David Laughton Consulting Ltd.
Draft: 05 Mar 2007

The Banff Taxonomy of Asset Valuation Methods: Lessons from Financial Markets for Real Asset Valuation in the Upstream Petroleum Industry

Abstract

The SPE (Society of Petroleum Engineers) held a workshop in Banff, Canada in 2003 to discuss the future of real asset valuation in the upstream petroleum industry. A taxonomy of valuation methods was developed to organise the work of that workshop. This paper describes the taxonomy and the reasons for its structure. It then reports on how it may be used to organise a discussion of:

1) the evolution of asset valuation:
   a) in the upstream petroleum industry since 1960; and
   b) in financial markets since 1970;
2) some lessons the industry might learn about valuation from developments in financial markets;
3) the foundations of, and similarities and differences among, the valuation methods used, or proposed for use, in the industry; .
4) some implications for the choice of valuation method arising from the interaction of strategic analysis and asset valuation as parts of the asset selection, design and management process;
5) organisational considerations in the design:
   1) of a valuation method for use in the asset design, selection and management process; and
   2) of a process to change valuation methods, if change is desirable;
6) several misconceptions about so-called "real options" approaches to valuation, in particular showing why two well-known proposals for change in the direction of dealing with "real options" are dead ends; and
7) potential future developments in real asset valuation.

The origin of the taxonomy

The managers of a public corporation have a fiduciary responsibility to maximise the value of its assets, as determined by their price in the financial markets of relevance to its investors. There are a wide variety of approaches that attempt to determine which management alternative in any given commercial situation will maximise asset value. Practice in the upstream petroleum industry reflects this diversity and has evolved over time, as organisations have sought processes not only to make better choices from among alternatives for the design and management of real assets, but also to generate a better range of alternatives to consider. An important part of many of these approaches is the estimation of asset value using discounted cash-flow (DCF) methods. There has also been exploration of some other approaches to estimating asset value, some of which are known in the industry as "real options analysis" (ROA) methods.

Unfortunately, there has been some confusion over the conceptual underpinnings of the various approaches to value estimation, the relationships among them and the implications of this for standards of best practice. Because of this, a group of upstream petroleum industry managers and consultants decided to use the technical forum and workshop process of the Society of Petroleum Engineers (SPE) to organise some pre-competitive investigations of these issues. Technical fora were held in 2000 and 2002. Participants in these fora requested a workshop on topics in "advanced" economic analysis, including probabilistic DCF analysis, decision tree analysis, and real options analysis.

In response to this, the SPE organised a workshop on "The Theory and Art of Asset Valuation: Building a Case of Change - Applying to the Oil and Gas Industry What Finance has Learned". The workshop was held on 15-17 September 2003 in Banff, Canada.

At this workshop, we addressed the following questions.

1) Can organisations in the upstream petroleum industry improve their asset design and management decision-making process?[1]
2) If so, is there a role in this for the use of "improved" methods to estimate individual real asset values?[2]
3) If so, are there insights to be had from developments in financial markets over the last 30 years about how better to estimate real asset values?
4) If so, is it worthwhile to consider changing valuation processes from those currently in use to gain these insights?
5) If so, how can this be done with most benefit at least cost?

One of the key outputs of the workshop was a taxonomy of valuation methods that has come to be known as "the Banff taxonomy". (See Figure 1.) It was designed to support discussions to clarify some of the confusion about valuation methods that has plagued the industry over the last couple of decades. Workshop participants and others since have found it to be very useful for this purpose.

---

[1] We defined an improved decision-making process in an organisation as one that is more likely to produce decisions about asset design and management that cause the value of assets of the organisation to be as large as possible.

[2] Because of asset interactions, acting to maximise the value of a single asset on a stand-alone basis may not maximize the value of the organisation as a whole. To make individual asset valuations more relevant, we can model some of the interactions among assets and include them in the individual asset valuation. These include some joint tax effects, some aspects of common infrastructure usage, and some informational externalities. However, we do not know enough about some other interactions to model them explicitly. These include the effects on value of the overall enterprise risk profile and of the financing needs of the organisation. These issues must be considered with more qualitative analysis.

The structure of the taxonomy

Financial markets are the markets where claims to future cash-flows are traded. There are two characteristics of cash-flow that help to determine financial market prices:

1) timing; and
2) uncertainty.

Timing is relatively simple, and the effects of time on value are determined in roughly the same way by most asset valuation methods, using prices determined from claims to risk-free cash-flows, frequently in government debt markets. These prices are usually represented by a risk-free interest rate or a term structure of such rates.

Fig. 1 The Banff taxonomy of asset valuation methods

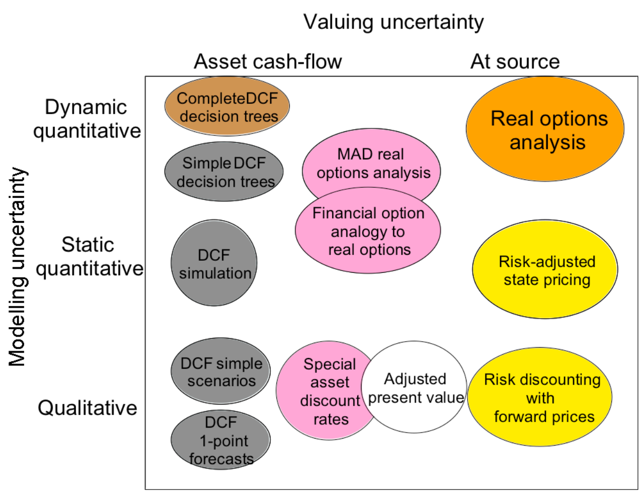

On the other hand, uncertainty is complex and multidimensional. It is modelled, and its effects on value are determined, very differently by different approaches to valuation.

Therefore, in organising our taxonomy of valuation methods, we decided to focus on how different valuation methods:

1) model uncertainty; and
2) determine the effect of uncertainty on asset value.

Models of uncertainty are usually based on possible realisations for the uncertain variables that are inputs into a model of the stream of asset cash-flows to be considered. These models of uncertainty can be either:

1) qualitative; or,
2) quantitative, if based on explicit probabilities for the realisations.

The quantitative models can be either:

2a) static; or,
2b) dynamic, if they model how the realisation probabilities are resolved over time in the face of the arrival of new information.

The effect of uncertainty on asset value is determined at the level of either:

1) the asset cash-flows themselves; or
2) the sources of uncertainty in the asset cash-flows (such as the term structure of oil or gas prices or geological uncertainties like oil-in-place).

If the valuation of uncertainty is done at the level of the asset cash-flow, this is usually accomplished by incorporating a risk premium, above the risk-free interest rate, into the discount rate used to discount some measure of each asset cash-flow.

The methods that use this approach are shown on the left side of the taxonomy in Figure 1.

If the valuation of uncertainty is done at the level of the sources of uncertainty, some risk-adjusted measure of the uncertainties underlying the asset cash-flow is used to determine a risk-adjusted measure of each cash-flow. These are discounted for time using risk-free interest rates to estimate asset value. Both the at-source risk adjustments and the time discounting in these methods are based directly on financial market prices, where those prices exist, or on proxies calculated as much as possible from financial market data where they do not exist. For this reason, the workshop called these "market-based valuation" (MBV) methods.

The methods that use this approach are on the right side of the taxonomy. The two approaches lower down on the vertical axis of the taxonomy are simpler implementations, that may be used in special circumstances, of the most general method of this type, which is shown by the bubble in the upper right corner of the taxonomy. More details on this are given below in the section on "Market-based valuation without decision tree analysis and the issue of planning vs market price forecasts".

There are also some methods (shown by the white and rose bubbles in the middle of the taxonomy) that are hybrid DCF/MBV approaches to the valuation of uncertainty. These are also discussed below in the sections on "Some partial moves to the right" and "Two dead end approaches to real options analysis".

A brief history of modern petroleum asset valuation

In the early 1960's, the upstream petroleum industry began to use asset value, estimated using DCF methods, as a metric in the process for making decisions about asset design and management.

DCF methods had been first applied correctly to asset valuation about 50 years earlier in the valuation of bonds, in situations where there is little cash-flow uncertainty. Prices for claims to future cash-flows were formulated as discount factors determined by the interest paid in the bond market as a result of the time value of money.

This was generalised to the valuation of assets with uncertain cash-flows, such as those in upstream petroleum assets, by discounting the forecast cash-flow, approximated by the cash-flow in the forecast realisation of the future, with discount factors based on a discount rate that includes the risk-free interest rate to take account of the time value of money plus a risk premium to reflect the generally lower value of claims to cash-flows that are uncertain. This is the method in the lower left corner of the Banff taxonomy, where the uncertainty is represented by the single forecast in Figure 2.

Fig. 2 The evolution of petroleum asset valuation

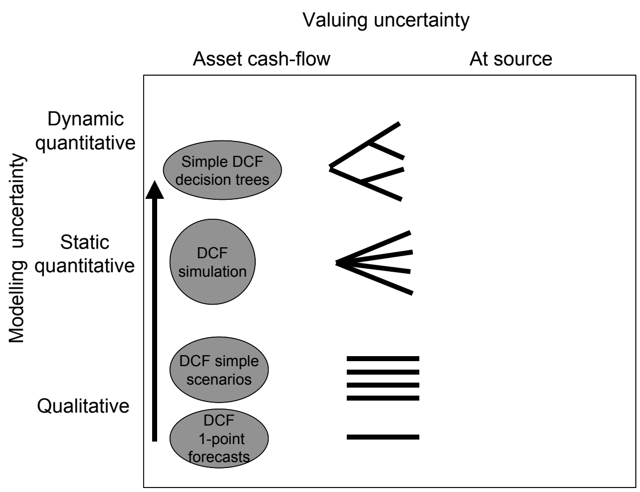

For the most part, the evolution in industry practice since then has been about the use of different ways of modelling uncertainty. This has roughly corresponded to a move up the left side of the Banff taxonomy, as shown in Figure 2. Three aspects of this have been:

1) the use of qualitative sensitivity analysis;
2) the introduction of quantitative probabilistic simulation based on random sampling (e.g. Monte Carlo simulation); and
3) the development of simple forms of decision tree analysis.

