Workshop on analytic optimisation methods

For business and field development planning

September 2000

Prepared by:

Peter Cunningham Serafim Ltd
Mark Cook Tracs

SERAFIM

info@serafimltd.com
P. +44 (0)2890 421106
www.serafimltd.com

Contents

Contents i

Foreword 1

Outline of the Workshop 2

Problem 1 - Injection water allocation 3

1.1 An answer to the fixed allocation problem 3

1.2 Two other capacity allocation problems 4

Replacing a watered-out well 5

Shutting in wells - capacity allocation varying with time 6

Problem 2 - Gas Field Development Planning 8

Background to the discovery 8

Volumetrics 8

Nearest infrastructure 9

Your overall task 9

Task 5 Well productivity versus cumulative production 12

2.1 - Influence diagram 14

2.2 Volumes of hydrocarbons in place 14

2.3 Material balance and the P/z plot 15

2.4 Inflow performance and tubing performance curves 16

2.5 Well productivity versus cumulative production 16

2.6 Production profile and well requirement 17

2.7 Interpretation of the above result 17

2.8 Reviewing the assumptions 18

Constructing mathematical models 19

Problem 3 - Optimising the producer:injector ratio 21

3.1 Influence diagram for the factor affecting the producer:injector ratio 21

3.2 Simplified mathematical description 22

3.3 Objectives 23

3.4 Achieving a target reservoir pressure 23

3.5 Strategies for calculating the optimal ratios 24

3.6 Expressing as a function of the producer injector ratio 25

3.7 Calculating the producer:injector ratio that maximises production 26
3.8 Calculating the economically optimal producer:injector ratio 28
Problem 4 - What is the optimal number of wells to drill on a field? 30
4.1 The effects of varying the number of wells drilled 30
4.2 Quantifying the effects of well numbers on the production profile .32
4.3 Simplifying economics 34
4.4 Finding the number of wells which maximises NPV 35
Analytically 35
In Excel 36
4.5 Incorporating value:investment hurdles 37
4.6 Results from gas field example 38
Problem 5 - Project optimisation under a fixed production limit 39
5.1 Effects of minimising project UTC 42
5.2 A general rule for optimising field development plans under a production limit 43
Using simple models for business decisions..44
Problem 6 - Optimising water injection on Catbert Alpha 45
6.1 Water injection optimisation 46
Forecasting infill well performance in a mature field 48
Non-exponential decline 53
Initial assumptions...54
Solving for cummulative oil 56
Appendix 1 – Mathematical proofs of the formulas in Problem 4 58
Assumptions and Definitions 58
Theorem 1 – Expressing NPV as a function of well numbers 59
Proof 59
Theorem 2 – Number of wells giving the highest NPV 60
Proof of Part I 61
Proof of Theorem 2, Part II 63
Theorem 3 – NPV for a development with the approximately optimal number of wells 65
Appendix 2 - Mathematical treatment of shadow NPV 66
Shadow NPV theorem 68

Proof 68

Foreword

Imagine that you were planning to take a walking holiday on the West Coast of Scotland and that your only source of information was a complete set of detailed 1:25,000 maps of the whole of Britain. You would probably find the task of deciding where to go quite laborious, involving lots of folding and unfolding of different maps, each only covering a small area. At the end of this, your choice of itinery would be reasonable, but not necessarily the best. You would be able to choose a mountaineous area, but it would be difficult to judge purely from maps whether the Cuillens in Skye, for example, were more suited to your tastes in walking than the Trossachs. It would also be difficult to plan how to get to the West Coast.

If, on the other hand, all you had were guidebooks describing walks in several areas, then it would be easy to choose the most suitable area. But, now, it would be difficult to do the actual walking without detailed maps.

In reality, when we are planning our holidays, we all use a combination of guidebooks and maps of different scales - detailed for walking, less detailed for going by road. This workshop argues that such a "maps and guidebooks" approach should be general practice in reservoir management. We should supplement the elaborate, reservoir simulation-based modelling with simpler models that give an overview, help us to generate development options and provide independent gross-error checks on reservoir simulation results.

Outline of the Workshop

The aim of this workshop is to encourage the use of simple mathematical models as a complement to more elaborate conventional models, such as reservoir simulation. Through examples, the workshop gives guidance on how to construct and optimise such simple models.

The first example deals with the question of allocating injection water between two fields. A simple solution is proposed, that serves to link together the results of two separate reservoir simulation models, and helps to overcome communication barriers between reservoir engineering and facilities engineering. This is an example of a capacity allocation problem. Two other forms, taken by such problems, are discussed.

The second example shows how simple methods can go a very long way in gas-field planning. It also illustrates the danger of allowing there to be gaps between commercial calculations and technical calculations.

Having gone through two examples, it is useful to try to describe a process for constructing good, simple models. The process is illustrated by looking at the most influential model in petroleum engineering-reservoir simulation.

The process of constructing mathematical models is described into five stages

1. Observe the problem in the real world.
2. Create a simplified mathematical description.
3. Carry out deductions/calculations.
4. Conclusions of deductions/calculations.
5. Relate the conclusions back to the real world.

The important question of using mathematical optimisation to help design development options is covered.

The importance of saving time and effort by, wherever possible, manipulating the results of existing models is illustrated by the fourth example, which deals with the question of what is the optimal number of wells to drill in a field.

