

INTRODUCTION TO FINANCIAL RISK ASSESSMENT USING MONTE CARLO SIMULATION

Robert A. Strong
Natalie M. Steiger

Maine Business School
5723 D. P. Corbett Business Building
University of Maine
Orono, ME 04473-5623, USA

James R. Wilson

Edward P. Fitts Department of Industrial
and Systems Engineering
North Carolina State University
Raleigh, NC 27695-7906, USA

ABSTRACT

The fundamental principles of financial risk assessment are discussed, with primary emphasis on using simulation to evaluate and compare alternative investments. First we introduce the key measures of performance for such investments, including net present value, internal rate of return, and modified internal rate of return. Next we discuss types of risk and the key measures of risk, including expected present value; the mean, standard deviation, and coefficient of variation of the rate of return; and the risk premium. Finally we detail the following applications: (i) stand-alone risk assessment for a capital-budgeting problem; (ii) comparison of risk-free and risky investment strategies designed merely to keep up with the cost of living; (iii) value-at-risk (VAR) analysis for a single-stock investment; (iv) VAR analyses for two-asset portfolios consisting of stock and either call or put options; and (v) VAR analyses for two-asset portfolios consisting of both puts and calls.

1 INTRODUCTION

The term *risk* encompasses not only the probability (likelihood) of various outcomes of interest but also the adverse or beneficial consequences resulting from the occurrence of those outcomes. Risk is not the same as uncertainty or random variation; that is, a truly risky situation must somehow involve a chance of monetary loss or some other type of adverse consequence as well as a chance of monetary gain or some other type of beneficial consequence. An individual who elects to take the afternoon off to play golf despite a 70% chance of rain may incur risk; but for that individual's office mate who remains in the office, the same uncertainty about the weather may not have an element of risk. Koller (2005) identifies two parameters associated with risk: (i) the probability of occurrence of a specific outcome; and (ii) the consequence (impact) associated with the outcome, which may be expressed in terms of the utility resulting from the outcome or in terms of the monetary loss or gain resulting from the outcome. In each application, both (i) and (ii) must be quantified.

Wharton (1992) makes a distinction between risk analysis and risk assessment. Risk analysis is the identification of possible outcomes, and risk assessment is the estimation of probabilities and the economic impacts that result from the corresponding outcomes. In many situations the required quantities are not observable, and estimates of these quantities must be based on someone's subjective judgment. In reality there may be little agreement between perceived and actual risk. Because all conclusions derived from any quantitative study will depend heavily on the inputs to that study, risk analysis and assessment must be preceded by as complete and accurate perception of the risk as is possible. Equally important is that risk analysis and assessment must include the following: (a) validation of the assumptions underlying the financial model used in the analysis; (b) thorough exploration of the sensitivity of the model results to changes in the underlying assumptions and changes in the input parameters whose values are uncertain or are subject to random variation; and (c) an honest representation of the inherent variability (lack of precision) in the final results that is due to uncertainty or random variation in the input parameters.

Complex problems in risk analysis and assessment arise in many application contexts, ranging from personal finance (Should I accept the terms of a contract offered for the purchase of my home?) to international banking (Should governments inject cash into their economies in an attempt to prevent a global depression?). Unfortunately, the pervasive nature of these applications seriously hinders an accurate perception of risk. William Schreyer, former chair and CEO of Merrill Lynch, once said of business decisions, "The key is not to eliminate risk; it is to measure it and manage it wisely." Koller (2005) details applications of risk analysis and assessment in the following areas:

- Legal actions—for example, evaluating the net benefit of settlement vs. the cost of litigation.
- Environmental health and safety—for example, evaluating the impact of a firm’s doing business in various foreign countries in terms of the firm’s profitability and the resulting effect on the environment and on the health and safety of the firm’s employees.
- Logistics systems engineering—for example, ranking and prioritizing various pipeline routes based on political, environmental, commercial, financial, technical, and other considerations.
- International economic development—for example, comparing countries as potential recipients of economic aid based on political stability, foreign investment conditions, operating environment, transportation infrastructure, etc.
- Capital project ranking and portfolio management—for example, comparing financial outputs using project safety and environmental aspects, cost estimates, incentives, discount rates, etc.

Regardless of the complexity of the problem, the essence of risk analysis and assessment is the same: to compare two or more options, where each option features probabilistic outcomes and the associated economic consequences. Figure 1 is a simple decision tree demonstrating the nature of the “basic risk paradigm” (MacCrimmon and Wehrung 1986). In this figure two decision alternatives are represented—a “sure” action with one outcome; and a “risky” action that has two possible outcomes, each having some probability of occurrence. This figure is a simplified ideal in the following respects: (i) in reality, the “sure” action may be only less risky than the one labeled “risky”; (ii) there may be more than two possible actions; and (iii) each action may have more than two possible outcomes. The sum of the probabilities of all outcomes must sum to one.

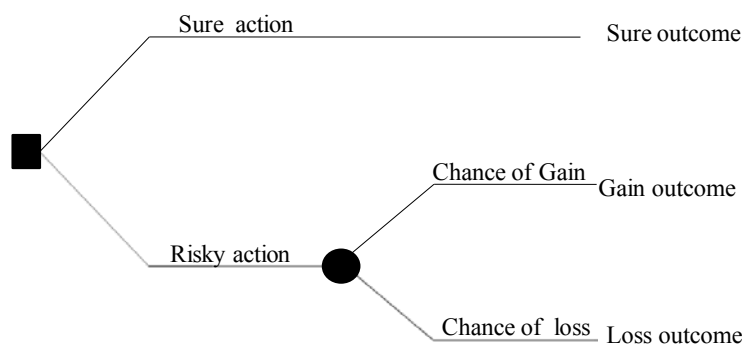


Figure 1: The basic risk paradigm

This paper deals primarily with risk assessment and techniques for processing and summarizing the probabilities and outcomes into information that can be used as the basis for evaluating alternative investments, with primary emphasis on the use of simulation for evaluating and comparing alternative investments. The rest of this paper is organized as follows. In Section 2, we discuss the main measures of performance for such investments, including net present value, internal rate of return, and modified internal rate of return. In Section 2 we also discuss types of risk and the main measures of risk, including expected present value; the mean, standard deviation, and coefficient of variation of the rate of return; and the risk premium. In Section 3 we discuss the following applications in detail: (i) stand-alone risk assessment for a capital-budgeting problem; (ii) a comparison of risk-free and risky investment strategies designed merely to keep up with the cost of living; (iii) value-at-risk (VAR) analysis for a single-stock investment; (iv) VAR analyses for two-asset portfolios consisting of stock and either call or put options; and (v) VAR analyses for two-asset portfolios consisting of both puts and calls. Finally in Section 4 we recapitulate our main conclusions. The slides for the oral presentation of this article are available online via www.ise.ncsu.edu/jwilson/files/wsc09itra.pdf, which also contains hyperlinks to the spreadsheet simulations for the applications discussed in Section 3.

2 ELEMENTS OF FINANCIAL RISK ASSESSMENT

With most financial investments, individuals or organizations spend money today with the prospect of earning even more money at some future time. So the first step in assessing risk is often to estimate the cash flows associated with the alternatives. As stated in the introduction, the estimation of the cash flows for each alternative investment is based on subjective evaluations by individuals. For the present we will assume that we have deterministic (fixed, nonrandom) estimates of all cash flows, which typically include the following:

1. Initial investment outlay—for example, (a) in the case of a project undertaken by a firm, this is the cost of acquiring fixed assets and increases to net operating working capital; or (b) in the case of a stock purchase by an individual, this is the cost of the stock.
2. Operating cash inflows and outflows over the investment's life—for example, in case (a) immediately above, this includes after-tax income and depreciation of equipment required for the project; and in case (b), this includes dividends paid to the stockholder.
3. Terminal year cash inflow and outflow—for example, in case (a), this is the salvage value of the project's fixed assets adjusted for any taxes and the return of the net operating working capital; and in case (b), this is the estimated value of a stock when it is sold.

2.1 Measures of Financial Performance

The financial performance of any investment may be measured by the investment's return, which attempts to relate the future cash flow to the expenditures required to obtain that cash flow. The simplest of the return measures is *dollar return*, which is given by the following simple relation: dollar return = (amount received) – (amount invested). This, however, gives no indication of the magnitudes of the amounts received and invested, the relative scales of the return and investment, or the timing of the return. *Rate of return*, given by the relation rate of return = (amount received – amount invested)/(amount invested), considers the return per unit of investment. Rates of return are typically expressed on an annual basis; for example, a 10% *annual rate of return* implies that after one year a \$100 investment will return \$10. The same \$100 investment which returns \$10 after five years results in a 1.92% annual rate of return because of the *time value of money*—the idea that a dollar received today is more valuable than a dollar received at a later time. More generally if P is a dollar amount received at the present time and R is the interest (or discount) rate per year, then n years into the future we see that the future value F of the present amount P is given by the relation $F = P(1 + R)^n$; conversely the present (today's) value P of the dollar amount F received n years into the future is given by the relation $P = F/(1 + R)^n$. Although the basic unit of time is a year in this discussion, we could have used a day, a week, a month, or some other time unit instead; and then the formulas given above relate the present value P and its future value F after n time periods, where R denotes the interest rate per period.

