

Using simulation to calculate the NPV of a project

Marius Holtan

Onward Inc.

5/31/2002

Monte Carlo simulation is fast becoming the technology of choice for evaluating and analyzing assets, be it pure financial derivatives or investments in real assets. Two of the main virtues of simulation are flexibility and simplicity. There is no constraint on the type of uncertainty that can be modeled and most any type of decision rule can be incorporated. Simulation is also easy to implement and models can easily be constructed in spreadsheet packages. Moreover, with the surge in computing power that has taken place over the last decade, run time is no longer an issue except perhaps in the most extreme applications. Wall Street now uses Monte Carlo simulation extensively when valuing mortgage backed and other asset backed securities as well as exotic options and other derivative securities. But Monte Carlo simulation is not only useful for Wall Street, professionals analyzing real investment problems can also greatly benefit.

Expanding production capacity, building new plants, investing in IT, harvesting natural resources, and investing in R&D are but some examples of real investment problems whose financial success can be efficiently analyzed with simulation. Typically, in such real investment projects we have an idea about the range values should fall within but we do not know the exact numbers by any means. Furthermore, project plans are often adjusted to reflect changes in the estimates of future prices, demand, or research outcomes. Simulation lets us generate any number of likely forecasts from a general specification of the overall cash flow distribution. Moreover, we can create management policies that define what should happen if for example price goes up (expand plant capacity) or what should happen if R&D is successful (launch new product).

This document provides a step by step introduction to project analysis with simulation. Most of the steps in such an analysis are no different from the steps underlying a regular NPV analysis. First, we create a cash flow model of the project. In this step we identify all relevant revenue and cost components as well as planning horizon and terminal value. Second, we identify the cash flow components that are heavily impacted by uncertainty and choose mathematical models that reflect the random characteristics of the

components. Next, we establish discount rates and discount method. The method and rates must reflect the types and forms of the uncertainties. When the previous steps are completed we are ready to implement and run a simulation. For those used to doing NPV analysis in a spreadsheet, the only extra implementation step required for a simulation run is to insert the random model for each uncertain forecast. As in any financial analysis, before communicating the results we should analyze the output in order to spot possible problems with our assumptions or the modeling. It is also important to perform sensitivity analysis on important parameters. Finally, the output should be checked for inconsistencies and sensitivity analysis should be performed on critical factors.

Perhaps the most interesting aspect of a real option analysis is the derivation of optimal policies. Simulation can successfully be utilized for this purpose as well. However, the methods available for deriving optimal policies is a rather lengthy topic by itself. We will therefore focus on the different steps underlying a NPV analysis with simulation and defer discussion of real option optimization to another time.

Cash flow model

The cash flow model consists of all current and future cash flows that result from undertaking the project. Cash flows for most projects include sales, cost of goods sold, taxes, and initial capital costs. For example, suppose your company wants to analyze the prospect of a development project. The project requires 6 months of development at a cost of about \$100,000, of which about \$50,000 is used for equipment. The sales department provides an expected sales forecast while the manufacturing group estimates that the cost of goods sold is about 30% of sales. The manufacturing group also estimates the capital equipment costs. Furthermore, the sales department provides an estimate of the general and administrative costs. Table 1 summarizes this information.

	2002:H1	2002:H2	2003:H1	2003:H2	2004:H1
R&D Cost	50				
Capital Investment	50	1000			
Sales		600	1000	600	200
Cost Of Goods Sold		180	300	180	60

Other Costs		100	100	100	100
Depreciation		12.5	345.83	345.83	345.83
Pretax Profit (Accounting)	-50	307.50	254.17	-25.83	-305.83
Tax at 30%	-15	92.25	76.25	-7.75	-91.75
After Tax Profit (Accounting)	-35	215.25	177.92	-18.08	-214.08

Table 1. Example R&D project description. The amounts are in \$1,000.

The project description in Table 1 is not a cash flow description of the project because capital costs are described by a straight line depreciation schedule rather than when they occur. The effect of the depreciation schedule is to reduce future tax payments. Although in Table 1 we used straight line depreciation for simplicity, it is important to be aware of the fact that other types of depreciation schedules exist. The choice of depreciation schedule can sometimes be an important factor of project value and care should be taken to ensure that the best schedule is chosen. For a proper NPV cash flow analysis we must work with cash flows as they occur. A cash flow analysis of the project is presented in Table 2.

