UNIVERSITY OF CAMBRIDGE

Judge Institute of Management

cambridge's business school

Risk-enhanced NPV analysis:

A Call for Computer Aided Investment Appraisal

by

Stefan Scholtes

Judge Institute of Management, University of Cambridge, Cambridge CB2 1AG, UK, s.scholtes@jims.cam.ac.uk

Summary

This paper presents a step-by-step process, which allows a decision maker to gradually enhance an NPV analysis by incorporating uncertainty and flexibility. Unlike most of the approaches in the literature we start and never deviate from the premise that the value of a project is, a priori, uncertain. The uncertainty is best depicted by a graphical representation of the distribution of the project value, such as a value-at-risk chart. This risk profile is an ideal representation of the trade-offs of the risks and opportunities inherent in a risky project. The risk profile is obviously not only affected by underlying exogenous uncertainty drivers but also by the project manager's skill in exercising managerial flexibility. Therefore a valid risk profile can only be determined on the basis of an a priori contingency plan that prescribes how key flexibilities will be exercised as the future evolves. Finding a suitable contingency plan is a matter of technical analysis and can employ various modeling paradigms, including decision tree analysis, real options valuation and dynamic and stochastic optimization. Different contingency plans induce different risk profile, whose graphical representations can be approximated by Monte Carlo simulation. The choice of the contingency plan that "optimizes" the risk profile depends on the decision maker's preferences in trading off risk and opportunity.

I. WHAT'S WRONG WITH NPV?

The value of a single project, a project portfolio or a company as a whole derives from the cash flows it generates over its lifetime or, in a practical setting, over a fixed planning horizon. Any evaluation must therefore begin with a *functionally* valid model of cash flow generation that explains how expenditure is incurred over time to generate income. This can be conveniently done in a spreadsheet. Once the model is functionally valid, all unknown value drivers, such as future demands, prices, costs, etc., are forecast and incorporated into the spreadsheet, which then automatically calculates the corresponding cash flows. Summing up the suitably discounted cash flows and subtracting the initial investment gives the net present value (NPV) of the project.

Albeit the industry standard for decades, NPV analyses are well known to be problematic in the presence of uncertainty about the future economic environment of the project. Indeed, an NPV analysis does neither reflect such uncertainty nor the managerial flexibility that can be used to cut losses if necessary or to exploit opportunities if and when they arise. The NPV spreadsheet represents a *fixed plan*, it is a snapshot of just one of the zillion paths the project may take when it is actually carried out. Even if the logic of cash flow generation is captured well in the model, it lacks precision since the crucial value drivers have been forecast and forecasts are bound to be wrong. Nobody can say what the demand for a new product will be and how it will evolve over time or how the prices of raw materials, energy, labour, etc. will change. Some risks or opportunities, such as the risk of technical failure or the chance of a negative outcome of a political regulation process, may be completely neglected in the NPV analysis. Since the analysis neglects uncertainties, it cannot account for managerial decisions that may be made over time, re-actively as circumstances unfold and pro-actively as the likelihoods of future scenarios change. These decisions will be made in an attempt to maximise the value of the project at the time they are taken and the decision maker will take all available information at the time into account. If we neglect this managerial flexibility we undervalue the project, possibly considerably. VCs and CEOs know this intuitively. Many projects that do not look promising on

the basis of an initial NPV analysis are nevertheless launched "for strategic reasons". This does often lead to reverse-engineered NPV spreadsheets. Numbers are massaged to make the project look good on paper. Such practices bear no rational or conceptual foundation and obviously undermine the credibility of quantitative analyses in general. What is needed is a quantitative analysis that allows the incorporation, and thereby rationalization, of the VC's or CEO's intuition.

The rationale behind the use of a single number, an average, to represent the value of a risky project is a tremendously powerful risk-management strategy: Diversification. This strategy is based on the law of large numbers, which states that the risk inherent in the return of a portfolio of independent projects is much lower than the risk of an individual project because underachieving on one project can be balanced out by overachieving on another and it is unlikely that all projects underachieve. In other words, a holder of a portfolio of independent projects can neglect individual project risk and can base decisions on expected return[1]. This is the basis for much of corporate finance theory. It assumes, however, that the decision maker holds a portfolio of independent projects. In the post-war period many companies developed into conglomerates, being engaged in many diverse businesses, not least to take account of their long-term investors' desire to minimize risk. In the more recent past, however, investors have realized that it is more efficient for them to take care of risk minimization through diversification themselves, through appropriately balanced and regularly re-balanced portfolios of stocks. They are primarily interested in companies' earnings potentials and are less willing to reduce return expectations for risk reduction due to corporate diversification. This development goes along with an increased emphasis on corporate strategy and its focus on core competencies. Obviously, there are still highly diversified and successful corporations. However, the rationale behind this diversification is not risk minimization but the desire to maximize return expectations by fully exploiting the core competencies of the company throughout the economy. The investors' focus on returns does of course not mean that they are blind for a company's risk; it only means that risk management through corporate diversification does not

---

[1] More precisely, if the projects are not fully independent and therefore the portfolio is not fully diversified, the decision maker can base the decision on the portfolio effect (the systematic risk) and can neglect the typically larger individual project risk (the unsystematic risk) since this is eliminated by the diversification effect.

add value to their portfolios. Risk management on the corporate level therefore has to focus on "active" risk management through good management and adaptive control of its operations. On the project level this requires an appropriate level of managerial flexibility to be able to react as uncertainties unfold. This has two consequences for project design and valuation: It requires corporate decision makers to take project-specific risk explicitly into account and it requires them to include managerial flexibility into project design and valuation.