Because the forecast cash-flows were subject to model uncertainty, the practice of doing sensitivity analysis using different plausible forecasts was frequently part of the original process. This is represented by the multiple unconnected realisations in Figure 2 just above the lower left corner of the taxonomy.

Unfortunately, this was conflated with the use of different realisations, not to take into account model uncertainty, but to take into account the uncertainty in the realisations themselves. Therefore, instead of having several estimates of asset value, which would be justified given our definition of value, the different "values" became summary statistics representing some measure of ex-post happiness or regret, dependent on which realisation occurred.

This tendency was reinforced by the introduction in the late 1960's of static quantitative models of uncertainty based on probability distributions for the relevant realisations. This is represented in Figure 2 by the multiple realisations tied together at the beginning by a static probability distribution. Typically, these distributions are sampled randomly, using what are commonly called "monte carlo" methods, to estimate statistics of the distribution of realisation "values". A whole technology has developed, with software to support it and a language to communicate it, around the use of such static quantitative representations of future realisation uncertainty. With the increase in the last decade of desktop computing power needed to run this software, the use of this approach to valuation has become more widespread. Typically, some central tendency, such as "expected value", is taken as the measure of value, with some measure of the spread being a measure of risk.

This use of expected "value" as a measure of value brings us back to the first DCF definition of value, but with expected cash-flow itself, as opposed to the cash-flow in the forecast realisation, being the cash-flow measure that is discounted.

However, the use of some spread in the realisation "values" as a measure of risk is problematic. The effect of realisation uncertainty on value is already notionally taken into account by the risk premium in the discount rate. While model uncertainty might be represented by a spread of values, it is almost certainly much less than the realisation uncertainty that is usually sampled.

The last move up the Banff taxonomy also began in the late 1960's, at about the same time as monte carlo methods were first used. Some people began to use decision tree analysis to examine simple exploration and appraisal problems. The focus of this method is on sequential decision-making under uncertainty, where earlier decisions must take into account their implications for later decisions.

To support this type of analysis, we need a dynamic quantitative model of how uncertainty is resolved over time as decisions are made. This is represented by a tree structure for the realisations of the underlying uncertainties, as at the top of Figure 2.

Branching on the tree represents the arrival of new information to differentiate among possible realisations. Each realisation is represented by a single line of "descent" across the tree to an end state, and may be labelled by that end state.

Note that each state on the tree consists of a set of as yet undifferentiated realisations, and that each realisation is associated with one and only one state at any given time.

A more detailed description of the structure of such trees is given in Appendix A, using, as an example, a simple tree for a short-term oil field development lease.

After an initial spurt of interest, the use of the method did not grow very much until a few years ago, where, once again with the development of appropriate desktop computational power, software and training, it began to be adopted more readily, at least for the examination of the class of exploration and development issues for which it was originally used.

A brief history of financial markets since 1970

Financial markets have changed dramatically since 1970, particularly in the period of the 1980's and early 1990's. These changes have occurred under the influence of both demand-and supply-side forces.

On the demand side, the floating of currency exchange rates in 1971 with the collapse of the Bretton Woods system, and the inflation rate and interest rate uncertainty triggered by this and the oil shocks of 1973 and 1979, increased uncertainty in the global economy in the 1970s and early 1980s. Increased economic uncertainty has continued since then with higher degrees of competition resulting from globalisation and greater rates of technological change. All of this has led to an increased demand for the risk management services that are supplied by financial markets.

On the supply side, increases in computational and telecommunications capabilities and infrastructure and the development of better asset valuation methods have facilitated the creation and trading of many new types of financial instruments. As the pricing history for these new assets has grown longer and the markets deeper, trading has been further facilitated producing positive feedback for more trading in a greater variety of assets.

Before 1973, methods of valuation used in financial markets were in the lower left corner of the Banff taxonomy, based on qualitative models of uncertainty and asset level valuation of that uncertainty. Bond market valuation focussed on patterns of yields, equity markets on price-earnings ratios or dividend growth models, and derivatives markets were small and undeveloped with no systematic approach to valuation available to them at all.

The new methods of valuation that changed this were first introduced by Black and Scholes and Merton in 1970-3 in their analysis of the value of equity options.

Contrary to popular belief, the treatment of optionality was not the most important aspect of what they did. The key innovations, which are explained below, were:

1) the use of dynamic models of how uncertainty is resolved over time;
2) the creative use of the Law of One Price; and
3) the valuation of the effects of uncertainty at the source of the uncertainty rather than at the level of the asset cash-flow.

The important result of these innovations is that the value of a complex asset may be estimated by examining a portfolio of simpler assets that either have observed prices or for which prices are easier to estimate than is the value of the complex asset.

To see this, consider this typical Black-Scholes-Merton (BSM) valuation. The asset to be valued has a single payoff: $1000 three months from now, if the closing price of a given stock, say Alcan, on that day is between two levels, say $20.00 and $20.25. The payoff is zero otherwise.

The BSM approach uses the fact that transactions costs and barriers for trading in financial markets are relatively small. They make the approximation that these costs are zero. In this approximation, the "Law of One Price" holds: two assets with the same structure of payoffs must have the same price. If this were not the case, trading would occur to lock in risk-free profits until the prices move into consistency. With costless trading this process would occur essentially instantaneously.

The first step in the BSM valuation is to determine the uncertain process by which the price of a share of Alcan will evolve over the next three months. The approximation is made that:

1) the possible future price realisations are continuous (no jumps);
2) the uncertainty in the price movements over small periods of time is bounded enough, so that a portfolio of Alcan shares and risk-free bonds can be rebalanced continuously at will, by trading the shares for the bonds; and
3) the cash-flow realised from holding the stock and the uncertainty in the stock price movements in each small time period depends only on the stock price at the beginning of the period.

The next step is to find a current portfolio of Alcan shares and risk-free bonds to begin a trading strategy that will result, no matter what Alcan price realisation is actually realised over the next three months, in a portfolio at the end of that period of time, consisting of one share of Alcan and no bonds, if the price of Alcan is between $20.00 and $20.25, and no shares or bonds otherwise. The theory that BSM and their successors have developed shows that there is a unique initial portfolio and trading strategy that accomplishes this task. The portfolio created at any particular time by this strategy is called the "replicating portfolio" for the asset being valued.

Finally, the Law of One Price is used. Because the replicating portfolio process creates an asset that has the same payoff as the asset we are trying to value, the two assets have the same value. Because we can determine:

1) the number of Alcan share and risk-free bonds in the replicating portfolio at any time (including now) in any realisation for Alcan prices; and
2) the share and bond prices at that time in that realisation,

we can determine the value of the replicating portfolio, and from this the value of the asset being considered.

Notice that, in this analysis, as promised:

1) the valuation of uncertainty is done at source of the uncertainty, in this case, the Alcan share price;
2) the analysis is based on a dynamic quantitative model of that uncertainty;
3) the valuation uses the Law of One Price, which is a result of the approximation that there are no transactions costs or barriers in financial markets; and
4) optionality, or the management of flexibility, is not an issue, while the valuation of a complex asset in terms of a combination of simpler assets is.

Analysis in the bond and derivative markets uses BSM methods almost exclusively, and has done so since at least 1990. This means that it has shifted from the lower left corner of the Banff taxonomy into the upper right. (See Figure 3.)

Most equity analysts still use simple rules of thumb to estimate value, although there has been some shift in the equities markets up and to the right in the taxonomy as well.

Fig. 3 The evolution of financial market valuation

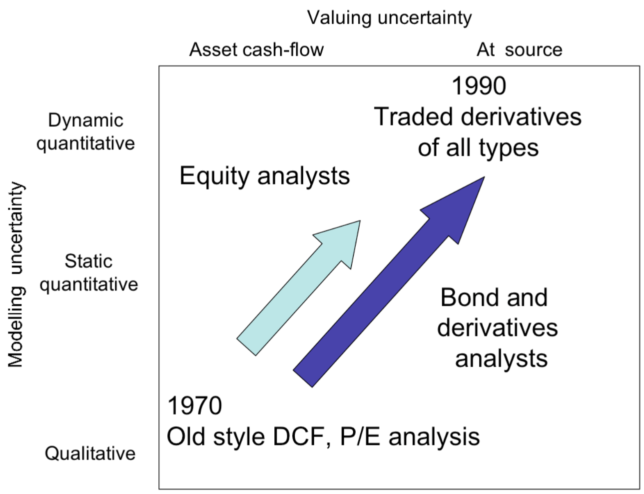

Where the petroleum industry might go with asset valuation

This financial market history gives mixed lessons about where the upstream petroleum industry might go in the future with real asset valuation.

Two facts must be kept in mind.

1) The relationship between the value of a single upstream petroleum asset and its sources of uncertainty is typically more complex than the similar relationship involving a bond or a derivative asset and its sources of uncertainty, and less complex than that involving the equity price for a corporation as a whole.

2) The valuation of assets in financial markets is used to support trading, where precision and speed are critical. The valuation of real assets in a commercial organisation is used to support decision-making where precision and speed are usually less important.

The overall lesson from derivative markets is that the industry would be well advised to follow those markets to the top right corner of the Banff taxonomy and begin to use dynamic quantitative models of uncertainty and to value uncertainty at its sources. There are two caveats to that lesson.

1) Analysis in equity markets has not yet made this leap completely, in part because of the greater analytical complexity involved.

2) The costs and benefits of such a move are different in different settings in which the valuation is done and used. The setting of an organisation trading in financial markets is very different from that of an organisation that explores for and produces oil or gas.

Taking these caveats into account, participants in the SPE workshop at Banff came, and others have come, to the conclusion that the industry should explore moves up and to the right in the taxonomy. If made, these moves could occur separately in either order or together. There are different organisations for which each of these three general patterns is likely to be best.