The fifth example aims to illustrate some of the more common tricks that can help in solving a problem. The example deals with optimising NPV when operating under fixed production limits, such as those imposed by gas supply contracts.The acceptability of simple models for business decisions is discussed. This is followed by the last example, a more complicated water injection problem.

The workshop finishes with looking at two aspects of what can be the very important relationship between recovery and number of wells drilled

- Actual recovery data from individual wells in a very large field is examined, to show that simple patterns of behavior can often be seen in actual field data.
- For fields showing non-exponential decline, a more general alternative to the hyperbolic curve is presented. This has been found to give good results for matching both field water-cut development and the field recovery vs number of wells drilled relationship (as predicted by ECLIPSE).

Problem 1 - Injection water allocation

This workshop aims to teach the methods of constructing simple mathematical models by going through a number of examples. Building such models is a skill that is best learnt through experience. Fortunately, we have all had considerable exposure to the construction of mathematical models of reality - this is, after all, the basis of physics and engineering.

The first example illustrates the simplicity of many models.

If you have a fixed amount of water available, what ways are there of calculating the optimal split between two fields?

You are the reservoir engineer putting together a field development plan for Griffith, a small offshore field, with an estimated 6 mmb recoverable reserves. It is planned to develop Griffith as a satellite of Rhame, a much large field that has its own platform with a range of production and injection facilities.

Let us assume that the Rhame platform's injection water capacity is a fixed 200,000 b/day and that the injection water from Griffith has to come out of this 200,000 b/day, and that the allocation, once determined, is fixed. (Let us imagine, for example that the value of Griffith is not sufficient to pay for an upgrade to the injection system on Rhame). There is ample capacity in the Rhame production facilities.

You have gone to the Rhame facilities engineer and asked him how much injection water would be available for Griffith. His initial answer was "How much do you want?"

Your task is to devise a method of calculating the optimal injection water allocation for Griffith.

It is suggested that you consider two important questions -

- how do the economics of the two fields interact?
- how can the problem be expressed in terms of a single decision variable?

Note - a "decision variable" is a variable that describes the decision. For example, if the decision was the location of a new well, the decision variables might be the x and y co-ordinates of the well bottom-hole location. Alternatively, it might be possible to reduce the question to the simpler one of "What is the optimal outstep from the existing well?", in which case the decision variable might be the distance, d, from the existing well.

1.1 An answer to the fixed allocation problem

It is likely that the tax and royalty regime acts in such a way that, to a good approximation, the NPV of the project can be expressed in a simple linear fashion as

NPV = PV(Rhame production) + PV(Griffith production) - PV(Costs)

In real problems, it is important to get confirmation about this from the petroleum economist. The confirmation can either come from direct knowledge of the way the various taxes and royalties interact, or it can come from running the economic model for a variety of different scenarios, and seeing whether the results match the equation above.

The next step is to express all these terms in the economics equation above as functions of the Griffith injection rate - a straightforward task with the individual field reservoir simulation models.

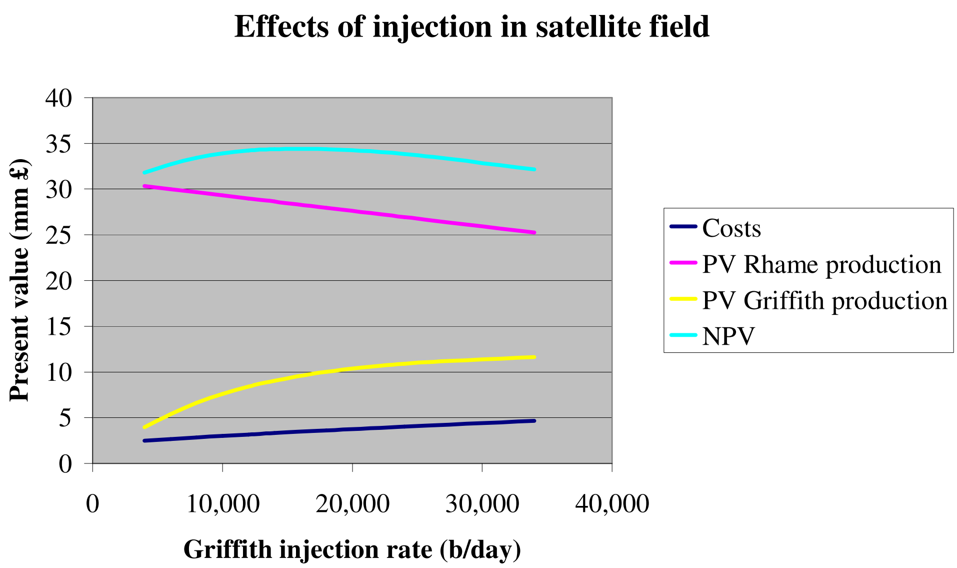

It is useful to note that for the larger field, Rhame, it probably suffices to make three reservoir simulation runs. Since the percentage changes in water injection in Rhame are relatively small, the effect of the changes is likely to be approximately linear. If this is the case, then all one needs do is run the Rhame reservoir simulation model with

- the extreme low value for the Griffith injection rate
- the extreme high value for the Griffith injection rate and
- a value in between.

This will probably be sufficient to establish the expected near-linear behaviour in Rhame.