	2002:H1	2002:H2	2003:H1	2003:H2	2004:H1
Sales		600	1000	600	200
Cost Of Goods Sold		180	300	180	60
R&D Cost	50				
Other Costs		100	100	100	100
Tax at 30%	-15	92.25	76.25	-7.75	-91.75
Operation Cash Flow	-35	227.75	523.75	327.75	131.75
Capital Investment	50	1000			
Profit	-85	-772.25	523.75	327.75	131.75

Table 2. Example R&D project expected cash flows

There are two additional points regarding cash flow identification worth mentioning. First, when considering cash flows on the cost side be careful that only costs attributable to the project going forward

are included. Do not include sunk costs and be careful with including overhead expense. The only overhead expense that should be included is the extra overhead resulting from initiating the project. Finally, some projects will have a horizon value, that is, a value associated with the ongoing business as proceeding from the end of the planning horizon. There are several ways of getting at a number for the horizon value. A popular one is to use the P/E ratio of the company times the projects earnings in the final year of the analysis. Regardless of the method, it is important to monitor the horizon value and verify that it is reasonable. It is often the case that most of the value of a project ends up at the horizon which should be a cause for concern. For more information on this topic one can refer to one of the many excellent finance/corporate finance texts that are readily available.

Modeling uncertain cash flows

Most cash flows are to some extent uncertain. The degree of uncertainty of cash flows, however, generally differ considerably. Moreover, the impact on the profit of the uncertainty of the cash flows usually also vary significantly. As a general rule, when creating a model it is best to focus on the one or two cash flows whose uncertainty has the most impact on the profit. Other uncertainties with little impact on profit will invariably also have little impact on the project value. In addition, if you try to model every minute detail errors will inevitably creep in, the model will become confusing, and simulation runs will take much longer than necessary to run.

For example, consider the extraction of a natural resource. The resource price and the resource amount are usually the two main uncertain variables. Although the expected growth of the price is low we expect it to fluctuate widely over a short period of time. Moreover, we never really know how much is left of the resource until we get to the end in which case the answer is obvious.

After identifying the uncertain cash flow components we choose models that match the uncertainty characteristics of the components. There are quite a number of different models of which many have been designed to fit a particular circumstance. Two models that are particularly useful because of their

versatility are the geometric growth model¹ and the mean reverting model². We will discuss these two models in more detail.

In the geometric growth model the uncertainty is gradually introduced in a multiplicative fashion. Figure 1 shows the 10/50/90 distribution plot for this model and also a typical sample path.




Figure 1. 10/50/90 distribution and sample path for a geometric growth model. The model's expected growth rate is 10% and the return standard deviation is 25%. The x-axis scale is in years while the y-axis scale is in the unit of the model.

In a 10/50/90 distribution plot there is a 10% chance that at a particular point in time the process could be below the red area, a 50 % chance that the process could be below the border between the red and green areas, and a 10% chance that the process could be above the green area. Thus, there is a 80% chance that the process at a point in time could be in either the green or the red areas.

¹ A technical name that is often used for this process is geometric Brownian motion.

² A technical name that is often used for this process is Ornstein Uhlenbeck model.

The geometric growth model is useful for capturing processes whose growth is typically thought of in percentage terms. Examples of such processes include stock prices, GDP, and demand for general product categories such as energy, cars, or computers.

In order to construct the geometric growth model it is necessary to specify a forecast of the expected rate of growth over time and the standard deviation of the rate of growth. For example, suppose in the project development example that the future sales is uncertain with expected sales growing according to the forecast but with a 20% standard deviation of the annual rate of growth. Figure 2 shows the 10/50/90 distribution under these assumptions.




Figure 2. 10/50/90 distribution for the sales forecast in the product development example when sales is modeled using the geometric growth model.

It is interesting to note that although in percentage terms the uncertainty grows with time, in absolute terms the uncertainty actually decreases for the last two periods. The reason for the decrease is that while the standard deviation of this process increases over time this standard deviation is specified in terms of the rate of change of the process while the distribution plot describes the uncertainty in terms of the actual value of the process.