In project risk analysis and capital budgeting problems, simulation practitioners find the following measurements of return useful in ranking and selection from a set of alternative investment options:

1. Net present value (NPV) is the present value of future net cash flows, discounted at some rate, minus the initial investment.
2. Internal rate of return (IRR) is the discount rate that equates the present value of a project's cash inflows to the present value of the project's cash outflows.
3. Modified internal rate of return (MIRR) is the discount rate that equates the following: (a) the present value of the magnitudes of a project's cash outflows discounted at the firm's cost of capital, as explained in the paragraph following Figure 2 below; and (b) the present value computed at the discount rate MIRR of the project's terminal value, where the latter is the sum of the future values of the cash inflows compounded at the firm's cost of capital.

NPV is defined by the relation

$$\text{NPV} = \sum_{t=0}^n \frac{\text{CF}_t}{(1 + R)^t},$$

where n is the number of periods; CF_t is the cash flow at time t (that is, the end of period t) for $t = 0, \dots, n$; and R is the discount (interest) rate per period. When using NPV to compare several alternatives, we should of course use the same value of R to compute all relevant NPVs. Typically the current cost of capital is used, as elaborated after Figure 2 below.

The definition of IRR can be stated mathematically as the rate that yields

$$\text{NPV} = \sum_{t=0}^n \frac{\text{CF}_t}{(1 + \text{IRR})^t} = 0, \quad (1)$$

where n and CF_t are defined as above; $CF_0 < 0$, indicating the initial investment outlay; and $CF_t > 0$ for $t=1, \dots, n$. Finding the IRR for a set of cash flows requires an iterative solution method, but most spreadsheet software packages include a built-in function for computing the IRR. Note that when using IRR as a financial performance measure, we usually assume that the initial outlay is the only negative net cash flow and all subsequent net cash flows are positive.

NPV and IRR are used to rank and select investment options. Table 1 shows the initial investment and subsequent cash flows for two potential projects. An investment (cash outflow) of \$100 at time zero (today) in project A or B would return a different set of net cash flows (NCFs) over the next four years, depending on which project is selected. We assume here that the cash inflow is received at the end of the each year. If we rank projects A and B based on NPV computed at the discount rate $R = 10\%$, then project A outranks (is preferred to) project B because $NPV_A > NPV_B$. With respect to IRR, project A also outranks project B because $IRR_A > IRR_B$.

Table 1: Examples of projects that have the same initial investment, but different cash flow streams, NPVs, and IRRs

Project A							
Time Period	0	1	2	3	4	NPV @ 10%	IRR
Net cash flow	-100	50	40	30	10		
Cumulative NCF	-100	-50	-10	20	30		
Discounted NCF (@ 10%)	-100	45.5	33.1	22.5	6.8	7.9	14.5%
Cumulative discounted NCF	-100	-54.5	-21.4	1.1	7.9		
Project B							
Time Period	0	1	3	4	NPV @ 10%	IRR	
Net cash flow	-100	10	30	40	60		
Cumulative NCF	-100	-90	-60	-20	40		
Discounted NCF (@ 10%)	-100	9.1	24.8	30.1	41.0	5	11.8%
Cumulative discounted NCF	-100	-90.9	-66.1	-36.0	5.0		

For selecting projects, NPV is a better comparison metric than IRR in some respects; however, IRR is widely used in industry. An *NPV profile*, a graph of the NPV of each alternative project as a function of the discount rate R , demonstrates why NPV may be a superior in some cases. For example, in Figure 2 notice that the curve for NPV_A crosses over the curve for NPV_B at the *crossover rate* of 7.2% so that when $R < 7.2\%$, we have $NPV_B > NPV_A$ and project B is preferred; on the other hand when $R > 7.2\%$, we have $NPV_A > NPV_B$ and project A is preferred.

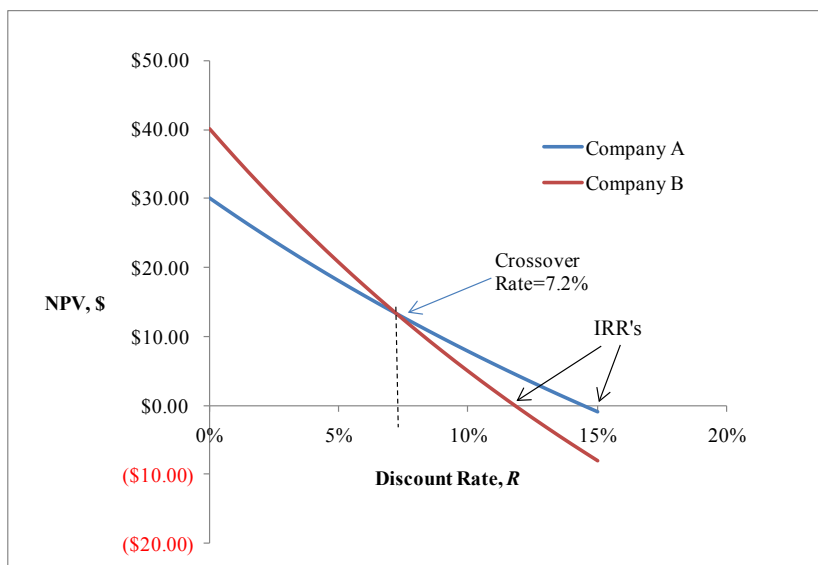


Figure 2: Net present value profiles of projects A and B

To continue our discussion of the merits of NPV *versus* IRR as measures of financial performance, we must consider the *cost of capital* k , which for simplicity we define as the interest rate at which the individual or firm must borrow to finance an investment; for more on this subject, see Chapter 9 of Brigham and Houston (2002). If a single independent project is being evaluated without considering other alternatives, then both IRR and NPV lead to the same decision—accept the project if $NPV > 0$ (or equivalently, $IRR > k$) and reject the project otherwise. However, in the case where A and B are mutually exclusive alternatives and k is used as the discount rate, then if $k < 7.2\%$, we see from Figure 2 that the NPV method chooses B, whereas the IRR method *always* chooses A regardless of the value of k , provided that $k < IRR_A = 14.5\%$. It follows that NPV is the preferred measure of financial performance because IRR can select the inferior alternative when the relevant discount rate is on the “wrong” side of the crossover point. Notice also that compared with the curve for NPV_A , the curve for NPV_B is more sensitive to changes in the discount rate R because of its steeper slope. The greater sensitivity to the discount rate is caused by the longer-term nature of project B compared with project A (that is, the bulk of the cash inflows for project B occur later than the bulk of the cash inflows for project A); and this phenomenon is sometimes interpreted as an indication that B is a higher risk than A.

Another complication in using IRR arises when a project has both inflows and outflows of cash over the course of the project so that the net cash flow changes sign at least twice. For example beyond the initial investment for construction of an offshore oil production platform (a negative cash flow at the beginning), the platform must be dismantled at the end of its productive life; and the dismantling cost will cause a negative cash flow after all positive cash flows have been realized. A cash flow is said to be *nonnormal* if it has two more changes in sign, which means that Equation (1) can have multiple real roots so that there are multiple possible values for IRR. To eliminate this problem, a *modified internal rate of return* (MIRR) can be calculated. From item 3 in the list of financial measures given at the beginning of this section, we see that MIRR is the rate that satisfies the equation

$$NPV = \sum_{t=0}^n \frac{CIF_t(1+k)^{n-t}}{(1+MIRR)^n} + \sum_{t=0}^n \frac{COF_t}{(1+k)^t} = 0 \quad \text{or} \quad \sum_{t=0}^n \frac{CIF_t(1+k)^{n-t}}{(1+MIRR)^n} = \sum_{t=0}^n \frac{|COF_t|}{(1+k)^t}, \quad (2)$$

where: n is the number of time periods; k is the cost of capital; COF_t is the cash outflow (a negative number) at time t so that $|COF_t| = -COF_t$ is the magnitude of the cash outflow at time t for $t = 0, 1, \dots, n$; and CIF_t is the cash inflow (a positive number) at time t for $t = 0, 1, \dots, n$. Equation (2) has a unique real root, which is the MIRR.

2.2 Types of Financial Risk

In addition to defining the metrics of evaluation, we must also identify the type of risk we are trying to quantify and manage. Most financial texts recognize three types of risk:

1. An investment’s *stand-alone risk* is the risk associated with a single investment opportunity, ignoring the fact that the investment may represent only one of several assets held by a firm or individual.
2. *Diversifiable risk* is the type of risk that can be diluted by viewing one investment as part of a group that may contain other investments, of which some may be less risky and some may be more risky compared with the single investment under consideration.
3. *Market, or beta, risk* is the type of risk that cannot be eliminated by diversification because it stems from pervasive factors that typically affect most investments—for example, inflation, recession, and high interest rates.

