of the option at that point.

In this section we will show how one can use dynamic programming to price an American option. The method involves two steps:

(i) Build a T stage tree of possible states. The states correspond to points visited by the underlying stock price process.

(ii) Use dynamic programming and backward recursion to determine the current value of the option.

12.1. Building the T stage tree

Cox, Ross, and Rubinstein, (1979) derived an exact options pricing formula under a discrete time setting. Following their analysis, we model stock price as following a multiplicative binomial distribution: if the stock price at the beginning of stage t is S then at the beginning of stage $t+1$ the stock price will be either uS or dS with probability q and $(1-q)$, respectively. Each stage has length Δt. We will build a "tree" of possible stock prices in any stage. When there is only one stage remaining, the tree looks like:

$$
S \begin{cases} uS & \text{with probability } q \\ dS & \text{with probability } 1 - q \end{cases}
$$

Suppose there are two stages. In each stage, the stock price will move by "up" by a factor of u with probability q and "down" by a factor of d with probability $(1-q)$. In this case, there are three possible final stock prices and the tree is given by:

u^2S with probability q^2

uS

S $\quad$ $duS=S$ with probability $2q(1-q)$

dS

d^2S with probability $(1-q)^2$

The tree continues to grow according to this method. In general, at stage t there are $t+1$ possible stock prices (states) that will appear on the tree. These are given by:

$u^j d^{t-j} S$, for $j=0,\ldots,t$.

The values of u, d, and q are determined based upon the assumptions of efficient markets, risk neutral valuation, and the fact that the variance of the change in stock price is given by $\sigma^2 \Delta t$ (from Section 10.3). These values are:

$$
\begin{aligned}
u &= e^{\sigma\sqrt{\Delta t}} \\
d &= e^{-\sigma\sqrt{\Delta t}} \\
q &= \frac{a-d}{u-d}
\end{aligned}
$$

where $a = e^{r\Delta t}$.

12.1. Pricing the option using the Binomial Tree

We now use backward enumeration through the binomial tree to determine the current stage 0 value of the option. We will illustrate the concept using the trees developed in Section 12.1 above. Let K denote the strike price. With one period remaining, the binomial tree had the form:

uS with probability q

S

dS with probability $1-q$

The corresponding values of the call at the terminal points of the tree are C_u=*max(0,uS-K)* with probability q and C_d=*max(0,dS-K)* with probability *(1-q)*. The current value of the call is given by the present value (using the risk free rate) of $qC_u+(1-q)C_d$.

When there is more than one period remaining, each node in the tree must be evaluated by comparing the intrinsic value of the option against its continuation value. The intrinsic value is the value of the option if it is exercised immediately; this value is determined by comparing the current stock price to the option strike price. The continuation value is the discounted value of the expected cash payout of the option under the risk neutral measure, assuming that the optimal exercise policy is followed in the future. Thus, the decision is given by:

$$g_t = \max\left\{\max(0, x_t - K), E[e^{-r\Delta t} g_{t+1}(x_{t+1}) \mid x_t]\right\}. \tag{27}$$

The expectation is taken over the risk neutral probability measure. x_t, the current state, is the current stock price. Notice that the action taken in the current state (whether or not to exercise) does not affect the future stock price. Further, this value function is not additive. However, its recursive nature makes a dynamic programming solution method useful.

12.2. A numerical example

We illustrate the approach using an identical setting as that used to illustrate the Monte Carlo simulation approach to options pricing. However, here we will assume that that we are pricing an American option. The investor wishes to determine the value of an American call option with a strike price of $40 that matures in one month. (We consider only one month to limit the size of the tree that we build.) The current stock price is

S=\$42. The stock return has standard deviation of $\sigma=20\%$ per year. The risk free rate is 8%.

We first build the binomial tree and then use the tree to determine the current value of the option. We consider weekly time intervals, i.e., $\Delta t=1/52$. Thus $T=4$ assuming, for the sake of simplicity, that there are four weeks in each month.

$$\begin{aligned} u &= e^{\sigma\sqrt{\Delta t}} &&= 1.028 \\ d &= e^{-\sigma\sqrt{\Delta t}} &&= 0.973 \\ q &= \frac{e^{r\Delta t}-d}{u-d} &&= 0.493 \end{aligned}$$

The tree of stock movements over the four week period looks like:

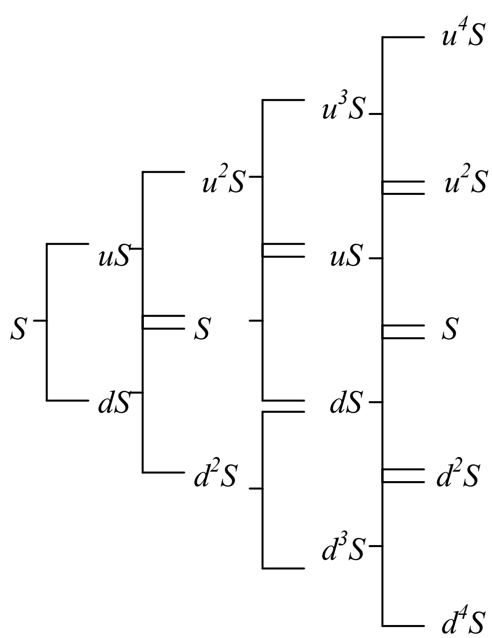

We will evaluate each node in the tree by backward evaluation starting at the fourth time period and moving backward in time. For each node we will use equation (write the equation number) to compare the intrinsic value of the option against its continuation value to determine the value of that node. The binomial tree for the value of the American call option is:

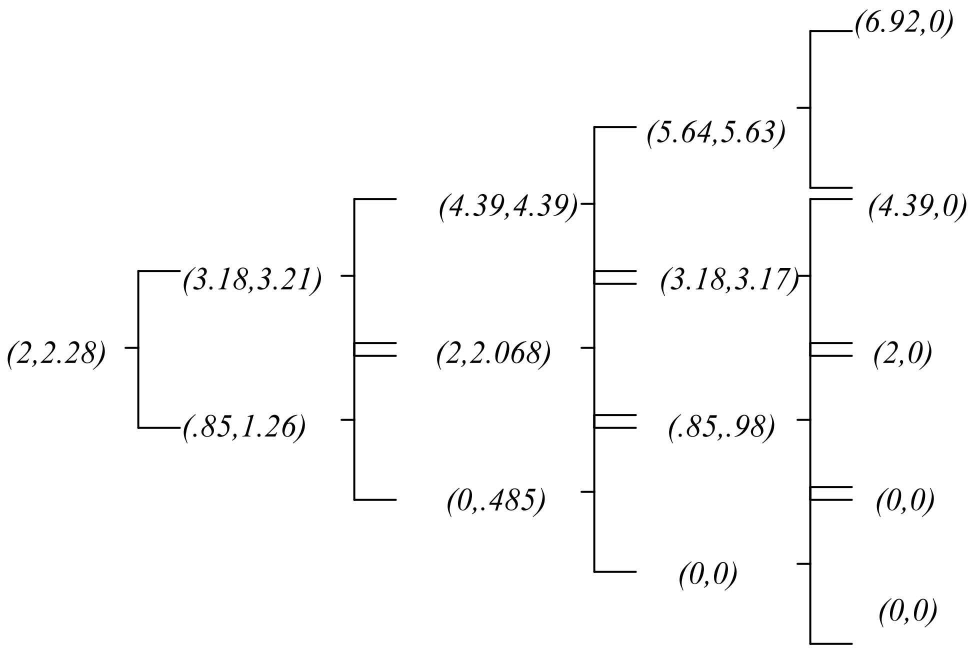

Every node in the tree contains two numbers in parenthesis. The first number is the intrinsic value of the option. The second number is the discounted expected continuation value, assuming that optimal action is followed in future time periods. The option value at time zero (current time) is 2.28. Note that although the option is currently in the money, it is not optimal to exercise even though it is an American option and early exercise is allowed. By using the binomial tree to evaluate the option we find that the expected continuation value of the option is higher than its current exercise value.

13. Comparison of Monte Carlo simulation and Dynamic Programming

Dynamic programming is a powerful tool that can be used for pricing options with early exercise features. However, dynamic pricing suffers from the so-called curse of dimensionality. As the number of underlying variables increases the time required to solve the problem grows significantly. This reduces the practical use of dynamic programming as a solution methodology. The performance of Monte Carlo simulation is better in the sense that its convergence is independent of the state dimension. On the other hand, as we have discussed, simulation has traditionally been viewed as inappropriate for pricing options with early exercise decisions since these require estimates of future values of the option and simulation only moves forward in time. However, recent research has focused on combining simulation and dynamic programming approaches to pricing American options to gain the benefits of both techniques. See, for example, Broadie and Glasserman (1997).