Time (Years)	0	0.5	1	1.5	2
Sales (\$1,000)		600	1,000	600	200
Rate of return			0.51	-0.51	-1.10
Standard deviation of the return		0.14	0.14	0.14	0.14
Normal random variable		0.88	0.40	-0.83	-0.47
Multiplicative growth random variable		1.12	1.74	0.53	0.31

Sample path	672	1,173	620	191
-------------	-----	-------	-----	-----

Table 3. Step by step calculation of sales sample paths using the geometric growth model to model sales.

Table 3 tabulates the steps necessary to obtain a sample path of the sales in the development project when assuming that the sales grows according to the geometric growth model. The rate of return for a particular period is calculated by taking the logarithm of the ratio of the current sales to the previous period's sales. The standard deviation of the return is the product of the annual standard deviation of the return which is 20% and the square root of 0.5 which is the time step. The row denoted normal random variable provides four samples of a standard normal random variable. The multiplicative growth random variable is the variable responsible for the geometric growth. It is calculated in two steps: First multiply the normal random variable with the standard deviation, add the rate of return, and subtract one half of the square of the standard deviation. Second, the growth multiplier is the exponential power of the result of the preceding calculation. The sample path is calculated as follows: The value for the second period is the product of the expected sales in that period and the multiplicative growth variable while the value for each of the other periods is the product of the value in the preceding period and the geometric growth variable.

When a cash flow or other time series display a lot of variability in the short term but generally stay in a band around a long term mean then we typically consider it to be mean reverting. Commodity prices are usually deemed to be mean reverting. Figure 3 presents a 10/50/90 distribution for a mean reverting model.

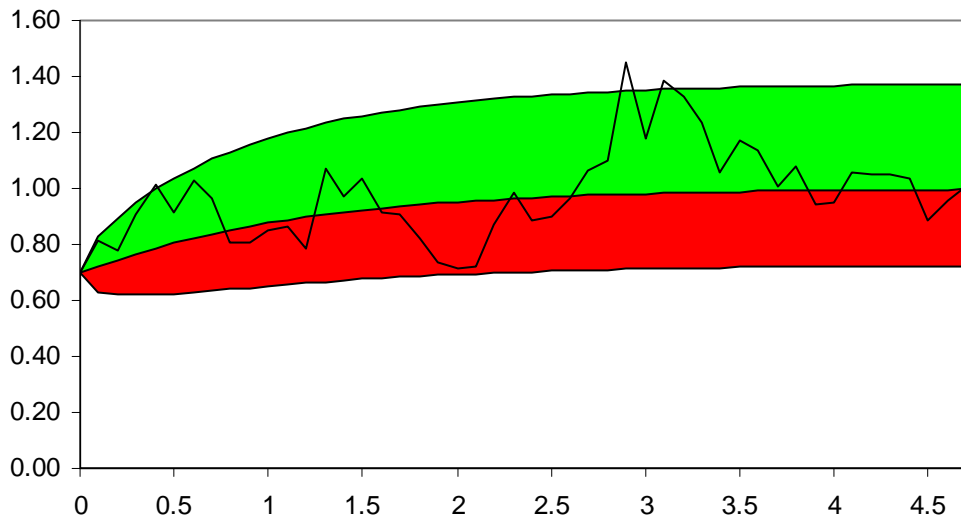


Figure 3. 10/50/90 distribution plot for the exponential mean reverting model. The mean reverting level start at 1 and grows at an annual rate of 10%. The process itself begins at 0.7. The long term standard deviation of its return is 25% and its half life is 0.7 years.

Contrast Figure 3 with the 10/50/90 distribution of the geometric growth process shown in Figure 1. After a few years the uncertainty of the mean reverting process is approximately constant while the uncertainty of the geometric growth model continues to grow. The reason is that as a particular sample path from the mean reverting model starts to deviate from the mean the growth rate adjusts with the effect of gradually pulling the path back towards the mean reverting level. No such adjustment occurs to the growth rate of the geometric growth model.