II. WHAT'S WRONG WITH ROV?

It has been suggested that a new paradigm should be used to overcome the pitfalls of fixed-plan NPV analyses: Real options valuation (ROV). This approach, akin to the Nobel-prize-winning financial options analysis, enhances the NPV calculation by taking account of managerial flexibility in a specific way. In its standard form, the real options approach assumes the current value of the "passive" project, i.e. the project without managerial flexibility, to be known. The value of this passive project is assumed to change over time and as the value moves up or down, managers make decisions, such as abandoning the project, extending the project, etc. The real options paradigm regards the movement of the value of the passive project as analogous to the movement of a stock price and managerial flexibility, the *real option*, as analogous to an option on the stock price. An option, whether real or financial, grants the right but not the obligation to exercise a certain action for a certain price at a certain point in time or at some point during a certain time interval. Because of the conceptual similarities between real and financial options it is tempting, and in some cases indeed plausible, to employ the well-developed concepts and procedures for financial options analysis also in the context of real options. In particular the formal use of financial options valuation mechanisms put a price tag on managerial flexibility. The overall value of the "flexible" project is then regarded to be the value of the passive project plus the obtained value of flexibility.

Whilst many practitioners agree that the standard NPV paradigm has pitfalls in the presence of uncertainty, they are reluctant to embrace the new real options paradigm. Indeed, the number of success stories,

e.g. with Merck and Shell, is relatively modest compared to the demand for risk-enhanced project appraisal. Many real options advocates argue that the reluctance of practitioners is due to the complexity of financial options valuation mechanisms and are optimistic that ROV will become a standard tool, just as NPV, once a generation of MBA students has been adequately trained to work with the paradigm. We are not as optimistic and believe that the reason for the lack of success lies deeper and is at least two-fold: Firstly, the assumptions underlying financial options valuations are very stringent in the context of real options and secondly, the approach provides the decision maker with little guidance to build up the intuition that is indispensable for the acceptance of results obtained by means of a complex decision support tool.

The limitation of financial options analysis, when applied to project valuations, begins with the notion of the "value" of the passive project. This value, unlike the value of a stock, is not a market value that is publicly listed but has to be determined by one or another method and any such method has elements of arbitrariness because the cash flows that will or would be generated by the passive project are not known a priori and typically depend on the unfolding of many uncertain value drivers. Secondly, the assumption that the project value, however determined, moves like a stock price may be tenable for certain operational, i.e. company-typical and recurring, projects but seems difficult to maintain for one-off strategic projects. Thirdly, the "narrow" view of real options assumes that the exercise of flexibility is triggered by the aggregate value of the project, which is difficult to maintain for many types of flexibility. The aggregate value may flourish because demand is higher than expected, which may trigger an expansion decision, or because the costs are lower than expected which, without the prospect of high demands, is not a sensible criterion for expansion. Finally, the main assumption of financial options pricing is the independent tradability of stock and option and fractions and multiples thereof. This implies a unique options price through a non-arbitrage argument. Can a project and the flexibility, or fractions and multiples thereof, really be traded independently, or traded at all? The list of limitations of a "narrow" real options view can be further extended.

We will not pursue a more detailed discussion of the conceptual limitations of ROV because we regard the second obstacle to its wide-spread dissemination as

more important and as a good starting point for a more mature analysis of risk: ROV lacks immediate intuition. A real options valuation is often seen as a black box, which, fed with a current project value and a volatility estimate, produces a number, which is claimed to reflect the value of flexibility. The real options approach does not allow for a gradual and gentle enhancement of the NPV analysis. Instead it requires the users to replace the NPV by a typically considerably larger value, without providing much guidance to build up the intuition behind this additional value. The aim of this article is to explain a process that starts from the NPV spreadsheet and enhances the analysis by gradually incorporating uncertainty and flexibility. It is hoped that this process will help practitioners to build up intuition at their own pace and thereby enhance their understanding of the risks involved in a project. As we will see later, standard real options ideas can be, but don't have to be, part of this process.

III. A STEP-BY-STEP VALUATION PROCESS

A. Setting the scene: Formulating the NPV model.

An NPV model is indispensable as a tool to capture the *logic* of cash flow generation. It describes the relation between expenditure and income and is as such a functional description of cash flow generation. To obtain the net present value, one assigns numbers to all the uncertain cells in the spreadsheet. The spreadsheet then calculates the corresponding stream of cash flows; the NPV of the project is obtained by summing up the period cash flows, discounted at the company's hurdle rate to reflect the company's cost of equity capital[2], and

---

[2] This is the return on equity capital that the company's owners expect to earn. In our setting, we do not want this discount rate to include a project-specific risk-premium, as is often done for projects that are perceived to be riskier than the company as a whole. We prefer to reflect the risk directly through an appropriate distribution of possible outcomes. Using a risk-premium in the discount rate would therefore double-count risk.

Traditially, cost of capital calculation are based on the capital asset pricing model (CAPM) or one of its relatives. The CAPM cost of capital is calculated as the risk-free rate of return (e.g on a 10 year treasury bill) plus a premium for the risk of the company as a whole. This premium is calculated as the historical difference between risk-free rate and the rate of return on the stock market as a whole (e.g. on the S&P 500 index), multiplied by a company-specific adjustment factor (called the stock's beta) which reflects the volatility of the stock as a percentage of the volatility of the market as well as the correlation with the market, i.e. the degree to which the stock's movements are aligned with movements of the market as a whole. CAPM cost of capital has been criticized in the literature. It is based on a diversification argument and assumes that company-specific, so-called unsystematic, risk is negligible since it can be reduced by diversification. An interesting alternative to CAPM has been recently suggested by McNulty et al. (What's your real cost of

subtracting the initial investment. Modern spreadsheets have the advantage that they allow a quick "what-if" analysis. Uncertain inputs can be changed and the corresponding changes of the NPV can be tracked and graphically represented as illustrated in the chart below, where the dashed lines represent the demand forecast and the zero-NPV line, respectively.
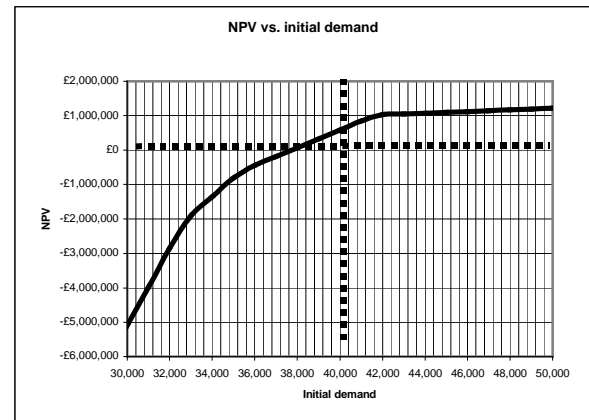

Such sensitivity charts are obviously very valuable and part of every professional NPV analysis. However, a major drawback of this type of analysis is that the uncertainties have to be dealt with one-by-one rather than simultaneously. We can draw a graph that shows how the NPV changes with initial demand, provided all other uncertainties, like demand in future years, prices, costs, etc. are fixed. However, such charts may be misleading since they typically neglect implicit dependences between variables, such as price and demand. Even if these dependences are incorporated, it is quite difficult to illicit sensible information about the risk inherent in the project from an inspection of many different charts. Our first aim is therefore to enhance the sensitivity analysis by producing a single chart that takes *simultaneous* changes of all inputs into account and thereby provides a better picture of the risk profile as a function of *all* underlying uncertainties.