14. Multi-period Asset Liability Management

The management of liability portfolios of relatively long-term products, like pensions, variable annuities, and some insurance products requires a perspective that goes beyond a single investment period. The portfolio optimization models of Sections 5 through 7 are short-term models. Simply rolling these models over into the next time horizon can lead to problems. First, the models may make an excessive number of transactions. Transactions are not free, and realistic portfolio management models must take trading costs into consideration. Second, the models depend only on second moments. Large market moves, such as during a market crash, are not part of the model assumptions. Finally, policies and regulations place conditions on the composition of portfolios. These are not part of the model assumptions.

Academic and finance industry researchers have, over the past few decades, been exploring the viability of using multi-period balance sheet modeling to address the issues of long-dated asset liability management. A multi-period balance sheet model treats the assets and liabilities as generating highly aggregated cash flows over multiple time periods. The aggregations are across asset classes, so that individual securities in an asset class, say stocks, are accumulated into a single asset class, say S&P 500. Other asset classes are corporate bonds of various maturity classes, and so forth. The asset class cash flows are aggregated over time periods, typically three or six months, so that cash flows occurring within a time interval, say, $(t-1,t]$, are treated as if they all occur at the end-point *t*. The model treats the aggregate positions in each asset category as variables in the model. There is a single decision to be made for each asset class at the beginning of each time period, which is the change in the holdings of each asset class. The asset holdings and liabilities generate cash flows, which then flow into account balances. These account balances are assigned targets, and the objective function records the deviation from the targets. The objective of the model is to maximize the sum of the expected net profits and the expected penalties for missing the targets over the time horizon of the model.

A simplified application of such a model to a property-casualty insurance problem is as follows. Let A_t denote a vector of asset class values at time *t* and i_t denote their cash flows (e.g., interest payment, dividends, etc.) at time t. Let x_t denote the portfolio of holdings in the asset classes. Cash flows are generated by the holdings and by asset sales:

$$C_t = A_t \Delta x_t + i_t x_{t-1}$$

where $\Delta x_t := x_t - x_{t-1}$. The cash flows are subject to market and economic variability over the time horizons of interest, say $t = 1,\ldots,T$.

Liability flows from the property-casualty portfolio are modeled by aggregating and then forecasting cash flows. The net liability flows L_t are losses minus premium income. Loss events are classified by frequency of occurrence and intensity given loss. These can be simulated over the time horizon T using actuarial models for insurance payments. The evolution of the liability portfolio composition (say, by new sales, lapses of coverage, and so forth) can also be modeled. The key correlation to capture in the liability model is the relationship between the state of economic activity and the asset markets. For example, high employment is indicative of strong economic activity, which can lead to increases in the insurance portfolio; high inflation will lead to higher loss payouts given losses; and so forth.

Various performance, accounting, tax, and regulatory measurements are computed from the net cash flows. For example, one measurement could be shareholder's equity at the time horizon S_T, another could be annual net income N_t, and yet another could be premium-surplus ratio P_t – a quantity used in the property-casualty industry as a proxy for the rating of the insurance company.

In these aggregated models, we try to model the change in performance measurements as linear computations from the cash and liability flows and the previous period quantities. For example, net income is computed as cash flow minus operating expenses. If $O_t \Delta x_t$ is a proxy for the contribution of portfolio management to expenses, for example the cost of trading, then net income can be modeled by the following equation

$$N_t = C_t - L_t - O_t \Delta x_t \, .$$

Shareholder's equity can change due to a number of influences; here we just capture the change due to the addition of net income

$$S_t = S_{t-1} + N_t .$$

Finally, premium-surplus ratio can be approximated by fixing premium income to a level L and assuming (this is a major simplification!) that the surplus is equivalent to shareholders equity:

$$P_t = L / S_t .$$

A typical objective for a multi-period balance sheet model/an asset-liability matching problem is to create performance targets for each quantity and to penalize the shortfall. Suppose that the targets are $\overline{N}_t$ for annual net income, $\overline{S}_T$ for shareholder's equity, and $\overline{P}_t$ for premium-surplus ratio. Then the objective function could be

$$\text{Maximize } E\{S_T - \lambda \sum_t [\overline{N}_t - N_t]^+ - \mu \sum [L - S_t \overline{P}_t]^+\}$$

$$\text{Subject to } \begin{array}{l} N_t = A_t \Delta x_t + i_i x_{t-1} - L_t - O_t \Delta x_t \\ S_t = N_t + S_{t-1} \end{array} \qquad (28)$$

where the parameters λ and μ are used to balance the various contributions in the objective function, the premium-surplus ratio target relationship has been multiplied through by the denominator to make the objective linear in the decision variables, and the objective is integrated over the probability space represented by the discrete scenarios.