To implement a mean reverting model it is necessary to define the mean reverting level, the current starting point, the long-term standard deviation, and the expected half-life. The expected half life is the expected amount of time it takes until a deviation away from the mean reverting level has returned half way back to the mean reverting level. The long-term standard deviation and the expected half-life are responsible for the width of the band around the mean reverting level and the speed at which the process returns to the mean reverting level. For example, if the expected half life is 0.7 then after 0.7 years we expect the process to be half-way back, in 1.4 years another half from where it was at 0.7 years which means 75% of its way back, and after 2.8 years we expect the process to be about 94% back to the mean reverting level. This is

significant also for the shape of the band as we can observe in Figure 3 where after 2.8 years the band only grows very slowly because we expect any deviation that occurred 2.8 years earlier to be 94% back to the mean. The maximum width of the band is defined by the long term standard deviation.

Suppose that in the development project you decide to model the sales forecast as being mean reverting with a half-life of 0.25 years and 20% long-term standard deviation. Furthermore, suppose that the sales forecast defines the mean reverting level³ and the expected starting point in the second half of 2002 is 600. Figure 4 shows the 10/50/90 distribution for this version of the sales forecast.

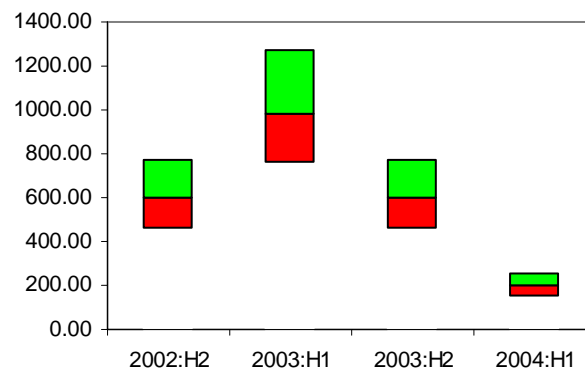


Figure 4. 10/50/90 distribution for the sales forecast in the product development example when sales is modeled using the exponential mean reverting model.

Comparing Figure 2 with Figure 4 we note that the uncertainty is about the same in the second halves of 2002 and 2003 for the mean reverting model while for the geometric growth model the uncertainty of the second half of 2003 is significantly larger than the uncertainty of the second half of 2002. Thus, we have verified that the uncertainty of the geometric growth model grows gradually over time while the uncertainty of the mean reverting model approaches a constant level of uncertainty after a few expected half-lives.

³ Strictly speaking, in the exponential mean reverting model the expected values and the mean reverting levels do not completely agree. Thus, if the sales forecast represent the expected values it is necessary to make an adjustment. Deriving the adjustment is a relatively straightforward exercise and will not be expanded upon at this point.

The exponential mean reverting model is somewhat more complicated to implement than the geometric growth process. You can find a detailed but brief description of this model in the appendix.

Discounting

Finance practitioners generally agree that a cash flow should be discounted at its opportunity cost of capital where the opportunity cost of capital of a cash flow is the expected return demanded by investors purchasing securities with risks comparable to the asset's risks. To obtain an estimate of the opportunity cost of capital many begin by estimating the company cost of capital, often utilizing the capital asset pricing model (CAPM) to estimate the company beta. Strictly speaking, the company cost of capital is only valid if the project cash-flows have the same risk characteristics as the company. Sometimes the project can be less risky and sometimes more risky. In addition, cash-flow characteristics often change over time further complicating the estimation problem. Advanced NPV practitioners know this and try to make the necessary adjustments in order to have a single discount rate for the project.

Another approach to discounting is to estimate a discount rate for each of the cash flow components. This approach is often simpler, particularly when cash flow characteristics change over time, and, as an added benefit, it is more convenient when we analyze options. For example, suppose a project's cash-flows are sales, variable costs, and fixed costs. Fixed costs are deterministic and do not change while variable costs are proportional to sales which are uncertain. The fixed costs should therefore be discounted at the risk free rate while variable costs should be discounted at the same rate as sales because they are proportional. It is sufficient to determine the discount rate for sales and the risk free rate where the latter can be derived from treasury bills. The project NPV is the difference of the sum of the present values of the revenue components and the sum of the present values of the cost components.