---

capital, HBR October 2002). They suggest that a company's cost of capital consist of three parts: the time value of money (taking care of macro-economic risks, such as inflation, and measured by the return on long-term government bonds), a debt-related component (taking care of default risk and measured by the difference between the company's credit spread, i.e. the difference between the return on government bonds and the return on corporate bonds), and a company-specific risk component (taking care of the risk associated with earnings volatility). The latter component is the price of an insurance against the event that the company's equity return will not meet the return on corporate bonds. The share price needed to achieve the return target of the corporate bonds can be easily computed. An insurance against the event that the share price will not exceed this target is therefore a put option with the latter share price as exercise price. The price of such a put option can be computed using standard financial options pricing formulas.

B. Incorporating uncertainty

Starting point for the incorporation of uncertainty in the model is the simple but important fact that the project value is known only ex post, when all uncertainties have been resolved. At present the value of a project is not a single number but a range of numbers. Therefore, whether managers like it or not, projects are gambles. No professional gambler would base the decision to engage in a gamble on a single number as an estimate of payoff. The gambler would want to have more information about the possible payoffs of the gamble and their likelihood. If this is so for a professional gambler, then it should be even more so for a professional business decision maker. We therefore start from the premise that the *value of a project is a distribution*, i.e. a range of numbers and associated likelihoods. Rather than attaching a price tag to a project we want to estimate its value distribution and then decide on the investment in the light of this distribution.

The beauty of the concept of a distribution is that, albeit a fairly complicated mathematical construction, it can be easily visualized and understood as a geometric *shape*. This allows the decision maker to grasp, after some habituation, the risk profile of the project at one glance. In fact there are several ways of depicting the risk profile of a project. The two most popular graphical representations are a histogram and a cumulative distribution function or value-at-risk chart.
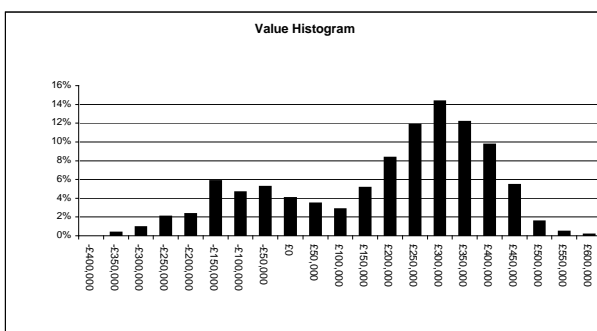
The histogram represents the chance of the value falling into a certain range. The range is typically indicated by a mid-point; the width of the intervals can be easily inferred from neighbouring mid-points.
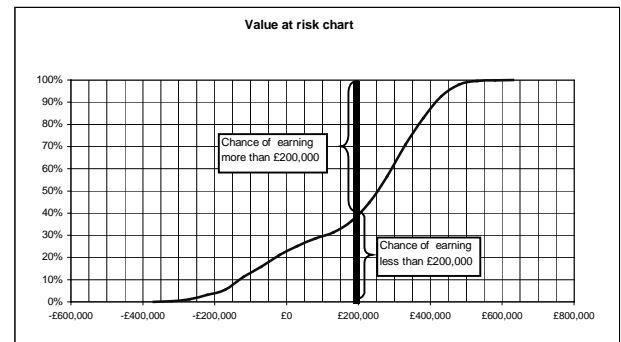
The curve on the value at risk chart gives the chance that the value lies below or above the target values given on the horizontal axis, as explained in the above figure. Which representation one prefers is a matter of taste and habituation. We prefer value-at-risk charts since they tell the decision maker at one glance the probability of earning at least £x pounds or loosing at least £x with the project.

Risk shapes are multivariate analogues of sensitivity charts; they provide a similar picture but incorporate all key uncertainties. A shape visualisation of the risk profile of a project is therefore the desired output of our analysis at every stage. We have no wish to condense these shapes to single numbers since numbers are misleading and give a false impression of certainty. Risk profiles are ideally suited to convey the risk inherent in a decision and they are, after some training, easy to understand and, in fact, easy to produce. Our experience shows that devices like value-at-risk charts, once the decision maker is used to them, are felt to be indispensable for rational decision making in an uncertain world.

So how can we produce risk profiles? The answer has been given decades ago by D.F. Hertz (Risk Analysis in Capital Investments, HBR 1964): Use Monte Carlo Simulation (MCS). The article dates back to a time when working with computers was difficult and setting up an MCS was a time-consuming task. Still the author considered it important and worth the extra effort. How much worthier is this in the days of powerful spreadsheet software that allows every reasonably experienced MBA to set up a meaningful MCS for many projects in less than an day. It seems irresponsible not to do it. We explain the workings of MCS in an appendix.

Aside: The flaw of averages.

