Issue 2

Making better use of reserves predictions when deciding where to drill an appraisal well

SPE Seminar : Reserves and Production Forecasting
Aberdeen, January 29 & 30, 1997

T J Allen & D J Franklin

AEA Technology plc

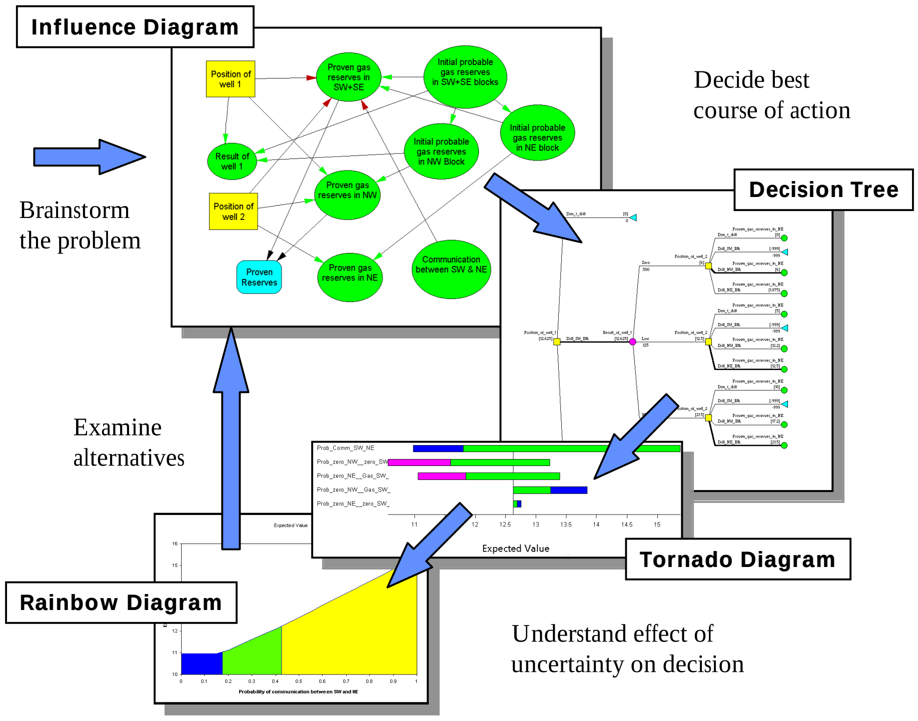

January 1997

Executive Summary

Decision analysis and risk assessment methodologies have been available for many decades, but although they have found some acceptance for ranking exploration prospects, they have not been widely used in appraisal or development drilling. This paper discusses simple methods to help the reservoir team to review their reserve estimates in more detail and optimise the appraisal strategy. The methods allow the team to account for the dependencies between various factors which influence the expected outcome of different strategies. What seems at first sight to be a fairly straightforward problem, is often revealed to be more complex, as all the options are fully considered. The sensitivity of a chosen drilling strategy to the inherent uncertainty in the input values can also be examined. This structured decision process makes it much easier to present the justification for the team's decision to management, since it shows that all the options have been fully considered.

A case study from the North Sea is presented, in which the optimum decision was not the obvious first choice. The prospect was a structural trap, well defined by seismic, and believed to contain several oil-bearing reservoirs, in close proximity to several productive fields. The structure was divided into several fault blocks, which geological mapping indicated may be isolated. Monte Carlo simulation had been used to examine the range and most likely reserves for each of the blocks.

The decision had to be made as to where to drill the first well. The outcome of this well would influence where the second well should be drilled. Because some of the uncertainties are inter-related, the expected value of reserves for the second well are dependent on where the first well is drilled. Following this through, the obvious location for the second well may not be the optimum one, since more data on the fault sealing may be gained by drilling a different fault block.

This case illustrates the value of these techniques, which can easily be adapted to a wide range of upstream and downstream problems. Although decision analysis techniques have been around for a long time, better software is now available to facilitate the application of these methods and to tie them into commonly-used spreadsheets for reserves estimation.

INTRODUCTION

Reserves estimation and production forecasting are full of uncertainty, and this carries through to the process of deciding how to develop a reservoir. Most oil companies these days use some form of probabilistic methods to try to address these uncertainties, but how do you decide on the best development strategy, whilst still keeping an open mind on the assumptions behind the probabilistic reserves numbers? Decision analysis tools can help you to develop the optimum strategy by handling the probabilities right through the decision making process.

Decision analysis and risk assessment methodologies have been used in the oil industry for several decades, although this has usually been for ranking exploration prospects, rather than in appraisal and development drilling. Nevertheless, in the last few years, the number of companies using the techniques to address a whole range of problems is growing rapidly.

Decision analysis is more than just building decision trees. Indeed, if the technique is to be used to its full potential, this can be considered a minor part of the analysis. It is a process for structuring a problem so that by handling the uncertainty more effectively, the basis for the decision can be better understood. Analysing the decision allows the input values and biases to be reviewed so that their impact on the decision policy is revealed. Alternative decisions can then be investigated in an iterative process. This approach also has the advantage of providing a framework for discussion, facilitating interaction between different professional disciplines, as well as between technical staff and management.

To illustrate some of the principles and methods of decision analysis, we present a case study which is somewhat hypothetical, but broadly based on recent work we have done to advise an operating company in the UK. The details of the study have obviously been changed to maintain the confidentiality of the client. The problem has also been simplified to some extent in order to make it easier to explain.

CASE STUDY

Our case study is an undrilled gas prospect in the North Sea which is scheduled for appraisal this year. The objective of the decision analysis was to find the appraisal drilling strategy which will maximise the proved reserves with the two wells budgeted. The decisions to be made were where to drill the first well, and given that result, where should the second well go?

This sounds as if it should be fairly straightforward: put the first well in the most favourable location, and once the result is known, optimise the location of the second well. However, if we examine the uncertainties in the prospect, we will find that it is not quite that simple. The prospect is a structural trap, reasonably well defined by seismic, which appears to be compartmentalised into four major blocks by faults, Fig 1. There are similar producing fields nearby and if the reservoir rock is present and it is gas-bearing, there is a very good chance that it will produce commercially.