Diversified risk can be managed in a properly chosen set of investments because by definition it is risk that can be reduced by combining investments that are negatively correlated or that at least are not perfectly positively correlated. (The phenomenon of correlation between investments is discussed in detail in Section 2.3 below.) Given proper diversification, market risk outweighs diversifiable risk as an issue of concern. In investment evaluations (whether projects or stock purchases), stand-alone risk may sometimes serve as a surrogate for harder-to-measure corporate risk and market risk. If the economy does well (as measured by market risk), then usually firms do well (as measured by corporate risk); and if a firm does well, then it can better control a project (as measured by stand-alone risk). In other words, the three types of risk are often closely related.

We note that most investors are *risk averse*, preferring to avoid risk and thereby requiring the promise of higher rates of return on risky investments. The *risk premium*, RP , is the difference between the expected rate of return on a given risky asset and the return on a less risky asset—that is, the additional return that investors require for assuming additional risk.

2.3 Measures of Risk

Because risk assessment must account for the value (gain or loss) resulting from each outcome as well as the corresponding probability of occurrence of that outcome, an appropriate performance measure must account for both of these quantities over the “sample space” of all possible outcomes. One such appraisal mechanism is the *expected present value*,

$$EV = \sum_{g=1}^m p_g v_g,$$

where the number of possible outcomes is denoted by m ; and for $g = 1, \dots, m$, the value of outcome g is denoted by v_g and the probability of outcome g is denoted by p_g . For a risky investment, the performance measure analogous to rate of return is the *expected rate of return*,

$$\hat{k} = \sum_{g=1}^m p_g k_g, \tag{3}$$

where k_g denotes the rate of return of outcome g for $g = 1, \dots, m$.

At this point we have to assign probabilities to each potential outcome and compute the expected rate of return as in Equation (3) via a *payoff matrix*. If we are computing NPV and IRR from a cash-flow analysis similar to that given in Section 2.1, then using this approach requires considering the probability distribution of all possible outcomes and the associated probability distribution of monetary gains and losses for each alternative investment. For an example in which there are three possible demand levels in the overall market for the competing products of companies A and B, see Table 2, which also gives the associated probabilities and the anticipated rates of return for each company and for each level of demand in the overall market.

Table 2: Computation of expected rates of return on two stock investments

	Company A		Company B
	Probability p_g of	Rate of Return k_g for	Rate of Return k_g for
Demand Level g	Demand Level g	Demand vel g	Demand Level g
Strong	0.3	100%	20%
Normal	0.4	15%	15%
Weak	0.3	-70%	10%
Expected Rate of Return \hat{k}		15%	15%
Standard Deviation of Return σ		66%	4%

Because the expected returns from these two alternative investments are equal, we are still in need of a way to compare the relative riskiness of these opportunities. If the alternative investments have the same rate of return, then the risk-averse investor will choose the alternative with the lower variability in the rate of return as measured by the standard deviation of the rate of return computed over all possible outcomes,

$$\sigma = \sqrt{\sum_{g=1}^m p_g (k_g - \hat{k})^2},$$

where \hat{k} is the alternative’s expected rate of return as given by Equation (3). The investment with the smaller standard deviation σ in its rate of return is usually regarded as being more dependable and hence less risky. If we do not know the probabilities associated the various outcomes, then we may estimate \hat{k} and σ using the sample mean and the sample standard deviation from historical rate-of-return data as follows,

$$\hat{k}_{\text{est}} = \frac{1}{n} \sum_{t=1}^n \bar{k}_t \quad \text{and} \quad \sigma_{\text{est}} = \sqrt{\frac{1}{n-1} \sum_{t=1}^n (\bar{k}_t - \hat{k}_{\text{est}})^2},$$

where \bar{k}_t denotes the past realized average rate of return in year t for $t = 1, \dots, n$ so that \hat{k}_{est} is the sample average annual return earned in the last n years and σ_{est} is the sample standard deviation of the annual return over the same n years.

If we are choosing between two investments with the same expected return but with different standard deviations as in Table 2, then we prefer the less-risky alternative, that is, the investment with less variability in its rate of return. If we are comparing two investments with the same riskiness (standard deviation) but with different expected returns, then we prefer the investment with the higher expected return. For the situation described in Table 2, we conclude that investing in company B is less risky than investing in company A, which exhibits substantial variability in its rate of return. For company B, the probability that the actual rate of return will be close to its expected value of 15% is much higher than for company A. In the case where one alternative has the higher expected return and the other alternative has the lower standard deviation, we use the *coefficient of variation* defined by $CV = \sigma / \hat{k}$, which gives the risk per unit of return, and provides a basis for comparison when two alternatives are otherwise not directly comparable.

Next we consider measures of diversifiable risk and market risk in a stock-investment context. In a diversified portfolio, we again begin by determining the expected returns of each asset. Then we compute the portfolio's expected return, which is simply the weighted average of the expected returns on the individual assets,

$$\hat{k}_P = \sum_{i=1}^a w_i \hat{k}_i,$$

where a is the number of assets in the portfolio; and w_i is the fraction of the portfolio's monetary value that is invested in stock (asset) i for $i = 1, \dots, a$. For asset i , we let \bar{k}_i denote the *realized rate of return* over some previous period; and we note that \bar{k}_i usually deviates somewhat from the expected rate of return \hat{k}_i . Similarly, the portfolio's realized rate of return,

$$\bar{k}_P = \sum_{i=1}^a w_i \bar{k}_i$$

(i.e., the return that the portfolio actually earned during the previous period under consideration), usually deviates from the portfolio's expected return, \hat{k}_P . However, the realized rates of return $\{\bar{k}_1, \dots, \bar{k}_a\}$ on the individual assets in the portfolio, although perhaps deviating substantially from their individual expected rates $\{\hat{k}_1, \dots, \hat{k}_a\}$, may tend to have deviations that offset (cancel) each other and result in a realized portfolio rate \bar{k}_P that is reasonably close to the expected portfolio rate \hat{k}_P .

Computing the standard deviation of a portfolio's return is slightly more complicated than the analogous computation for a stand-alone investment because of the possible existence of correlation between the returns on the different assets in the portfolio. The covariance between the returns X and Y on two investments is given by $\sigma_{XY} = \text{Cov}(X, Y) = E[(X - \mu_X)(Y - \mu_Y)]$, where μ_X and μ_Y denote the expected value of the returns X and Y , respectively. The correlation between the two returns is defined to be $\rho_{XY} = \sigma_{XY} / (\sigma_X \sigma_Y)$, where σ_X and σ_Y denote the standard deviations of the returns X and Y , respectively. For a portfolio in which w and $1-w$ are the fractions of the portfolio's value invested in X and Y , respectively, the variance of the portfolio's return is

$$\sigma_P^2 = \text{Var}[wX + (1-w)Y] = w^2 \sigma_X^2 + (1-w)^2 \sigma_Y^2 + 2w(1-w) \sigma_{XY}. \quad (4)$$

If the two returns X and Y were perfectly positively correlated (so that $\rho_{XY} = 1.0$), then diversification would *increase* the risk of the resulting portfolio compared with investing everything in the stock with the smaller standard deviation. If the returns X and Y were perfectly negatively correlated (so that $\rho_{XY} = -1.0$), then theoretically diversification could lead to the

elimination of all risk, at least in the case that $w = \sigma_Y / (\sigma_X + \sigma_Y)$ so that $\sigma_P^2 = 0$. Brigham and Houston (2000) state that “on average the correlation coefficient for the returns on two randomly selected stocks would be about +0.6, and for most pairs of stocks, ρ_{XY} would lie in the range of +0.5 to +0.7.” Table 3 displays the average annual returns over a five-year period for two stocks C and D whose correlation coefficient is 0.67. Notice that although both stocks have not only the same expected return but also the same standard deviation, a portfolio containing both stocks has a smaller standard deviation than either of the individual stocks. This example shows the benefit of diversification—although risk cannot be eliminated, it can be reduced by diversification.

Table 3: Comparison of expected returns and standard deviations of two stocks and a portfolio consisting of both equally

Year	Stock C	Stock D	Portfolio 50% Invested in both C and D
2004	40.0%	28.0%	34.0%
2005	-10.0%	20.0%	5.0%
2006	35.0%	41.0%	38.0%
2007	-5.0%	-17.0%	-11.0%
2008	15.0%	3.0%	9.0%
Average Return	15.0%	15.0%	15.0%
Standard Deviation	22.6%	22.6%	20.6%

If we have a stocks with returns $\{X_1, \dots, X_a\}$ where the values of the stocks respectively represent the fractions $\{w_1, \dots, w_a\}$ of the portfolio’s total value, then the variance of the portfolio’s overall return is an extension of Equation (2),

$$\sigma_P^2 = \text{Var} \left[\sum_{i=1}^a w_i X_i \right] = \sum_{i=1}^a w_i^2 \text{Var}(X_i) + 2 \sum_{i=1}^{a-1} \sum_{j=i+1}^a \text{Cov}(X_i, X_j).$$

As the number of stocks in a portfolio increases, the risk in the portfolio generally decreases to some limit—namely, the risk in a portfolio consisting of all stocks in the market. The return k_M of a portfolio consisting of all stocks (a *market portfolio*) has a standard deviation σ_M approximately equal to 20.4%; and this is much smaller than the standard deviation of an average stock, which historically has been approximately 35% (Brigham and Houston 2000). This market risk, also called *systematic risk*, is the risk that all stock investors must bear. For an individual security, the systematic risk is measured by its “ β coefficient,” which describes how the stock’s return moves in relation to the return of the general market as gauged by some market index such as the Dow Jones Industrial Average. A stock’s coefficient $\beta = 1.0$ implies that if the market increases by 20%, then the value of the individual stock will also increase by 20%. See Figure 3 for theoretical examples of high, average, and low volatility stocks.