An important positive side-effect of MCS is that it avoids the "flaw of averages". This term was coined by Sam Savage (HBR, Nov. 2002) to describe the subtle fact that values calculated on the basis of average inputs are not necessarily average values. In other words, even if we were able to give precise estimates of the averages of the value driving uncertainties, like prices, costs, demands, etc., we could not guarantee that the NPV calculated on the basis of these averages is the average NPV of the project. The flaw of averages is typically a consequence of non-linearities or of statistical dependencies in the model; both are abound in business spreadsheets. To illustrate the flaw, suppose the only uncertain element is the demand for a product and that this demand is symmetrically distributed about a mean. Suppose also that there is a capacity constraint in the system. If the realised demand falls below the mean demand, your NPV suffers but the chance of that happening is balanced out by the chance that your demand falls above the mean in which case your NPV flourishes. A naïve person would now argue that the NPV should balance out as well and therefore the average NPV should be the same as the NPV you calculate on the basis of the average demand. Upon reflection, the fallacy in this argument becomes obvious: Very low NPV scenarios due to low demand are not balanced out by very high NPV scenarios due to very high demand because high demand cannot be fully realised if it exceeds the capacity constraint. Therefore the NPV based on average demand, which is the number calculated in your NPV analysis, is larger than the average NPV of the project. Savage gives several interesting examples of the flaw of averages, which is lurking in almost every business spreadsheet. This is important as it shows that the use of projections incorporates bias into an NPV analysis; even if the projections are the expected quantities, the corresponding NPV is typically not the expected NPV. MCS avoids this problem; the expected value of the produced sample of NPVs is, by the very nature of MCS, an unbiased estimator of the expected NPV.

C. Incorporating flexibility

Just as the latter step, this third step starts again from an obvious premise: Flexibility has value. The more flexibility a project incorporates, the more robust it is against adverse effects in the future and the more it will allow the exploitation of positive opportunities if and when they arise. If we have the flexibility to extend the project then this can increase the value of the project substantially if the demand for our product or service is larger than expected. This can, for example, be an argument to build a plant in a more expensive area, if this allows easy hiring of qualified staff if we wish to

expand. Every project incorporates such flexibilities and these flexibilities have value which, however, is not incorporated into an NPV analysis. It is not only not incorporated in the analysis, it is also often overlooked when crucial decisions are being made during the design phase of the project. Indeed, whilst most projects incorporate certain generic flexibilities, such as expand, abandon, contract, defer, etc., it will sometimes be worthwhile to spend extra money at the start of a project to allow for more flexibility later on. A decision maker needs a way to incorporate flexibility into the analysis, to see how flexibility changes the risk profile of the project and to gauge whether it is worthwhile to spend additional money to add more flexibility.

Our starting point is the risk profile of the "passive" project obtained in foregoing step. On the basis of this risk profile, the project design team will have to think carefully and creatively about the types of flexibility that may have a substantial impact on the risk profile of the project, i.e. shift the value-at-risk curve to the right. The tails of the risk profile are particularly interesting: Which scenarios are responsible for detrimental effects and which types of flexibility will allow the project manager to cut losses when these detrimental effects occur? Which scenarios are responsible for positive values and how can these opportunities be amplified if and when they arise? The former types of flexibility have an insurance character and are analogous to put options, which hedge the owner of a stock against a drop in share price; the latter types are analogous to call options, which allow the owner to buy cheap if share prices increase.

To begin with, we incorporate decision points into our model. At these points in time we revisit the project and ask ourselves, in the model world, which of several alternative actions we should take. At the time of the decision, some of the underlying uncertainties will have been resolved and the decision will be based on the current information. For example, the decision to change the scale of the project can be taken after the first year, when crucial initial demand or other market information is available. The decision therefore depends on how the first year evolves; at present we can only device a contingency plan or decision rule, which specifies which decision will be taken under which circumstances. We could, for example, say that we abandon the project after the first year if the first year's demand is below a failure threshold and expand if demand exceeds a success threshold. Structurally, such decision rules can be easily

incorporated in a spreadsheet. However, just as in the case of the passive NPV spreadsheet, the decision rules need to be populated with numbers. In the above case we need to specify the failure and success thresholds as well as the size of expansion in case of success. But how should we determine these parameters in a decision rule? Obviously, we should choose those parameters which maximise the value of the project. However, we had said that the value of the project is *not* a number but a distribution. Changing the parameters of our decision rules will change the distribution and whether one distribution is preferable to another is not always clear-cut and may depend on the decision maker's attitude towards trading off risk and opportunity. Let us stress this simple observation: *Whilst numbers can be optimised, distributions cannot.* Our preference for distributions is not as clear-cut as "*x* is larger than *y* therefore *x* is preferable to *y*".
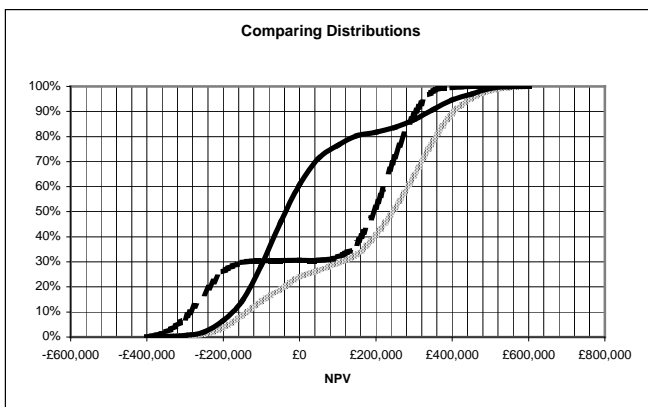
In the above graph, the gray distribution is clearly preferable to both the black and dashed distribution. Indeed, the further to the right the value-at-risk curve, the better the risk profile of the project. However, if we had to choose between projects with the black and the dashed risk profile, respectively, then the decision would be less clear-cut: The dashed profile is preferable in the middle range values, but has larger downside risks and less up-side opportunity.