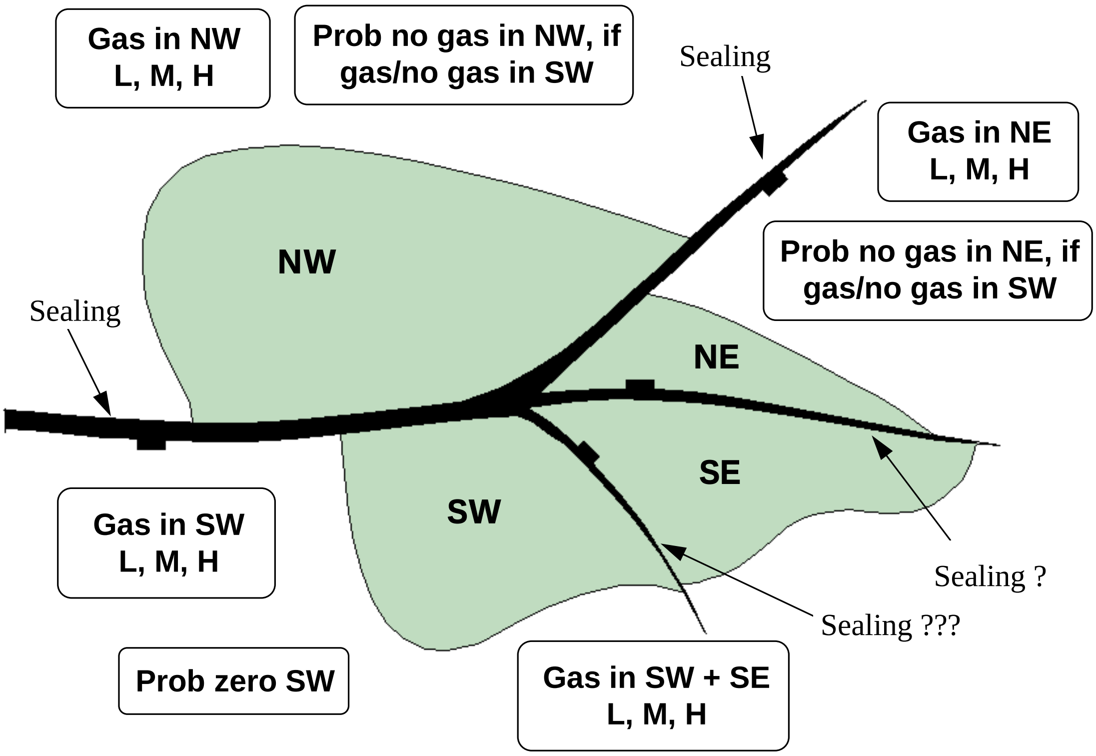

Fig 1 *Diagram of prospect area showing fault blocks*

The northwest (NW) fault block is the largest individual block, and on strong geological evidence, it is assumed to be totally isolated by the major faults running east-west and northeast-southwest. The fault dividing the SW and SE blocks is considered to be less sealing than the fault dividing the NE and SE blocks. Therefore, if the SW and NE fault blocks can be shown to be in pressure communication, all the blocks apart from the NW are assumed to be in communication, and the SW and SE blocks can be treated as a single block.

So the main uncertainties are the probabilities of gas reserves being present, and the communication between the SW and NE blocks.

A decision had already been made to drill the first well in one of the larger blocks, either the NW or the SW. What complicates the decision is that, not only does the result of the first well determine where the second well should be drilled, but that the potential result of the second well influences the decision of where to locate the first well. After the first well is drilled, the results may alter the probabilities and the values in the undrilled blocks.

Why is this so? Simply, because the first well will only provide limited data on the prospect - considerable uncertainty will remain. Remember that the objective is to maximise the proved reserves after the two wells. Any blocks which are undrilled, and not proved to be in communication, will not count towards our reserves' total. Any information which the wells provide on the fault sealing may, therefore, be critical.

The Decision

Our initial estimate of the distribution of probable gas reserves is shown in Table 1. These are derived from a Monte Carlo simulation in the usual way, sampling from distributions of porosity, net thickness, etc, to generate a probability distribution of gas-in-place. Discrete values at the 10, 50 and 90 percentiles have then been extracted (as per the oil company's policy) to give low, medium and high volumes. The probability of these volumes being present is also shown. You can then calculate the initial probabilistic reserves expected in each block, which is illustrated in the pie chart, Fig 2.

Block	Initial Probable Reserves, (MMSTBOE)			Probability	Initial Probable Reserves
	Low	Med	High		
NW	3	15	40	0.55	10.0
SW	5	10	20	0.50	5.6
NE	3	6	10	0.55	3.4
SE	5	10	15	0.50	5.0

Table 1 *Initial Understanding*

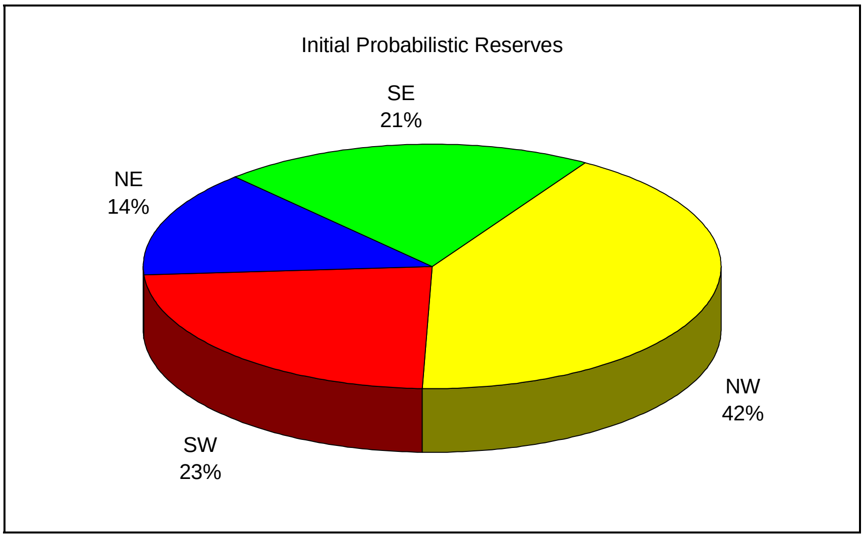

Fig 2 *Initial Probabilistic Reserves*

Clearly, the NW block contains the largest expected volume. So it seems fairly logical that we would want to drill it, probably as our first well. What are our options? These are listed in Table 2, but note that the company has set a condition that the first well must be in the NW or SW blocks.

