SPE 78329

A Stepwise Example of Real Options Analysis of a Production Enhancement Project

William Bailey, SPE, Schlumberger-Doll Research, Johnathan Mun, Decisioneering Inc. and Benoît Couët, SPE, Schlumberger-Doll Research

This paper was prepared for presentation at the SPE 13th European Petroleum Conference held in Aberdeen, Scotland, U.K., 29–31 October 2002.

Abstract

Practical solution to a Real Options analysis of a somewhat generic production enhancement problem is presented. First the problem is properly framed, suitable data is gathered, the analysis is performed and finally the results are communicated. Ancillary issues concerning the most appropriate solution methodology and the quantification of a representative volatility component are also touched upon. The narrative is constructed in such a way as to, hopefully, enable usage beyond the specifics of the problem discussed. A form of visual representation of results is proposed that may facilitate communication, improve understanding and help gain acceptance of the benefits of such an analysis from management and decision-makers' points of view.

Introduction

It is understandable that Real Options can be perceived as not being the clearest, most intuitive of financial analysis techniques. The apparently abstruse concepts involved, new jargon and our limited experience have conspired to limit their application somewhat. Although several books and articles have been published that remove the veneer of complexity and demonstrate applications to various problems[1-7], they are not always readily translated to specific petroleum-related problems. While they indeed provide insight there is still a need to provide a step-by-step framework for a Real Options (RO) analysis for just such issues. Furthermore, there remains a genuine difficulty in effectively and concisely communicating the insights derived from such an analysis, in a lucid and convincing manner, to decision makers. This has been an impediment to RO's gaining universal management acceptance.

We present a step-by-step guide that the reader may be able to apply to other problems. We are proposing a new way to convey the findings of such an analysis (in graphical form) that may make communicating the results a little easier (although not foolproof). It should be noted, however, that the framework *cannot* be universal. Certain classes of problem may just not be amenable or may require specialized or custom-made solutions. Rough guidelines to recognize such cases are discussed. Apart from two introductory passages (overviews of Real Options and a general problem description), the paper organization more or less follows a 5-step process:

Step 1: Frame the Problem
Step 2: Quantify Internal Uncertainties
Step 3: Quantify External Uncertainties
Step 4: Do the Numbers
Step 5: Communicate Results & Decide

A brief glossary is provided (Appendix A) which maps the oft-cited financial engineering terminology to their petroleum engineering equivalents.

Options & Real Options: What They Are

The classic definition of an option is "the right, but not the obligation, to execute a particular strategy or initiative." While Luerhman[5] provides a succinct and lucid introduction to Real Options that may be difficult to better, it is still worth recapping: an option value is simply the amount you should be willing to pay for the right to acquire a stock, a stream of cash flows or commodity at some point in the future – but whose ultimate value is uncertain – for a predetermined amount today. Similarly, RO's represent the value one would be willing to pay now for the right to an uncertain stream of cash flows in the future. As such Real Options only really apply when one is faced with uncertainty.

For example, today we have a daily cash flow of $1,000 and wish to have the same one year hence. However, future cash flows are uncertain, subject to a known volatility. Fortunately we have an option to purchase information that will provide much better estimates of future cash flow (furnishing greater certainty). Do we then accept fate, see what happens and possibly get less than $1,000? Or do we purchase the information? If so, how much should we pay? This is effectively what Black-Scholes[8] and Merton[9] (BS-M) provided in their classic, Nobel prize-winning formulation, simplified as:

$$V = SN(d_1) - Xe^{-r(T-t)}N(d_2) \tag{1}$$

where V is the option value, S is the underlying asset value and X is the strike price and

$$d_1 = \frac{\ln(S/X) + (r + 0.5\sigma^2)(T-t)}{\sigma\sqrt{T-t}} \tag{2}$$

and

$$d_2 = d_1 - \sigma\sqrt{T-t}\ . \tag{3}$$

$N(.)$ is the cumulative probability distribution function for a standardized Normal variable[*]. This equation is, however, loaded with assumptions[10] and must be applied with care (it is not applicable in all cases). Conceptually it states that the value of the option ($\$V$) for the right to buy (at price $\$X$) at time T a future benefit is the value of the probable benefit (first term on right hand side) minus its cost if exercised (the last term on the right hand side of Eq.1).

If volatility is 25% p.a., the risk free rate, r, is 1.75%, the time period, $(T–t)$, is one year and both S and X = $1,000, the BS-M values this option at $107.71. Simply put RO's are extensions of conventional NPV such that:

NPV = Benefits – Cost (4)
Option = Benefits$\cdot P(x_1)$ – Cost$\cdot P(x_2)$ (5)

$P(x_2)$ represents the probability that our option will be positive ("in-the-money", i.e. the benefit is greater than the cost). Assuming that the option itself has been exercised (costs sunk), $P(x_1)$ represents the probability that subsequent benefits are also positive. The problem is knowing the value of $P(.)$ which is an indirect objective of this article.

Real Options are then:

- Useful when we are faced with future uncertainties,
- Useful when we have the flexibility to respond in some way to new information, conditions or event outcome, where these uncertainties bring with them valuable information and management can execute contingencies (or flexibilities) when required,
- Useful when a classical Discounted Cash Flow (DCF) analysis indicates marginal project viability,
- Essentially advanced DCF calculations that consider uncertainty, flexibility and new information
- Extensions of, but not fully beholden to, existing financial option theory (which form the basic conceptual foundations),
- Valuable as they confer the right, but not the obligation, to capture future benefits depending on prevailing conditions (e.g., values an additional exploration well to provide new information),
- Contingent decisions. Depending on the learning obtained from some future event, RO's help us to decide whether to instigate, defer or curtail some action,
- Everywhere. From the decision to buy a new house to developing a new field, RO's are prevalent. In many ways they are inescapable in an uncertain environment.

Real Options: What They Are Not

Real Options are, however:

- Not the answer to all our valuation needs. They still require imprecise inputs, such as volatility and Weighted Average Cost of Capital (WACC) that are subject to uncertainty. (Some issues concerning WACC are discussed briefly in Appendix B),
- Not suitable when project is highly lucrative or highly uneconomic (known as being *deep in-* or *deep out-of-the-money* respectively),
- Not just the blind application of a BS-M type model taken from classical financial options theory[a]. Such models are bound by a host of assumptions and caveats (Appendix C) and are defined by unambiguous contracts. RO's have less (or no) contractual structure *per se* but instead are loaded with opportunities and impactful decisions,
- Not applicable when there is no uncertainty or doubt as to future cash flows (an option model would then simply emulate a deterministic DCF),
- Sometimes *non-trivial* to perform and frame.

Other contemporary issues and concerns are discussed by Luerhman[11]. Contrary to conventional wisdom RO's do not always guarantee higher value on a project. Laughton[12] has demonstrated how RO's can actually reduce value – albeit an exception, not the norm. Although an option's value is always zero or positive, the cost of instigating the option may exceed its value, therefore the entire option strategy may take on a negative economic position – even if the option itself is positive.

[*] Mean of zero and a variance of one. `NORMSDIST()` is the related *Excel* function.

[a] Readers should consult Wilmott[10] for a detailed financial guide, Armstrong & Galli[2], Leslie & Michaels[4] and Luerhman[5] for examples of application to Real Options.

General Problem Outline

We shall apply a Real Options analysis to a problem involving production enhancement of a single well. The analysis narrative is generalized so as not to be unduly specific.