For the i th stock in a portfolio, its β coefficient may be expressed in terms of the mean and standard deviation of its return as well as its covariance with the market return and the standard deviation of the market return as follows,

$$\beta_i = \rho_{i,M} \sigma_i / \sigma_M = \text{Cov}(k_i, k_M) / \sigma_M^2,$$

where: k_i is the return on stock i ; σ_i is the standard deviation of the return on stock i ; k_M is the market return; σ_M is the standard deviation of the market return; and $\rho_{i,M}$ is the correlation between the return of stock i and the market return. A portfolio consisting of stocks with all coefficients $\beta_i = 0.5$ can be only half as volatile as the market, provided that the correlations between the returns of the stocks in the portfolio are not too large. If we plot the return of stock i as a function of the market return, then we see that β_i is the slope of the regression line for which the return of stock i is the dependent variable (plotted on the vertical axis) and the market return is the independent variable (plotted on the horizontal axis). The β coefficient

cient of a portfolio is the weighted average of its individual β_i 's. Therefore, since the coefficient β_i for stock i determines how the stock affects the riskiness of a diversified portfolio, it is the most relevant measure of the stock's risk.

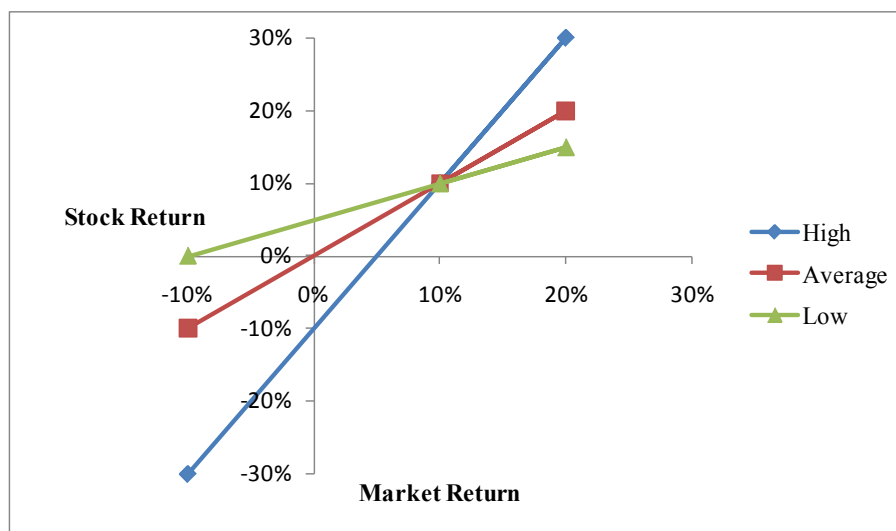


Figure 3: Examples of high ($\beta = 2.0$), average ($\beta = 1.0$), and low ($\beta = 0.5$) volatility stocks

The riskiness of an investment as measured by its β coefficient is related to a level of return that we would require for compensation of the risk assumed. We measure the premium that investors require for bearing the risk of an average stock as

$$RP_M = k_M - k_{RF},$$

where: k_M is the required rate of return on a portfolio consisting of all stocks and is also the required rate on an “average” stock (that is, a stock with $\beta_{avg} = 1.0$); and k_{RF} is the risk free rate of return, usually measured by the return on long-term treasury bonds. The risk premium RP_M measures the additional return above the risk-free rate that is required to compensate for the “average” (that is, market) risk. Unfortunately, it is hard to estimate k_M .

Assuming that we have an estimate for RP_M , then an expression for the risk premium for investment i is

$$RP_i = RP_M \beta_i.$$

It follows that the required return on investment i is

$$k_i = k_{RF} + RP_M \beta_i. \quad (5)$$

The Security Market Line (SML) given by Equation (5) specifies the relationship of the required rate of return on a risky investment and the investment's β coefficient. The SML is depicted in Figure 4. The SML and an investor's position on it may change over time because of changes in interest rates, the investor's aversion to risk, and individual investments' β values. Inflation will also affect the risk-free rate, k_{RF} . In fact, we may think of k_{RF} as consisting of a “real” inflation-free rate, k^* , and an inflation premium, IP, that is the additional premium required by the investor to compensate for inflation so that we have $k_{RF} = k^* + IP$. The “real” inflation-free rate k^* is typically between 2% and 4%.

This discussion leads us to the following questions:

1. What do we use as a discount rate when computing NPV?
2. What is the threshold value for the rate of return that is necessary for acceptance of an investment option?
3. How is risk related to the rate of return required to make an investment attractive?

As mentioned in Section 2.1, the cost of capital to the firm or individual is often used as the discount rate for the calculation of NPV and is referred to as a “hurdle rate” in that an investment’s rate of return should exceed this rate in order to be accepted. For more detail on how to arrive at an appropriate cost of capital to use as the hurdle rate, see, for example, Brigham and Houston (2000). In this paper, we will proceed on the assumption that this rate has already been determined. If the IRR exceeds the cost of capital, then of course $NPV > 0$ when computed at the cost of capital. We will also assume that in the case of diversified investments, we have determined a risk-adjusted rate of return.

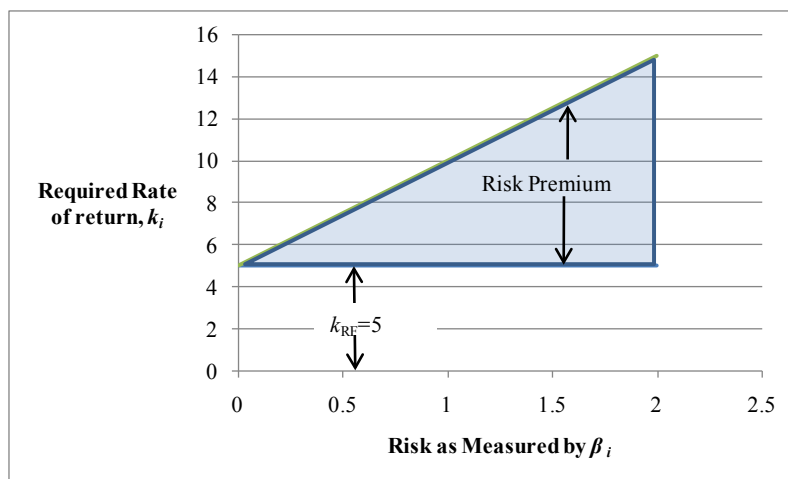


Figure 4: Security market line

3 APPLICATIONS

3.1 Capital Budgeting Problem in Stand-Alone Risk Context

A typical application of stand-alone risk assessment is found in the case of a firm that is considering investing in a project. This may be the development of new product, entrance into a market previously untapped, the building of a new manufacturing facility, or the expansion of an existing facility, etc. In order to analyze whether an investment is worthwhile, the company must estimate many of the input parameters. Gathering this data may involve the use of historical data (costs, sales prices, inflation rates, etc.) as well as other sources (market research, etc.); and some of the required inputs may be only someone’s best guess. Of course, all these input parameters are subject to uncertainty and variability. Typically, the analysis will begin with some sort of “base case.” This case may represent one possible outcome on a decision tree and is often the “most likely” scenario.

Table 4 shows a base case for a hypothetical company considering the construction of a new production facility. The input parameters are the best guesses at current costs, prices, demand, etc.; demand growth rate is zero; and inflation of costs and product price is moderate. Of course, the outcome is uncertain, and we must be compensated for the risk assumed. Therefore we must modify the discount rate used for NPV and the threshold rate used for IRR. In essence, we require a risk-adjusted discount rate for the NPV analysis and risk-adjusted threshold for IRR.

The *certainty equivalent* approach to this problem involves scaling down cash flows that are not certain. This is based largely on subjectively judging which CFs to scale down and by how much. We could also use a *risk-adjusted* rate of return—average risk projects are discounted at an average cost of capital; less risky ones, at a lower-than-average cost of capital; etc. But we are again forced to choose somewhat arbitrarily how much adjustment must be made.

After estimating and evaluating a base case, the next step is to run a worst-case scenario where, for example, initial demand is lower, fixed and variable costs are higher, and high inflation affects costs but not sales price. We also ran a best-case scenario where demand is initially higher and grows over time, initial fixed cost is lower, there is no inflation on costs, and there is moderate inflation on product price. Table 5 displays the output from a “scenario manager” typically available with spreadsheet packages. Again, this approach requires changing all input parameters simultaneously for some limited number of possible outcomes.