Option	Action	Potential Knowledge Gain	Expected Proved Gas
1	Drill NW, then drill SW	Prove reserves in NW & SW, change probable reserves in NE & SE	15.6
2	Drill SW, then drill NW	Prove reserves in NW & SW, change probable reserves in NE & SE	15.6
3	Drill NW, then drill SE	Prove reserves in NW & SE, change probable reserves in NE & SW	15.0
4	Drill SW, then drill NE	Prove reserves in SW & NE, change probable reserves in NW, may prove reserves in SE	14.0
5	Drill NW, then drill NE	Prove reserves in NW & NE, change probable reserves in SW & SE	13.5
6	Drill SW, then drill SE	Prove reserves in SW & SE, change probable reserves in NW & NE	10.6

Table 2 *Options*

We can eliminate some of these options with a little thought. The right hand column gives us the total expected value for the drilled blocks (ignoring for the moment that the blocks are not independent). Options 5 and 6 can obviously be eliminated, so then we are left with the following options:

- drill the NW first and then drill either the SW or SE, or
- drill the SW first and then drill NW or NE.

At first glance, it does seem as though our best option would be to drill in the NW first. However, as we have already suggested, the blocks are not independent, and the probabilities used to calculate the expected values are linked. One of the complications is the communication between the SW and NE blocks which depends on the fault sealing. It is not too difficult to convince oneself, therefore, that it is usually better to drill the SW block first, since you still have the option of drilling the NW (Option 2), or you can try to solve the communication unknown by drilling the NE (Option 4).

We are starting to see some of the complexity of the problem, which is largely due to the way in which the assumptions are linked. As an example, if we drill the SW block and discover gas, the probability of there being gas in the other blocks goes up. However, because of the variable fault sealing, the change in the probabilities is different in each block, so that there is a larger increase in the probability of gas in the NE and SE, than in the isolated NW block.

If we drill the SW block first, the result will influence the expected value which we calculate for the second well options. These values are shown in Table 3. Note how the expected values vary according to the outcome of the first well and that now the highest expected volumes come from Option 4.

Option		Zero in SW	Low in SW	Med/high in SW
2	Drill SW, then NW	9.125	15.95	24.28
4	Drill SW, then NE (& SE)	1.875	10.0 (17.5)	18.33 (25.83)
6	Drill SW, then SE	2.500	12.5	20.83

Table 3 *Expected reserves for second well options, when accounting for dependencies*

To account for the dependencies in a more rigorous way, we will need to do a full analysis using the various decision analysis techniques which are available. Note, however, that in our example, we will not be able to give a full explanation of the various steps, since this would take several hours. The aim of this paper is to give you some idea of the complexity of the effects of dependency and how the techniques can be used to handle this.

Influence Diagram

We started the study by holding a brainstorming session with the oil company appraisal team, bringing together key personnel from different professional backgrounds to identify the critical issues which affect the decision.

A useful tool to help facilitate this session is to construct an influence diagram. This has been constructed in two parts, Figs 3 and 4. The issues which have been identified can then be classified as decision (▭), chance (◯), or value (▢) nodes. This is a method for laying out the problem which is easy for everyone to see and understand. Decisions and uncertainties can be discussed openly, with the whole team contributing their views. It is effective because it can be used to help divide the uncertainties into smaller factors so that experience/intuition can function effectively. If the uncertainties are too complex, the probabilities will include factors which are beyond individuals' experience.

The influence diagram is also used to consider how the uncertainties might be related. If new data are obtained to resolve one of the uncertainties, will this affect the uncertainty of other factors?

The value nodes can be either input parameters or variables to be calculated as the outcome of the process. In our example, the input parameters are laid out adjacent to the main area of the influence diagram, Fig 3, and are either values or probabilities.

The input values are the volumes of gas estimated in each fault block. On Figure 3, these are shown as high, medium, low and zero gas value nodes for the NW, SW and NE fault blocks. There are two nodes for each SW case, depending on the fault communication. The probability value nodes shown on Figure 3 represent the uncertainty about the presence or absence of gas in each fault block, and also the uncertainty of communication between the SW and NE blocks.