Description: A faulted (but not compartmentalized), mature, offshore reservoir comprises several interbedded reservoir units. Interbedding shales are not always continuous. The throws and connectivity of the smaller faults are not certain. Most of the 30 production wells are completed over several pay-zones. Pressure support is provided by water injection and a modest aquifer. Field water cut averages 65% but varies between 20% and 85% for individual wells. The hydrocarbons are sweet (no solids content). A detailed reservoir simulation model is not available but material balance and NODAL analysis models are. Significantly, there is a paucity of recent production log data – the last major logging run was conducted a number of years previously. Most producers started with near-zero water cuts but this has steadily increased over time.

In order to improve production (and/or reduce water cut) a high-level field-wide diagnostic study was initiated. It identified and prioritized candidate wells for further detailed analysis. Lack of relevant PLT data cast doubt on the individual contributions of each producing layer (and associated water cut). Considering the single primary candidate well alone, a DCF indicated that economic benefit was sensitive to the effectiveness of the proposed intervention. The lack of PLT measurements meant that an incorrect diagnosis was certainly possible and the economic impact of such could be significant. A small-scale sensitivity test on the DCF (using perceived low- and high-side values along with a best guess) indicated that cash flows straddled the break-even point. The decision to go-ahead with the intervention itself was (economically) not clear-cut nor was the viability of running a new production log that would ensure a more accurate prognosis. As such management wanted more insight into the available alternatives, hence a RO analysis was performed.

For this article the following abbreviations are defined: J_0 refers to the no-intervention case, J_1 refers to well intervention but with *no* additional PLT data, J_2 refers to well intervention but with the additional PLT data (which is assumed to cost $125,000 plus lost production and rig time).

Step 1: Frame the Problem

Ostensibly this step qualifies and bounds the problem. It helps define which parameters are significant (includes only what is relevant) and clarifies what are the actual options embedded in the project. It can also create alternative solutions not previously considered. The activity can, however, range in complexity from trivial to confounding. It all depends on the nature of the problem itself.

Semi-formal structures upon which one may compose a detailed frame have been proposed[13,14]. These can consider multiple factors, inputs and conditions using a consistent methodology. Effectively this activity is a self-questioning qualitative assessment of the problem conducted by a multi-disciplinary, cross-functional team covering the full knowledge chain (our problem stake-holders). Ultimately, however, the decision-makers' perspective is the one that really counts. The framing activity for stake-holders can be summarized as follows:

a. Stake-holders define a problem statement,
b. Agree on 2 to 5 major strategic decision types: for our example they will be "which well to intervene," "type of intervention" and "timing". (Note: a decision is an irreversible allocation of resources),
c. List alternatives pertaining to each strategic decision (e.g., our candidate well(s) under column A),
d. For each alternative obtained in step c, brainstorm a list of associated issues and label each as either: *fact*, *minor decision* or *uncertainty*,
e. Construct a strategy table and use it to trace a series of realistic decisions based on available options[*]. Figure 1 shows our basic strategy table and illustrates a trace representing a realistic, but "*Cautious*" decision route: A1→B2→C2 equates to J_1. Other decision routes considered are: J_2 = A1→B3→C2 and

		A	B	C
Strategic Theme →		Which Well to Intervene	Type of Intervention	Timing of Intervention
"Negative"	1	Primary candidate	Do nothing, natural decline	Act now regardless and charter vessel at current rates
"Cautious"	2	Secondary candidate	Shut-off without PLT data	Act in 6 months time (vessel already chartered)
"Adventurous"	3	Third-placed candidate	Shut-off with PLT data	Act only when net revenue to $20/bbl

Fig. 1. Basic strategy table. Issues requiring resolution head columns A to C. Listed under each are the different options associated with them. A decision then comprises a combination of alternatives (a "*Cautious*" decision trace shown: A1→B2→C2). *Note:* we could add to, and refine, this table by adding other issue-related columns, e.g. one describing intervention method (CT, slickline etc). The labels used under strategic themes ("*Negative*", etc.) are optional but are useful for classifying our strategic decisions and reporting purposes.

[*] The word 'options' here is used to mean 'alternatives', not financial options

J_0 = A1→B1→C2. For each item listed in the table there are related facts, uncertainties or minor decisions (not shown in the figure). This information is used in Step 2.

The above outline is highly stylized. More rigorous approaches utilize influence diagrams, decision boards, decision hierarchies and other conceptual tools (refer to Coopersmith *et al.*[15] and Bailey *et al.*[16] for details).

Step 2: Quantify Internal Uncertainties
We have now revealed a number of realistic strategic decisions for which we have a list of associated facts, uncertainties and minor decisions. In so doing we have stated the main factors that will be influential in our RO analysis. It is now necessary to identify the nature of the uncertainties and how they may influence one another (interdependencies). We should now define which of our uncertainties are private (internal) or market (external)[17]. Market uncertainty (e.g. oil price, risk-free rate) is outside the control of the decision maker and is covered in Step 3. However, our stakeholder can reasonably expect to know more about private uncertainty. We should now establish spread, distribution shape, correlation and certitude of our 'internal' parameters (access to databases, experience, experts, or by choosing to acquire new data).

In the strategic decision traced in Fig. 1 the following private post-intervention uncertainties were identified: initial production, water cut, job duration and OpEx (lift, process and transport). Rate of decline is uncertain for all *pre*- and *post*-intervention scenarios.

Probability Density Function's (PDF's) should now be generated for all private uncertainties. The impact on production profiles for jobs J_1 and J_2 are shown in Fig. 2. It often transpires that formulating the necessary PDF's is a time-consuming task that requires searches and possibly substantial investment.

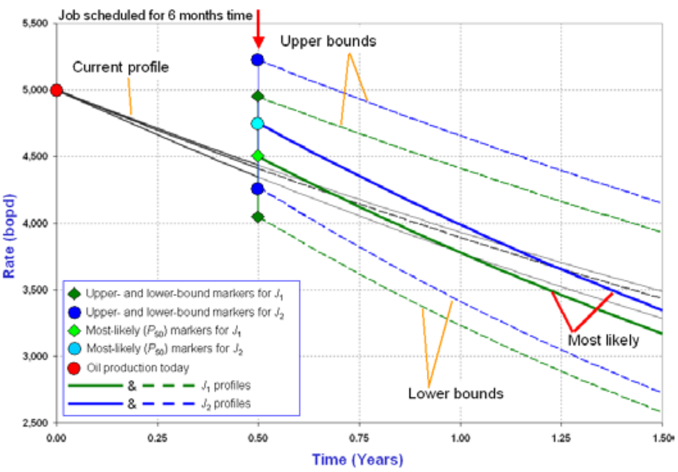

Fig. 2. Uncertainties with respect to initial oil production. Plot shows median (P_{50}) along with upper- and lower-bounds initial production forecast with their respective declines. Option 1 refers to intervention without new PLT data, option 2 refers to one with this important data.

Step 3: Quantify External Uncertainties
It is just not possible to portend the future price of oil (or any other commodity subject to open worldwide markets). All we can do is make a 'best guess'. A number of approaches have been adopted to provide such a prediction: using company policy, pure guess work, trend forecasting, stochastic forward pricing[18] and so on. While some are more sophisticated than others, intricacy alone does not guarantee a better forecast. The approach suggested here is, therefore, no better or worse than any other. It is, however, consistent. The reader is at liberty to apply any model they deem suitable to quantify these exogenous uncertainties.