Combining the three component PVs gives total NPV as a function of the Griffith injection rate. It is then easy to choose the injection rate that gives the maximum NPV.

This problem may seem very simple, but it illustrates the importance of deriving the correct framework for efficient communication between different disciplines. Otherwise, we are back at the "How much water can I have?" "How much water do you want?" stage.

1.2 Two other capacity allocation problems

The example above is deliberately simple and so deals with the possibly slightly artificial situation in which the allocation of capacity is fixed, and cannot be changed later in field life. In real life, there are a number of situations in which other capacity allocation problems arise, usually with individual wells. It is interesting to examine briefly two of them -

- Abandoning a watered-out well and using its slot to drill an additional infill well - a one-off replacement
- Given facilities constraints on both oil and gross liquids rates, deciding on which wells to shut in - i.e.capacity allocation that can vary with time.

Replacing a watered-out well

Here, the decision variable is the time at which the old well is abandoned. It is straightforward to calculate the value derived from the old well (the PV of its production up to abandonment) and the value derived from the new well (its project NPV, discounted for the time prior to abandonment of the old well). This is illustrated in the plot below for an old well producing 800 stb/day at time t=0 and declining at 25% per annum, and for a new well with a project NPV of £8 million.

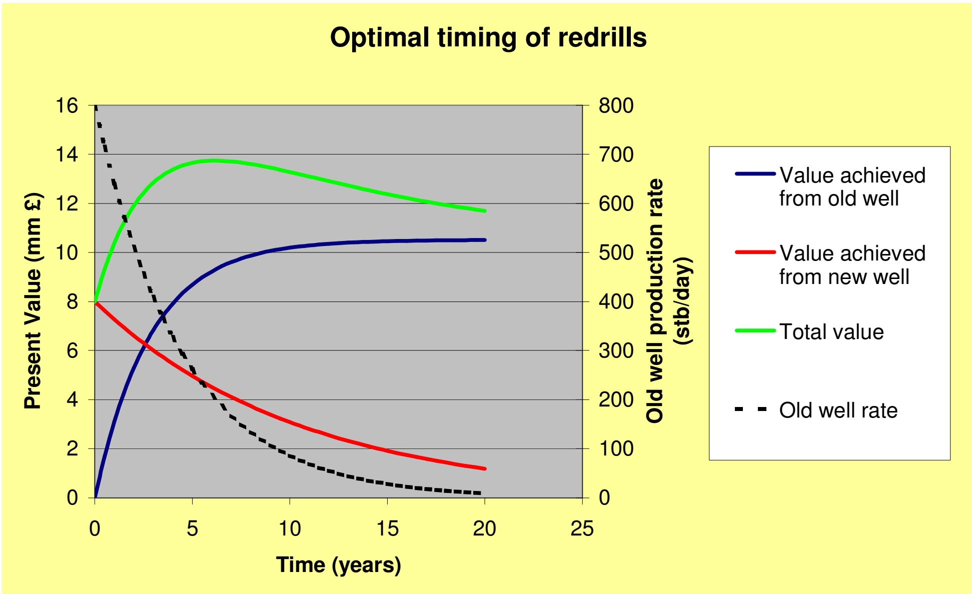

It can be seen that the maximum total value is achieved if the old well is kept in production for a further 6 years, until its production rate drops to 200 stb/day.

The criteria for selecting abandonment time - maximising total value - can be expressed in terms of derivatives, by setting the derivative of total value (with respect to abandonment time) to zero. Such an equation can also be derived by considering revenue (which is the derivative of value).

In this case, the old well should be abandoned when

Net revenue from old well < Cost of delaying new well (i.e. Discount rate x NPV of new well)

Shutting in wells - capacity allocation varying with time

It is often the case that production facilities are constrained and so some wells need to be shut in or choked back. Usually, such settings can be changed, almost on a day by day basis.

If a field is constrained by water or gross liquids handling capacity, then the optimal strategy is very straightforward - open up the wells in order of increasing water-cut until the constraint is hit. This maximises oil production for each moment in time and so maximises field value (assuming that any cost implications are minor).

The problem is more complicated if the field is limited both by both oil and gross liquid facilities constraints. We will argue that the optimal strategy is to maximise both oil and liquid production i.e. to choose the well choke settings so that both constraints are simultaneously hit.

An oil-field can be considered to have an overall life-time that can be expressed in terms of both cumulative oil production and cumulative liquid production. By maximising both oil and liquid production, the field is aged as quickly as possible – in other words, the overall production occurs as soon as possible. Consider a period when the well capacity is above the facilities capacity, both in oil rate and liquid rate. If one maximises both oil and liquid production, then some of the drier wells are shut-in. Their production is then conserved until later, enabling one to prolong the plateau period of maximum oil production.

This argument can be illustrated by an idealised example. Consider a field (similar to the Alba field) in which 30 wells enter into production at time 0 and will give an ultimate recoverable of 250 mm barrels. Let the facility constraints be 400,000 stb/day liquid and 100,000 stb/day oil.

The field unconstrained cumulative oil vs cumulative liquid relationship is illustrated below.