Moving up the taxonomy

As we have already noted, many exploration and production organisations currently use decision tree analysis, based on dynamic quantitative models of uncertainty, to help with decisions made during the initial exploration and appraisal phases of the asset life cycle. These are usually decisions about whether and how to expend resources to resolve, or partially resolve, some geological uncertainty. Typically, the only uncertainty that is modelled to be dynamic is the geological uncertainty about which information is being sought. Price uncertainty is usually treated using a 1-point forecast model or a static probability model.

A move up the taxonomy is to models that treat as dynamic all uncertain variables that actually have dynamic uncertainty, including prices and other commercial variables. It is also a move to models of dynamic uncertainty throughout the life cycle rather than only in the initial phases of the life of a typical petroleum asset.

More complete dynamic uncertainty models may be required for three purposes:

1) the use of decision trees, which we call "complete decision trees", to analyse future flexibility in response to all dynamic uncertain variables throughout the life cycle of the asset;
2) to support BSM replication, as noted above, in any shift to the use of valuation methods on the right side of the taxonomy; and
3) the modelling of how different uncertainties are resolved over time in different ways

This last use for dynamic models can have important effects on valuation, independent of the effects that stem from a decision tree analysis of future flexibility. While static models of uncertainty can be constructed to reflect different temporal patterns of uncertainty, it is generally easier to think of such matters within the context of a dynamic model. As the importance of all of this is revealed when we look at methods that are on the right side of the taxonomy, we shall defer further discussion of this until we look at the issues arising from a shift to market-based valuation (MBV) methods on the right of the taxonomy.

The shift up the taxonomy to more complete dynamic models of uncertainty has significant implications for the complexity of models to be considered. Moreover, because more detailed decision trees over longer times periods can rapidly become decision bushes or even decision forests, computational intensity increases dramatically with a shift to more complete decision tree analysis, and interpretation and presentation of the results can be more difficult. A good technical analogy would be a shift from using 2-D to 4-D seismic.

Moving to the right of the taxonomy

Figure 4 below shows a schematic for the mechanics of asset valuation on the left and right side of the taxonomy.

Fig. 4 Valuation Mechanics: Left and Right

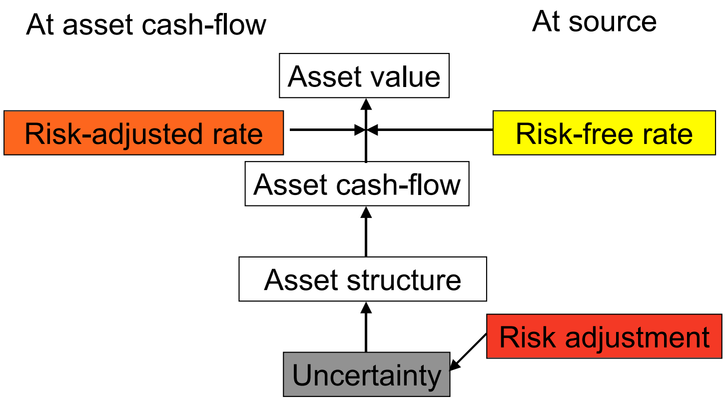

To value any asset, realisations of the relevant underlying uncertain variables are fed through a model of the asset structure to produce realisations of asset cash-flow. Some measure of the asset cash-flow at each time is adjusted or discounted for time and risk and summed over time to produce an estimate of the asset value.

On both the left and right, time is taken into account by discounting for time a measure of the cash-flow already adjusted or discounted for risk. This is done, in principle, by using the pricing of claims to risk-free cash. Frequently, an approximation to this pricing is used, based of a single constant discount rate, which is called the risk-free rate.

Uncertainty is taken into account on the left of the taxonomy by discounting an estimate of the expected cash-flows for risk with discount factors built typically using a single discount rate for risk. Therefore, on the left, time and uncertainty are taken into account together through a term structure of joint discount factors for time and risk, usually based on a single risk-adjusted discount rate, which is a sum of the risk-free rate to discount for time and a risk premium, which is another name for the discount rate for risk. The risk-adjusted discount rates used by an organisation are almost always constant across broad classes of assets. Usually, one rate is used to value all of the assets that an organisation considers.

Uncertainty is taken into account on the right of the taxonomy by "risk adjusting" the structure of the underlying realisations before they are fed into the asset structure model. As a result, the measure of asset cash-flow that is to be discounted for time is already adjusted for risk in a manner that depends on the asset structure and on the structure of the input realisations. BSM technology is used to determine how the realisation structure is adjusted for risk.

We shall examine how this works first in a simple set of situations and then in general.

Moving to the right: A simple case

Consider a producing oil field with known production and costs, which has the following asset cash-flow structure, linear in the uncertain oil price:

cash-flow amount at time t =
production_t * oil price_t - cost_t .

The 1-point forecast DCF estimate of value is:

sum over times t (forecast cash-flow amount_t
* corporate risk-adjusted discount factor_t)

= sum over times t
((production_t * forecast oil price_t - cost_t)
* corporate risk discount factor_t
* time discount factor_t).

The market-based value (MBV) estimate can be determined by using two corollaries of the Law of One Price.

First, the Principle of Value Additivity states that the value of a portfolio of assets is the sum of the value of the assets in the portfolio. The allows us to partition the cash-flows of an asset at will, and add the value of the claim to each resulting part to determine the original asset value.

The most convenient partition of the cash-flow of this simple oil field is into parts proportional to each oil price (i.e., the revenue at each time) and the parts that are without uncertainty (i.e., the cost at each time). The resulting value estimate is:

sum over times t
(production_t * value of the claim to oil price_t
- cost_t
* unit price of a claim to risk-free cash-flow at time t)

The unit price of a claim to risk-free cash-flow at time t is just the time discount factor for time t.

Second, Salahor (1998) shows, using a version of the Law of One Price, that the value of a cash-flow claim for a cash-flow at time t can be split into two factors:

the forward price for the cash-flow
* unit price of a claim to risk-free cash-flow at time t.

The forward price is defined as follows. A forward contract for a cash-flow is a contract between two parties that requires the parties to exchange, at the time of the cash-flow, the uncertain cash-flow amount for a fixed known cash-flow amount. This fixed known amount is called the forward price of that cash-flow.

Intuitively, the forward price takes the uncertainty out of the cash-flow, and must reflect the effect of uncertainty on the value of the claim to that cash-flow. Therefore, this forward price need only be discounted for time to produce the value of the cash-flow claim.

Moreover, given the forecast cash-flow amount, we can define the risk discount factor for the cash-flow as

the forward price for the cash-flow
= forecast cash-flow amount * cash-flow risk discount factor

When we use this to relate the value of claim to oil price at time t to the forward price for oil at time t and insert the result into the oil field valuation, we get:

sum over times t
((production_t * forecast oil price_t * oil risk discount factor_t
- cost_t)
* time discount factor_t).

Compare this to the DCF valuation:

sum over times t
((production_t * forecast oil price_t - cost_t)
* corporate risk discount factor_t
* time discount factor_t).

Notice that, as promised, risk adjustment (in this case direct risk discounting of the oil price forecasts) takes place at the source of the uncertainty, which lies in the individual oil prices, not at the level of the asset cash-flow as is the case in the DCF valuation.

Moreover, for oil forward markets that exist, oil forward prices can be observed and used directly in determining the MBV estimate of asset value. Recall that MBV methods use financial market data as much, and as directly, as possible to inform value estimation.

If we look at the taxonomy in Figure 1, we have shown the analogy between the 1-point forecast DCF method on the left and the risk discounting with forward prices method on the right at the bottom of the matrix. Moverover, we have shown the special circumstances where this specialised MBV method may be used: assets with linear models of cash-flow.

One last point. In the MBV method, the temporal pattern of uncertainty in the oil prices (which is the underlying pattern of uncertainties in this case) and the production and cost profiles (which give the asset structure in this case) both influence the implicit discounting for risk at the asset level. Moreover, this discounting will vary across assets in systematic ways depending on the production and unit cost profiles and the pattern of oil price uncertainty. Therefore, the use of a single discount rate for risk in the DCF approach to value estimation introduces systematic biases into asset valuation.

In other work (Salahor 1998), it has been shown that, as a result of ignoring the effects on risk discounting of both asset structure and the temporal pattern of underlying uncertainties, the DCF approach to asset valuation results in at least two important systematic biases.

1) If costs are less risky than revenues, DCF tends to undervalue both future costs and future revenues. Therefore it tends to undervalue investments in cost reducing technology, overvalue contracts to outsource at fixed prices activities like gathering, processing and transport, undervalue the negative effects of revenue-based, as opposed to profits- based, taxation, undervalue assets with lower than average costs, and overvalue assets with higher than average costs.
2) If there are short-term shocks and long-term equilibrating forces in the petroleum (or more generally the revenue) markets, DCF undervalues long-term revenues. Therefore it tends to undervalue long-lived assets, and overvalue production capacity in general and acceleration projects in particular.

Moving to the right: The general situation

The general situation is one where the cash-flows have non-linear dependencies on underlying uncertain variables. This is the case, for example, if there is uncertainty in both output prices and production quantities or input prices and input amounts. There may also be non-linearities if there is uncertainty in potential production in the face of finite processing or transportation capacity. Input and output prices may be related through market forces in ways that are non-linear. Finally tax systems are inherently non-linear if the effective tax rate depends on the output price level, as is the case in some royalty or windfall profit tax regimes, or on the amount of some definition of profit or loss, as is the case in resource rent tax and most production sharing agreement regimes.

The general situation is also one where there are multiple asset management decisions to be made in the future, some of which are sequential and some made in response to the arrival over time of new information about potential asset cash-flows and their valuation. These decisions may concern, for example, the timing of various activities related to the asset, or they may be about a choice of technology, or about the location, intensity or capacity of those activities. Because different choices may be made in different realisations of the future, management flexibility also results in non-linearities in the asset cash-flows.

For now, we shall consider only situations where the resolution of uncertainty is not influenced by any decision that we can make. We shall deal with the completely general situation later. If the resolution of uncertainty is not influenced by any decision, the resulting tree of realisations is the same for all possible management alternatives or policies that we might want to consider.