So how shall we design a decision rule if we find it difficult to decide which of two risk profiles is preferable? Designing good decision rules is obviously all but simple. Notice, however, that whatever decision rules we incorporate into our analysis, they will lead to a conservative valuation. It may well be that we find better decisions in real-time as the project unfolds. The call is therefore for tools, which enable us to design sensible decision rules that shift the value-at-risk curve

"as far to the right as possible" in order to make sure that we do not turn down good projects. The right-shift of the value-at-risk curve stipulates the value maximisation effort of good project management. But how can we "optimise" a decision rule? Initially, it may be possible to find parameters that shift the curve as a whole to the right but at some point changes in the parameters will only move one part of the curve to the right whilst another part will pop out to the left. This is reminiscent to changing the shape of a balloon that contains air and water, but whilst air can evaporate, water can't. Suppose we want the points on the surface to be as close as possible to some target point in the middle of the balloon. By letting the air out of the balloon we can make each point on the surface getting closer and closer to the target point. When there is no air left, we can still press some points closer to the target but only at the expense of some other points on the surface moving away from the target. If we are at this point in our risk profile design, we may still be able to push the left-side of the value-at-risk curve to the right, thereby improving the downside risk, but only at the expense that the right-side of the curve pops to the left, having a detrimental effect on the positive opportunities, and vice versa. A push on the left side makes the range of possible outcomes smaller, a push on the right side makes it larger. Now it depends on the decision maker's attitude to risk which risk profile he or she prefers.

In technical terms, we will invariably apply some sort of an optimization process to find a suitable shape of the curve. Whatever can be said about multi-criteria optimization and trade-offs: If we wish to optimise, we somehow need to condense the objective to a single number. We do not suggest going back to the "numerical project value"-paradigm but we must acknowledge that specifying a decision rule through an optimisation process will require us to specify a numerical valuation mechanism. There are various levels of sophistication in designing decision rules. Tools that are available include decision tree analysis and more general dynamic programming approaches, real options analysis, and stochastic optimisation. Decision trees are particularly valuable in setting up the timing and structure of the key decision rules, because they allow a visualisation of the interplay between uncertainties and flexibilities. It is interesting, and not a coincidence, that their use was promoted by J.F. Magee in the Harvard Business Review shortly after Hertz's article on Monte Carlo simulation (Decision trees for decision making, HBR July 1964). Decision trees can, to some extend, also be

used to optimize the parameters in the decision rules, either by a standard expected utility based valuation or by a replication based real options valuation. Indeed, it is at this stage of "optimising" the use of flexibility, where the risk-neutral valuation process underlying standard real options valuations can be very helpful indeed, either within a decision tree or as a stand-alone procedure to determine a sensible exercise rule for flexibility. However, in contrast to ROV, the obtained numerical "value of flexibility" is not so important; what is important is the contingency plan that ROV provides. Simulation-based stochastic optimisation is another alternative technology to design contingency plans. It allows us, for example, to determine those parameters of a decision rule that maximise the expected value or a certain percentile of the value, or to determine the parameters that shift the risk profile as close as possible to a target profile. We will not go into the technical details here. The more tools are available, the better it is, and more research in this area will hopefully expand and improve the tools available to practitioners. What is important is that all these techniques, including real options valuation, are part of a toolbox that helps to produce sensible decision rules rather than project values. Which decision rule we prefer in the end is up to us; the optimisation tools do not prescribe an "optimal" rule but are merely heuristics that help us to find candidates for decision rules. We need to incorporate the most promising rules into the risk enhanced spreadsheet to see what effect they have on the risk profile and then choose a rule according to our risk-opportunity tradeoffs. It is important to stress again that this process is conservative as we may find more suitable decisions in real time than the ones we anticipate in our contingency plan.

D. Producing a risk profile for the flexible project

To test a decision rule, we incorporate it into our risk-enhanced NPV model. For simple rules this can be directly done in a spreadsheet and does not require sophisticated programming. Once that has been done and the parameters of the rule have been chosen, the risk profile will have changed, hopefully to a more favourable shape, pushing the value-at-risk curve further to the right. Some manual tuning of the parameters may further help to improve the risk profile of the project. Once we have decided on the decision rules, we have produced the final risk profile of the project, which now replaces the NPV as a basis for the investment decision.

IV. AN ILLUSTRATIVE EXAMPLE

Consider the decision to launch a new product, requiring an investment into new production facilities.

Step 1: Setting the scene: Formulating the NPV model

We begin by formulating the cash flow generation mechanism of the project. A simple initial NPV analysis, based on a 5 year time horizon, could look like the following self-explanatory spreadsheet, where uncertain cells are marked in gray.

Year		1	2	3	4	5
Initial demand		40,000				
Demand change			40%	30%	-20%	-70%
Demand		40,000	56,000	72,800	58,240	17,472
Production		55,000	55,000	55,000	55,000	0
Units sold		40,000	56,000	69,000	55,000	0
Inventory position		15,000	14,000	0	0	0
Price		£300.00	£300.00	£200.00	£180.00	£150.00
Unit cost		£150.00	£140.00	£130.00	£120.00	£110.00
Production fixed costs		£950,000	£950,000	£950,000	£950,000	£400,000
Marketing and sales cost		£2,500,000	£2,000,000	£1,500,000	£250,000	£100,000
Initial investment	£9,500,000					
Salvage value						£2,500,000
Cash flow	-£9,500,000	£300,000	£6,150,000	£4,200,000	£2,100,000	£2,000,000

| DCF Analysis						
Discount rate	10%					
Discounted Cash flow	-£9,500,000	£272,727	£5,082,645	£3,155,522	£1,434,328	£1,241,843
NPV	£1,687,065					

Step 2: Incorporating uncertainty

In the second step we implement distributions for each uncertainty. Here, we use triangular distributions, which are determined by a range and a most likely value (see the appendix for details). The distributions cover a range "projection +/- spread" with the projection being the most likely value. The key uncertainty is initial demand which we allow to fluctuate by 50% about its projection.