The approach utilizes actual market data to define what the current exchange considers to be commodity price volatility and drift. New York Mercantile Exchange (NYMEX) sweet crude futures and options data were used. Options data is used to compute forward implied volatilities while futures contracts provide drift. A mean reverting factor could also be extracted from the options data. Figure 3 shows the drift based on current market expectations (contracts) and is around minus 8.3% for the first year and less after that (see inset to Fig. 3). Figure 4 shows the variation of implied volatility with strike and time to expiry using the NYMEX market data (for July 5th, 2002) to solve the generalized BS-M model. Solution of implied volatility requires robust root-finding algorithms that necessitate proper bracketing. The Brent method[19] was found to be the most stable and rapid algorithm considered. Figure 4 also shows long-term mean-reversion in action (an empirical observation when studying market trends) and the implied volatility smile for short-term (near-term) contract expiration dates (the deadline when the contract is either exercised or left to die). Implied volatility frowns for low-strike medium-term options are also observed.

Implied volatility 'smiles' are expected as in order for the price of sweet crude to achieve these more-extreme values it must vary by a significant amount (exhibit high volatility). However, when the near-term future price is

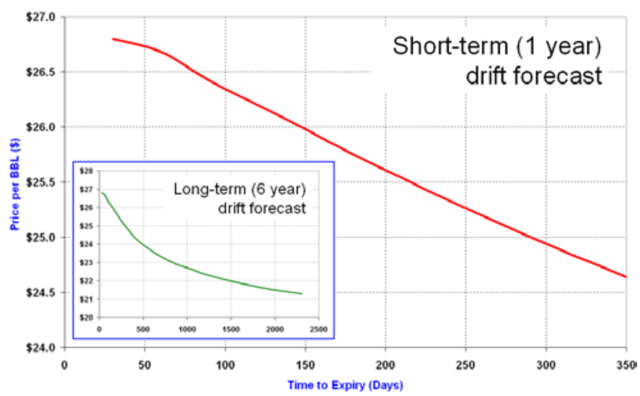

Fig. 3. Short-term (one year) drift forecast (main) and the long-term oil-price futures drift (inset) based on market expectation from NYMEX sweet crude contracts dated 5 July 2002.

close to the current price of crude our oil price does not need to vary greatly, hence a smaller implied volatility (refer to Wilmott[10] for elaboration).

Step 4: Do the Numbers
This is divided into 5 parts for ease of application.

Part 4.1: Construct a DCF model
For each of the strategic decisions realized in Step 1 construct a dedicated DCF model. Our model also applied a mean-reverting stochastic model for oil prices over time (using drift and volatilities from Step 3). The crucial activity here is to identify what is the main underlying involved. In this example we wish to establish the viability of a particular well intervention proposal (do we, don't we? If so, how?). We also want to quantify the benefits (if any) of collecting additional production log data (the "value of information"). As such the primary underlying (denoted as S) considered is the difference between the revenue streams from the well with and without any intervention. *Note:* J_0, J_1 and J_2 now refer to the net revenue streams (post tax and OpEx) for their respective job types (defined earlier) *but* excluding any direct job-related costs. In other words our option represents how much one is willing to pay for the right to the underlying (S), namely the incremental net present revenue streams such that $S = J_1 - J_0$ *or* $J_2 - J_1$).

A simple Monte Carlo model was built into the DCF model so that all uncertainties (defined in step 2) were considered.

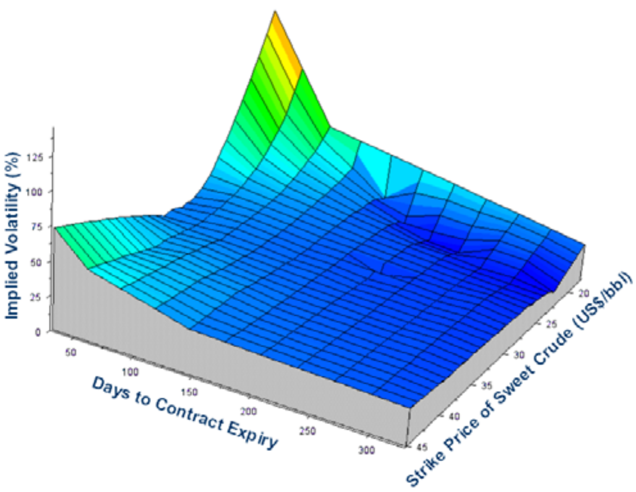

Fig. 4. The solution surface for implied volatility (vertical axis) of sweet crude options (for contracts dated July 5th, 2002) computed from the generalized BS-M model. The remaining axes show expiry (days until the option contract must be excercised or left to die) and the strike price (the "guess" for sweet crude prices on the expiry date). The "volatility smile" is observed at early expiry times (roughly 0 to 75 days). A mild "volatility frown" is observed at about 200 day expiry.

Volatility. Volatility is defined as the standard deviation of the forecast distribution of the underlying at time T. It is a critical issue that effectively impels a whole industry in the finance community. Here we need to determine the volatility that adequately reflects the behaviour of our underlying (S). Several approaches are available. One such is GARCH[20,21] (Generalized AutoRegressive Conditional Heteroskedasticity) but insufficient data exists to proceed in that direction. Another is to use a logarithmic ratio which states that:

$$\sigma = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(x_i - \bar{x}\right)^2} \ , \tag{6}$$

where x is the natural log ratio of the underlying for periods i and $i-1$ such that

$$x = \ln\left(\frac{Jx_i}{Jx_{i-1}}\right) \tag{7}$$

and Jx is a net present revenue stream: J_1 or J_2. This method is easy to implement, provides a single value for σ and is widely used in the financial markets. However, it has caveats, namely the possibility of negative cash-flows over particular time periods. This would mean taking the logarithm of a negative number – which does not exist. As such σ will not fully represent *all* possible cash flows.

Another method collapses all future cash flow estimates into two sets of present values[22]: one for the first time period and another for the present which are then used to compute a log ratio G:

$$G = \ln\left(\frac{\sum_{i=1}^{n} Jx_i}{\sum_{i=0}^{n} Jx_i}\right). \tag{8}$$

The process is summarized in Table 1 with the volatility being the standard deviation of G. Values for σ obtained using this approach were 12.22% for J_2 and over 200% for J_1 (unrealistically high). This is due to the presence of strong negative cash flows for certain simulation trials that causes model misbehaviour. Volatilities from Eq.(6) were (as expected) on the low-side (again due to the presence of negative cash flows) and were 11.3% and 7.3% for J_1 and J_2 respectively. As such these approaches are considered unreliable for cash flows where portions of them go negative. Another observation of this approach is that in a RO the variability in the cash flow (our underlying) is the key driver to value and not the variability of discount rates, which have a significant effect in this method.

The volatility used here was taken as being the standard deviation of the resultant forecast for the underlying obtained from a 10,000-trial Monte Carlo simulation. Figure 5 shows this for the underlying for J_1 (intervention without new PLT data). Volatility for J_2 was less (18.69%) and can be partially attributed to the increased certitude of increased production provided by the additional PLT data (hence more positive cash flows). This approach does not suffer the drawbacks demonstrated earlier.