In this type of situation, the DCF value is:

max over policies p (expected "value"(p))

where expected "value"(p) is the expectation of the realisation "value" that would occur if the asset is managed according to policy p.

Those readers who are not familiar with this formulation of decision tree analysis (DTA) or the concept of a management policy should consult Appendix A where two examples of a oil field development lease are used to illustrate them.

We can expand the expected realisation "value" for each policy considered to show that the DCF value has the following form:

max over policies p
(sum over realisations r
(probability$_r$
* (sum over times t
(cash-flow$_{r,t}$(p)
* corporate risk-adjusted discount factor$_t$))))

where cash-flow$_{r,t}$(p) is the cash-flow from the asset that will occur at time t in realisation of the future r if the asset is managed according to policy p.

Finally, we can separate the risk-adjusted discount factors into separate discount factors for time and for risk:

max over policies p
(sum over realisations r
(probability$_r$
* (sum over times t
(cash-flow$_{r,t}$(p)
* corporate risk discount factor$_t$
* time discount factor$_t$)))).

We now turn to the market-based value (MBV) estimate in this type of situation. If we look at the taxonomy in Figure 1, we are up in the upper right hand corner doing real options analysis (ROA). Therefore, according to this definition, real options analysis (ROA) is just another term for complete decision tree analysis (CDTA) using market-based valuation (MBV) to value payoffs on the decision trees.

As in the simple situation of the last section, the MBV estimate can be determined by using corollaries of the Law of One Price.

First, we use the Principle of Value Additivity again. The key innovation is to use the states on the realisation tree to partition the cash-flow for valuation.

Let us suppose that, for each state n on the realisation tree, we can determine the value, V(n), of the claim to a pattern of cash-flow that is zero if that state is not realised and a unit amount at the time of the state if it is. Then, for each management alternative or policy, p, that we wish to consider, the value of the asset under that policy is:

sum over states n (cash-flow(n,p) * V(n)).

where cash-flow(n,p) is the cash-flow in state n under policy p. If we act to maximise asset value, then the asset value is:

max over policies p
(sum over states n (cash-flow(n,p) * V(n))).

V(n) is called the state price for the state n. How can these state prices be determined?

Let us consider a situation where the relevant realisations are realisations of oil prices, and the cash-flow to be valued depends only on the contemporaneous oil price (as was the case in our simple oil field in the last section). The relevant states will consist of all realisations where the oil price at the time of the given cash-flow is in a given small range within which the cash-flow may be considered to be essentially constant, for example the range from $20.00 per bbl to $20.01 per barrel.

This looks very much like the situation we considered to illustrate BSM valuation technology, only in that situation, instead of the oil price, we had the Alcan share price, and instead of a range of $20.00 per bbl to $20.01 per bbl, we had a range of $20.00 to $20.25.

If we can determine the process for the evolution over time of the value of the claim to the oil price we are considering, just as, in the Alcan example above, we had a process for the evolution of the Alcan share price, we can use BSM valuation techniques to determine the state price for this state, or any other state, of the given oil price. Moreover, we can determine the process for the value of a claim to this oil price if we know the process for the evolution of its forward price. Finally, we have market information about oil forward prices that we can use to estimate the process for their evolution.

This approach to valuing state prices can be generalised to states where more than one variable is in play, including, for example, the realisation of oil prices over time.

It is useful to deconstruct state prices into factors: the probability of the state occurring, the time discount factor and a risk adjustment, so that:

V(n) = probability$_n$ * risk adjustment$_n$
* time discount factor$_{t(n)}$

where t(n) is the time at which state n occurs.

It is also useful to recall (see Appendix A for details) that each state on the realisation tree has associated with it a given set of realisations such that:

probability$_n$ =
sum over the realisations r in state n (probability$_r$).

Moreover, each realisation r is associated at a given time t with a single state n(r,t).

Given these observations, we can return to our state price valuation of an asset, and reformulate it so as to compare it to the DCF value arising from a complete decision tree analysis.

Recall that the asset value is:

max over policies p
(sum over states n (cash-flow(n,p) * V(n)))

If we use our deconstruction of state prices, this becomes

max over policies p
(sum over states n
(cash-flow(n,p) * probability$_n$ * risk adjustment$_n$
* time discount factor$_{t(n)}$)).

If we now use our relationship between states and realisations, we obtain

max over policies p
(sum over times t
(sum over states n at time t
(sum over realisations r in state n
(probability$_r$ * cash-flow(n,p) * risk adjustment$_n$
* time discount factor$_t$))))

= max over policies p
(sum over realisations r and times t
(probability$_s$ * cash-flow(n(r,t),p)* risk adjustment$_{n(r,t)}$
* time discount factor$_t$)).

With appropriate definitions of the cash-flow and risk adjustment at any given time t in any given realisation r, this may be rewritten as

max over policies p
(sum over realisations r
(probability$_r$
* (sum over times t
(cash-flow$_{r,t}$(p)
* risk adjustment$_{r,t}$
* time discount factor$_t$)))).

Now compare this with the DCF value estimate:

max over policies p
(sum over realisations r
(probability$_r$
* (sum over times t
(cash-flow$_{r,t}$(p)
* corporate risk discount factor$_t$
* time discount factor$_t$)))).

The decision tree aspects of the valuation are handled in both approaches by a search, if more than one possible management policy is being considered, to find the policies that result in the most value. The methods for conducting this search are the same for the two approaches, DCF and MBV, to value estimation.

The probabilities, cash-flows and time discounting are also the same.

The only difference between the two approaches is that, for any given policy, the contribution, under that policy, to the asset value of the cash-flow at any given time in any given realisation is weighted in the DCF calculation on the left side of taxonomy by a corporate risk discount factor that depends on the time but not the realisation, while it is weighted in the MBV on the right by a risk adjustment that depends on which realisation is being considered.

Let us examine some properties of the risk adjustments before comparing DCF risk discounting with MBV risk adjustment.

First, under MBV, a cash-flow that has no uncertainty is not discounted for risk. We can show that this means that, for each time t:

1 = sum over the realisations r
(risk adjustment$_{r,t}$ * probability$_r$).

As a result, the product of the probabilities and the risk adjustments has all the properties of a probability distribution: They are a set of positive numbers, one for each future realisation, that sum to 1. We call this product the risk-adjusted probability distribution, and we call any statistic, such as an expectation, with respect to this distribution, a risk-adjusted statistic.

Second, we can show, by another application of the Law of One Price, that the risk-adjusted expectation of any variable realised at a given time is its forward price.

Fig. 5 True (black), risk adjusted (red) and risk discounted (blue) probabilities

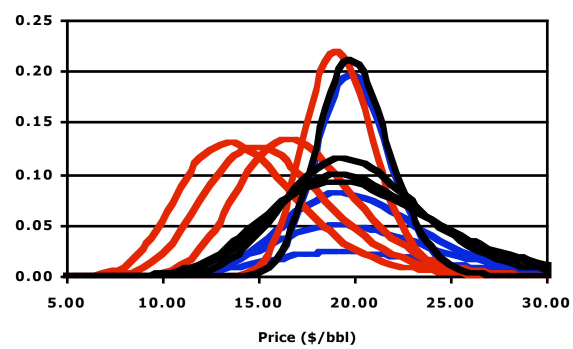

We can illustrate how DCF risk discounting differs from MBV risk adjustment by using an example where the underlying uncertainties are in oil prices. Figure 5 illustrates, in a particular commonly used type of oil price model, the actual and risk-adjusted marginal probability distributions of the oil price 1, 5, 10 and 20 years in the future (in black and red respectively), along with the DCF risk-discounted probabilities for these prices (in blue). The risk-adjusted probabilities are constructed from forward oil prices using BSM technology.

DCF methods account for uncertainty by weighting the contribution of all realisations, at a given time, to asset value with a weight that is less than and proportional to their actual probability and that decreases the further into the future we look. Thus all cash-flows are discounted in the same way no matter what their actual uncertainty is.

MBV methods account for uncertainty, in this example, by weighting the contribution to asset value of high oil price realisations with a weight that is less than their actual probability, and of low oil price realisation by a weight that is more. Thus, cash-flows that increase with the oil price will be discounted for risk. Moreover, the steeper the dependence, the more the discount. Cash-flows that decrease with increases in the oil price will be marked up for risk and the steeper the dependence the more the markup. Cash-flow that have no dependence on the price are not discounted for risk.

Notice also in this model of oil prices, there is long-term equilibrium in the sense that the amount of uncertainty in oil prices increases with the term of the price at a decreasing rate, as does the amount of risk-discounting in the forward prices. Cash-flows that depend on long-term prices inherit this lower per period risk discounting in the long term.

Does all this make sense?

It certainly makes sense to discount less for risk if there is less risk. In particular, if implemented, it would clear financial analysis of the often laid charge that it discounts the long-term too much.

What about the pattern of risk adjustments for realisations with different price levels?

We know that most investors are risk averse in the sense that they prefer to get a bit of extra cash more in states where they are otherwise poor than in states where they are otherwise rich. This is the fundamental reason for risk discounting in financial market prices. If oil forward prices are less than oil price expectations and thus discounted for risk, it is because high/low oil price states are more likely to be states where the economy is doing well/badly. In this situation, cash-flow to be received in a high oil price state that occurs at a given time with a given probability is worth less than cash-flow in a low oil price state at the same time with the same probability. As we can see in Figure 5, risk adjustment in MBV honours this. Risk discounting in DCF does not.

As a result, DCF methods are systematically biased to favour asset designs where value stems from cash-flows received in future states were the economy is doing well and against designs that increase cash-flow in states where the economy is in bad shape. Thus it systematically mishandles the analysis of many types of non-linear cash-flow patterns. For example, it typically suggests that development occur too early and abandonment too late. It also overvalues wind-fall profit taxes, or any other taxes or contractual terms where most of the value comes from cash-flows received in high price environments.

We have one final point to clean up, and that is the possible dependence of the realisation tree on the management policy chosen, if the policies can include actions to reveal information about the uncertainties represented by the tree.