If we wish to arrive at an expected NPV or expected rate of return from these three scenarios, then we must somehow assign probabilities to the worse, base, and best cases and compute the expected outcomes based on these probabilities.

However, the estimates are crude, because we may not have a good idea of the probability that the worst-, base-, or best-case values of *all* input parameters will be achieved *simultaneously*.

Table 4. Base case NPV analysis of potential production facility

Discount Rate	10%	Initial Variable Cost/Unit	\$1.00			
Initial Annual Prod. Qty.	21,500	Variable cost inflation Rate	3%			
Demand Growth Rate	0%	Initial Product Sale Price	\$50			
Initial Fixed Cost	\$200,000	Product Price Inflation Rate	3%			
Fixed Cost Inflation Rate	3%	Salvage	\$100,000			
Year	0	1	2	3	4	5
Capital Exp.	\$3,000,000					
Prod. Qty. (units)	21,500	21,500	21,500	21,500	21,500	21,500
Fixed Costs	\$200,000	\$206,000	\$212,180	\$218,545	\$225,102	\$225,102
Variable Costs	\$21,500	\$22,145	\$22,809	\$23,494	\$24,198	\$24,198
Total Costs	\$221,500	\$228,145	\$234,989	\$242,039	\$249,300	\$249,300
Revenue	\$1,075,000	\$1,107,250	\$1,140,468	\$1,174,682	\$1,209,922	\$1,209,922
Salvage						\$100,000
Cash Flow	-\$3,000,000	\$853,500	\$879,105	\$905,478	\$932,642	\$1,060,622
Disc. CF	-\$3,000,000	\$775,909	\$726,533	\$680,299	\$637,007	\$658,563
NPV @ 10% Discount Rate =	\$478,311		IRR = 15.8516%			

Table 5: Sample output from Excel's Scenario Manager for the production facility example

Scenario Summary	Current Values:	Worst	Best
Changing Cells:			
Discount Rate	10%	10%	10%
Initial Annual Production Quantity	21500	21500	15000
Demand Growth Rate	0%	0%	5%
Initial Fixed Cost	\$200,000	\$220,000	\$180,000
Fixed Cost Inflation Rate	3%	10%	0%
Initial Variable Cost/Unit	\$1.00	\$5.00	\$0.75
Variable cost inflation Rate	13%	10%	0%
Initial Product Sale Price	\$50	\$30	\$75
Product Price Inflation Rate	3%	0%	3%
Salvage	\$100,000	\$90,000	\$110,000
Initial Capital Expenditure	\$3,000,000	\$500,000	\$2,700,000
Result Cells:			
NPV @ 10% Disc Rate	\$461,162	-\$3,079,172	\$8,114,226
IRR	15.6585%	<<0%	93.4354%

Additionally we might want to examine the sensitivity of NPV and IRR to various input parameters, considered one at a time. There are several ways to do this. We can use the scenario manager to generate the information in a tabular format similar to Table 5. Alternatively, after Excel's scenario manager is used to generate NPV and IRR at various changes to the input parameters, we can display the results as shown in Figure 6.

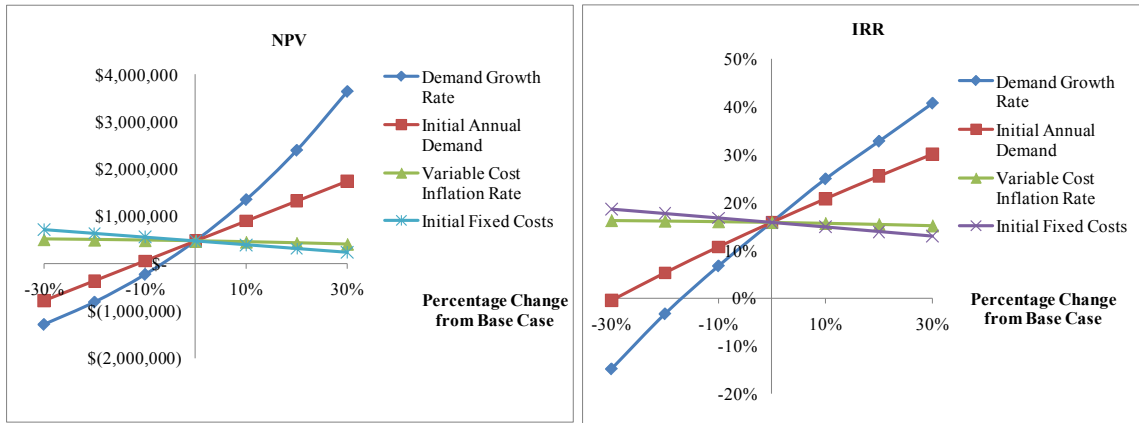


Figure 6: Sensitivities of production facility NPV and IRR to input parameters

Simulation enables us to obtain substantially more information than can be obtained from a small number of user-defined scenarios such as those depicted in Figure 6, where the input parameters (demand growth rate, initial annual demand, variable cost inflation rate, and initial fixed costs) are varied one-at-a-time with all other input parameters fixed. Using simulation, we may generate a large number of observations of the NPV and IRR metrics computed from a set of randomly generated scenarios in which all input parameters subject to uncertainty are simultaneously varied. We achieve this by defining appropriate probability distributions for those input parameters based on whatever information is available about those parameters. Then on each replication of the simulated project, we randomly sample new values of those input parameters from their corresponding distributions. For example, in the new production facility problem, we obtained subjective estimates (that is, estimates based on expert opinion) of the minimum, most likely, and maximum values for each input parameter listed as a “Changing Cell” in Table 5; and then we fitted a beta distribution to these three estimates using the approach detailed in Kuhl et al. (2009a, b). (Note that each resulting beta distribution is *not* the so-called “beta-PERT” distribution or the “RiskPert” distribution; for more on this issue, see Kuhl et al. (2009a, b).)

We implemented the fitted beta distributions using the “Define Assumption” option in Crystal Ball (Oracle Corporation 2008); and we performed 100,000 independent replications (runs) of the project simulation. This simulation experiment yields a multiway probabilistic sensitivity analysis that reveals the uncertainty (lack or precision) in the final results. Cumulative frequency histograms for NPV and IRR from the Crystal Ball simulation of 100,00 independent replications are shown in Figures 7 and 8.