Year		1	2	3	4	5	Spread
Initial demand	forecast	40,000					
	realised	30,978					50%
Demand change	forecast		40%	30%	-20%	-70%	
	realised		37.9%	29.1%	-20.6%	-72.1%	10%
Actual demand		30,978	42,729	55,166	43,782	12,214	
Production		55,000	55,000	55,000	55,000	0	
Units sold		30,978	42,729	55,166	43,782	12,214	
Inventory position		24,022	36,293	36,127	47,345	35,131	
Price		£300.00	£300.00	£200.00	£180.00	£150.00	
Unit cost	forecast	£150.00	£140.00	£130.00	£120.00	£110.00	
	realised	£149.80	£142.93	£139.77	£125.51	£108.84	10%
Production fixed costs	forecast	£950,000	£950,000	£950,000	£950,000	£400,000	
	realised	£865,876	£979,942	£1,019,420	£973,545	£414,603	10%
Marketing and sales cost		£2,500,000	£2,000,000	£1,500,000	£250,000	£100,000	
Initial investment	£9,500,000						
Salvage value	forecast					£2,500,000	
	realised					£2,543,293	10%
Cash flow	-£9,500,000	-£2,311,408	£1,977,469	£826,444	-£245,936	£3,860,824	

DCF Analysis						
Discount rate	10%					
Discounted Cash flow	-£9,500,000	-£2,101,280	£1,634,272	£620,920	-£167,978	£2,397,268
NPV	-£7,116,798					

Notice that the *realized* uncertainties in the spreadsheet are just *one* sample. We can use the spreadsheet to re-sample from the distributions by pressing the re-calculate button. Each time we do this, we obtain new realizations of the uncertain quantities and a corresponding new NPV sample. If we do this many, say 5,000, times and collect these NPVs in a list, we obtain a good idea of the distribution of the NPV. The list of NPVs can be graphically represented either as a histogram or as a value-at-risk chart like the one below.
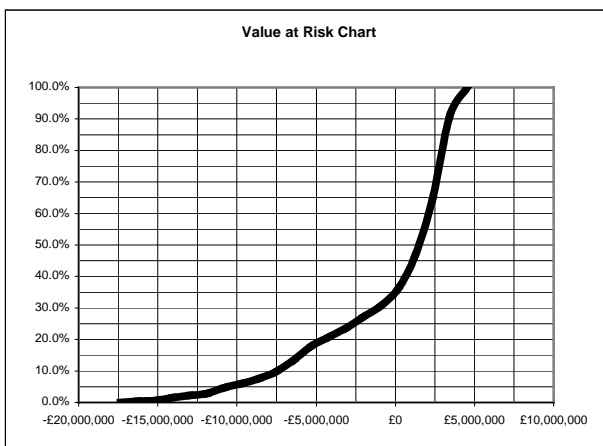
Value at Risk Chart

The value-at-risk chart shows quite a risky project. The chance of a negative NPV is estimated at roughly 35% and there is a 10% chance of loosing more than £7,500,000. On the other hand, the upside opportunities seem bounded. On the upside, the chance of an NPV above £2,500,000 is roughly 30%.

How does the obvious riskiness of this project compare to our initial NPV analysis, which had resulted in a juicy NPV of more than £1,600,000? The reason for the discrepancy is the flaw of averages that we mentioned earlier. Losses due to potential negative demand are not balanced out by large profits due to positive demand because the constraint on production capacity of 55,000 units does not allow us to capture the full demand if the realized demand is above the capacity constraint of 55,000 units. In fact, a 95% confidence interval for the average NPV, on the basis of a sample of 1,000 NPVs, gives a *negative* expected NPV estimate of £ -450,000 +/- £ 300,000. Compare this with the originally estimated NPV of £ 1,687,065! This is a stark example of the flaw of averages at work: The NPV calculated on the basis of average uncertainties is far from the average NPV of the project.

Step 3: Incorporating Flexibility

Presented with the risk-profile of the "passive" project, we are now challenged to incorporate flexibility to improve the risk profile. One may for example incorporate an abandonment option if the demand is low in order to avoid large losses that can occur in the left tail of the distribution. As for the right tail, we should search for flexibilities that allow us to exploit opportunities if and when they arise. We may, for example, start with a smaller production plant to begin with and expand the plant after the first year, provided the crucial initial demand is large enough. Such decisions can be easily implemented in the spreadsheet. A simple expansion decision, for example, depends on three parameters: the initial production capacity, the extension capacity and the observed demand size in year 1 that triggers the extension. It is easy to implement such a decision rule in the spreadsheet. Each set of parameters then leads, after simulation, to a different risk profile and it is quite a challenge, even for only three parameters, to find satisfactory parameter settings. It is here, where sophisticated technical analysis can help tremendously, in particular with complicated composite options that can depend on many parameters. For the simple expansion option in our example we

found, by trial and error, that an initial capacity of 50,000 units and an extension by 10,000 units if demand in year 1 exceeds 40,000 units improves the risk profile quite significantly.

Step 4: Value-at-risk chart for the flexible project

A Monte Carlo simulation of the spreadsheet with this decision rule results in the risk-profile below, with the gray profile corresponding to the original "passive" project.
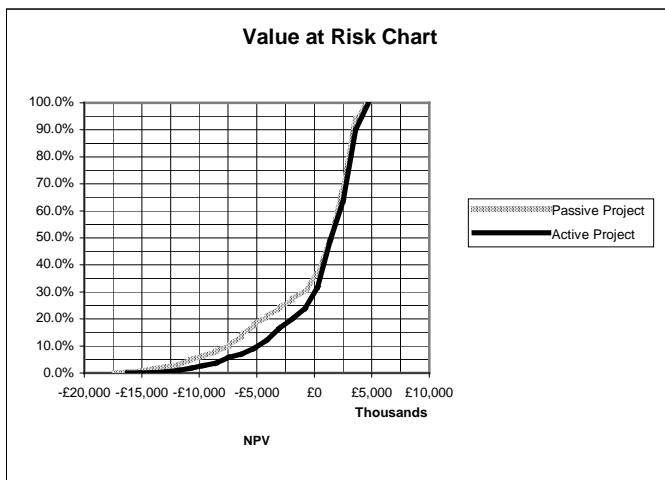
As can be seen, the flexibility has impact mainly on the downside risk. The chance of a negative NPV has been reduced from 35% to 25% and the 10% value at risk is now £5,000,000 as opposed to £7,500,000 earlier. The average NPV of the flexible project has improved significantly and is now, again on the basis of a sample of 1,000 NPVs, estimated to be positive at £ 450,000 +/- £ 240,000 with 95% confidence.

The project can now be further inspected and additional flexibility can possibly be added to further improve the risk profile.