First, the uncertainties that are resolved by such local action cannot be directly correlated with overall economic uncertainty. Therefore, we can split up the state structure on any realisation tree into:

1) a realisation tree representing the resolution of uncertainty in the overall economic state and anything correlated with it, which is independent of our control; and
2) within each overall economic state, a set of local asset states, the structure and tree linkages of which may depend on the policy under consideration.

Because of the lack of correlation, the probability of any joint local and overall economic realisation is the product of the unconditional probability of the local and overall economic realisations separately. As a result, DCF value estimate may be rewritten as:

max over policies p
(sum over local realisations l
(local probability$_l$
* (sum over overall economic realisations e
(overall economic probability$_e$
* (sum over times t
(realisation cash-flow$_{e,l,t}$(p)
* corporate risk discount factor$_t$
* time discount factor$_t$)))). (1a)

Because the local uncertainty is not correlated with their overall economic well-being, investors are indifferent about how that uncertainty is resolved. If this is the case, we can show, using the Law of One Price, that risk adjustment in state prices depends only on which overall economic realisation states are realised, and not on local states. Therefore, risk adjustment is also independent of the policy variables. As a result, the formula for calculating market-based value estimate is:

max over policies p
(sum over local realisations l
(local probability$_l$
* (sum over overall economic realisations e
(overall economic probability$_e$
* (sum over times t
(cash-flow$_{e,l,t}$(p)
* risk adjustment$_{e,t}$
* time discount factor$_t$)))). (1b)

This analysis of the comparison between complete DCF decision tree analysis and real options analysis (ROA) demonstrates that from an asset modelling and a computational point of view, a shift to the right represents no or little change respectively. On the other hand, there is a major conceptual change required by the different approach to valuing uncertainty, and an increase in the analytical effort required to compute the structure of risk adjustments instead of the risk premium in a single corporate risk-adjusted discount rate.

Once again, an analogy in the geological/engineering sphere is telling. Estimating the effect of uncertainty in the cash-flows of an asset on its value by assuming that uncertainty is set to the corporate average uncertainty is like setting the potential well productivity, in the analysis of a development plan for a particular asset, to the average productivity of all the production wells in which the corporation has an interest.

We have observed that most people think this treatment of a key technical variable like well productivity in an asset study would be a career-limiting move. Should this be so for a similar treatment of the effects of uncertainty on value?

The science and art of value estimation

We complete our consideration of the technical issues in real asset valuation by reviewing the science and art involved.

1) The general MBV formula for asset value (equation 1b) is strictly valid only in the limit where financial markets have no transactions costs or barriers. While low on average, these costs and barriers do exist and affect the validity of the calculations and usefulness of the results. For example, enterprise risk management and the detailed management of corporate financing are both all about managing the effects of different costs and barriers in financial markets. There is as yet no good quantitative science that lies behind these activities, and so, as we noted above, the link between an asset-level valuation and the effects of asset decisions on corporate value is not yet complete. Therefore, because our interest is in the effect of decisions on corporate value. there is an element of art to using the asset-level value estimates obtained through equation 1b to aid in decision-making.

2) The DCF formula for asset value (equation 1a) is strictly speaking justified only under the same conditions as the MBV formula, with the additional proviso that the use of a single average risk discounting structure suffices to treat the effects of uncertainty on value. For most organisations, the use of a single discounting structure is likely not to be sufficient. However, as we have noted, most organisations use DCF computations not to determine an estimate of asset value, but rather to determine a qualitative or quantitative distribution of realisation "values". Statistics of this distribution (including the expectation computed in a version of equation 1a without the search for a best policy) are used to provide ad hoc summary measures of potential regret or happiness at making a particular decision. There is no science behind this procedure.

2) There is art in the choice of underlying uncertain variables: fewer variables are better than more for computational and presentation reasons, while more variables are better than fewer if we wish to capture more of the reality of the world. There is some controversy at the current time over this issue, to which we refer below in our discussion of what we consider to be two "dead end" approaches to real asset valuation.

3) There is also science and art:

   1) in the specification of the model of the uncertainty in these variables;
   2) in the specification of the model of risk adjustment (or of corporate risk discount factors if DCF methods are used[3]); and
   3) in the use of data to parameterise these models.

   Most E&P organisations do not currently have the expertise to formulate and parameterise these models by themselves, but there are contractors available who can assist them, including some with experience in working with mining and upstream petroleum corporations.

3) There is art in the choice of the decision alternatives to be considered. Again, fewer is better than more for computational and presentation reasons, while more is better than fewer if we wish to capture more of the reality of the world. As we have noted, most organisations have not gone beyond some preliminary steps to examine, in any formal way, sets of decision alternatives that include sequential contingent responses to the resolution of uncertainty in the future.

4) There is art in cash-flow modelling, primarily in choosing the level of detail to be modelled.

5) Finally there is science and art in the manner in which the search and sums in equation 1a (or 1b) are accomplished. There are two aspects of this:

   1) the choice of overall method; and
   2) the choices made in implementing the method(s) chosen.

   There are several criteria for choosing an overall method of computation:

   1) generality (restrictions, if any, on types and scale of problems or models that can be considered);
   2) efficiency (accuracy and precision for different given amounts of computation applied to problems at different scales of complexity and "difficulty");
   3) verifiability (how sure are the accuracy and precision estimates);
   4) transparency (how easy it is for users to understand the method at the level they need to understand it); and
   5) progammability (what software tools are needed, how much they cost, and how easy it is to use them to construct and maintain applications).

   Again there is controversy at the current time about what general computational framework to use. We discuss one of these controversies below when we look at what we consider the "dead end" approaches to valuation just mentioned above.

Linking asset valuation and strategy development

Any asset selection decision should be required to pass a two-part test.

1) *The strategy test.* Unless an organisation has at least a temporary competitive advantage in owning and operating an asset, the value of that asset cannot be positive. Therefore the source of the competitive advantage must be identified before a positive value is to be believed.

2) *The valuation test.* Unless the potential positive value identified through the strategy test can be confirmed by an actual valuation, it should not be believed.

Moreover, because strategic decision-making is the art of taking action to create and then use competitive advantage, it involves imagining sequences of:

1) actions by the organisation to attempt to create the advantage;
2) what might happen partially as a result of those actions;

[3] Some managers may be familiar with exercises where the capital asset pricing model (CAPM) is used to determine the so-called corporate "beta" and from this the equity premium in a weighted average cost of capital (WACC) for a corporation

and
3) how the organisation should act again in response to use and extend any advantage created.

Therefore, the valuation methods that are most useful for supporting decision-making are those that:

1) support the analysis of sequential decision-making that seeks to create advantage and then exploit and renew it; and
2) focus on the sources of value.

Analysis of sequential decisions requires a method based on dynamic models of uncertainty structured to encourage consideration of a wide range of evolving alternative futures. A focus on sources of value requires methods that treat the effects of uncertainty at their source.

These are the methods respectively at the top and on the right-side of the taxonomy. Therefore, of all the valuation methods available to us, real options analysis (ROA) provides the best interaction with strategic analysis to support strategic decision-making.

Organisational considerations

Decision-making in a commercial organisation is a social process. Because of this, organisational considerations are important in determining the best valuation methods for any prticular organisation to use as part of its decision-making process. These considerations can be divided into:

1) those that influence the choice of methods if there is no process already in place; and
2) those that influence whether and how to change any existing process.

The structuring of any decision-making process in a complex organisation must recognise two sets of facts.

1) The information and expertise needed to define an appropriate set of alternatives to consider in any given situation, and to choose wisely from among them, is widely distributed among the people involved. Therefore:
    1) there must be a clear and efficient language to communicate this information among these people so that there is consistency on the treatment of different asset decision alternatives; and
    2) it is desirable for these people to have a shared understanding of the implications of the process, including its limitations, preferably based on extensive joint experience in using it.
2) The decision-making process and the resulting decisions have implications for the well-being of the participants in the process through several channels:
    1) compensation tied to the decisions;
    2) changes in job security and career prospects;
    3) changes in power within the organisation; and
    4) changes in the working environment more generally.

    Therefore:
    1) it is difficult to create conditions where the information is communicated is free of bias, whether conscious or unconscious; and,
    2) in particular, it is difficult to create conditions where the goals of the people involved are aligned with those of the organisation as a whole, and to mitigate any problems caused by a lack of goal congruence.

    Dealing with these issues takes:

1) a significant amount of senior management attention;
2) education and training;
3) specialised communication and analytical tools and protocols;
4) shared experiences and intuition;
5) organisational structures for:
    1) coordination;
    2) quality control;
    3) the maintenance of consistency across analyses; and
    4) the maintenance of system integrity .

These are costly and take time to build, develop or undertake. Therefore an incumbent process always has an advantage, and the benefits of adopting a new system may be outweighed by the costs of doing so.

Moreover changing the decision-making procedures in an organisation is a social process itself, akin to the diffusion of an innovation. Such processes are more likely to succeed (Rogers 2003):

1) if the perceived advantage of the innovation is large;
2) the more compatible with current practice the innovation is perceived to be;
3) the less complex the innovation is perceived to be;
4) the greater the frequency with which the innovation can be experimented with limited consequences; and
5) the more the results of using the innovation can be observed or described.

Because of all of these considerations, the detailed design of any process of change will strongly influence its net benefit.

1) Each step in the process should be big enough to be worth taking (size of perceived advantage). Moreover the process should be designed so that there are some important benefits from taking the first step.

    In a shift from DCF to MBV methods, these initial benefits come for many organisations from the greater understanding of tradeoffs between long- and short-term production and from a better understanding of operating and fiscal leverage.

    The initial benefits from a move up to more complete decision trees will depend on the flexibility issues faced by the organisation. For some organisations, this has come from a better understanding the tradeoff between economies of scale that might come from a megaproject development design and value of the flexibility to respond to market conditions that comes from phased development.

Others have benefited initially from a better understanding of their development options in brownfields situations. Others have benefited from a better understanding of the value of slack in the infrastructure they are constructing in a new area with possible follow-on developments.