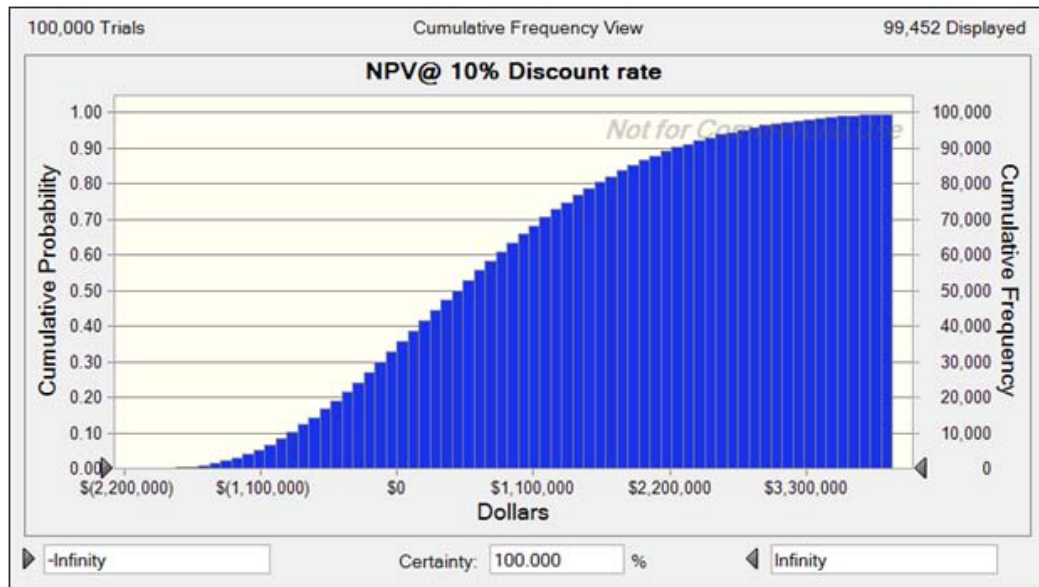


Figure 7: Crystal Ball frequency distribution of NPV for the production facility

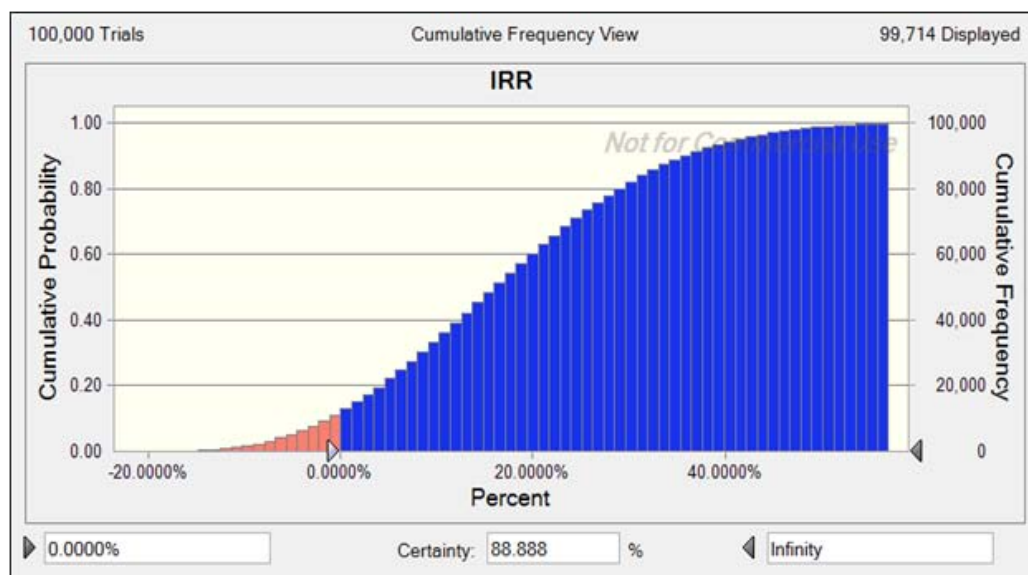


Figure 8: Crystal Ball frequency distribution for the IRR for the production facility

The probability that the output measure of interest exceeds the value in the lower left-hand side box (in Figure 8, cutoff value IRR = 0.0%) is displayed in the “Certainty” box (in Figure 8, Certainty = 88.9%). In other words, over 100,000 runs of the simulation, we observed that on 88.888% of those runs, condition $IRR > 0$ occurred. We may change the values in the lower far left- and right-hand side boxes to view the relative frequencies of IRR values below or above any desired cutoff value.

3.2 Risky vs. Risk-Free Investment Comparison

This example (Strong 2008) illustrates how simulation using Crystal Ball (Oracle Corporation 2008) can facilitate complex stock market scenario analysis. Suppose there are two available asset classes, large-capitalization common stock and long-term U.S. Government treasury bonds. A risk-averse investor wants to distribute his money across these two asset classes so as to maximize the likelihood that he at least earns the rate of inflation over a twenty-year time horizon. The objective is *not* to achieve the largest terminal portfolio value; it is to keep up with the cost of living.

Table 6 shows historical means and standard deviations for the annual rates of return on the asset classes and for the annual inflation rate. As is standard practice, we assume that these rates are distributed as normal random variables in the simulation. Table 7 gives the correlation coefficients for the three rates. The simulation also includes the historical serial correlation of 0.65 for the inflation rate. These statistics are from Ibbotson Associates (2005).

Table 6: Asset class statistics

	Mean	Standard Deviation
Return on Large-Cap Stocks	12.4%	20.3%
Return on Government Bonds	5.8%	9.3%
Inflation	3.1%	4.3%

Table 7: Correlations

	Large-Cap Stocks	Government Bonds	Inflation
Return on Large-Cap Stocks	1.00		
Return on Government Bonds	0.12	1.0	
Inflation	-0.02	-0.14	1.0

The simulation investigates 101 different asset allocations of an initial \$10,000 investment, ranging from 0% stock and 100% bonds to 100% stock and 0% bonds, with allocations changing by 1% increments. The distribution surrounding the

stock return, bond return, and actual inflation are modeled according to the data in Tables 6 and 7. The Crystal Ball model is set to run 100,000 replications of each portfolio. Each replication covers a 20-year time period, and therefore replication requires the sampling of 20 stock returns, 20 bond returns, and 20 inflation values. A total of 6 million random samples are generated from the underlying probability distributions for each of the 101 scenarios. On a 1.6 GHz computer with 2 GB of RAM, each replication takes approximately 28 seconds to complete. Figure 9 is an example of the Crystal Ball output from one portfolio combination.

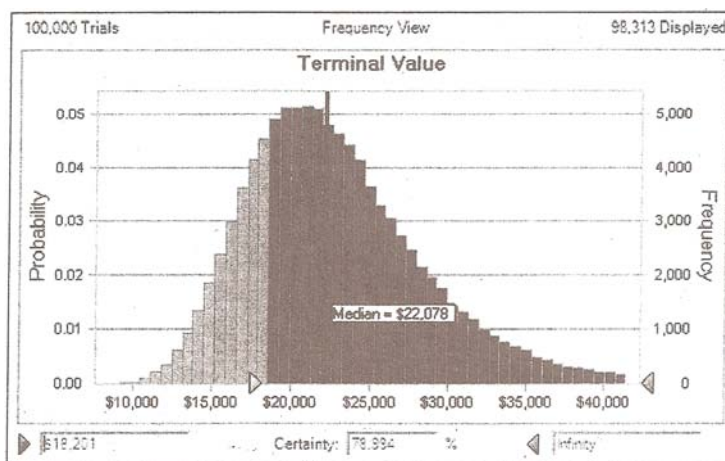


Figure 9: Sample Crystal Ball output of one asset allocation scenario

Table 8 shows a portion of the results for various asset allocations. We make the following observations about column (e) in Table 8: with the introduction of a relatively small percentage in stock, there is a dramatic increase in the probability that the terminal portfolio value (column (c)) exceeds inflation (column (d)); moreover, this probability levels off once the portfolio is about one-third invested in stock. Figure 10 shows the complete scenario results graphically.

Table 8: Sample simulation results for varying asset allocations

Percentage in stock (a)	Percentage in bonds (b)	Terminal Value median (c)	Inflation Threshold (d)	Probability Terminal Value > Inflation (e)
0	100	\$18526	\$18164	53.8%
5	95	22066	18189	79.0
10	90	24963	18184	87.0
15	85	27628	18189	90.1
20	80	30330	18158	91.9
25	75	33097	18183	92.9
30	70	35926	18208	93.4
35	65	38679	18190	93.8
40	60	41458	18137	94.1
45	55	44145	18173	94.3
50	50	46932	18176	94.4
55	45	49701	18166	94.5
60	40	52521	18175	94.5
65	35	55008	18175	94.5
70	30	58168	18169	94.5
75	25	61066	18179	94.6
80	20	63775	18151	94.7
85	15	66599	18118	94.7
90	10	69405	18194	94.5
95	5	72208	18174	94.6
100	0	74824	18181	94.7




Figure 10: Probability of return on portfolio exceeding the inflation rate

3.3 Determination of Value at Risk (VAR) Using Simulation

Another important use of Monte Carlo simulation in financial risk assessment lies in the calculation of *Value at Risk* (VAR). This concept has become increasingly important as the financial markets and their regulators pick up the pieces after the subprime disaster and subsequent market meltdown.

In any kind of investment strategy, a very reasonable question is “What is the most we can lose?” For years a useful answer to this question was elusive. Theoretically, if we buy a U.S. Treasury bond we could lose 100% of our money, but this is highly unlikely and not a particularly helpful answer. It would not be informative to use “all our money” as the default answer to the question regardless of the asset in question. Losing all our money on a penny-stock speculation is a much more likely outcome than losing all our money with a government bond investment.

A better answer would incorporate an estimate of what *realistically* can go wrong over a *particular time horizon*. If we spend \$1,500 on 100 shares of Harley-Davidson stock today, we may lose all our money, but it probably will not happen this week or this month. A more precise question might be “What is the most we can lose in one day?” Again, the answer is 100%, but in the 5,455 trading days between November 5, 1987, and July 6, 2009, there were only 15 days in which the stock fell by 10% or more. There were only 122 days when it fell by 5% or more. In fact, the average daily change was 0.099% with a standard deviation of 2.51%. Such an investment becoming worthless overnight would be a $40\text{-}\sigma$ event (that is, an extremely rare event).

A convention has developed within the finance industry to quote VAR on the basis of the distribution of daily price movements. When we ask “What is our value at risk?” we are asking for the size of the loss associated with a 5% tail of daily price changes. In other words there is a 95% probability that the loss in one day will not exceed the VAR. In the Harley Davidson example, 1.96 standard deviations of 2.51% each results in a 4.92% price movement. Therefore on a \$1,500 investment, the VAR is about \$74. A 100-share investment certainly has down days when it loses more than this, but over the past twenty plus years they have not been common. VAR has become the standard currency in answering the question about what can go wrong.

3.3.1 Single-Stock Investment VAR Using Simulation

Consider an investment position consisting solely of 10,000 shares of a stock, currently selling for \$55 per share, for a total value of \$550,000. We can simulate performance of this portfolio over a one-day period using the same historical input parameters as given in Table 6; that is, average annual return equals 12.4% and average annual standard deviation is 20.3%. There are typically about 252 trading days during the year, so to convert these annual values into daily values we divide the

mean by 252 and the standard deviation by the square root of 252, resulting in 0.0452% and 1.2851%, respectively. Again, assuming return rates are normally distributed and performing 100,000 replications yields the results in Table 9.