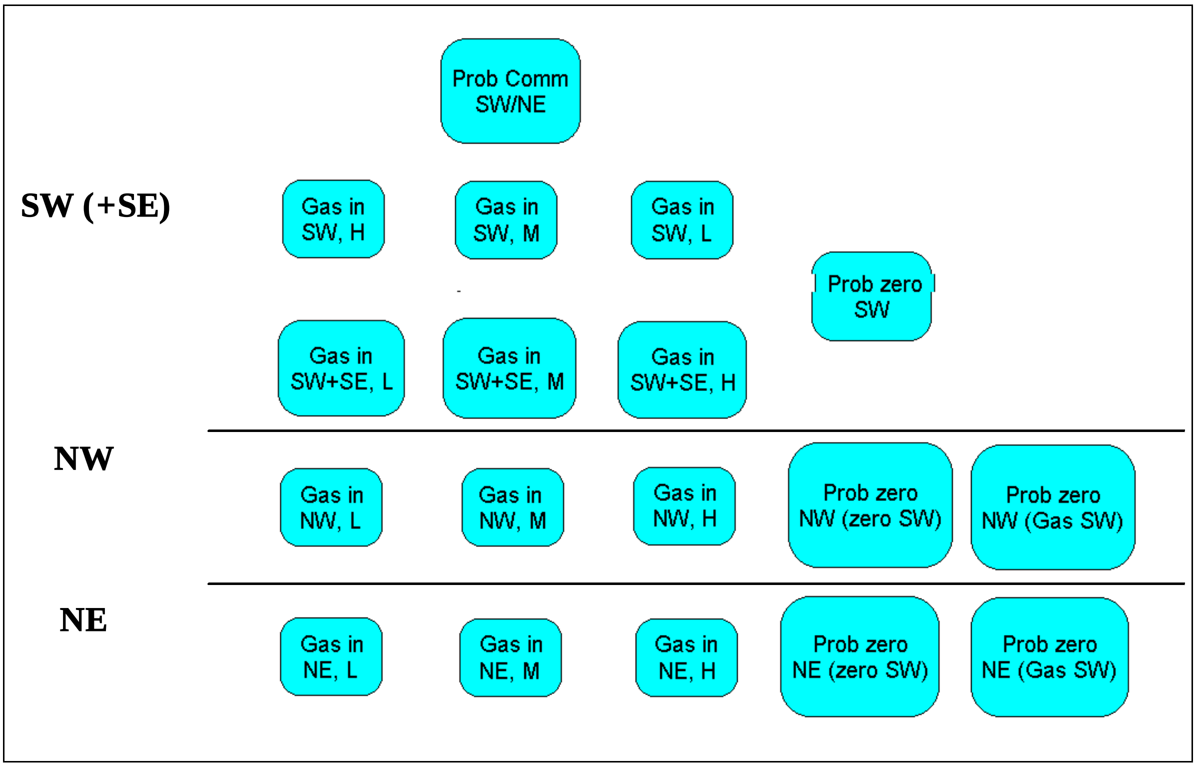

Fig 3 *Influence Diagram 1 - Value Nodes*

Each of the above value nodes has a single number associated with it. This is either:

i) a probability of a particular event, between 0 and 1, e.g. the node "Prob zero SW", representing the probability that there is no gas in the SW block has the value 0.5, or

ii) the value associated with a particular event, e.g. the value of gas defined as being "Low" in the NW is 10 MMSTBOE.

The numbers represented here are the numerical input to the problem.

The chance nodes shown on Figure 4 represent the uncertainty about the volume of gas reserves in each fault block, the **proven** gas reserves in each block, the uncertainty of communication between the SW and NE blocks, and the result of Well 1. In our example, the gas reserves are not the same as the proven gas reserves, since the latter must be confirmed by drilling.

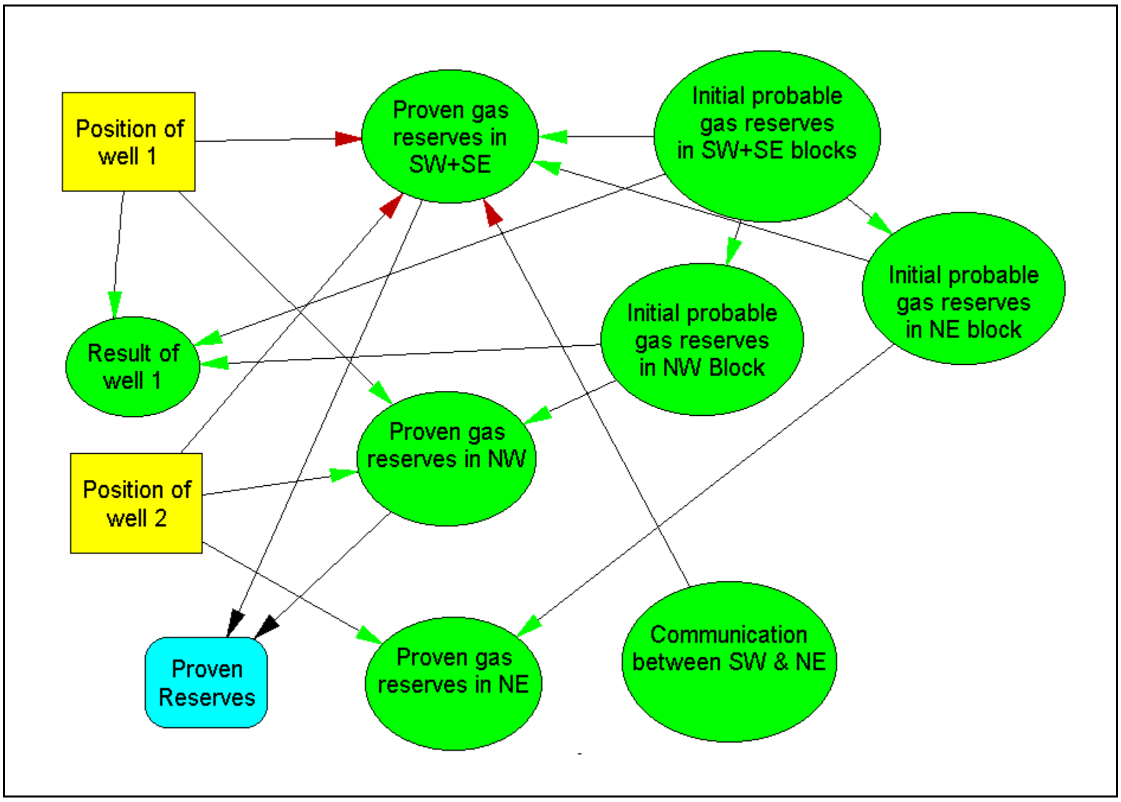

Fig 4 *Influence Diagram 2 - Chance and Decision Nodes*

The influence diagram shows the various factors affecting the problem. Chance nodes (e.g. "Result of well 1") are considered variables, with associated probabilities and values for different outcomes (see Figure 5 for an example). The arrows shown represent influences which lead to different probabilities and/or values, depending on the outcome of the influencing event. E.g. the presence or otherwise of gas in the SW block (represented here by "Initial probable gas reserves in SW block") directly affects the probability of gas in the NE block (represented here by "Initial probable gas reserves in SW block"), and so an arrow indicates this influence. The arrangement of the influence arrows is not necessarily unique, as indirect influences can be represented in different ways. For instance, clearly the "Result of well 1" will influence the "Proven gas reserves in SW", but there is no arrow on the diagram. This is because this influence is catered for indirectly by other arrows. Describing the influences in a consistent and correct manner is a non-trivial task, requiring many iterations, and we do not suggest that the reader attempts to understand all of the influences shown above.

These are obviously judgements, based on the experience of the professionals familiar with the geology and the producing characteristics of the neighbouring fields. Later on, sensitivity cases can be run to see how critical these judgements are to the decision strategy. The last nodes to add to the influence diagram, Fig 4, are the decision nodes (where to drill) and the end result (in this case, the total proved reserves).

Finally, the influence arrows can be added, to show the dependencies between the various factors. This is often quite difficult to get right, and thinking through the dependencies will, in itself, often help considerably in the understanding of a decision problem. A straightforward example of a dependency, see Figure 4, is that the volume of probable gas in the NW block depends on the probable volume of gas in the SW. Figure 5 is a probability tree which illustrates this relationship in more detail.