Figure 6 (back of paper) summarizes the whole calculation process using the input for one trial. Part 6.1 of this figure shows the oil price multiplier applicable to this trial based on a mean-reverting (MR) random walk generated using the drift and volatility obtained in Step 3 (as well as a mean-reversion parameter, α not covered in this text, refer Clewlow & Strickland[18]). The associated (highly simplified) DCF sheet for 5 discrete points in time is also shown. The various cash flows (rows E and H in part 6.2 of Fig. 6) are determined and the difference (row S) is computed and recorded for one trial. The process is then repeated until sufficient trials have been recorded. The standard deviation of the resultant distribution of the underlying is the volatility we require. Part 6.3 of Fig. 6 shows J_0 (blue) and J_1 (red) revenue streams for 16 such

Period, t	Undiscounted Underlying	PV for time $= t$	PV for time $= t-1$
0	x_0	$\frac{x_0}{(1+r)^0}$	–
1	x_1	$\frac{x_1}{(1+r)^1}$	$\frac{x_1}{(1+r)^0}$
2	x_2	$\frac{x_2}{(1+r)^2}$	$\frac{x_2}{(1+r)^1}$
3	x_3	$\frac{x_3}{(1+r)^3}$	$\frac{x_3}{(1+r)^2}$
⋮	⋮	⋮	⋮
N	x_N	$\frac{x_N}{(1+r)^N}$	$\frac{x_N}{(1+r)^{N-1}}$
Sum		$\sum_{i=0}^{N} J_i$	$\sum_{i=1}^{N} J_i$

Table 1. Methodology for the log-ratio volatility evaluation technique.

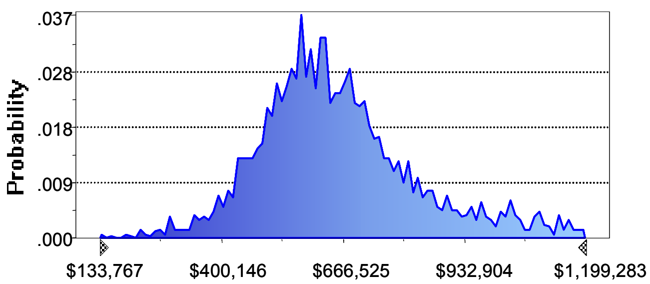

Fig. 5. Frequency distribution forecast (for 10,000 trials) for the underlying asset, S, ($J_1 - J_0$: incremental net revenue streams excluding job cost). Standard deviation is $208,000 (30.9% of the mean).

trials. Their corresponding differences (the underlying, S) is shown beneath on the same scale. The intervention itself is scheduled for 6 months time, hence the underlying must be zero from t=0 to t=0.5 then finite as the new post-intervention production takes effect. The resultant forecast frequency distribution plot for $J_1 – J_0$ from Monte Carlo simulation yields a standard deviation of the underlying of $208,000 or 30.9% of the mean (see Fig. 5). By the same token the volatility for the underlying involving $J_2 – J_0$ was found to be 18.69%.

Part 4.2: Is it worth doing RO's at all?

This stage utilizes a simple graphical tool comprising two vertical axes: y_1 is Net Present Value (NPV) and y_2 is PV revenue frequency count (dimensionless). PV revenues (*excluding* job costs) are plotted on the x-axis. First the linear payoff function for the job is plotted. Next, a "first-pass-only" option calculation is plotted against y_1. Unless the payoff structure is complex and/or conditional, the Bjerksund & Stensland[23] approximation of an American call is probably sufficient (refer Haug[24] for code). Figure 7 demonstrates the Payoff-Frequency plot for two cases: where option values exist (PV revenue forecast 'D') and one where they do not (PV revenue forecast 'E'). The only purpose of this plot is to screen-out - via visual

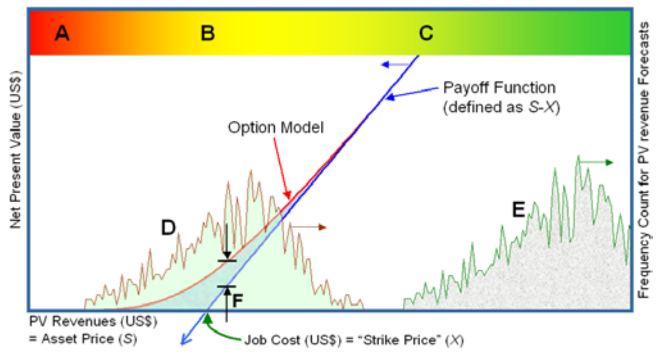

Fig. 7. Payoff-Frequency plot. Upper coloured band signifies whether project has negative value ("out-of-the-money" – point **A**), marginal ("at-the-money" - point **B**) or profitable ("in-the-money" – point **C**). Project is said to be "deep-in-the-money" if located on the far right. The linear payoff function is PV revenue (S) minus job cost (X). Forecast '**D**' of the underlying asset (obtained from Monte Carlo simulation) straddles the break-even point. The value of the option is – roughly – the difference between the option model and the linear payoff function $S–X$: margin '**F**' roughly estimates the option value at the P_{50} for forecast '**D**'. The option model becomes asymptotic with the payoff function as we go deeper in-the-money. Any underlying (PV revenue difference) in this region has, therefore, negligible option value. Forecast '**E**' shows this. Forecast '**D**' are the true values – hence visual inspection indicates that there is indeed value to the options. σ = 30.9%, X = job cost, S = underlying (difference between J_1 and J_0 net PV revenues - excluding job cost), T = 1 year r = 5%.

inspection - *deep* in-the-money (and possibly *deep* out-of-the-money) projects. When deep in-the-money the option model is asymptotic with the linear pay-off function and implies that there is essentially no option value to be gained. The plot should *not* be used when performing the analysis proper (*note:* this option model is not suitable for all cases). It is, however, considered accurate enough at the asymptote to enable visual screening.

We have now established that our problem has a forecast PV revenue stream that straddles the break-even point and that some option value does exist so it is appropriate to proceed with the analysis.

Part 4.3: Available Solution Options

Numerous methodologies can be used to calculate a financial option's value. These range from closed-form models (e.g., BS-M and variants), Monte Carlo path-dependent simulation, lattices (bi-, tri- and multi-nomial trees), variance reduction to partial differential equations (PDE) and so on.

Closed-form models are exact, rapid and easy to program - but difficult to explain (stochastic calculus is not usually amenable for management summary). They are also highly specific in nature with limited modeling flexibility. Appendix C summarizes the assumptions inherent in the BS-M equation. PDE's are also not transparent, apply to a specific valuation function and can be difficult to explain. The use of a closed-form equation used in 4.2 above was justified for purposes of screening but no further. If, however, the closed-form model (or PDE) applies perfectly to the particular project then they should be used. However effective communication and explanation of results may become an issue.

For our needs a bi-nomial lattice is preferred as it is flexible, easy to use and relatively transparent. The main drawback is the need for small time-steps to obtain good approximations. As the time-step decreases the expanse of the lattice becomes significantly greater and significant computer power is needed to solve the model. Nevertheless, at the limit, results obtained through such lattices tend to approach those derived from closed-form solutions. For example: consider a case where S = X = \$100, T = 1 year, r = 5%, σ = 25% and no dividends returns, the generalized BS-M model yields an option value of \$12.3360. A binomial lattice constructed using the same data provides the following: 10 lattice time-steps: \$12.0923; 50 steps: \$12.2867; 1,000 steps: \$12.335 with convergence to closed-form achieved at 50,000 steps: \$12,3360. 1,000 time-steps is sufficient for reasonable approximations. Tri-nomial lattices (and higher) have also been proposed[25] – however at very small time-steps such intricate lattices will provide the same result as a bi-nomial lattice.