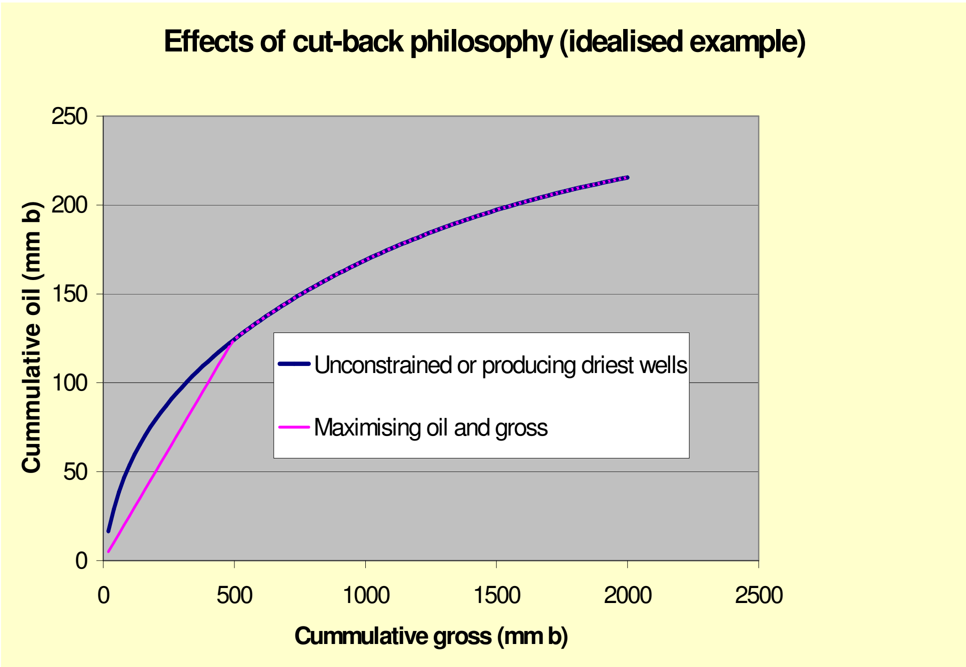

A cut-back strategy of producing the driest wells will lead to all the wells having similar water-cuts (if some are drier, they are produced preferentially until their water-cuts build up. Hence, if the wells are similar, the cumulative oil vs cumulative gross relationship with a driest wells cut-back strategy is very similar to the unconstrained relationship.

A strategy of maximising oil and liquid will, for the period of time in which this is possible, cause the field cumulative oil vs cumulative gross relationship to go in a straight line. This is illustrated in the plot for the extreme case in which it is possible to maximise both oil and liquid from time 0 up to the point of intersection with the unconstrained case[1].

In reality, it would not be possible to maximise both oil and liquid over this period. For example, at the start of production, prior to the build up of water-cut, the oil constraint will dominate. However, the straight line provides a simple extreme case from which it is easy to calculate the corresponding extreme case production profiles and PV oil.

The production profiles from the above cumulative relationships are given in the plot below. It can be seen that the strategy of maximising both oil and liquid rates enables the plateau oil production to be prolonged by more than a year. The total PV oil with this strategy is 5% higher than with the strategy of producing driest wells first.

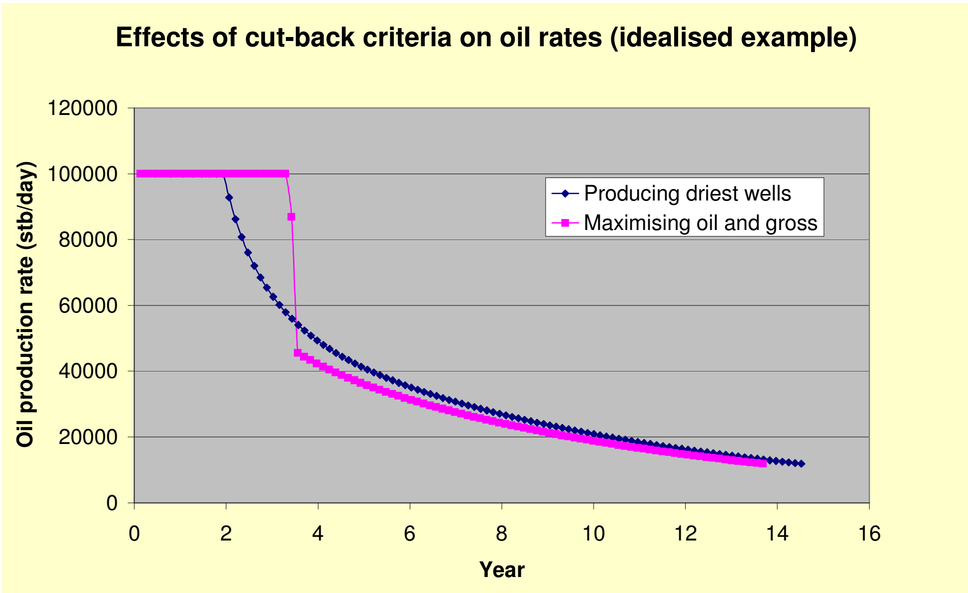

NB - For the strategy of maximising both oil and liquid, there is a sudden drop in oil production. This is an artificial feature that is the result of the artificial assumption made in constructing this extreme case – the assumption that maximum oil and liquid rates can be maintained up to

---

[1] It is possible to demonstrate that with a fixed number of identical wells and with water-cut increasing with cumulative gross, one cannot get above the unconstrained line.

the moment of hitting the unconstrained cumulatives curve. With more realistic assumptions, the difference between the two production profiles would be slightly less, and the difference in PV oil would be less – maybe 4%.