Regardless of the discounting approach it is still necessary to estimate one or more discount rates. The CAPM-based approach mentioned earlier is often used. Other approaches include estimating a return from historical data and estimating a return from fundamental data.

Monte Carlo simulation

During simulation the computer generates many sample paths of the values of the uncertain cash flow components. A single sample path represents one possible set of values for a particular component, one value for each time period. It is typical to let iteration denote a single set of sample paths. When we calculate the NPV using simulation we proceed by generating some number of iterations, each iteration representing a different set of sample paths. The cash flow components of each iteration are discounted using the corresponding discount rates, a net amount is calculated for each iteration, and the average of all the net amounts provides an estimate of the NPV⁴.

Simulation Algorithm

Set *estimatedValue* = 0

For i = 1 to *maxIterations*

 Generate a new sample path for each uncertain process

 Discount each cash flow component's for each period

 Calculate the discounted profit for each period

 Calculate the sum of the discounted profits, divide the sum by *maxIterations*, and add result to *estimatedValue*

Next i

Exhibit 1. Sketch of simulation algorithm for a straight project valuation

To illustrate the simulation approach let us again consider the development project. Table 1 and Table 2 provide a starting point for our analysis. As the first step in the implementation we input the cash-flow information into a suitable analysis package. Spreadsheet applications are good for this purpose. Next, we create an additional section for each of the uncertain components. In each section we input the formulas that describe how the corresponding component evolves over time. For each iteration each section generates a sample path. In a direct valuation we proceed by writing a new cash flow statement where references to the sample path formulas replace the forecasts of the uncertain components. The net cash flow in a single period usually looks similar to the following formula:

⁴ The process becomes a little more complicated when we include decisions. I will address this later.

$$\text{net cash-flow} = (\text{revenue} - \text{variable costs} - \text{fixed costs} - \text{depreciation}) \cdot (1 - \text{tax}) + \text{depreciation}.$$

If we choose to use a different discount rate for each cash flow component then we should rewrite the above into

$$\text{net cash-flow} = (\text{revenue} - \text{variable costs}) \cdot (1 - \text{tax}) - \text{fixed costs} \cdot (1 - \text{tax}) + \text{depreciation} \cdot \text{tax}.$$

Revenue and variable costs are discounted at the same rate while fixed costs and depreciation are discounted at the risk free rate. Note that the product of the tax rate and the depreciation amount ends up as a positive addition to the net cash flow. We can therefore interpret depreciation as a tax shield.

In the development project we group Sales and COGS together, multiply the difference by $(1 - \text{tax})$ and discount by the sales discount rate. Moreover, we sum R&D Costs and other costs and multiply by $(1 - \text{tax})$. To this we add capital investment costs and subtract the product of the tax rate and the depreciation cost. The result is discounted by the risk free rate. Next, we calculate the net cash flow of the iteration. Finally, the average of the iteration results gives us an estimate of the NPV.

To illustrate this approach, suppose the sales discount rate is 10% per year and the risk free discount rate is 5% per year. Table 4 shows the discount factors over time.

Discount factors	2002:H1	2002:H2	2003:H1	2003:H2	2004:H1
Sales (10%/yr)	1	0.95	0.90	0.86	0.82
Risk free (5%/yr)	1	0.98	0.95	0.93	0.90

Table 4. Discount factors over time for the development project. The discount factors are calculated using the continuously compounded formula.

Suppose further that the future sales is modeled as a geometric growth process. Figure 2 shows the distribution and provides the parameters of the model. Table 5 presents an example iteration of the project cash flows.

	2002:H1	2002:H2	2003:H1	2003:H2	2004:H1
Sales		616.78	1126.71	681.80	271.87
COGS		185.03	338.01	204.54	81.56

(Sales-COGS)(1-tax)df		287.48	499.55	287.55	109.07
R&D Costs	50				
Other Costs		100	100	100	100
Fixed Costs(1-tax)df	(50.00)	(68.27)	(66.59)	(64.94)	(63.34)
Depreciation*tax*df		3.66	98.69	96.25	93.88
Capital Investments*df	(50.00)	(975.31)			
ProjectCash Flow	(100.00)	(752.44)	531.65	318.86	139.61
Net Cash Flow	137.67				

Table 5. Example iteration for the development project.