Part 4.4: Constructing the Lattice

No matter what type of option is being considered, at least two lattices need to be constructed. The first shows the evolution of the underlying, S, while the second is the actual valuation lattice. Time-steps are defined as the number of branching events in a lattice. Figure 6 shows a 5-step lattice with corresponding up and down movement factors (u and d) computed using the volatility of 30.9% established in part 4.1 earlier.

If we know for certain what our pre- and post-intervention cash flows will be then our underlying in Part 6.3 in Fig. 6 will be a single straight line with zero volatility. As such the option value will be exactly the same as our deterministic DCF model. This is a good test of the model.

The lattice follows the evolution presented in Part 6.4 in Fig. 6. The lattice is combining at all nodes due to the presence of a single volatility throughout. If multiple volatilities were used over the period, (say σ_1 for t=0 to 30 days, σ_2 from t=30 to 90 days and so on, (reasonable considering the data shown in Fig. 4) - a non-recombining lattice would result which would grow nodes exponentially: 2^1 nodes at time-step 1, 2^2 nodes at time-step 2, and so forth until we achieve 2^{1000} nodes at time-step 1,000 (about $2{\times}10^{301}$). Such a daunting construct emerges when $S_0ud' \neq S_0du'$. (Note that S_0 denotes S at t=0). In our analysis we circumvented the need for such complexity by establishing, for each trial, a unique oil price path for the period that incorporates time-dependent volatility and drift explicitly and is applicable to all job scenarios (parts 6.1 and 6.2 in Fig. 6). The result is an underlying that can be represented by a single volatility.

The valuation function defined in the valuation lattice is likely to be a function of S and X. If X is paid up-front (at t=0) then the valuation lattice would use S–X. In this case we are essentially valuing an 'alternative' (J_1 or J_2) as the 'option' to pay X at t=T has already been taken (so there is effectively no option available). Conversely, if X is to be paid only when the project has proven economically successful (at t=T) then MAX(S–X, 0) should be used in the valuation lattice. This type of valuation (shown in Figure 6) implies that project revenue risk is being underwritten by the service provider. This, from the asset-holder's perspective, can now be considered to be a real option (only pay X if 'in-the-money'). This option then behaves like a simple BS-M European call. In reality we are likely to encounter more complex valuation structures. For example an internal hurdle rate for return on investment of 25% may be demanded. Such a criterion would then be defined in the valuation lattice as $S-1.25X$, if X is paid at $t=0$.

Rutherford[26] used exotics to value the options available for a farmout opportunity. The question was: should a marginal discovery be farmed-out right now, later (after further delineation) or not at all. The problem was resolved using standard barrier and cash-or-nothing binary options. However, unless one is well versed in the parlance of such exotic financial instruments it may not always be clear which exotic is required. Not so when using a lattice as it is merely a matter of conditioning the lattices properly. Table 2 provides a results summary of our analysis.

	J_1	J_2
Option Value	\$151,075	\$207,181
$S - P_{50}$ value	\$626,734	\$1,023,201
$X - P_{50}$ value	\$524,083	\$869,305
σ	30.9%	18.7%
Values from deterministic DCF		
NPV	(\$303,369)	(\$307,308)
Incremental NPV*	\$119,946	\$116,007

Table 2. Results summary. Options and deterministic analysis. *Incremental NPV is $J_{1,2}$ NPV minus J_0 NPV (at one year). Note that deterministic net revenues for J_0 (no intervention) were negative (\$423,315) (due to higher water cut and the associated lifting cost). Alone this would result in shutting-in the well. r=5%. See Step 5 for further interpretation.

Consideration of the aforementioned leads to an apparent contradiction: we know that our future incremental revenue is not a single value but rather a range of values. However, we are given just one value for our option – how can this be, shouldn't it be a range of values? This can be resolved by accepting the appropriate definitions: namely our option value is the difference between the linear pay-off function, ($S-X$, which assumes a constant X) and the value of the option *at the median (P_{50}) value of the underlying*. (This definition excludes any additional income obtained from investment deferral and such like). Our $J_{1,2}$ options are valued at these P_{50} values of S. This is acceptable as we cannot presage what S will be in the future. We then hedge our chances by accepting a 50-50 probability of being higher, or lower, than the stated value. The next section discusses this in more detail.

Step 5: Communicate Results & Decide

This can easily be the hardest part of any RO analysis. It is also a serious weakness. We propose a graphical vehicle that visualizes the options themselves against the underlying. This is shown in Fig. 8. So as not to clutter the picture, a fixed value of X is assumed. First a linear pay-off function $S-X$ is generated. The first y-axis (y_1) will represent NPV. Second, super-impose the PDF of the underlying computed earlier (against the second y-axis, y_2). Third, generate a range of option values by varying S and plot against y_1. This visualizes the relationship between our expected benefit (S), its cost (X) and the option value. The difference between ($S-X$) and the option curve is the value of the RO (this assumes zero additional income due to decision deferral and other off-DCF analysis incomes).

For J_1 this "gap" equals \$48,424 [Option Value minus ($S-X$)] and \$53,285 for J_2. In other words: *in order to gain the right to the cash flow stream generated by the job I would, theoretically, be willing to pay this additional amount*. (This considers jobs as discrete entities with single service provider). In other words, the service provider, by effectively underwriting job-related revenue risk, could justifiably charge this premium over and above the job cost X. If, however, $S-X$ was used then the 'option' value would represent a reasonable upper job cost bound. The analysis also allows one to evaluate the relative benefits of competing service providers: not just on price but also on (the more elusive) additional value added. We are now in a somewhat stronger position to respond to the following:

Q: Should we intervene in the well at all?
A: Yes. While deterministic DCF indicates a negative in all cases (justifying shutting-in the well), the fact that these cash flows are volatile means that there is an upside which could be captured. Our analysis (and Fig. 8) shows that it is worthwhile to intervene.
Q: Which job-type should we go for?
A: J_2 (intervene and collect new PLT data). Although J_2 costs more there is a differential benefit of \$56,106 - the difference between J_2 RO value and that for J_1. If this was negative we would opt for J_1.
Q: Theoretically, could I pay more for the PLT data and by how much?
A: Yes. The RO indicates that the current PLT charge is less than its value (recall it was assumed to be \$125,000 plus additional lost production and rig time). Parity between J_1 and J_2 benefit is achieved when PLT cost is increased to about \$247,000 (a difference of \$122,000). This value then represents the latitude for contract negotiation, assuming a single possible service provider (i.e. the upper threshold at which the additional information is not worth acquiring).
Q: What insight has RO brought that a DCF could not?
A: The deterministic DCF model is effectively a slave to perception. The negative NPV's for the job (Table 2) were obtained from using "best guesses". A crude sensitivity analysis indicated that there was certainly an upside but significant revenue downsides too. We would have, ordinarily, rejected the proposal to intervene in the well and shut it in. The fact that positive future cash flows

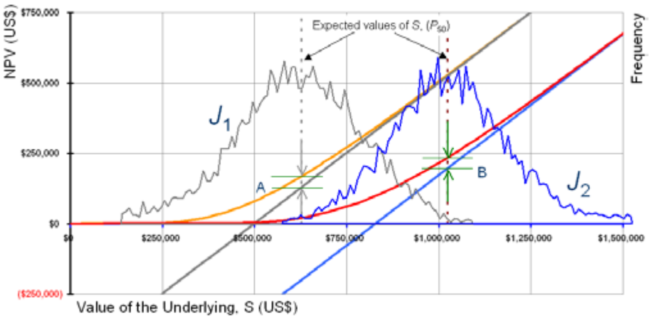

Fig. 8. Payoff-Frequency plot of J_1 (left, in grey) and J_2 (right in blue) options. Margin ‘A’ equals $48,424. Margin ‘B’ equals $53,285. These points mark the gap between linear pay-off and the option curve at the P_{50} value of the underlying (S).

were possible, even probable, was lost in the deterministic approach. For example a Monte Carlo model for non-intervention job-type J_0 NPV’s revealed the following: P_{50} $1,733,148, P_{10} ($3,137,831) and P_{90} $7,078,060. This demonstrates the considerable variability inherent in this DCF analysis. The stated NPV of minus ($423,315) using “best guess” values indicates that these values were not reasonable (in fact they represent P_{37} values – not the P_{50} as thought). Such is the inescapable bias inherent in deterministic models.