Problem 2 - Gas Field Development Planning

This case study will take you through a procedure for gas field development planning, using the approach of defining the problem as a set of simple relationships, and solving graphically to determine the optimum solution. It is based on a discovery in the Central North Sea in the UK Continental Sector.

Background to the discovery

The discovery well found gas condensate in Palaeocene sandstones at around 10,000 ft tvss. The well was tested with a DST, and an interpretation is available. The sands are low permeability (2 mD average), and the maximum DST flow rate was 10 MMscf/d.

Reservoir fluid is a gas condensate with a dew point of 5490 psig at the reservoir temperature of 280°F. Reservoir pressure is 5500 psig at 10,100 ft tvss, so retrograde condensate behaviour is expected during depletion. The initial condensate gas ratio is 130 bbl/MMscf, which represents a medium-rich gas. PVT data is available for the tested fluid.

Volumetrics

Based on the information from the 2-D seismic shot across the Block, and the well information, the following estimates have been made:

Geology/Petrophysics	P85	P50	P15
Gross Rock Volume (MMm^3)	380	565	1550
Net/Gross	0.3	0.5	0.7
Porosity	0.16	0.19	0.22
Hydrocarbon Saturation	0.25	0.5	0.75
Gas Expansion Factor (sm^3/rm^3)	240	250	260
Dry gas/wet gas shrinkage factor (stm^3/rm^3)		0.9	
Condensate/Gas Ratio (bbl/MMscf)		130	

Nearest infrastructure

The nearest infrastructure is a production platform 20 km away, which has spare capacity to treat up to 100 MMscf/d of wet gas, for export into a pipeline which runs to shore in the UK. Preliminary engineering calculations indicate that the only commercially viable development option would be a subsea tie-back. The minimum tubing head pressure required for natural flow from the wellhead to entry at the first stage separator on the platform is 600 psig.

Your overall task

Using the guidelines provided, create a development plan for the field, including well numbers and a production profile. Combined with well and facilities costs, this would be sufficient information to run the economics of the development. The example will be broken down into a series of stages.

Task 1

Draw an influence diagram of the parameters that will have an impact on the production forecast. Indicate some of the relationships that link these parameters.

Task 2 Volumes of hydrocarbons in place

Using the data provided, estimated the volume of

- wet GIIP
- dry GIIP
- condensate initially in place (CIIP)

Task 3 Material balance and the P/z plot

The following table has been generated, based upon the PVT report available.

Table 1

Component	mol %	Pressure (psig)	Viscosity (cP)	z factor	Bg (rcf/scf)	m(P) $(psig^2/cP)x10^6$	CGR (b/MMc
CO_2	2.80	5500	0.0380	1.0431	0.0040	1280	130
N_2	1.35	5000	0.0354	0.9979	0.0042	1130	115
C1	68.37	4500	0.0327	0.9552		992	100
C2	9.20	4000	0.0299	0.9162		845	85
C3	5.44	3500	0.0271	0.8824		696	70
i-C4	0.82	3000	0.0243	0.8563		551	60
n-C4	2.28	2500	0.0215	0.8413		410	50
i-C5	0.68	2000	0.0190	0.8414		278	45
n-C5	0.97	1500		0.8814		164	40
C6	1.21	1000		0.9517		75	35
C7+	6.87	500		0.9951		19	30

The z factor is the gas deviation factor, and is predicted from equation of state (EOS) simulation by the company performing the PVT study.

m(p) is the gas pseudo pressure which is useful in the linearised form of the transient and semi-steady state inflow equation for gas flow. This is calculated by integration as follows.

$$m(p) = 2\int_{p_b}^{p} \left(\frac{p}{\mu_g Z} \right) dp$$

This effectively lumps the viscosity, compressibility and non-ideal gas behaviour terms into a pseudoised "pressure" term and allows the radial diffusivity equation for gas to be linearised and solved. The resulting equations for transient and semi-steady state flow are:

m(p) -m(p_{wf}) = [1422 QT/kh]* {ln (re/rw) - 0.75 +s} **[equation A]**

where

p	=	average reservoir pressure	(psig)	
p_{wf}	=	wellbore flowing pressure	(psig)	
Q	=	gas flowrate	(Mscf/d)	BEWARE UNIT!
T	=	reservoir temperature	(°R = 460+°F)	
kh	=	permeability*thickness	(mD.ft)	
s	=	skin		

The material balance equation for gas can be reduced to

$$\frac{P}{z} = \frac{P_i}{z_i}\left[1 - \frac{G_p}{G}\right]$$

Where

P	=	current pressure	(psig)
Z	=	gas deviation factor	
P_i	=	initial reservoir pressure	(psig)
G_p	=	cumulative gas production	(scf)
G	=	GIIP	(scf)

Complete the following plot of P and P/z against Gp/G (which is of course the recovery factor, RF)