Table 9: Example of single-stock investment VAR

Initial Purchase			Results after 100,000 replications of one simulated day			
Quantity (Shares)	Price	Value	Average Price	Price Range	5 th Percentile Value	VAR
10,000	\$55	\$550,000	\$55.03	\$46.09–\$63.86	\$515,041	\$34,959

The mean stock price after one day being slightly above the starting price is not surprising given that we input a positive mean return. The stock price range, however, is more than \$8 on either side of the initial \$55. Price swings of this magnitude do not happen very often, but most investors can recall instances when they experienced such an event. It is a simple matter with Crystal Ball to find the portfolio value beyond which 5% of the adverse distribution lies. Note that the VAR concept is one-tailed; good days in the market (i.e., a big price rise if you own the stock) are not cause for concern.

3.3.2 Portfolio VAR

The plot thickens when we consider that we seldom hold investments in isolation; they are normally part of a portfolio. Sometimes the portfolio components have market risk sensitivities that are opposite in sign, such as Harley Davidson’s put and call options. A *call option* gives an investor the right (but not the obligation) to buy the underlying asset at a set price for a set time period. A *put option* is similar except it is the right to sell. Certain factors, such as the price of the underlying asset, that would push a call price *up* would *depress* the price of a put. Other factors, such as the passage of time, affect puts and calls in a similar direction; both types of options lose value as their expiration approaches.

Those factors which may be negatively or positively correlated to each other have implications in a VAR context. In a slow market with little volatility, there may be days in which both the puts and the calls decline in value because they are close to expiration and some of the “time value” has evaporated. On an active day, if the market moves sharply up the call will *increase* in value while the put will *decrease*. Conversely, with a big down day the puts will rise in value while the calls will decline. Therefore a portfolio holding some of each type of asset cannot lose on both when there is a significant change in the value of the underlying asset. Logically, the answer to the question “What is the worst that can happen?” should recognize this fact rather than simply aggregating the individual VARs.

A merger arbitrage hedge fund is one type of institutional player for which the VAR concept is especially relevant. This investor seeks opportunities arising because of the planned acquisition or merger of two (or more) companies. Often the stock of the acquiring firm declines while the stock of the acquired firm rises. The hedge fund may have an array of stock option holdings (puts and calls) on both of these firms as well as positions in the actual shares. The merger process is filled with uncertainty pending the resolution of regulatory concerns, shareholder approval, potential for competing bids, etc. Using Monte Carlo simulation, we can test various scenarios that might develop and see their consequences in a VAR framework.

Simulation results provide very helpful information to those involved with risk management and strategy oversight. Given a value for the underlying asset, the option valuation is relatively straightforward (although complex) and is likely to be determined from the Black-Scholes option pricing model. If C is the current value of the call option, S is the current stock price, K is the exercise (or strike) price of the option (that is, the set price at which the owner of the call has the right to purchase the underlying asset), and T is the remaining time until the option expires (in years), then the Black-Scholes model is given by

$$C = S\Phi(d_1) - K\exp(-k_{RF}T)\Phi(d_2), \tag{5}$$

where $\Phi(z) = \int_{-\infty}^z \exp(-u^2/2) du$ denotes the cumulative distribution function for the standard normal distribution, and

$$d_1 = \left\{ \ln(S/K) + \left[k_{RF} + \left(\sigma^2/2 \right) \right] T \right\} / (\sigma\sqrt{T}), \left. \begin{array}{l} \\ d_2 = d_1 - \sigma\sqrt{T}, \end{array} \right\}$$

and as usual k_{RF} denotes the risk-free interest rate and σ denotes the estimated standard deviation of the stock's rate of return over the life of the option. See Black and Scholes (1973) for a derivation of Equation (5).

VAR becomes an increasingly important concept when derivatives are present in a portfolio. We will now introduce call options into the portfolio. There are a variety of reasons why an investor might do this, but suppose we sell call options on the full 10,000 shares giving someone else the right to purchase those shares at \$55 by a set future date. This is the most common use of options and is referred to as *writing a covered call*. Options are valued using the Black-Scholes option pricing model, Equation (5). Once we have the value of the call option, we can use the put-call parity model,

$$C - P = S - K\exp(-k_{RF}T),$$

to determine the value P of the put. Complete coverage of option pricing is beyond the scope of this paper, but these two models help demonstrate the mathematical complexity of complicated derivative portfolios that scenario analysis can help mitigate.

In our example, we let $S = K = \$55$, $k_{RF} = 6\%$, $T = 87$ days or 0.2384 years, and $\sigma = 34\%$. Solving (5), we obtain $C = \$4.00$ and $P = \$3.25$. To calculate the VAR with this two-asset portfolio, we also describe a probability distribution around the estimate of σ . Future volatility is the one variable in the pricing model that is not observable and must be estimated. The model is quite sensitive to the input used. We initially chose a triangular distribution with a mean of 34% and a min and max 10% (or 3.4 percentage points) on either side of 34% to model the distribution of σ . Results of this simulation are in Table 10.

Note that the VAR of this portfolio, which contains options, is *less* than the stock-only portfolio. Table 10 has an additional column in the far right indicating the portfolio delta. Delta is the first derivative of the Black-Scholes model with respect to the underlying stock price. Another convenient interpretation of position delta is the common stock equivalency of the aggregate portfolio. In Table 10, a position delta of 4,338 means this covered call position will experience price changes equivalent to 4,338 shares of stock, *provided the price changes in the underlying asset are not large*. In the options world, second derivatives are important; VAR helps mine the information they carry.

Table 10: Covered call VAR

Position	Quantity	Price	Value	Delta
Stock	10,000	\$55	\$550,000	10,000
Calls	-10,000	\$4	-40,000	-5,662
		Total	\$510,000	4,338

Simulation Results:

Stock price range: \$45.88 – \$64.17

Mean stock price: \$55.03

95% threshold: \$491,050

VAR = \$510,000 – \$491,050 = **\$18,950**

Market volatility in the past year was substantially higher than the historical average. A risk manager may want to incorporate the potential for renewed high volatility in the future, and consequently model σ with a much higher maximum than normal. Suppose we alter the triangular distribution for σ to specify a maximum value of *twice* the base case mean, or 68%. Higher expected future volatility will increase the value of both the puts and the calls, but this does not have a clear influence on the stock price. If we rerun the simulation, the results change dramatically: VAR is more than twice as much as the base case.

Stock price range: \$46.05 – \$65.63

Mean stock price: \$55.03

95% threshold: \$470,870

VAR = \$510,000 – \$470,870 = **\$39,130**

Stock with Puts Now consider another two-asset portfolio, this time with stock and put options. Table 11 shows this portfolio has fewer shares of stock, contains no calls, and has a long position in puts. Its initial value and position delta are

very similar to the values for the covered call portfolio in Table 10. The VAR shows, however, that this portfolio is much safer than either of the previous two. The 95% threshold is \$506,165 meaning the VAR is \$509,903 - \$506,165 = **\$3,738**.

Table 11: Stock and Puts VAR

Position	Quantity	Price	Value	Delta
Stock	8,680	\$55	\$477,400	8,680
Puts	10,000	\$3.25	32,500	-4,338
		Total	\$509,900	4,342

Straddle Writer A speculator writing a straddle contemplates a very different market direction than the previous examples. This speculative position involves selling both put *and* call options on the same underlying asset, and with similar terms. Suppose we have the two-asset portfolio in Table 12.

Table 12: Writing a Straddle VAR

Position	Quantity	Price	Value	Delta
Calls	-65,800	\$4.00	-\$263,200	-37,243
Puts	-75,900	\$3.25	-246,675	32,925
		Total	-\$509,875	-4,318

As time passes, the value of an option deteriorates if the underlying asset does not change in price. This is what the straddle writer wants. The negative dollar value in Table 12 reflects the fact that the writer does not own the options; he created them and sold them to someone else. To close out the position, the writer would have to buy them back. If their value deteriorates over time, this is good from the perspective of the writer.

However, sharply rising *or* falling prices can result in big losses from this strategy. The 100,000-run scenario analysis with Crystal Ball gives a 95% threshold of -\$893,035 for a whopping VAR of $(-\$509,875) - (-\$893,035) = \mathbf{\$383,160}$.

Straddle Buyer If the strategy of writing the straddle seems likely to result in a large loss, logically the opposite strategy (buying the straddle) should be attractive. The danger here is that the market may reduce its estimate of future volatility and return to more normal input parameters. If we change the distribution for σ back to the original triangular distribution of 34% mean plus or minus 10%, the threshold becomes \$484,333 for a VAR of **\$25,664**. The ability to test the “what if” scenarios is the primary advantage of Crystal Ball in a VAR Analysis. The following three figures are a portion of the Crystal Ball report showing the 95% threshold with the increased future volatility assumption (Figure 11), the “normal” volatility assumption (Figure 12), and the sensitivity analysis showing that the real driver of the results is the market’s assessment of future volatility (Figure 13). These simulation results give the risk manager plenty to think about.