2) On the other hand, each step should small enough to be done at as little cost as possible, so that:

   1) intuition is not lost, but transformed; and
   2) processes, language, power and culture can gradually adjust.

   Because the taxonomy separates the issues of modelling and valuing uncertainty, and clarifies and makes as small as possible the steps involved in moving up or to the right, its use can help different organisations with different needs to explore, in a more organised way and with greater insight, the different possible changes in their approach to valuation most suited to their circumstances. It can also help organisations to avoid paths that lead to dead ends, an issue to which we return below in the sections entitled "Some partial moves to the right" and "Two dead end approaches to real options analysis".

   We have designed a shift to MBV valuation so that an organisation need absorb only the change from risk discounting to risk adjustment. Moreover, there is a path where the initial step requires an understanding only of the role of forward prices in asset valuation as price forecasts already discounted for risk. We discuss the implications of this approach for perceived compatibility and complexity of the MBV innovation below in the section entitled "Market-based valuation without decision tree analysis and the issue of planning vs market price forecasts".

   The shift up the taxonomy to more complete decision trees can also be done in small steps as an increasingly broad range of dynamic sequential issues are formally treated as such.

3) Perceived advantage, compatibility and complexity depend on the context in which an innovation is being made. If the whole process is designed well, each step should:

   1) be in the "right" direction toward the possibly moving target; and
   2) lead naturally to taking the next desirable step toward that target.

   For example, the first use of dynamic price models may involve some relatively simple models of output price reversion. Typically, this leads to the demand of a more thorough treatment of price uncertainty, not only for outputs but also for inputs.

4) Many large E&P organsiations have a small corporate group that is responsible for the maintaining the analysis manual, and upholding the integrity of the analytical process that feeds into decision-making. Moreover, in many organisations, a few large assets contribute a large amount of value to the organisation. New valuation methods are trialable with limited consequences if tested initially within this small group of people or on this small group of assets (or both) by being run in parallel with current forms of analysis. If managed properly, this will help with building new intuition, and revising language and processes.

5) The benefits of a change in valuation method are difficult to observe directly. This is a major obstacle for the diffusion of this type of innovation. A change will have been beneficial if its leads to better decisions (i.e. decisions that increase the value of the assets of the organisation as much as possible) more often than would have been the case otherwise. This can be observed only if we can distinguish good/bad decisions from good/bad luck. This is hard to do especially if the outcomes of decision-making under uncertainty are not observed until many years, if not a few decades, have passed. Moreover, even if a good/bad decision can be distinguished from good/bad luck, it is difficult typically to pinpoint a specific role for differences in valuation methods in influencing decision-making. Attempts have been made to do this by comparing realised returns to asset value of organisations that use different types of methods. The results have been controversial at best.

   A more limited criterion for a beneficial change might be that a better valuation method is one that give better estimates of value. Tests of this for long-lived assets of at least moderate complexity have been few and far between and do not have much statistical power.

   This leaves casual empiricism and the inchoate sense that the changes being considered are actually more accurate models of the real world:

   1) that lead to important insights that once stated become obvious; and
   2) that evoke a better range of decision alternatives to consider.

Whether and how an organisation should change what it is doing will depend on its needs and current capabilities and culture.

1) Organisations that deal with a wide variety of situations, none of which is a norm, will benefit, more than organisations that face a very narrow set of decisions, from using valuation methods on the right hand side of the taxonomy that can tune valuations to this variety of circumstances.

2) Organisations that face situations where future flexibility is important will benefit from using methods at the top of the taxonomy more than organisations that face a constrained decision-making environment.

3) Organisations will be able to move up the taxonomy or to the right more easily if they have a senior management or key staff with experience respectively in using:
   1) probabilistic models or formal decision analysis for real asset evaluation; or
   2) financial markets for risk management.

Consistency and control

Many managers have expressed a concern about how a shift up and to the right in the taxonomy will affect the consistency of analysis throughout an organisation and the control of the decision-making process by senior management.

There are several issues that arise under these general headings.

1) There is the issue of capital (or, more generally, resource) rationing. If financial markets were free of transactions costs or barriers, capital rationing would never be necessary: assets with positive value could and should be created, and financing could be obtained on fair terms if necessary to do so. In the real world with costs and barriers in financial markets, this is not always the case. There are several models of why and how this is the case, none of which is complete or usable in a general approach to asset valuation. The temptation to over-quantify the analysis for dealing with this may lead to a masking of the real situation in which a organisation finds itself. At this point, it would probably be better to calculate asset values without taking these constraints into account, and use simulation information in the rollup procedure to inform the portfolio choice, and possibly overrule the asset value indicators. A next step might be some sort of formalised two-stage shadow pricing process to take into account resource constraints. This ought to be a topic of informed applied research.

   Capital rationing is also used as a control device in some organisations to control for "over-optimism" in asset teams and to force choices down into the organisation. This is costly. As far as we know, it is an open question as to what circumstances, if any, make this a desirable control tactic.

2) Many organisations delegate responsibility for asset management through the use of deterministic or quasi-deterministic targets on secondary variables, such as production or reserves addition, that are correlated with value creation. This is difficult to square with the use of asset value as a decision metric, and the contingent nature of future decision-making that is the focus particularly of valuation methods in the upper part of the taxonomy.

   The effectiveness of deterministic target-setting as a control device is, as far as we know, also still an open question. Indeed, we have not resolved the whole issue of providing incentives for good decision-making in situations with exogenous uncertainty outside the control of the managers involved.

3) It is a common concern in the industry that a shift to market-based valuation (MBV) on the right of the taxonomy will decrease the consistency of analysis across assets, by allowing different discounting structures for different assets. There is also a concern that this will give more "dials" to asset teams to use to "game" the system, increasing the agency problems that are inherent in the allocation process.

   Both these concerns are based on a false premise. An examination of equations 1a and 1b shows that the treatment of economy-wide uncertainty is the only difference between the DCF and MBV approaches to value estimation. In both approaches, this treatment - risk discounting in DCF and risk adjustment is MBV- is specified centrally, not by the asset team. Under both approaches, the asset teams provide the same information (i.e. the details of the decision alternatives to be considered, the probabilities of realisations of local asset variables such as those defining the reservoir, the quantities of inputs and outputs involved in the asset, and possibly some local price spreads) and do the same basic kind of analysis. Therefore, there is no loss of consistency or control due to difference in the locus of control over information provided or analytical techniques used.

   MBV actually provides a more consistent treatment of value because it values each commodity (e.g. oil, gas, steel, drilling mud, different types of labour) of a given grade, bought or sold in a given place and at a given time, in the same way no matter how it is to be used or how it has been produced. DCF does not do this. DCF implicitly values two barrels of Brent crude to be sold at Rotterdam on 01 Jan 2007 differently if one is extracted offshore in Indonesia and the other in Angola. MBV gives these two barrels the same value.

   Because the asset teams provide the same information and do the same kind of analysis under both approaches, they do not in fact have more "dials" to turn nor any greater opportunities to game the system under the MBV approach.

Market-based valuation without decision tree analysis and the issue of planning vs market price forecasts

An organisation that moves up the taxonomy to use complete decision tree analysis (CDTA) would be able to model formally many more types of business situations for the purposes of asset valuation. However. making this move would come with costs associated with developing the expertise to create the appropriate asset models, the inputs needed to support the analysis and the computational tools needed to perform the required calculations. Therefore, most organisations that wish to explore market-based valuation (MBV) should do so first in settings without the complications that arise with the use of CDTA. Indeed we have used this above by first showing how MBV works if there are no cash-flow non-linearities let alone future management flexibility to be analysed within the valuation. We showed that the only new inputs required in such situations are forward prices for the underlying variables and that the required computations parallel closely those needed for a simple 1-point forecast DCF analysis.

There is a more general parallel between DCF and MBV methods that allows for general non-linear cash-flow models, but pre-specifies the management policy to be used rather than searching for it. With the management policy fixed, the search in the first line of equations 1a and 1b becomes moot and we are reduced to the computations in 1a and 1b without this search:

sum over local realisations l
(local probability$_l$
* (sum over overall economic realisations e
(overall economic probability$_e$
* (sum over times t
(cash-flow$_{e,l,t}$
* corporate risk discount factor$_t$
* time discount factor$_t$))) (2a)

and

sum over local realisations l
(local probability$_l$
* (sum over overall economic realisations e
(overall economic probability$_e$
* (sum over times t
(cash-flow$_{e,l,t}$
* * risk adjustment$_{e,t}$
* time discount factor$_t$))) (2b)

This is a parallel in the middle of the taxonomy between what we call "DCF simulation" on the left[4] and "risk-adjusted state pricing" on the right.

If the same model of uncertainty is used, both left and right, a comparison of the value estimates made with these two methods puts the focus completely on the difference between single-rate DCF risk discounting and market-based risk adjustment in the context of assets with general cash-flows but little or no management flexibility.[5]

However, most companies that perform DCF simulation do not use the sort of market-based dynamic models of uncertainty in economic variables that would be used in an MBV approach. Rather they use a "planning price deck" approach to modelling this uncertainty. A set of price realisations are used in the analysis. In each realisation, the economic variable are prices and each price path is smooth, usually constant in real terms.

Those organisations that do valuation based on the qualitative models of uncertainty at the bottom of the taxonomy typically also use planning price decks for their price forecasts.

In both cases, the planning price deck(s) involved in the analysis may not be closely connected to the market consensus on price expectation and uncertainty than informs the formation of the term structure of forward prices and the risk-adjusted probability distribution of prices.

For example, since 2004, there has been a significant increase in the oil price forecasts implicit in financial market price data. The planning oil price decks used by most organisations have increased but are much lower that any possible implicit market forecast.[6]

Therefore, an organisation that uses planning price forecasts in its DCF valuations will have to deal with two changes when contemplating a shift to the use of MBV methods:

1) the change from corporate risk discounting to market risk adjustment that inherent in the switch; and
2) the switch from planning price forecasts to market price models that goes along with it.