For the iteration shown in Table 5 the net cash flow turned out to be \$137,670. That was, however, a particularly optimistic figure as the estimated value after 10,000 iterations was -\$9,840. The 90% confidence interval of this estimate was \$3,150. Thus, the project is profitable but only barely.

For this straightforward application we do not actually need to resort to simulation to calculate the NPV. Because the project characteristics never change, there are no decisions involved, a simulation with enough iterations will simply return the expected values. However, although not necessary, a straightforward NPV calculation is a good modeling test before you turn to valuing a project with decisions or options.

The actual expected NPV value came out to -\$10,340. We see that the simulated value came very close to the actual value.

Simulation with decisions

The general approach to valuing projects with decisions is perhaps best illustrated by way of an example. Suppose, therefore, that the management of the development project can choose to abandon the project after the R&D stage but before money is spent on manufacturing equipment. In particular, suppose that management proposes that no investment will be made if sales is projected to be less than \$500,000 in the second half of 2002.

There are now two possible versions of the project. If sales turn out to be greater than \$500,000 then project cash flows are modeled as before but otherwise the only cash flows from the second half of 2002 is

depreciation of the initial capital investment. What is the effect on the NPV of the project? If we knew that we would abandon the project we would discount at the risk free rate while if we knew we would continue then we would discount as before. What is the correct way when we have some of both?

The right way to calculate the NPV is to first discount and then consider the impact of the decisions on the discounted cash flows. This initially seem counterintuitive to most people as many would first sort out the impact of the decisions and then discount. To understand the right way we must delve deeper into the relationship between NPV and discounting.

We only want to initiate projects with positive net present value. Discounting is the process by which we place future cash flows onto an equal footing with current cash. Nobody believes that we will actually get the discounted value of a future cash flow, yet we treat this discounted value as we would cash when calculating NPV. Thus, we see that the discounted cash flow is a representative at the present time of a future cash flow. From an NPV perspective we consider the representatives of all the cash flows that make up a project, in fact we calculate the net of all these representatives and treat that net as the value of the project. Thus, from an NPV perspective we are only interested in the representatives of the future cash flows. We see that there is a decoupling between the NPV of a cash flow which is a present time concept and the future realized amount of the cash flow.

Therefore, in terms of deriving the NPV of a particular decision rule we must calculate the net when the decision rule is applied to the representatives of all the possible cash flows because when we consider the NPV we are only viewing the representatives of the cash flows.

Let us now turn back to the problem facing the management of the development project. We need to make two modifications to the straight valuation procedure. The first modification is to include a new section in the spreadsheet where we calculate the value of ending the project after the development stage. This value only includes the initial \$100,000 investment in research and capital and the value of the depreciation tax shield from the capital portion of the initial investment. The second modification is to include an if-then-else statement in the simulation algorithm.

Simulation Algorithm with Decision

Set *estimatedValue* = 0

```

For i = 1 to maxIterations
    Generate a new sample path for each uncertain process
    Discount each cash flow component for each period
    If discounted sales is greater than $500,000 then
        Calculate the discounted profit for each period
        Calculate the sum of the discounted profits, divide the sum by
        maxIterations, and add result to estimatedValue
    Else
        Calculate the value of ending the project after the development stage, divide
        by maxIterations, and add result to estimatedValue
Next i

```

Exhibit 2. Simulation algorithm with abandonment decision rule.

Figure 5 compares the project value corresponding to the proposed decision rule to the values corresponding to a few other decision rules and the project value with no abandonment. We observe that while the proposed decision rule yield a marked improvement in value the project is even more valuable if we decide to abandon when the sales figure becomes less than \$550,000. Thus, we see that the abandonment option, even in the rudimentary form examined here, can have a significant impact on a project's value.