When this process ends, and no additional flexibility can be thought of that will substantially improve the project, then we have obtained the final risk profile of the project. It is this risk profile that the company's investment decisions should be based on. This device seems infinitely better suited as a basis for an investment decision then a single number, in however sophisticated a manner obtained. Also, and possibly more importantly, the project design and valuation team will have learned a lot about the risks and opportunities inherent in the project by trying to improve the risk

profile through the incorporation of flexibilities. In this way, the process not only improves the quantitative understanding of risk and opportunity but also provides a guide to pro-active rather than re-active risk and opportunity management during the project design phase by focusing the view on the interplay between uncertainty and flexibility.

Aside: Valuing flexibility

It is worth mentioning that MCS can also be used to separately value flexibility, here of the extension option, and thereby help to decide whether or not to spend money to allow for additional flexibility. The value of flexibility is of course uncertain: It will be equal to the cost of flexibility if the flexibility is never exercised and even if it is exercised, its value depends on just how good the opportunity is that it allows to exploit or how bad the loss would have been that it allows to avoid. To estimate the distribution of the value of flexibility we simulate, simultaneously, the NPVs of the passive and the flexible project and subtract the former from the latter. Thus we generate many, say 5000 scenarios of the exogenous uncertainties and subtract the corresponding value of the "passive" project from the value obtained by exercising the flexibility according to our contingency plan. In the example, this results in the value-at-risk chart below.
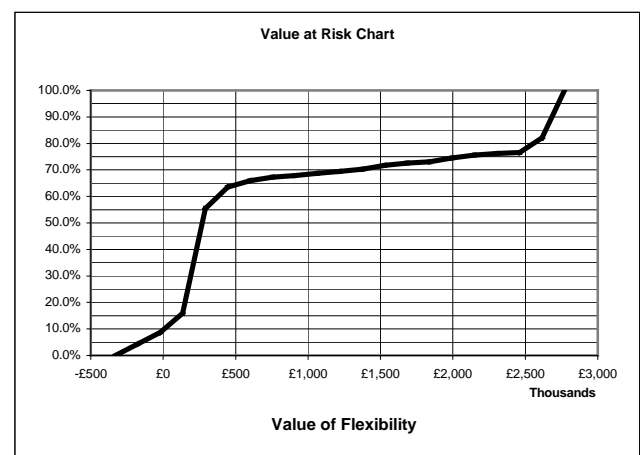
This value at risk chart for the expansion flexibility exhibits a rather typical bi-modal distribution. There is a roughly 65% chance that the value of the expansion flexibility is in the range of -£300,000 to £500,000, corresponding to scenarios where no expansion occurs, and a roughly 25% chance that the flexibility adds value in the range of £2,500,000 - £2,8000,000, coming from

scenarios when the option is exercised. Notice that the small chance that the value of the extension flexibility is negative is not apparent on the foregoing chart because its resolution is not fine enough. The negative value of flexibility can occur because we have decided to start the flexible project with a smaller capacity of 50,000 units, instead of 55,000 units for the project without the extension. This lack of initial capacity can lead to foregone profits, for example if initial demand is high but demand growth over time is low or even negative.

V. FROM SINGE PROJECTS TO PROJECT PORTFOLIOS

Projects are seldom considered in isolation. The value of a project is, in fact, the value that this project adds to an already existing portfolio of projects. In an uncertain world, project values are not additive. Indeed, running many small projects whose fates depend on different and fairly independent uncertainty drivers is much preferable to a portfolio of a few large projects because the chance of many of the small projects going wrong is rather low. A failure of one project is likely to be balanced out by a success of another. This effect, known as diversification or the law of large numbers, can be easily replicated by setting up a master NPV distribution which is obtained from the individual NPV distributions by sampling from these distributions and adding up the sampled values to obtain a single scenario for the portfolio NPV. Repeating this in a Monte-Carlo fashion gives the risk profile of the overall portfolio.

Unfortunately, this simple process is very likely to give a wrong representation of the risk profile of the portfolio. The reason is that the assumptions for the law of large numbers are not satisfied. This law is only valid if the value drivers underlying the individual projects, and therefore their NPVs, are independent. In the real world, the value drivers of projects often depend on one-another. In order to deal with this phenomenon we need to make sure that the dependent value drivers of the individual projects are sampled appropriately. One way of doing this is to estimates the correlations between the NPVs of the individual projects and to use this information to sample appropriately from the NPVs, taking their dependence structure into account. However, any such methodology, being based on direct sampling from NPV profiles, assumes implicitly that the decision rules of the individual projects are independent and can therefore remain unchanged in the portfolio

context. Obviously, this assumption is difficult to maintain in a world of budget constraints.

Portfolio managers have to allocate limited financial resources to the various projects and will change these allocations over time, depending on how the projects unfold. The allocation of project budgets has obvious consequences on the admissible managerial flexibility, i.e. the possible contingency plans, for each project. What is needed, for example to inform budget allocation decisions, is an integrated model, which takes account of the interdependences between the projects and allows a simultaneous design of decision rules, taking account of the common constraints and statistical dependencies. One approach could be to build a regression model, which specifies certain independent value drivers and how the remaining value drivers of all projects depend on these independent values. The value drivers for all projects are then obtained by sampling the independent values and sampling, independently for each value driver, from a suitable error term. There are alternative ways of dealing with the issues of dependence between projects but their explanation goes beyond the scope of this article. In any case, an effort to model the total portfolio value is much preferable to sampling from value distributions of individual projects.