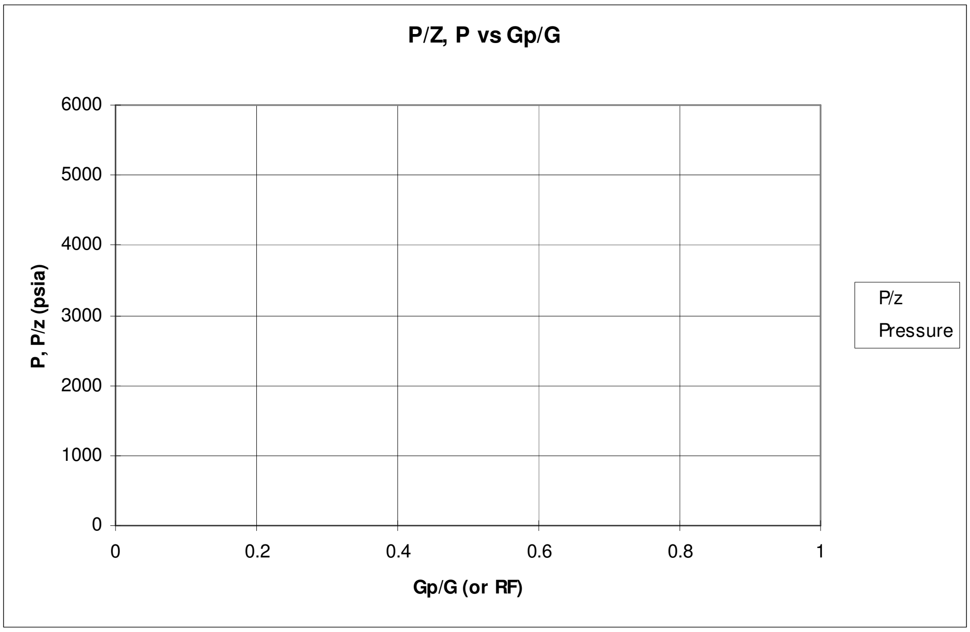

Task 4 Review of Inflow Performance Relationship and Tubing Performance Curves

Inflow Performance Relationship (IPR)

Equation A and the psuedo pressure vs pressure relationship in Table 1 can be used to create the IPR for various average reservoir pressures. This has been done for you in the plot below, assuming

T	=	280°F = 740°R
S	=	0
ln(re/rw)	=	8
kh	=	2 mD * 54m (177 ft)*N/G (0.5) = 177 mD.ft

The **tubing performance curve** below was created using a vertical lift performance package PROSPER, assuming a 3.5” tubing.

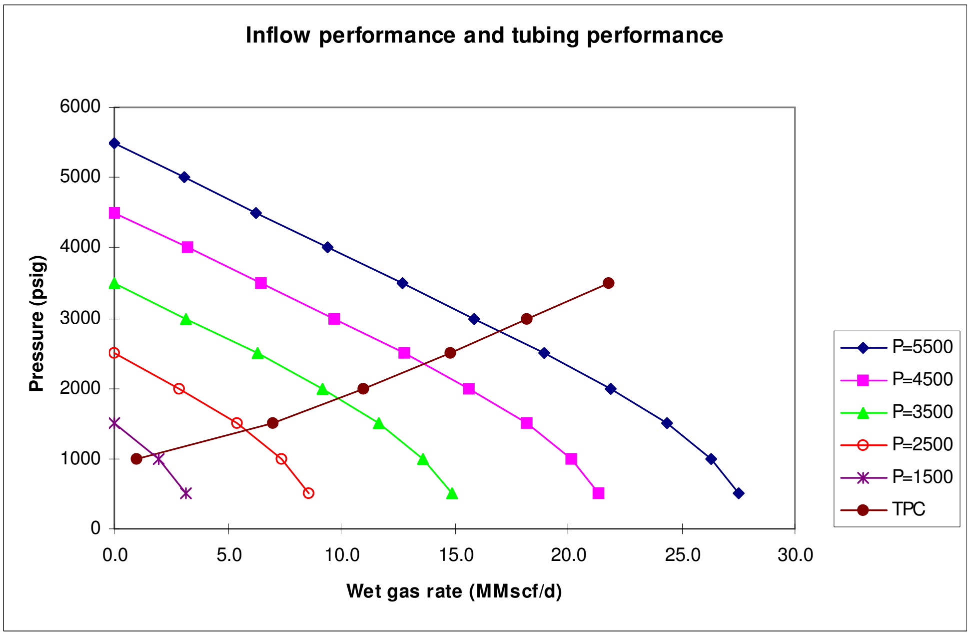

This plot will be used to determine well productivity for any current reservoir pressure.

What does this plot tell you about the appropriateness of the tubing size selected?

What abandonment pressure do you predict for the field?

What final recovery factor do you predict?