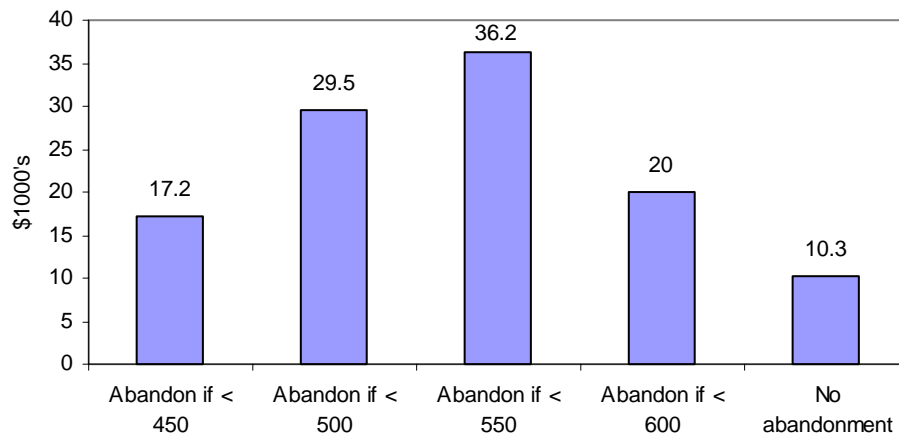


Figure 5. Comparison of estimated project values for different decision rules. The project is abandoned after the development stage if the estimated sales at that time is less than the stated amount.

Another interesting output is the distribution of future earnings. Figure 6 shows the 10/50/90 distribution of earnings after taxes but not including the capital investment costs. This type of graph gives a useful picture of the likely level and amount of earnings over time for a project. In this case we note that the majority of the cash flow is earned by the first half of 2003.




Figure 6. Distribution of earnings over time after taxes before inclusion of capital costs.

We can also estimate the probability of abandonment by observing how frequently the actual sales figure in the second half of 2002 would be less than \$550,000 (and not by estimating how frequently the discounted sales figure would be less than the discounted value of \$550,000). This figure came out to about 29%. Moreover, in the even of abandonment we would still earn the depreciation tax shield for the initial investment. The value of the depreciation tax shield is \$28,200.

We can summarize the conclusions of the development project as follows: The initial investment is \$100,000. The estimated value of this investment is \$36,000. Our rough analysis showed that we should abandon the project if the expected sales in the second half of 2002 is estimated to be less than \$550,000. We estimated the probability of abandonment to be 29%. In the event that we abandon the project we will receive the depreciation tax shield for the initial capital investment. The value of this shield is \$28,200. Thus, while \$100,000 is the amount necessary to begin the project, we know we can get the depreciation tax shield back which carries a value of about \$28,200. The best initial cost figure to use is therefore

perhaps \$71,800. With a probability of 71% we will choose to continue the project. The expected value at the present time of the investment is \$36,000 which is about a 50% return on the initial capital in excess of the market estimated rate of return.

Appendix

This section describes a procedure for implementing the exponential mean reverting model. Let $f(t)$ denote the mean reverting level in period t , let $s(0)$ denote the current starting point and $s(t)$ the value of the process in period t , let v denote the long term standard deviation, let h denote the expected half life, and let

Δt denote the time step (in fraction of years) between periods. Define $\kappa = \frac{1}{h} \ln(2)$, $\sigma = v\sqrt{2\kappa}$,

$$\beta = \sqrt{\frac{\sigma^2}{2\kappa} (1 - e^{-2\kappa\Delta t})}, \text{ and } m(t) = \ln(f(t)).$$

We proceed by modeling the return process of s . Let x denote this return process. The relationship between x and s is captured by the equation $s(t) = e^{x(t)}$. First, the initial value of the return process is $x(0) = \ln(s(0))$. Furthermore, for the mean reverting model the expected growth rate depends on the distance between the current value and the mean reverting level. Let $\bar{x}(t + \Delta t)$ denote the expected growth rate from time t to time $t + \Delta t$. The formula for $\bar{x}(t + \Delta t)$ is

$$\bar{x}(t + \Delta t) = (x(t) - f(t))e^{-\kappa\Delta t} + f(t).$$

Because the process is uncertain we must also include a random term. Let \mathcal{E}_t denote a standard normal random variable whose value we know at time t . Then

$$x(t + \Delta t) = \bar{x}(t + \Delta t) + \beta\mathcal{E}_{t+\Delta t}.$$

As mentioned above, the value of the exponential mean reverting process is given by the formula

$$s(t) = e^{x(t)}.$$