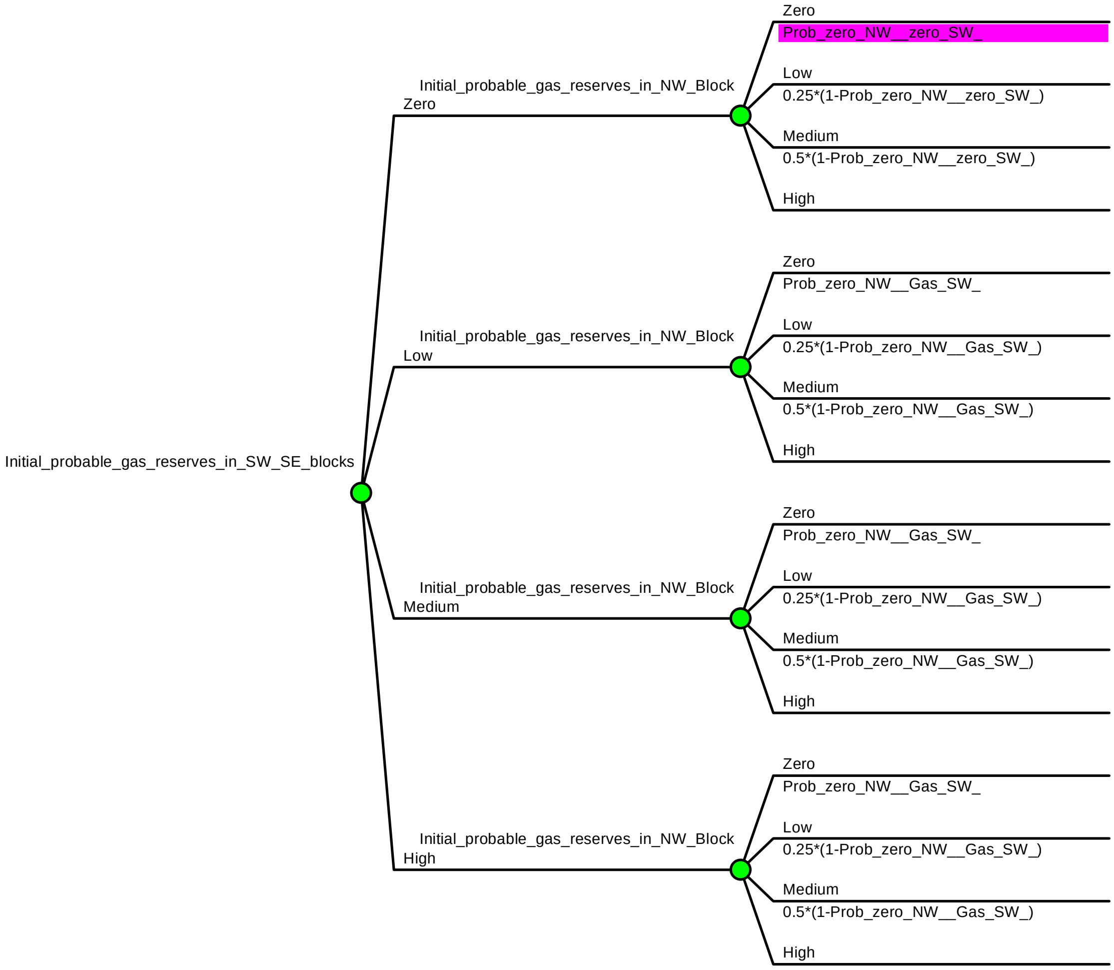

Fig 5 *Probability Tree*

This figure shows the data for probabilities stored "behind" the chance node "Initial probable gas reserves in NW block". There are four possible values for this node - Zero, Low, Medium and High. The probabilities are dependent on the node "Initial probable gas reserves in the SW block" and so the four cases for this node (also Zero, Low, Medium and High) lead to four separate branches on the probability tree shown above. For each of the four branches, the probabilities must sum to one. In this instance, we cause the software to automatically calculate the probabilities for the High branches by leaving these blank. The other probabilities are entered using the variables set up in the value nodes (see Figure 3).

Decision Tree

The next step in structuring the decision is to consider the timing of events, and to do this, we constructed a decision tree, Fig 6. Obviously, the first decision is where to drill the first well. The choices are the NW block, the SW block, or don't drill at all (this might be a relevant option if the economics of drilling were to be considered as part of the analysis). The result of drilling is entered for the possible locations as zero, low, or medium/high proved gas volume proved, as defined on the value nodes of the influence diagram. (Medium and high results are sometimes combined for geological reasons specific to this study).

We then construct the tree to consider the location of the second well. The only restriction this time is that we do not drill the same block as before (which results in an asymmetrical tree). Once the tree has been constructed in this way for all the relevant branches, the values and probabilities can be incorporated, with the end nodes being the proved gas reserves for each block. Note that the tree shown in Figure 6 is a compressed version; the full number of end nodes is actually 2311, although this is very hard to display!