Finally, if we only let the job price (X) vary, but keep the underlying the same, we can observe how the option value varies to shifts in the linear pay-off function. Daily rig rate and job duration are considered variable here. This is shown in Fig. 9 for J_2 and also for the difference between J_2 and J_1 option values (where a negative value indicates a preference for J_1 over J_2).

Conclusions and Process Summary

A Real Options analysis was conducted on a somewhat generic production enhancement problem. It utilized market-expectations for oil price volatility and drift in a

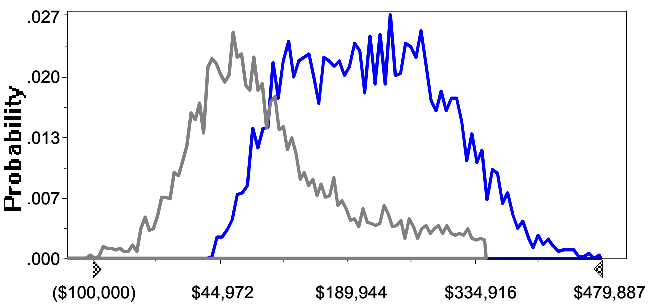

Fig. 9. Variation in J_2 option value (dark blue curve) from varying job cost (X) alone (equivalent to shifting the linear payoff function. The lighter grey curve is the difference between J_2 and J_1 option values (J_1 was subject to the same job cost changes except J_2 takes longer). Cost variables considered were job duration and rig rate. Product and service charges were kept constant.

mean-reverting stochastic oil price model. The process summary is as follows:

- Step 1: Frame the problem. Identify strategic decisions and main cash flow drivers in a qualitative framework.
- Step 2: Quantify the internal uncertainties for each of the strategic decisions defined in step 1. (Here they are J_0, J_1 and J_2). Our uncertainties were OpEx, initial production rate, decline, job duration, job cost and water cut but others may be possible.
- Step 3: Quantify the external uncertainties. An optional step as we can only presage, not define future oil prices. Our approach utilized actual market data.
- Step 4: Constructed a Monte Carlo DCF model. Identify the correct underlying, S (not always a simple task). Here it is incremental revenues at one year. Obtain a forecast of the underlying and note the underlying volatility (standard deviation as a percentage of the mean). Evolve a lattice and use P_{50} values for the underlying to calculate the RO value.
- Step 5: Communication (interpretation and make a decision). Generate a Payoff-Frequency plot to visualize the results for demonstration and communication purposes. Possibly use results to define contract terms and evaluate potential vendors. Luehrman[11] expands on this step with additional follow-up activities (monitoring, managing, marketing and evaluating success of the final decision).

While the above process should be applicable to a variety of real problems there will be exceptions. More likely than not the underlying will be a cash flow stream of sorts and the option will reflect the premium one would be willing to pay to obtain the right to that cash flow. Financial options act in exactly the same way. Nevertheless there are a number of unresolved technical issues (WACC, volatility). More importantly though is the need to find a means to communicate and market these principles to the decision makers, otherwise some of the value embedded in our uncertain projects will be lost.

Nomenclature

$d_{1,2}$	=	Black-Scholes-Merton parameters Eqs.(2 & 3)
J_0	=	Job type and/or net revenues (after tax and OpEx) for no intervention
J_1	=	Job type and/or net revenues (after tax and OpEx) after well intervention without any new production logging data
J_2	=	Job type and/or net revenues (after tax and OpEx) after well intervention but with new production logging data
$N(.)$	=	cumulative distribution function for the standardized (zero mean, unit standard deviation) Normal distribution

r = discount rate (specify whether risk-free rate or WACC, %p.a.)
S = underlying asset (US$)
x = undiscounted cash flow for period, Table 1 (US$)
X = job cost (equivalent to the strike price, US$)
δt = time step used in lattice
σ = volatility of the underlying (% of mean)

Abbreviations

BS-M Black-Scholes-Merton
CT Coiled Tubing
DCF Discounted Cash Flow
NPV Net Present Value
OpEx Operating Expenditure (US$/bbl)
PDF Probability Density Function
PLT Production Logging Tool
PV Present Value
RO Real Options
WACC Weighted Average Cost of Capital

References

1. Amram, M. and Kulatilaka, N.: *Real Options - Managing Strategic* Investments *in an Uncertain Worl*d, Harvard Business School Press, Boston MA, 1999
2. Armstrong, M. and Galli, A.: "Option Pricing, A New Approach to Valuing Mining Projects," *CIM Bulletin*, (April 1997) **90**, No. 1009, pp. 37-44.
3. Copeland, T. and Antikarov, V.: *Real Options – A Practitioner's Guide*, Texere, New York, NY, 2001
4. Leslie, K.L. and Michaels, M.P.: "The Real Power of Real Options," *The McKinsey Quarterly*, (1997) No. 3., pp.97-108.
5. Luehrman, T.A.: "Investment Opportunities as Real Options: Getting Started on the Numbers," *Harvard Business Review*, reprint number 98404, (July-August 1998), pp.3-15
6. Luenberger, D.G.: *Investment Science*, Oxford University Press, Oxford, 1998
7. Trigeorgis, L.: *Real Options – Managerial Flexibility and Strategy in Resource Allocation*, The MIT Press, Cambridge, MA, 1996
8. Black, F. and Scholes, M.: "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* (May/June 1973) **81**, pp. 673-659.
9. Merton, R.C.: "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, (Spring 1973) **4**, pp. 141-183.
10. Wilmott, P.: *Paul Wilmott on Quantitative Finance*, Vols 1 & 2, John Wiley & Sons, 2000
11. Luehrman, T.A.: "Extending the Influence of Real Options: Problems and Opportunities," paper SPE 71407, SPE Annual Technical Conference and Exhibition, New Orleans, Louisiana (30 Sept. 3 Oct, 2001)
12. Laughton, D.: "The Management of Flexibility in the Upstream Petroleum Industry," *The Energy Journal*, IAEE (1998) **19**, No. 1, pp.83-114.
13. Bodily, S.E. and Allen, M.S.: "A Dialogue for Choosing Value-Creating Strategies," *Interfaces* (the journal of INFORMS), (November-December, 1999), 29, pp. 16-28
14. Claeys, J. and Walkup, G.: "Discovering Real Options in Oilfield Exploration and Development," paper SPE 52956, SPE Hydrocarbon Economics and Evaluation Symposium, Dallas, TX (20-23 March, 1999)
15. Coopersmith, E., Dean, G., McVean, J. and Storaune, E.: "Making Decisions in the Oil and Gas Industry," *Oilfield Review* (Winter, 2000), pp. 2-9.
16. Bailey, W.J., Couët, B., Lamb, F., Simpson, G. and Rose, P.: "Taking A Calculated Risk," *Oilfield Review* (Autumn 2000), pp. 20-34.
17. Faiz, S.: "Real-Options Application: From Successes in Asset Valuation to Challenges for an Enterprisewide Approach," *JPT* (Jan 2001), pp. 42-47 & 74. Paper SPE 68243 originally presented at the SPE Annual Technical Conference and Exhibition, Dallas, TX (1-4 Oct, 2000).
18. Clewlow, L. and Strickland, C.: *Energy Derivatives Pricing and Risk Management*, Lacima Publications, London, UK, 2000
19. Press, W.H., Flannery, B.P., Teukolsky, S.A. and Vetterling, W.T.: *Numerical Recipes in C. The Art of Scientific Computing*, Cambridge University Press, 1990.
20. Engle, R.: "Autoregressive Conditional Heteroskedacity with Estimatyes of the Variance of the United Kingdom Inflation," *Econometrica*, (1982) **50**, pp.987-1007
21. Bollerslev, T.: "Generalized Autoregressive Conditional Heteroskedacity," *Journal of Econometrics*, (1986) **31**, No. 2, pp.307-327
22. Mun, J.: *Real Options Analysis: Tools and Techniques for Valuing Strategic Investments and Decisions*, Wiley Finance, John Wiley & Sons, 2002
23. Bjerksund, P. and Stensland, G.: "Closed-Form Approximation of an American Options," *Scandinavian Journal of Management*, (1993), **9**, pp. 87-99.
24. Haug, E.G.: *Option Pricing Formulas*, McGraw-Hill, New York, NY, 1997
25. Clewlow, L. and Strickland, C.: "Valuing Energy Options in a One-Factor Model Fitted to Forward Prices," working paper, University of Technology, Sidney, Australia (1999)
26. Rutherford, S.R.: "Farmout Valuation Using Exotic Options," paper SPE 71409, presented at Annual Technical Conference, New Orleans, Louisiana (30 Sept - 3 Oct, 2001)