For some organisations, this disconnect between their planning prices and the market environment in which they find themselves is an extra incentive to consider the use of MBV. For others, the fact that they would have to change the type of price forecast that they use presents another reason not to consider a change from their current analytical practice.

Some partial moves to the right

There are some valuation methods in the middle of the taxonomy that are partial shifts from valuing uncertainty at the asset level to valuing it at the source of the uncertainty. In this section we shall discuss the two at the bottom of the taxonomy. We discuss the two that are higher up in the next section.

The first method, the "special asset discount rates" approach, uses DCF methods but applies different discount rates to assets of different types, presumably to represent their different average types and levels of risk. Several methods of asset classification have been used by different organisations to implement this type of approach.

Some organisations have used different discount rates for assets in different countries. A so-called "country risk premium" is meant to reflect different levels of political risk in different countries. Typically, the cash-flows to be discounted are modelled ignoring political risk, so that the country risk premium is a mixture of a true return premium for this political risk and a measure of the expected proportional loss of asset cash-flow resulting from political risk.

Some organisations determine "divisional costs of capital" so that the discount rate depends on which division of the organisation "owns" the asset. For example, there may be a

4 This is also known in the industry as monte carlo analysis or expected net present value.

5 Some flexibility may be considered if a few pre-specified policies are considered and compared with a separate calculation being done using equation 2a and b for each policy considered.

6 Some organisations use lower planning prices because they believe that the current oil market prices are a price bubble and they are speculating against the market. This is not usually stated clearly, and it is also not clear why these organisiations do not speculate directly in the financial markets, where it is typically easier to unwind speculative positions than those positions taken with the structure of their real assets.

Others say that they do this to be "conservative" and to ensure that their assets work in the low price environments that might occur. They are not trying to estimate asset value when doing the analysis but are doing a form of low price asset simulation.

different discount rate for assets in the onshore oil division and those in the natural gas and power division.

Finally, some organisations use different discount rates for assets in which different activities are taking placeFor example, assets that are acceleration projects will be discounted at a different rate from assets that are projects to be undertaken to meet environmental regulations.

This approach of using different discount rates for different classes of assets might be useful if it were a stepping stone for a move more completely to the right. It does get at the idea that different types and amounts of uncertainty ought to have different structures of risk discounting. However, it still has many of the problems with DCF, by not dealing directly with, for example, the bias against long-term assets, unless the term of the asset is an explicit part of the classification scheme. Moreover, it would be difficult to implement and the inevitable political fights over the classification scheme and the discount rates to be used means that it would be likely to ossify the organisation into a state where it uses this scheme that results with little possibility for any further changes. For this reason its bubble in the taxonomy is coloured rose, indicating we do not consider it to be a desirable change.

The second approach, the "adjusted present value" approach, separates asset cash-flows into two components: a component that stems from corporate obligations (e.g., contractual pipeline charges) that look like corporate debt payments, and the rest. The first component of cash-flow is modelled ignoring the possibility of default and discounted at the yield on the equivalent debt securities of the corporation. The rest is discounted at the standard risk-adjusted corporate discount rate. The two values are added, using value additivity as the justification, to estimate the asset value.

A shift to the use of adjusted present value would be useful if it were a stepping stone for a move more completely to the right. It does get at the notions that:

1) asset cash-flow can be split up into bits with different patterns of uncertain to be valued separately; and
2) value additivity can be used to bring the bits back together.

For some organisations, this might be a useful first step, if expectations for change are managed appropriately. For this reason its bubble in the taxonomy is coloured white, indicating we do consider it a possibly useful first step in a process of change.

Two dead end approaches to real options analysis

The term "real options analysis" has been defined by many people to mean one of two related approaches to valuation,[7] either:

1) the "financial options analogy" approach to real options; or
2) the so-called MAD ("marketed asset disclaimer") approach to real options (Copeland and Antikarov 2003).

Both of these methods were created as approximations to full real options analysis in an attempt to satisfy the second of the design criteria for a change in valuation method (i.e., to be as close as possible to what the organisation involved has been doing). Both methods do this by treating real options analysis as an "add-on" to the type of DCF simulation that many organisations are currently using. The general idea is:

1) to use DCF simulation to analyse a version of the asset without "flexibility"; and
2) if there is flexibility in the future management of the asset, to use the results of the DCF analysis in a Black-Scholes-Merton (BSM) estimation of the value of the asset with the flexibility.

First, we shall outline briefly why these approaches are dead ends (and thus coloured rose in the taxonomy). Then we shall expand on these observations by giving more details about how the analysis is done in each approach, and about the problems encountered.

Unfortunately, although these approaches do incorporate a partial shift toward using Black-Scholes-Merton methods of valuation, it is in the context of an underlying DCF analysis. As a result, organisations that adopt these methods will not gain the insights that are available from an MBV approach, but they do pay some of the costs of shifting in that direction. Moreover, they are not led naturally to contemplate more steps toward a complete MBV framework.

Finally, their analysis of flexibility is constrained dramatically by treating it as an add-on to a static DCF valuation, as opposed to making it an integral part of the basic analytical framework from the start. As a result, organisations that adopt these methods are not lead naturally to deal with general forms of flexibility beyond the simplistic situations that these "static DCF with BSM add-on" methods can address.

The financial-options-analogy approach implements the general "static DCF with BSM add-on" strategy by searching for a financial option that "looks like" the real asset being considered. The financial option chosen must have a value formula, constructed using BSM methods. This approach then demands that the analytical team find the real asset equivalents to the input parameters for the financial option value formula, and use them in this formula to estimate the real asset value.

The canonical example of this approach is the analysis of a non-renewable natural resource development lease (i.e. the right to initiate, during a given period of time, development of the resource for production). The financial option analogy is a call option on an equity (i.e. the right to buy the equity during a given period of time, called the exercise period, at a fixed price, called the exercise price). In this analogy,

1) the claim to the production cash-flows of the developed resource takes the place of the equity;
2) the value of the claim to the development costs takes the place of the exercise price; and

[7] Some people also refer to complete DCF decision tree analysis as real options analysis, ignoring the whole right hand MBV side of the taxonomy

3) the length of the lease takes the place of the length of the exercise period.

The original Black-Scholes formula gives the value of an equity call option as required, but only under some very restrictive circumstances. Among other things, it requires that:

1) there be no possibility of dividends issued from the equity during the life of the option;
2) the change over any small period of time in the logarithm of the value of equity be normally distributed with a variance proportional to the duration of the time period and known with certainty; and
3) the rate of return on risk-free assets at any time in the exercise period be known with certainty.

The input parameters for the Black-Scholes formula are:

1) the exercise price;
2) the length of the exercise period;
3) the value of a risk-free bond with a single unit payoff at the end of the exercise period;
4) the current value of the equity; and
5) the variance of changes over the life of the option in the logarithm of the value of the equity.

For the development lease, the analogous input parameters are:

1) the value of the development costs, presumed to be risk-free and constant with respect to the time at which the development begins;
2) the length of the lease;
3) the value of a risk-free bond with a single unit payoff at the end of the exercise period;
4) the expected realisation "value" of the claim to the production cash-flows, found using DCF simulation; and
5) the variance for changes over the life of the lease in the logarithm of the value of the claim to the production cash-flows, found somehow possibly using DCF simulation.

Unfortunately, very few real assets are strictly analogous to a financial option with a simple value formula. In our example of a development lease:

1) there are the likely to be the equivalent of dividends that accrue from having developed the resource and putting it into production;
2) changes in the logarithm of the value of the claim to the production cash-flows are not likely to have a variance rate that is known with certainty;
3) there are likely to be choices that must be made, not only about the timing of the development, but also about its design, as well as about any appraisal activities or works requirements during the life of the lease; and
4) a DCF analysis of the value of the claim to the production cash-flows would suffer from all the failings of DCF analysis in general.

Typically, the features of the real situation that cannot be fit into any simple analogy are material to the valuation and the decisions that are being analysed. Moreover, even in the few situations where a sufficiently close financial option analogy can be found, the valuation formulae that do exist are almost always obscure "black boxes".

The MAD approach to real options is a little bit more general in that it does not require the existence of a financial option analogy with a valuation formula.

The real asset being considered is stripped of managerial flexibility by pre-specifying a possibly suboptimal management policy. We shall call this the "underlying asset".

Its value and some measure of the volatility of changes in that value are determined using DCF simulation methods. From these parameters, a realisation tree is set up for realisations of this value, and state prices for states on the tree are determined, assuming the value evolves in a specific way (usually as a geometric brownian motion).

A decision tree analysis of the asset is then done using state prices on the realisation tree of underlying asset values.

Several problems arise from using an underlying asset value to describe the state of the asset with flexibility, rather than using directly the major uncertain determinants of the asset cash-flows. These are problems for the financial-option-analogy approach as well.

1) The value of an asset is determined in reality by the structure of its cash-flows, not, as in this approach, the other way around. The payoffs on the decision trees to be analysed are typically cash-flows for the asset being considered. It is only in very special circumstances that these payoffs can be accurately determined as a function of an underlying asset value and that value alone. This greatly restricts the types of assets and flexibility that might be accurately considered using this method.

2) The risk-adjusted process, by which the value of an underlying asset changes over time, is usually much more complex, especially over the long time scales that are relevant for most real assets, than the process by which the determinants of the asset cash-flows evolve. The underlying asset process typically cannot be expressed well using the underlying asset value itself as the only input. Even if it could, the appropriate specification and parameterisation of the process would not be well understood from a simple static DCF simulation of that asset, because of the limitations of DCF analysis discussed above. Once again this greatly restricts the types of assets that might be accurately considered using this method.

3) An underlying asset value must be calculated rather than observed directly. As such, it does not provide as good a signal for decision-making as do other variables that can be observed directly. This makes the method hard to use in practice.

A shift to using either of these approaches is not likely to be a good first step in improving the asset selection, design and management process. The one aspect of the DCF approach currently used that should be changed - constant discounting for risk - is either not changed, or is only partially changed in an obscure and almost surely inaccurate way. Moreover, these methods do not address the issue of valuation of flexibility in a way that:

1) is a natural extension of what is currently being done; or
2) is naturally extendable beyond their inherent limitations.