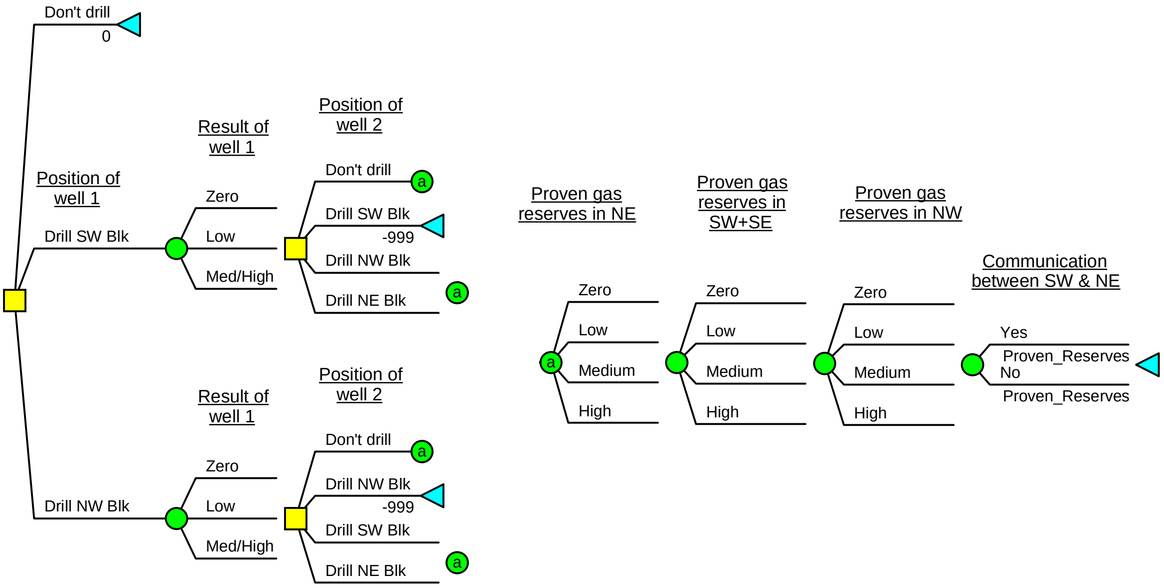

Fig 6 *Decision Tree*

The decision tree arranges the nodes from the influence diagram in the order that they occur. Certain routes down the tree may not be permissible, for instance drilling the same block twice. Here this is handled by stopping calculation of the tree at those points, leading to an asymmetric tree. In the figure, each route from left to right through the tree is calculated. In order to aid efficient representation of the data, branches which are repeated in different parts of the tree do not need to be shown twice. At the most simple level, this is illustrated by all four outcomes (Zero, Low, Medium, High) from "Proven gas in NE", all leading into the same node "Proven gas reserves in the SW+SE". When the same group of nodes are repeated, this is also simplified, by using alphabetic characters to indicate links. In the tree above, the node "Proven gas in NE" has the letter "a" at its left hand end. Wherever the letter "a" appears on the outcome at the right hand end of a node (for instance "Position of well 2", "Don't drill"), this outcome is followed by the node "Proven gas in NE".

To illustrate the effect of dependency, consider the following case. If the SW block is drilled first and finds no gas (at a 50% probability), the probability of finding no gas in the NE block increases to 70%, whilst the chance of finding no gas in the NW block (which is isolated) increases to 60%. If, instead, the first well finds a low gas volume in the SW, the probability of there being no gas in the NW drops to 40%, but in the (possibly connected) NE block, the probability of no gas drops to 20%. Note that, although the gas volume probabilities are changed according to the results of the first well, the probability of the SW and NE blocks being in communication stays the same. Until both the SW and NE blocks are drilled, there is no new information to change this probability.

In some cases, events may take place which will affect the probabilities of other events. As an example, if there is no gas in the SW block, the probability of zero gas being found in the NW has been assessed and entered in the influence diagram (Figs 3 and 4). But the reverse is also true: the probability of gas being present in the SW block is affected by the presence or absence of gas in the NW. Therefore, in the case where the NW block is drilled first (the lower branch on the decision tree, Fig 6), the probability for the presence of gas in the SW block needs to be revised to account for the probability of zero gas in the NW. This is done using Bayes' Rule and can, in more complicated situations, be a fairly tedious process involving reversal of the order of nodes to manipulate the probabilities. Thankfully, this can usually be handled automatically by software.

Once its construction is completed, we can back-calculate through the decision tree, to obtain the expected value of proved reserves for any decision policy. The optimum decision is highlighted (bold line) as the one which will give the highest expected value, Fig 7.

The Decision Policy

The optimum strategy appears to be to drill the SW block first and then, depending on the result, drill the NW or NE blocks with the second well.

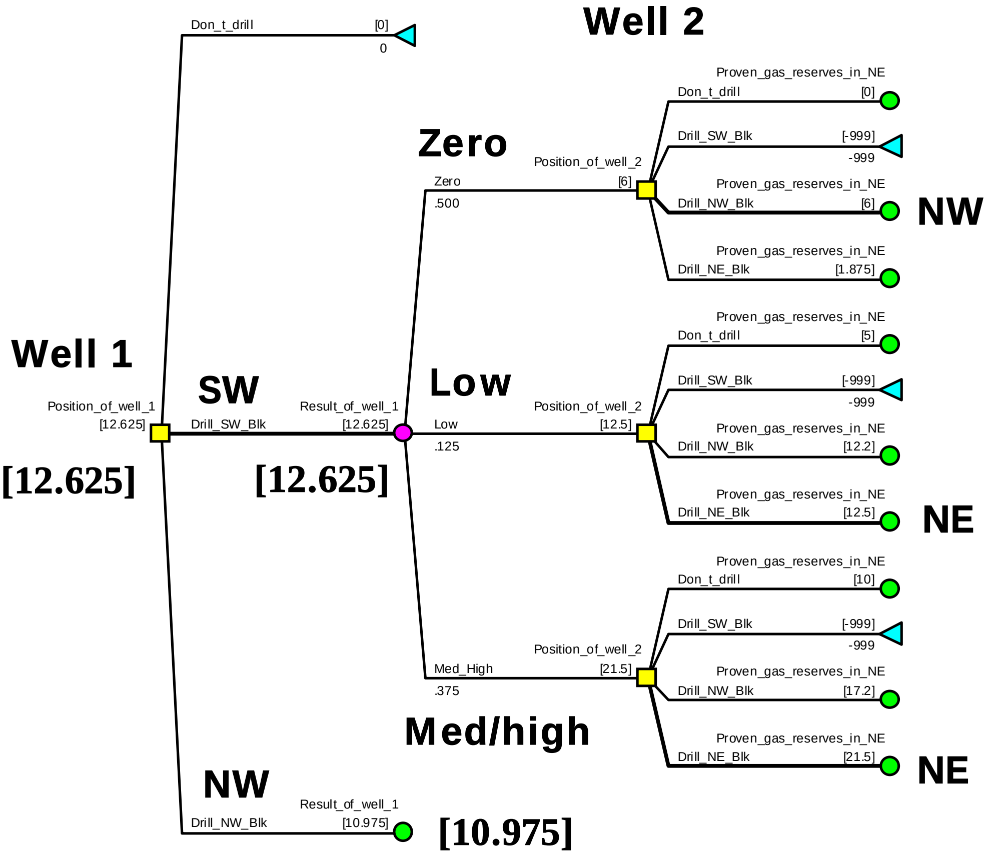

Fig 7 *Optimum Decision Policy*

The previous figure represents the summary decision tree once the software has calculated the various probabilities and values. The solid black lines represent the optimum decision policy, i.e. the policy to follow to maximise the expected value. The numbers immediately to the left of a node (whether a square decision, oval chance or triangular value node) show the expected value at that point in the tree. So here the overall expected value, before any drilling is carried out is 12.6251. On drilling the first well, if the result is zero, the expected value is 6, whilst if the result is low the expected value would then be 12.5. The numbers to the right of oval chance nodes represent the probability of that particular event. In the above figure, the only such chance node is the "Result of well 1" node, for the which the probabilities are 0.5, 0.125 and 0.375 respectively for the three outcomes Zero, Low and Medium/high, when the first well is drilled in the SW block.