Appendix A: Terminology Equivalents

Financial Parlance	Equivalent
Underlying Asset, S	Project value; present value of future cash flow streams
Strike, X	The cost (i.e. job cost) to have the right to the desired underlying (S).
Expiration, Expiry, T	The time when we are evaluating our future cash flow *or* time from now when we need to make the decision.
(Plain) Vanilla	Simple cost structure with no complicating factors, i.e. $\mathrm{MAX}(S-X,0)$
Exotic	Complex cash flow/cost structures - for example internal hurdle rates for ROI, abandonment threholds, etc.
European	An option that can only be exercised at a specific time, T
American	An option that can be exercised at *any time* up to T

Appendix B: Issues with WACC

The discount rate that is used is usually calculated from WACC, Capital Asset-Pricing Model (CAPM), Multiple Asset-Pricing Model, Arbitrage Pricing Theory (APT), set by management as a requirement for the firm, or as a hurdle rate for specific projects. In most circumstances, if we were to perform a simple DCF model, the most sensitive variable is usually the discount rate. The discount rate is also the most difficult financial variable to correctly quantify. Hence, this leaves the discount rate to potential abuse and subjective manipulation. A target NPV value can be obtained by simply massaging the discount rate to a suitable level. In addition, certain input assumptions required to calculate the discount rate are also subject to question. For instance, in the WACC, the input for cost of common equity is usually derived using some form of the CAPM. In the CAPM, the infamous beta *(*β - a measure of systematic risk) is extremely difficult to calculate. In financial assets, we can obtain beta through a simple calculation of the covariance between a firm's stock prices and the market portfolio, divided by the variance of the market portfolio. Beta is then a sensitivity factor measuring the co-movements of a firm's equity prices with respect to the market. The problem is that equity prices change every few minutes! Depending on the time frame used for the calculation, beta may fluctuate wildly. In addition, for non-traded physical assets, we cannot reasonably calculate beta this way. Using a firm's tradable financial assets' beta as a proxy for the beta on a project within a firm that has many other projects is ill-advised.

There are risk and return diversification effects among projects as well as investor psychology and overreaction in the market that are not accounted for. There are also other more robust asset-pricing models that can be used to estimate a project's discount rate, but they require great care. For instance, the APT models are built upon the CAPM and have additional risk factors that may drive the value of the discount rate. These risk factors include maturity risk, default risk, inflation risk, country risk, size risk, non-marketable risk, control risk, minority shareholder risk, and others. Even the firm's CEO's golf score can be a risk hazard (*e.g.,* rash decisions may be made after a bad game or bad projects may be approved after a hole-in-one, believing in a lucky streak). The issue arises when one has to decide which risks to include and which not to include. This is definitely a difficult task, to say the least.*

One other method that is widely used is that of comparability analysis. By gathering publicly available data on the trading of financial assets by stripped-down entities with similar functions, markets, risks and geographical location, analysts can then estimate the beta or even a relevant discount rate from these comparable firms. For instance, an analyst who is trying to gather information on a research and development effort for a particular type of drug can conceivably gather market data on pharmaceutical firms performing only research and development on similar drugs, existing in the same market, and have the same risks. The median or average beta value can then be used as a market proxy for the project currently under evaluation. Obviously, there is no silver bullet, but if an analyst were diligent enough, he or she could obtain estimates from these different sources and create a better estimate. Monte Carlo simulation is most preferred in situations like these. The analyst can define the relevant simulation inputs using the range obtained from the comparable firms and simulate the discounted cash flow model to obtain the range of relevant variables (typically the NPV and IRR).

Appendix C: BS-M model limiting assumptions

In order to fully understand and use the BS-M model, we need to understand the assumptions under which it was constructed. These are essentially caveats that go into using RO's in valuing any asset. These assumptions are violated quite often, but the model should still hold-up to scrutiny. The main assumption is that the underlying's asset price structure follows a Geometric Brownian Motion with static drift and volatility parameters and that this motion follows a Markov-Weiner stochastic process. The general derivation of a Markov-Weiner stochastic process takes the form of

$$dS = \mu S dt + \sigma S dz \qquad \text{(A.1)}$$

where

$$dz = \varepsilon \sqrt{dt} \qquad \text{(A.2)}$$

and dz is a Weiner process, μ is the drift rate, σ is the volatility measure and ε is a Normal random variable. The other assumptions are fairly standard, including a fair and timely efficient market with no riskless arbitrage opportunities, no transation costs (frictionless market) and no taxes. Price changes are also assumed to be continuous and instantaneous. The risk-free rate is known as a function of time and no dividend payments. Delta hedging is performed continuously and instantaneously (an impossibility). Refer to Wilmott[10] for a detailed appraisal of BS-M model assumptions from a financial risk management perspective.

---

* A multiple regression or 'principal component' analysis can be performed but probably with only limited success for physical assets as opposed to financial assets because there are usually very little historical data available for such analyses.

6.1: Oil Price Definition (external uncertainty): Any suitable model can be used. Here NYMEX data for 5 July 2002 provides market expectations for implied volatility (from options) and drift (from futures) and is used in a Mean Reverting stochastic price model.