The objective function in formulation (28) can be viewed as a variation of the Markowitz style, where we are modeling "expected return" through the shareholder's equity at the end of the horizon, and "risks" through the shortfall penalties relative to the targets for net income and premium-surplus ratio.

14.1. Scenarios

In multi-period asset liability management the probability distribution is modeled by discrete scenarios. These scenarios indicate the values, or states, taken by the random quantities at each period in time. The scenarios can branch so that conditional distributions given a future state can be modeled. The resulting structure is called a "scenario tree''. Typically there is no recombining of states in a scenario tree, so the size of the tree grows exponentially in the number of time periods. For example, in the property-casualty model, the scenarios are the values and cash flows of the assets and the cash flows of the liabilities. The values of these quantities a each time point *t* and scenario *s* is represented by the triple (A_t^s, i_t^s, L_t^s). The pair (s,t) is sometimes called a "node" of the scenario tree. The scenario tree may branch at this node, in which case the conditional distribution for the triple $(A_{t+1}, i_{t+1}, L_{t+1})$ given node (s,t) is the values of the triples on the nodes that branch from this node.

It is important to model the correlation between the asset values and returns and the liability cash flows in these conditional distributions. Without the correlations, the model will not be able to find positions in the asset classes that hedge the variability of the liabilities. In property-casualty insurance, for example, it is common to correlate the returns of the S&P 500 and bonds with inflation and economic activity. These correlations can be obtained from historical scenarios, and conditioned on views as discussed above in Section 7.

The scenario modeling framework allows users to explicitly model the probability and intensity of extreme market movements and events from the liability side. One can also incorporate "market crash" scenarios in which the historical correlations are changed for some length of time that reflects unusual market or economic circumstances – such as a stock market crash or a recession. Finally, in these models it is usual to incorporate the loss event scenarios explicitly rather than follow standard financial valuation methodology, which would tend to analyze the expected value of loss distributions conditional on financial return variability. Such methodology would ignore the year-to-year impact of loss distribution variability on net income and shareholder's equity. However, from the ALM perspective, the variability of the liability cash flows is very important for understanding the impact of the hedging program on the viability of the firm.

14.2. Multi-period Stochastic programming

The technology employed in solving an asset liability management problem such as this is multiperiod stochastic linear programming. For a recent survey of stochastic programming, see Shapiro and Ruszczynski (2003).

The computational intensity for these models increases exponentially in the number of time periods, so the models must be highly aggregated and strategic in their recommendations. Nevertheless, the models do perform reasonably well in practice, usually generating 300 basis points of excess return over the myopic formulations based on the repetitive application of one period formulations, primarily through controlling transaction costs and because the solution can be made more robust by explicitly modeling market crash scenarios. A recent collection of this activity is in the volume edited by Ziemba and Mulvey (1998). See also Ziemba's monograph Ziemba (2003) for an excellent survey of issues in asset liability management.

15. Conclusions

In this chapter we saw the profound influence of applications of Operations Research to the area of finance and financial engineering. Portfolio optimization by investors, Monte Carlo simulation for risk management, options pricing, and asset liability management, are all techniques that originated in Operations Research and found deep application in finance. Even the foundations of options pricing are based on deep applications of duality theory. As the name financial engineering suggests, there is a growing part of the body of financial practice that is regarded as a subdiscipline of engineering in which techniques of applied mathematics and operations research are applied to the understanding of the behavior and the management of the financial portfolios underpinning critical parts of our economy in capital formation, economic growth, insurance, and economic-environmental management.

16. Bibliography

1. Ashford, R.W., R.H. Berry, and R.G. Dyson, 1988, "Operational Research and Financial Management," *European Journal of Operational Research,* 36 (2), 143-152.

2. Bengtsson, Christoffer, 2004, "The Impact of Estimation Error on Portfolio Selection for Investors with Constant Relative Risk Aversion." Working paper, Department of Economics, Lund University.

3. Bertsekas, D.P. and J.N. Tsitsiklis, 1996, *Neuro-Dynamic Programming*. Athena Scientific, Belmont (MA), 1996.

4. Bevan, A. and K. Winkelmann, 1998, "Using the Black Litterman Global Asset Allocation Model: Three Years of Practical Experience," Fixed Income Research, Goldman, Sachs, and Company, December.