It is important to stress that this is no "black box" which cranks out a ready-made decision for you. The whole point of the analysis is that you are trying to understand the basis for any decision, and so the whole process is generally iterative. Misunderstandings of the problem are often revealed by examining which factors control the decision policy. Further sensitivity analyses will help in this examination.

Sensitivity Analysis

Clearly, the decision policy for the second well, given a low gas volume result in the SW block, is somewhat problematical. Therefore, it is worth examining more closely the uncertainty factors which control this decision. By constructing a tornado diagram, Fig 8, we can identify which of the variables has the greatest impact on the EV. The diagram in this case shows that the greatest uncertainty is associated with the probability of communication between the SW and NE blocks.

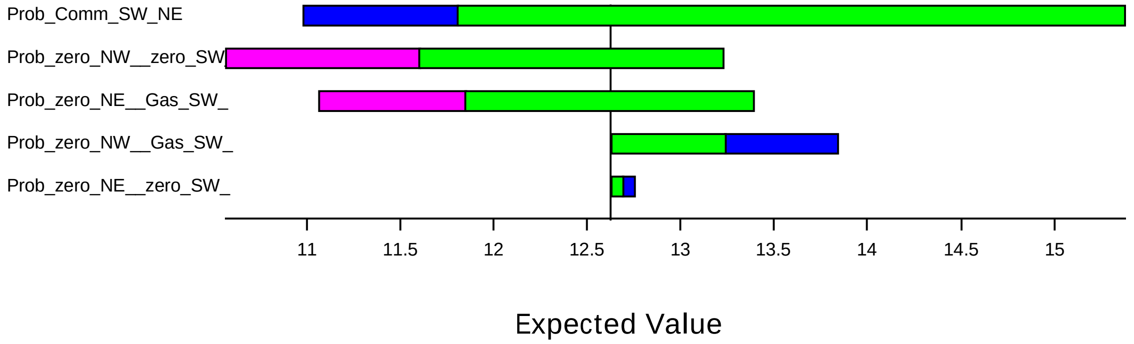

Fig 8 *Tornado Diagram*

The tornado diagram has as its x-axis the expected value. On the y axis are shown different variables (in this case probabilities of different events). Each bar represents the range of the expected value if the variable is varied between user-specified values. So for instance, the impact of varying the probability of communication between the SW and NE can be varied from 0 to 1 (as has been done here) revealing that this gives a range of expected value from 11 to 15.5. The vertical line represents the nominal expected value, and a change in shading between this nominal value and either end-point indicates a change in the optimum decision policy.

In order to investigate this further, we can use a rainbow diagram to plot, Fig 9, the EV as a function of the variable, and to show the impact of the variable on the decision policy. The x-axis is the value of the variable, in this case the probability of SW/NE communication. The y-axis indicates the EV for the optimum decision policy and the different shaded regions represent changes in optimum decision policy.

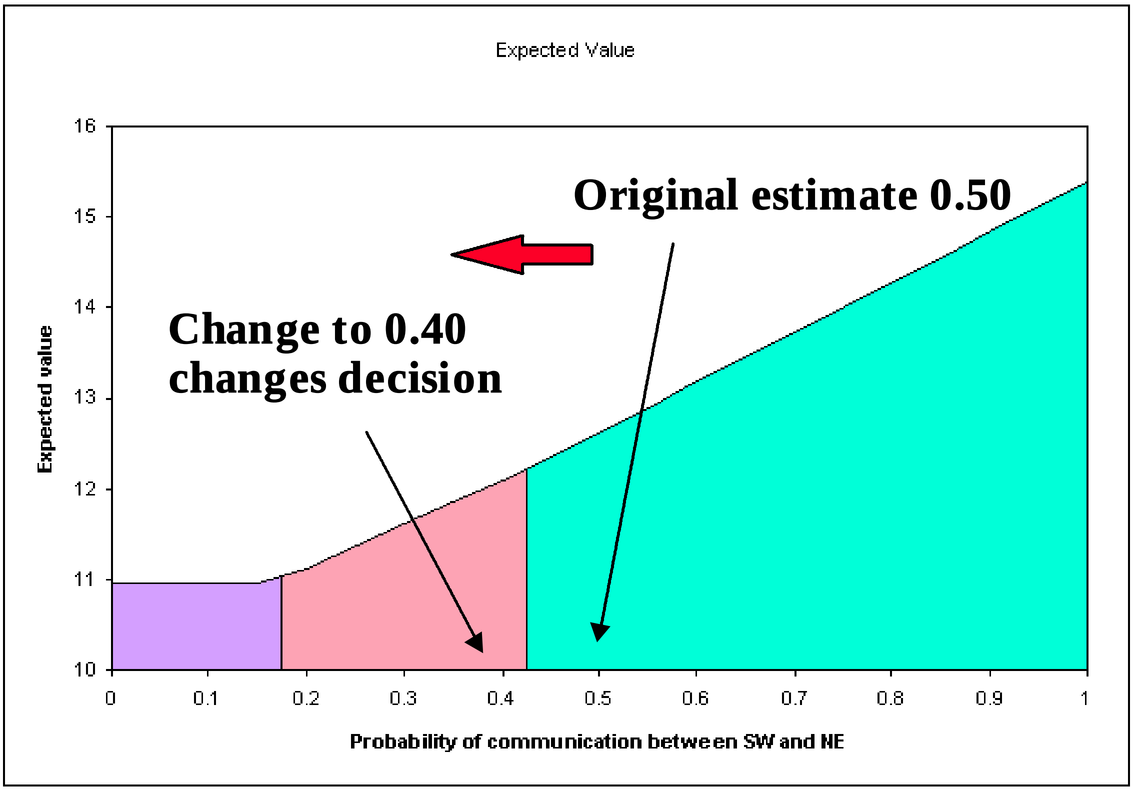

Fig 9 *Rainbow Diagram*

For our base case decision, the input estimate for the probability of communication was 50%. The figure shows that if this probability is lowered to 40%, the decision policy would change. If we were to rerun the analysis with this lower estimate, the optimum location for the second well (given a low gas volume in the SW) would switch from the NE to the NW block. The other change in decision policy indicated in Figure 9 is where the probability of SW/NE communication is so low (< 15%), that whatever the result of a SW well, the other well should always go in the NW block.