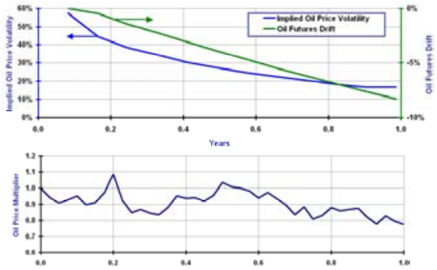

6.2: DCF Model: Highly simplified DCF model (model includes uncertainties).

t =	0	0.2	0.4	0.6	0.8	1
A	1.000	1.086	0.938	0.939	0.880	0.777
B	0.000	0.200	(0.054)	(0.149)	0.187	(0.080)
Job 0: Rates and Net Present Earnings without any intervention						
C	0	47,708	45,382	43,168	41,063	39,060
D	0	143,125	136,145	129,505	123,189	117,181
J0	$0	$175,052	($9,183)	$234,511	$39,109	($423,315)
Job 1: Post intervention rates and Net Present Earnings BUT without any new production log data						
F	0	47,708	45,382	45,012	44,211	40,105
G	0	143,125	136,145	127,661	120,041	116,136
J1	$0	$175,052	($9,183)	$356,332	$413,201	$203,419
Difference between Job 1 and Job 0 and represents the underlying, S						
S	$0	$0	$0	$121,821	$374,092	$626,734

Key
A Oil Price Multiplier
B "dz" for stochastic price model
C Oil production, bopd
D Water production, bwpd
J0 Net Revenue (post-tax and OpEx)
F Post-job oil production, bopd
G Post-job water production, bwpd
J1 Net Revenue (post-tax and OpEx)
S Revenue difference (underlying, S)

Internal uncertainties defined in DCF model

6.3: Quantify Uncertainty in Underlying, S: This is the difference in net revenues between intervening (J_1) and not intervening (J_0) in the well. Top plot shows revenue uncertainty while lower shows the difference (S). Just 16 separate trials are shown here.

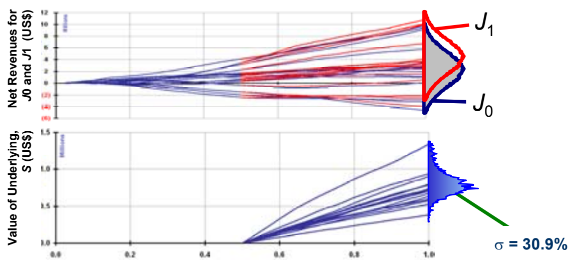

6.4: Lattice Construction: Construct two lattices: one for the underlying asset, S, and one for the option's valuation. The 5-step lattice shown was not used for actual pricing, but is presented for demonstration purposes only. The actual analysis used a 1,000-step lattice. Parameters: δt = 0.2 yrs, σ = 30.9%, r = 5% (this risk-free rate is known)

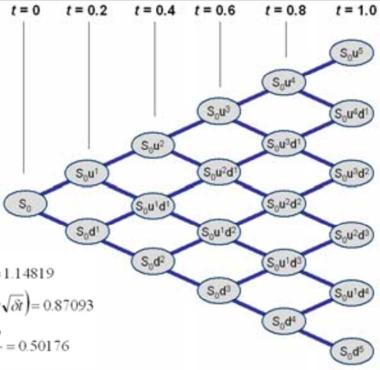

$$u = \exp\left(\sigma\sqrt{\delta t}\right) = 1.14819$$

$$d = \frac{1}{u} = \exp\left(-\sigma\sqrt{\delta t}\right) = 0.87093$$

$$p = \frac{\exp(r\delta t) - d}{u - d} = 0.50176$$

6.5: Number Generation: Lattice evolution of the underlying is straightforward once σ is known. Of note is how $f(S, X, \ldots)$ in the valuation lattice is defined. The valuation function here is simply MAX $(S\text{-}X, 0)$. Complex (exotic) option structures (barriers etc) are defined through $f(.)$ but will have more involved structures. Here (for J_1): S_0 = \$626,734 and X = \$524,083. The 5-step option value is \$150,529. The 1,000 step lattice yields \$151,075.

Lattice Evolution of the Underlying, *S*

t=0	t=0.2	t=0.4	t=0.6	t=0.8	t=1.0
					$1,250,710
				$1,089,206	
			$948,696		$948,696
		$826,252		$826,252	
	$719,611		$719,611		$719,611
$626,734		$626,734		$626,734	
	$545,844		$545,844		$545,844
		$475,394		$475,394	
			$414,037		$414,037
				$360,599	
					$314,058

S_0ud^4

S_0d^5

MAX $(S_0u^nd^n - X, 0)$

Valuation Lattice (Backward Induction)

t=0	t=0.2	t=0.4	t=0.6	t=0.8	t=1.0
					$726,627
				$570,418	
			$434,991		$424,614
		$317,658		$307,384	
	$222,594		$205,906		$195,528
$150,529		$131,349		$107,866	
	$80,993		$58,916		$21,761
		$31,916		$10,810	
			$5,370		$0
				$0	
					$0

Intermediate value \$222,594 given by backward induction:
$[p\times(317{,}658)+(1-p)\times(131{,}349)]\times\exp(-r\delta t)$

Fig. 6. Stepwise process for calculating option value using a binomial lattice. ***Part 6.1*** (optional) takes the external uncertainty (considered to be oil price) and uses the time-dependent volatility and drift (from *Step 3*) to compute a mean-reverting stochastic oil price path over the period of interest. ***Part 6.2*** sets up the actual DCF model which incorporates all internal uncertainties (illustrated as pdf's for oil rates, job time, cost over-runs, water cut and other uncertainties identified in *Step 2* earlier. ***Part 6.3*** determines the uncertainty attributable to our underlying asset, S – which is the incremental revenue one would obtain from intervening in the well. X (the strike price in financial parlance) represents the job cost, or more formally "the amount one if willing to pay to acquire the right to this incremental net revenue stream". The frequency distribution of S is shown and its standard deviation is the volatility we need to proceed with the actual option pricing. ***Part 6.4*** demonstrates the evolution of a 5-step lattice with equi-probable up and down behaviour (u and d). Rigorous analysis requires a much more refined lattice (we used a 1,000 time-step construct). ***Part 6.5*** populates the lattices. The lattice of the underlying asset (upper) is straightforward. The lower valuation lattice utilizes our valuation function, here we apply MAX $(S\text{–}X, 0)$, and through backward induction (shown) computes the option value, which is \$150,529. Note the difference between this value and that presented in Table 2 (\$151,075). This is due to the latter value being computed using 1,000 time-steps which is more accurate. More complex pay-off structures (such as knock-out barriers, minimum R.O.I. and so on) can be accommodated in these lattices. The distinction in interpretation when applying $S\text{–}X$ and MAX $(S\text{–}X, 0)$ is important. The former assumes that X is paid up-front - (sunk at t=0). This means that we are effectively valuing an 'alternative' (scenario J_1 or J_2) because we do not have the ability to exercise an 'option' at time T. The function MAX $(S\text{–}X, 0)$, however, assumes that the asset holder has an option at $t=T$ to either pay X (because $S\text{–}X$ is positive) or pay nothing at all. This transfers the risk burdon onto the service provider. In this case the option value represents the premium the risk taker should charge for accepting this risk – over and above the stated job cost. When using $S\text{–}X$ instead of MAX $(S\text{–}X, 0)$ the option value for the 5-step lattice was computed to be \$128,211. This value is smaller because it accounts for the possibility of negative outcomes (refer Fig. 8). (For J_2 this option value was \$196,293).