The interrelations between projects in a portfolio are quite complex and it is therefore unrealistic to assume that they can be captured by simple models or rules-of-thumb. Nevertheless, trading off simplicity and accuracy is important in this context. A fully integrated model of a large project portfolio is unlikely to be acceptable as a decision aid, not only because of its inherent complexity but also since projects are managed individually and therefore the individual project managers need to be able to recognize the boundary and interplay between their project and the remaining projects in the portfolio. What is needed is a suitably decomposed representation of the portfolio, capturing, on the one hand, the common as well as the individual value drivers of the projects and, on the other hand, the common resource constraints and the impact of resource allocations on the available managerial flexibility within each project. Clearly, making sense of the risks and opportunities inherent in a project portfolio subject to resource constraints requires a substantial effort.
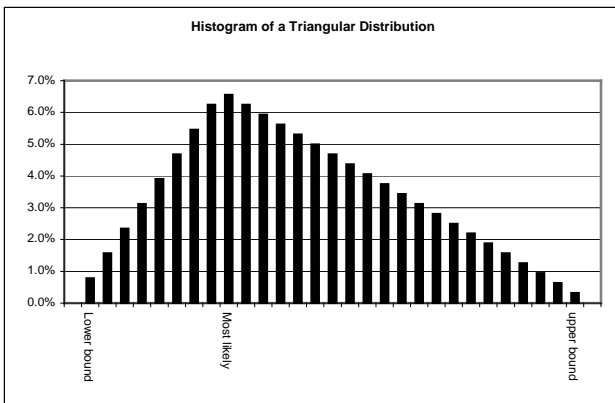
VI. CONCLUSIONS

The real options lesson is that uncertainty and flexibility have to be taken into account to get a valid picture of the project value. Uncertainty implies that the project value is not a number but a distribution, whilst flexibility implies that this distribution can be improved, provided the manager is skillful in using the flexibility to avoid losses and to exploit opportunities as the future unfolds. There are a variety of techniques for project appraisal and they are all limited in scope. It is therefore important to follow a robust valuation process that, being based on the two premises that the project value is a distribution and that flexibility improves the distribution, allows the incorporation of many techniques and thereby leads to an improved understanding of the risk and opportunities. We have suggested such a process, which can incorporate many state-of-the-art quantitative methodologies, including real options valuation, Monte Carlo simulation, and stochastic optimisation. The process allows for a smooth transition from NPV to a mature recognition of risk and opportunity. Whilst the quantitative information that is gained during the process can be very helpful, we wish to stress that it does not add information. It only makes information more transparent. One should never take the risk profile as a precise valuation. After all it has been generated through a sampling process and, more importantly, it is likely that some of the input distributions are based on limited data or include subjective judgments. However, this does not devalue the process. Indeed, the process itself is arguably its most important aspect. The desire to produce sensible risk profiles forces us to make risk, opportunity and managerial flexibility an integral part of our valuation process. The basic understanding that the journey is the goal in business modeling remains true.

We hope that the use of sensible risk- and opportunity-enhanced project design and valuation processes like the one presented in this paper will enable companies to gradually develop a "real options culture". Such a culture will constitute a move away from the still prevalent backwards looking "project control" perspective which focuses on plan achievement, and will provide a more forward looking "project management" perspective with an emphasis on value maximisation. In such a culture, uncertainty is not an evil but is embraced as an opportunity for those who are prepared to deal with it better than their competitors. Pro-active project design with a recognition of potential risks *and* opportunities is an important aspect of this culture. However, in order to develop a real-options culture, companies need to move away from "number-based" performance measures. They are a major obstacle to sensible strategic decision making. Numerical performance indicators do not embrace what is common wisdom in decision making: *Good decision can have bad outcomes in an uncertain world*. If a manager's pay or even job depends on a particular number then his or her first goal is to make sure that this number is achieved, thereby possibly foregoing risky but potentially very beneficial opportunities. Numerical performance indicators are only sensible in a business environment, where the law of large numbers governs the exogenous effects on overall performance. They are particularly misleading when important strategic decisions are being taken. Companies should recognize this: Whilst single decisions can have bad outcomes, a culture of making informed decisions and taking calculated risks will, in the long run, lead to good outcomes and is very much in the long-run interest of the company.

Appendix: How does Monte Carlo Simulation work?

Whilst traditional NPV analyses take projections of the uncertainty drivers as input and produce a numerical estimate of the value of the project, a Monte Carlo simulation takes distributions of the uncertainty drivers as inputs and produces a distribution of the value. For example, instead of asking the marketing department for a single estimate of the demand for the new product, we may ask for a distribution of this demand. More simply even, we may ask the marketing manager for a range and a most likely value, construct a distribution, and check with the marketing manager that the distribution is sensible. A very simple but effective distribution, based on a range and a most likely value, is the triangular distribution depicted below.

Histogram of a Triangular Distribution

quantity is higher than expected, then the other is likely to be higher as well, or negatively, i.e., when one quantity is higher then expected then the other is likely to be lower. Such dependencies have to be taken into account when the spreadsheet is randomised. We will not go into technical details here but alert the reader that caution is required in such an analysis and that the advise of an expert should be sought in case of doubt. Nevertheless, a reasonably experienced user, far from being an expert, can validly randomise NPV spreadsheets of many projects.

In other words, the likelihood of demands close to the ends of the range is small while the likelihood of demands close to the most likely demand is large. The triangular distribution is uniquely determined by the lower, upper, and most likely value.

The underlying technology for Monte Carlo simulation is that of random number generation. A computer can be programmed to produce sequences of numbers that look very much as if they were randomly generated, even to the testing eye of a statistician. Triangular and many other random number generators are included with spreadsheet MCS add-ins. If such a generator is invoked in a cell then each time we press the re-calculate button, the computer samples a new value from this distribution; if we sampled many, say 5,000, values with a triangular generator, collected them in a list and constructed the histogram of these data then the shape would resemble the graph above.

In this way, the NPV spreadsheet can be "randomised" by replacing the number in each uncertain cell by a random number generator which samples a number from an appropriate distribution. Each time we press the re-calculate button, the computer fills all uncertain cells simultaneously with sampled values. In other words, the computer creates a scenario from appropriate distributions and calculates the corresponding NPV. We can ask the computer to do this many, say 5,000 times and to gather the corresponding NPV values. This is easily accomplished with appropriate spreadsheet software. The risk profile of the project is now obtained by setting up a histogram or value-at-risk chart from this list of NPV values. This shape reflects the spread of NPVs obtained from 5,000 or more scenarios, each sampled from the appropriate distributions.

There are many technical aspects to consider when an NPV spreadsheet is randomised. We mention here only the problem of statistical dependence. This relates to the fact that uncertain values, such as the demand in year 1 and the demand in year 2 or interest rates and inflation rates, can depend on one-another, either positively, i.e., when one