5. Birge, J. and V. Linetsky, *Financial Engineering, Handbooks in Operations Research and Management Science*, Elsevier, Amsterdam, Forthcoming, 2007.

6. Black, F. and R. Litterman, 1992, "Global Portfolio Optimization," *Financial Analysts Journal*, September/October, 28-43.

7. Black, F. and M. Scholes, 1973, "The Price of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637-654.

8. Board, J., C. Sutcliffe, and W.T. Ziemba, 2003, "Applying Operations Research Techniques to Financial Markets," *Interfaces*, 33, 12-24.

9. Broadie, M. and P. Glasserman, 1997, "Pricing American-Style Securities Using Simulation," *Journal of Economic Dynamics and Control*, 21, 1323-1352.

10. Campbell, J.Y., A.W. Lo, and A.C. MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press, Princeton, N.J.

11. Chopra, V.K. and W.T. Ziemba, 1993, "The Effect of Errors in Mean, Variance, and Co-variance Estimates on Optimal Portfolio Choice," *Journal of portfolio management*, 19, no. 2, 6-11.

12. Cox, J. C. and S. A. Ross, 1976, "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, 3, 145-166.

13. Cox, J.C., S.A. Ross, and M. Rubinstein, (1979), "Option Pricing: A Simplified Approach", *Journal of Financial Economics*, 7, 229-263.

14. Fama, E.F., 1965, "The Behavior of stock prices," *Journal of Business*, 38, 34-105.

15. Fusai, G. and A. Meucci, 2003, "Assessing Views," *Risk*, March 2003, s18-s21.

16. Glasserman, P., 2004, *Monte Carlo Methods in Financial Engineering*, Springer-Verlag, New York Inc.

17. Harrison, J. M. and S.R. Pliska, 1981, "Martingales and Stochastic Integrals in the Theory of Continuous Trading," *Stochastic Processes and their Applications*, 11, 215-260.

18. He, G. and R. Litterman, 1999, "The Intuition Behind Black-Litterman Model Portfolios," Investment Management Research, Goldman Sachs & Company, December.

19. Holmer, M.R., 1998, "Integrated Asset-Liability Management: An Implementation Case Study," in *Worldwide Asset and Liability Modeling*, W. T. Ziemba and J. R. Mulvey (eds.) , Cambridge University Press.

20. Jarrow, R. V. Maksimovic and W. Ziemba, eds., *Finance, Handbooks in Operations Research and Management Science*, Volume 9, North-Holland, 1995.

21. King, A.J. 1993, Asymmetric risk measures and tracking models for portfolio optimization under uncertainty, *Annals of Operations Research*, vol. 45, 165-177.

22. Michaud, R.O., 1989, "The Markowitz Optimization Enigma: is 'Optimized' Optimal?" *Financial Analysts Journal*, 45, no. 1, 31-42.

23. Mills, T.C., 1999, *The Econometric Modelling of Financial Time Series*, 2nd ed., Cambridge University Press, Cambridge, U.K.

24. Pedersen, C.S. and S.E. Satchell, 1998, "An Extended Family of Financial Risk Measures," *Geneva Papers on Risk and Insurance Theory*, 23, 89-117.

25. Rockafellar, R.T. and S. Uryasev, 2000, Optimization of Conditional Value-at-Risk, *The Journal of Risk*, vol. 2, No. 3, 21-41.

26. Rockafellar, R.T. and S. Uryasev, 2002, Conditional Value-at-Risk for General Distributions, *Journal of Banking and Finance*, vol. 26, No. 7, 1443-1471.

27. Samuelson, P.A., 1965, "Rational theory of warrant pricing," *Industrial Management Review*, 6, 13-31

28. Shapiro, A. and A. Ruszczynski, 2003, *Stochastic Programming,* Handbook in Operations Research Volume 10, Elsevier.

29. Sharpe, W. F., 1964, "Capital Asset Prices: A Theory of Market Equilibrium." *Journal of Finance*, September, 425-442.

30. Uryasev, S. 2001, Conditional Value-at-Risk: Optimization Algorithms and Applications, Financial Engineering News, 14, 1-5.

31. von Neumann, J. and O. Morgenstern, 1944, *Theory of Games and Economic Behaviour*, Princeton University Press.

32. Ziemba, W.T. 2003, *The Stochastic Programming Approach to Asset, Liability, and Wealth Management,* AIMR Research Foundation Publications, No. 3: 1-192.