A direct shift either toward MBV, toward complete decision tree analysis (CDTA) or toward both at the same time is more likely to lead to successful long-term continuous improvements in valuation as part of the decision-making process.

Future developments in valuation methods

Work is continuing on several fronts to develop MBV and CDTA methods further.

1) The analysis of future flexibility in the face of multiple sources of uncertainty is a serious issue for CDTA.

   Increasing the number of underlying uncertainties and the complexity of the decision environment dramatically increases computational intensity. It also makes results more difficult to comprehend and communicate. On the other hand, uncertainty and flexibility in the real world are both complex, and incorporating that complexity explicitly can increase accuracy and believability.

   There is an art to choosing the set of underlying uncertainties and decisions to be analysed that makes the best tradeoff among these considerations, as well as connecting to available data.

   This art is being developed as more complete decision tree analyses are done.

   There is also work being done to develop more efficient computational methods, that combine random sampling (e.g. monte carlo) methods with specialised forms of search algorithms over different types of policy sets so that computation time and precision are less of a constraint on the choice of underlying uncertainties to be modelled.

   Work has yet to begin on making results from models with many underlying uncertainties easier to understand and communicate.

2) Work is being done to develop better specifications and methods of parameterisation of models of underlying uncertainties and their risk adjustments.

   The modelling of oil and gas prices is moderately well developed from work that has been done to support financial market trading, but research is continuing particularly on:

   1) the parameterisation of long-term price models; and
   2) the relationship between oil and gas and light and heavy oil.

   There is great deal of work to do to determine best practice in the modelling of input prices, including:

   1) how detailed the cost models should be;
   2) what costs should be considered together as having the same effective price uncertainty;
   3) what the structure of various types of input prices is; and
   4) how best to determine the parameters of the models of input price uncertainty.

   Finally, we do not yet have a good feel for potential best practice in the art of pruning down the massive amounts of geological and technical data, associated with any real potential or actual petroleum production asset, to a state where it can be used in manageable decision trees over the whole of the asset life cycle. Work is just beginning on this topic.

3) Tools to implement the modelling and computations needed by CDTA are not yet available in commercial form. There has been some preliminary work on this, but more may not be done before there is a greater indication of the likely demand.

4) Better training materials are needed for both MBV and CDTA and are gradually being developed. These would benefit from more examples of actual use of the methods. Good insights into how best to handle organisational issues will also come only with use.

Summary

In this paper we have introduced a taxonomy of valuation methods, organised by how uncertainty is modelled and valued, and showed how to use this taxonomy to describe the past evolution of valuation in the upstream. petroleum industry and in financial markets.

Both the industry and financial markets have evolved to use more quantitative models of uncertainty.

However, financial markets have for some time typically supported asset valuation with complete dynamic models of all uncertainties over the whole of the relevant asset life cycle, while most organisations in the industry still treat uncertainty as being static, except possibly for some geological variables during the exploration and appraisal phases of the asset life cycle. As a result, most organisations in the industry have yet to use what we call "complete decision tree analysis" (CDTA) to examine management flexibility in response to all uncertainty throughout the asset life cycle.

Moreover, financial markets determine the effects of uncertainty on value by using what we call "market-based valuation" (MBV). In this approach to valuation, financial market data is used as directly as possible to adjust the contribution to asset value over time of different possible realisations of the future in a manner that reflects the temporal and correlative structure of the uncertainty. Meanwhile, most of the upstream petroleum industry is still using "one-size-fits-all" constant rate risk discounting to account for the effects of uncertainty on value

We then described what is involved technically if an organisation were to shift to using compete decision trees or to market-based valuation, outlining some of the benefits and the technical process costs of each type of move.

We have also discussed how the interaction of strategic analysis and asset valuation, as parts of the overall asset selection, design and management process, favour the use of ROA methods for asset valuation because of their focus on sequential decision-making and the analysis of the effects of the sources of uncertainty on value.

We then described the organisational considerations in the design:

1) of a valuation method for use in the asset management process (to complement the technical considerations and the issue of compatibility with, and support for, strategic analysis raised earlier);
2) of a process to change valuation methods, if change is desirable.

We focussed on the issues of information transmission and incentive compatibility, and came up with some criteria for managing the process of change. We also looked at issues of control and consistency, including capital rationing and subsidiary goal setting.

Next, we suggested how market-based valuation (MBV) might be explored better in setting without an imbedded search for a best management policy. We also discussed briefly the problems and opportunities arising from a disconnect, as is currently the case for oil, between the planning price decks that many organisations use in DCF simulation and price forecasts implicit in market forward prices

Within this overall context, we showed why a change to use different discount rates for different classes of assets and why two well-known proposals for change in the direction of dealing with "real options" are dead ends that should not be considered. We also discussed why some organisations might want look at adjusted present value as a stepping stone to the use of full market-base valuation (MBV) methods.

Finally, we outlined some potential future developments in this field.

Appendix A: Complete decision tree analysis

In this appendix, we set up the analysis of two models of oil field development leases to illustrate what we mean by a realisation tree, and by the set of management policies to consider in an asset valuation.

To keep things simple, we presume that the only uncertainty that we need to consider is oil price uncertainty.

In the first example, oil price uncertainty may be represented approximately by a model where, in any given year, the oil market in the following year may in be in one of two states, where the oil price is either more or less than previously expected. Our time horizon is two years.

Fig. 6 The oil market realisation tree

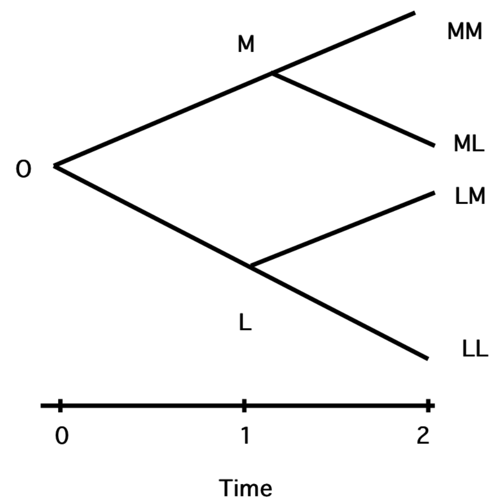

The relevant tree is shown in Figure 6. There is the original oil market state O at time 0. There are two possible oil market states at time 1: M, if the realised oil price at time 1 is more than price at time 1 was expected in state O to be, and L, if it is less. If the state realised at time 1 is M, there are two possible states at time 2: MM, if the oil price at time 2 is more than was expected in state M, and ML, if it is less. Similarly, if the state realised at time 1 is L, there are two possible states at time 2: LM, if the oil price at time 2 is more than expected in state M, and LL if it is less. Therefore, viewed from state O at time 0, there are four possible states at time 2: MM, ML, LM and LL.

In this situation there are four possible realisations of the future which may be labelled by the end states on the tree: MM, ML, LM and LL. The state at time 0 associated with each of the realisations is O, which is the only state at time 0. The state at time 2 associated with each realisation is the state used to label it: e.g, state MM is associated with the realisation MM. The state at time 1 associated with the realisations MM and ML is M, and with LM and LL is L.

What policies need we consider for managing this lease?

In each state on the tree we have two possible decision alternatives, if the field had not already been developed: develop now (D) or wait (W). Note that waiting in a state at time 2 is the same as walking away from the lease. If, in any state, the field has already been developed, no choice is possible (N).

In each state on the tree we have two possible decision alternatives, if the field had not already been developed: develop now (D) or wait (W). Note that waiting in a state at time 2 is the same as walking away from the lease. If, in any state, the field has already been developed, no choice is possible (N).

A possible management policy would be a pattern of D's and W's for each state on the tree where development is possible, and N's on all other states. Development is possible in a state, if the policy does not specify development in any predecessor state.

To make this more concrete, suppose that, for each state, we have calculated the value in that state of developing the field beginning at the time of that state. Suppose these values are all positive, except for those in states L and LL. Therefore we need not consider any policy that contemplates development in either of these two states. This leaves three possible policies to consider:

1) develop at time 0 (D in state O and N in all other states)
2) develop at time 2 if state MM, ML or LM is realised (D in states MM, ML and LM, and W in all other states)
3) develop at time 1 if state M is realised or at time 2 if state LM is realised (D in states M and LM, W in states O, L and LL, and N in states MM and ML).

For a second example, consider a development lease of finite duration with a somewhat more realistic oil market model. As in the example just discussed, the decision to develop or not will be taken annually during the life of the lease, but the price realised at any given time may take on any non-negative magnitude. Suppose that this price is sufficient to specify the state at that time of the oil market and its future, and that the expectation in any state of all then future prices increases with the price realised in that state. Suppose that the value of beginning to develop the field now is negative, and that the value at any given time in the future of beginning development at that time is an increasing function of the price at that time, and is positive for high enough prices. In this situation, each policy in the set of possible management policies may be represented by an annual future term structure of critical oil prices. The policy rule for such a policy would be to begin to develop the field at the first time where the realised oil price is greater than the relevant critical price. The optimal policy would set the critical price at any time to the price for which the value of beginning development at that time would be equal to the value of waiting.

References

Copeland, T. and V. Antikarov (2003) *Real Options: A Practitioner's Guide* Thomson-Texere

Laughton, D. (ed.) (1998) "The Potential for Use of Modern Asset Pricing methods for Upstream Petroleum Project Evaluation" *The Energy Journal* **19(1)** pp. 154

Laughton, D., J. Sagi and M. Samis (2000) "Modern Asset Pricing and Project Evaluation in the Energy Industry" *Western Centre for Economic Research Information Bulletins* **56** pp. 76

Rogers, E.M. (2003) *Diffusion of Innovations, 5 Ed.*

Salahor, G. (1998) "Implications of Output Price Risk and Operating Leverage for the Evaluation of Petroleum Development Projects" *The Energy Journal* **19(1)** 13-46