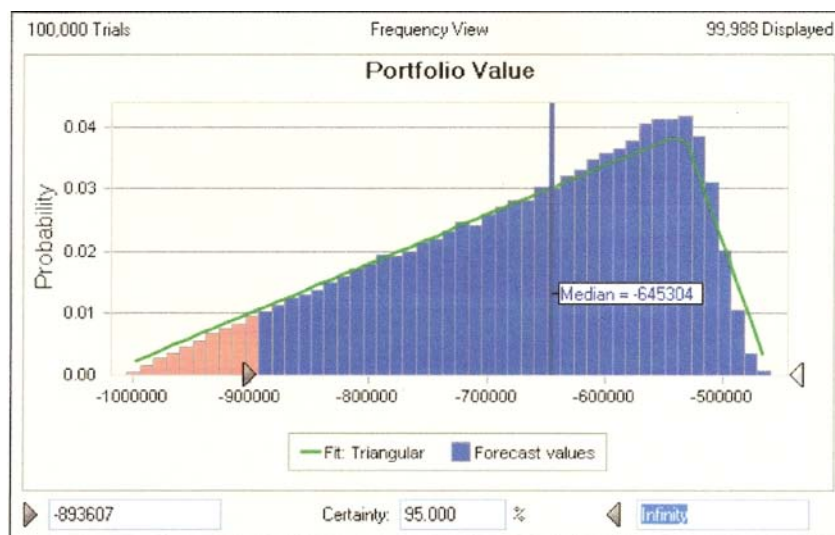


Figure 11: Increased volatility assumption

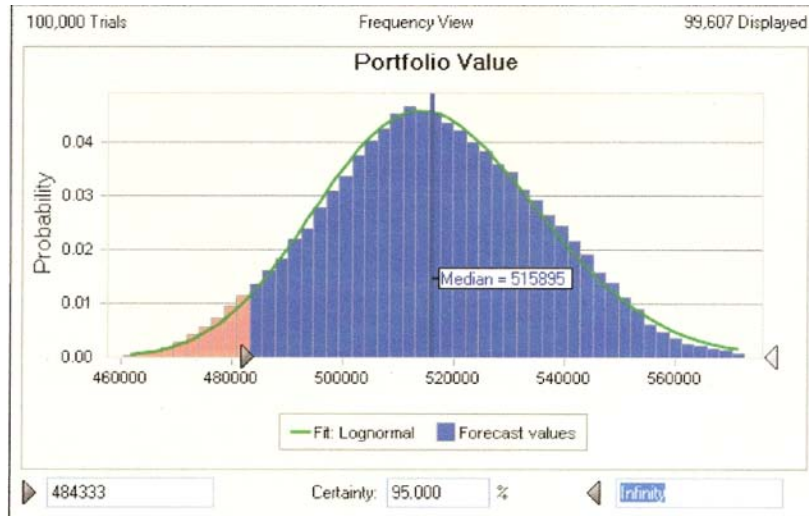


Figure 12: Normal volatility assumption

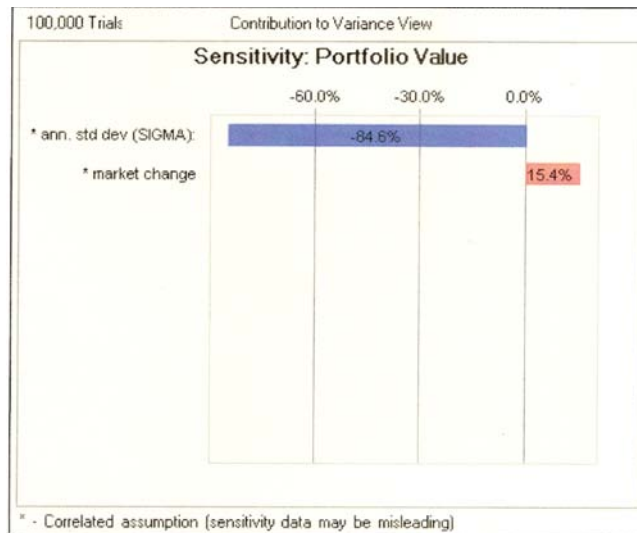


Figure 13: Sensitivity analysis

4 CONCLUSION

Several themes run through this tutorial on financial risk assessment. First, spreadsheet-based Monte Carlo simulation can provide the user with a powerful tool for implementing financial models to perform risk assessment in complex applications. Second, Monte Carlo simulation enables the user to do the following: (i) check the validity of the assumptions underlying a financial model; (ii) explore the sensitivity of the model results to the input parameters whose values are uncertain or are subject to random variation; and (iii) honestly represent the inherent variability of the final results. For example, in the Black-Scholes model used for VAR analyses of portfolios containing stocks, call options, or put options, the standard deviation of the relevant rate of return per year is unknown and must be estimated either from subjective information (expert opinion) or historical data; and the results in Section 3.3 demonstrate the sensitivity of the simulation-based results to changes in the assumptions made about this key input parameter. Clearly the proper accounting for uncertainty and random variability in the inputs and outputs of a financial model is a critical element of financial risk assessment; and simulation provides an effective mechanism for doing this.

REFERENCES

- Black, F., and M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81:637–659.
- Brigham, E. F., and J. F. Houston. 2000. *Fundamentals of financial management*. 3rd ed. Mason, OH: South-Western.
- Ibbotson Associates. 2005. *Stocks, bonds, and inflation 2005 yearbook*. Chicago: Ibbotson Associates.
- Koller, G. 2005. *Risk assessment and decision making in business and industry: A practical guide*. 2nd ed. New York: Chapman & Hall/CRC.
- Kuhl, M. E., J. S. Ivy, E. K. Lada, N. M. Steiger, M. A. Wagner, and J. R. Wilson. 2009a. Introduction to modeling and generating probabilistic input processes for simulation. In *Proceedings of the 2009 Winter Simulation Conference*, ed. M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, forthcoming. Piscataway, NJ: Institute of Electrical and Electronics Engineers, Inc.
- Kuhl, Michael. E., Julie S. Ivy, Emily K. Lada, Natalie M. Steiger, Mary Ann Wagner, and James R. Wilson. 2009b. Univariate input models for stochastic simulation. *Journal of Simulation* forthcoming. Available online via www.ise.ncsu.edu/jwilson/files/kuhl09ajos.pdf [accessed August 8, 2009].
- MacCrimmon, K. R., and D. A. Wehrung. 1985. *Taking risks: The management of uncertainty*. New York: The Free Press.
- Oracle Corporation. 2008. *Oracle® Crystal Ball getting started guide*. Redwood Shores, CA: Oracle Corporation.
- Strong, R. A. 1998. *Practical investment management*. Cincinnati, OH: South-Western.
- Wharton, F. 1992. Risk management: Basic concepts and general principles. In *Risk: Analysis, assessment and management*, ed. J. Ansell and F. Wharton, 1–14. Chichester, UK: John Wiley & Sons.

AUTHOR BIOGRAPHIES

ROBERT A. STRONG is professor of finance in the Maine Business School at the University of Maine. He has received the following recent professional recognition: the 2005 Maine Distinguished Professor Award; appointment as the University Foundation Professor of Investment Education and Professor of Finance; and the 2007 Maine Professor of the Year by the Carnegie Foundation. He has written more than seventy-five articles for business journals, many of them about his teaching methods, and books on a variety of financial topics. His three textbooks are used in more than one hundred universities around the world. In 1993, Dr. Strong, working with the University of Maine Foundation, started a club that has allowed thirty to forty students a year to act as real-money portfolio managers for UMaine Foundation funds. The Student Portfolio Investment Fund, or SPIFFY, has since increased its original \$200,000 allotment to more than \$1.2 million today. His e-mail address is Robert_Strong@umit.maine.edu.

NATALIE M. STEIGER is an associate professor of production and operations management in the Maine Business School at the University of Maine. Her research interests include simulation modeling and analysis, especially techniques to improve simulation analysis in practical applications. In 2000 she received the IIE Pritsker Doctoral Dissertation Award (Second Place) for her dissertation research, which forms the basis for the ASAP3 procedure for steady-state simulation output analysis. She is a member of IIE and INFORMS. Her e-mail address is nsteiger@maine.edu and her Web address is www.mbs.maine.edu/steiger_natalie/.

JAMES R. WILSON is a professor in the Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University, where he served as director of graduate programs (1995–1999) and department head (1999–2007). His current research interests are focused on probabilistic and statistical issues in the design and analysis of simulation experiments. He has published ninety-five archival journal articles, seven book chapters, two books, and seventy-five papers in the proceedings of conferences. He has advised twenty-five doctoral students and thirty-seven master's (thesis) students. Currently he serves as Editor-in-Chief of *ACM Transactions on Modeling and Computer Simulation*. He is a member of AAUW, ACM, and SCS; and he is a Fellow of IIE and INFORMS. His e-mail address is jjwilson@ncsu.edu, and his Web page is www.ise.ncsu.edu/jwilson.