As we have already seen, the optimum location for the first well to maximise the EV is the SW block. Figure 10 shows the cumulative probabilities for different decision policies. Note that at various points the lines cross, indicating where the optimum decision would change, if a certain volume of proved reserves was needed.

Clearly, examining the sensitivity of the decisions to the input values is very important, and this will lead to many iterations as the analysis proceeds. These insights give the professionals on the team the opportunity to investigate and reconsider their assumptions, and the effect that they have on the appraisal strategy.

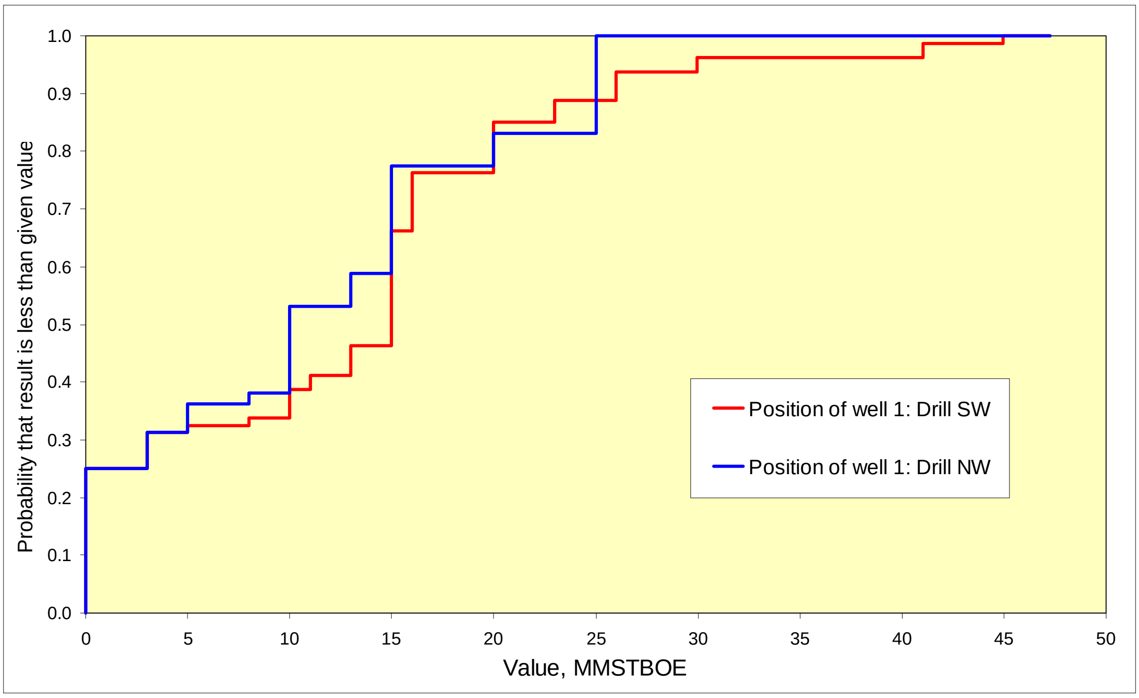

Fig 10 *Value of Policy Decisions for Well 1*

Software

As mentioned in the introduction, risk assessment and decision support techniques have been around for several decades. A development drilling example is given in Newendorp's classic textbook published in 1975. What has changed in recent years is that the techniques have become more accessible, due to the availability of easy-to-use, PC-based software.

A search on the internet for software in this area using Yahoo! found about 50 entries at the site ***http://www.yahoo.com/Business_and_Economy/Companies/Computers/Software/Business/Decision_Risk_Analysis***. Many of these companies offer commercially available decision analysis tools.

Some of the software products will link directly to the spreadsheets in which many companies generate their reserves numbers. Since some of the commercial Monte Carlo simulation programs also link into the spreadsheets, you are able to link the decision analysis directly back to the probabilistic reserves generation, and examine how changes to the Monte Carlo input might affect your decision policy.

CONCLUSION

Reserves estimates are inherently very uncertain. A probabilistic approach to reserves estimation is a positive step to recognise this uncertainty and to attempt to quantify it. However, when decisions on appraisal and development options need to be made, it is important to remember where the numbers came from, so that the underlying assumptions can still be questioned and re-examined.

Decision analysis provides us with the means to put together all these factors to come up with an optimum strategy and examine how sensitive that strategy is to the uncertainties. Finally, it can help us with the important task of justifying the decision to management and partners.

What we hope to have illustrated, is that there are well established techniques available for decision analysis which can help you examine your decisions in more detail, and handle the probabilistic numbers in a relatively simple, easy-to-understand way. We have tried to explain, without recourse to too much jargon, how some of these techniques have been applied to a recent appraisal study.

Obviously, decision analysis methods can be applied to a wide range of problems, and we are aware of their use in such diverse issues as well design, gas contracts, and platform abandonment. We hope that you may have been stimulated to explore the use of these methods for yourselves.

REFERENCE

Decision Analysis for Petroleum Exploration, Paul D Newendorp, PennWell Books, 1975

ACKNOWLEDGEMENTS

The authors would like to thank Laurence Wickens and Bob Foulser, AEA Technology, for reviewing this paper and making some useful suggestions to improve it. However, the ideas and statements made in this paper are the responsibility of the authors alone.

Contacts:

Tim Allen
Principal Reservoir Geologist

Laurence Wickens
Decision Risk Management Consultant

AEA Technology plc
Winfrith
Dorchester
Dorset DT2 8DH
Telephone 01305-203610
Facsimile 01305-202110
e-mail tim.allen@aeat.co.uk
laurence.wickens@aeat.co.uk

Lawrence Roberts
Principal Reservoir Engineer

AEA Technology plc
Exploration House,
Offshore Technology Park, Bridge of Don
Aberdeen AB23 8GX
Telephone 01224-254207
Facsimile 01224-254201
e-mail lawrence.roberts@aeat.co.uk