Task 5 Well productivity versus cumulative production

As depletion takes place, the reservoir pressure drops and the well deliverability declines. To meet gas nominations it is important to know how many wells are required. By combining the material balance P/z plot with the above IPR/TPC plot, construct the following graph

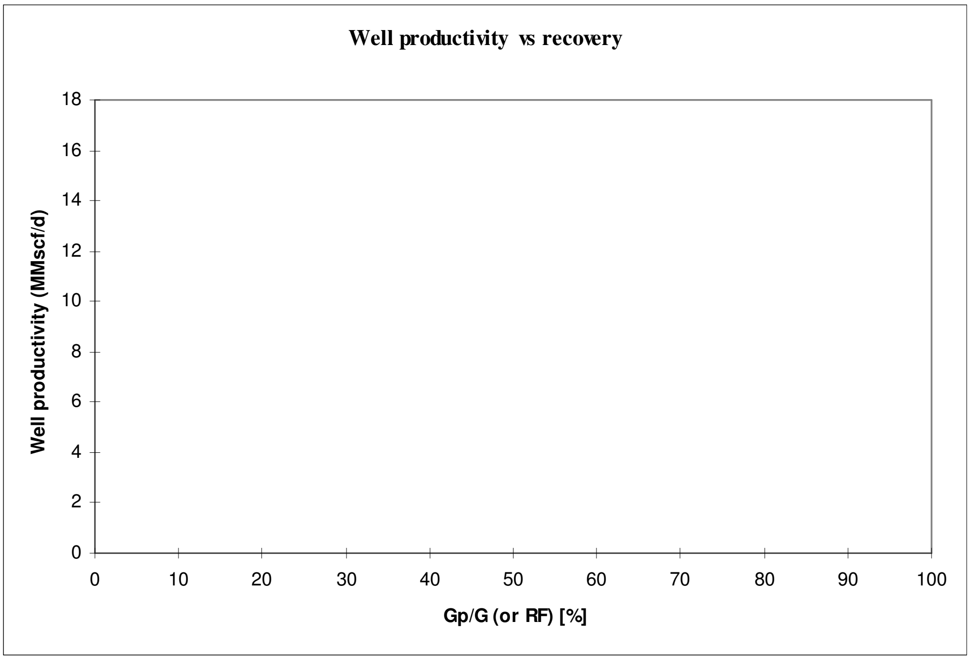

Task 6 Production profile and well requirement prediction

Assume that it is proposed to produce the field at 75 MMscf/d for 5 years, followed by a decline period. Complete the following table to estimate the well requirement. Assume Gp and Gp/G are calculated to the end of each year

Production year	Wet gas rate (MMscf/d)	Gp (wet gas) (Bscf)	Gp/G (or RF)	Rate/well (MMscf/d)	No. Wells required
1	75				
2	75				
3	75				
4	75				
5	75				
6	60				
7	45				
8	30				

The wet gas sales forecast would then be split out into dry sales gas plus condensate rate. In fact the condensate is an important part of the value of the product stream, and should not be neglected in the economics of the project.

Task 7 Analysis of the proposed production forecast

From the results of task 6, comment upon the validity of the proposed forecast, and suggest how the field development could be optimised.

2.1 - Influence diagram

The attached slide shows some of the main parameters of influence in production forecasting for a gas field. The main linkages between these are the material balance, inflow performance and tubing performance relationships.

Influence diagram - Gas Field Development Planning

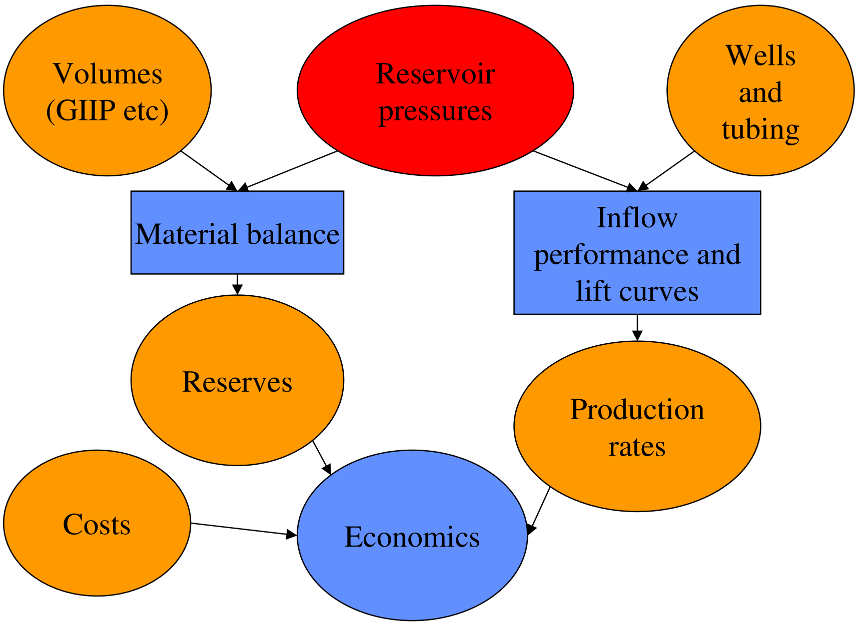

In many respects, reservoir pressure is the key parameter, since it links the material balance and the calculations of production rates.

2.2 Volumes of hydrocarbons in place

For the p50 case

Wet GIIP = GRV*N/G*Sg*φ*E

= $565\ 10^6\ m^3 * 0.5 * 0.5 * 0.19 * 250\ sm^3/rm3 *$
$35.3\ scf/sm^3$

= $6.7\ Bsm^3$ = 237 Bscf

Dry GIIP = Wet GIIP * dry gas/wet gas shrinkage factor

= 237 Bcf * 0.9

	=	213 Bscf	(approximately 35 MMboe)
CIIP	=	237 Bscf * 130 bbl/MMscf	
	=	30.8 MMstb	

2.3 Material balance and the P/z plot

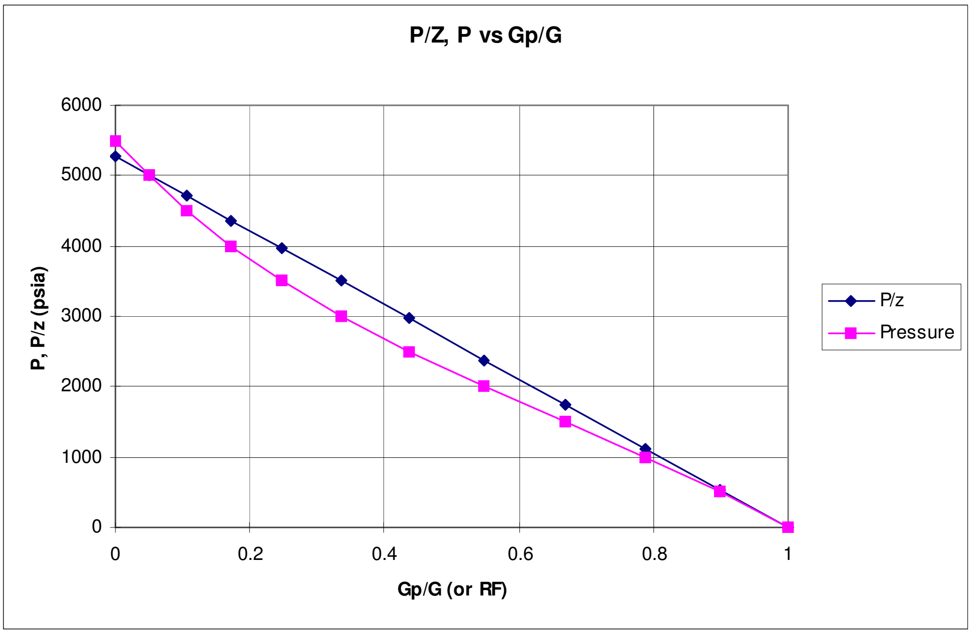

The P/z plot is a straight line between P_i/z_i and $G_p/G=1$.

2.4 Inflow performance and tubing performance curves

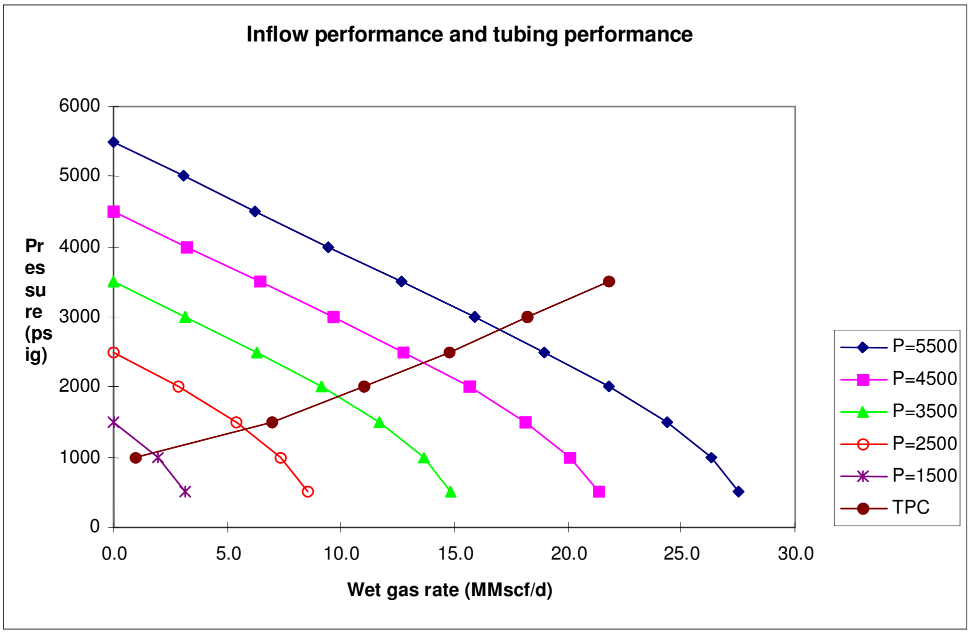

What does this plot tell you about the appropriateness of the tubing size selected?

The 3.5" tubing will be able to lift the produced fluids right down to the abandonment pressure. There will be no requirement for a tubing replacement during the field lifetime (at least not for lift reasons).

What abandonment pressure do you predict for the field?

If the tubing curve cannot be extrapolated for lower rates, then the abandonment would occur at 1500 psig average reservoir pressure, 1000 psig minimum flowing bottom hole pressure. If the tubing curve can be extrapolated, then it might be possible (given enough time) to take both reservoir pressure and flowing bottom hole pressure down to 1000 psig.

What final recovery factor do you predict?

From the P/z plot at a reservoir pressure of 1500 psig, RF should be around 67%. At a reservoir pressure of 1000 psig, RF should be around 80%.

2.5 Well productivity versus cumulative production

The initial well rate of 17 MMscf/d is in excess of the DST on the discovery well. This may be because the discovery well had a positive skin, or that it was actually choked back due to flaring constraints, or separator capacity. The DST interpretation should be checked to ensure that the difference can be explained.

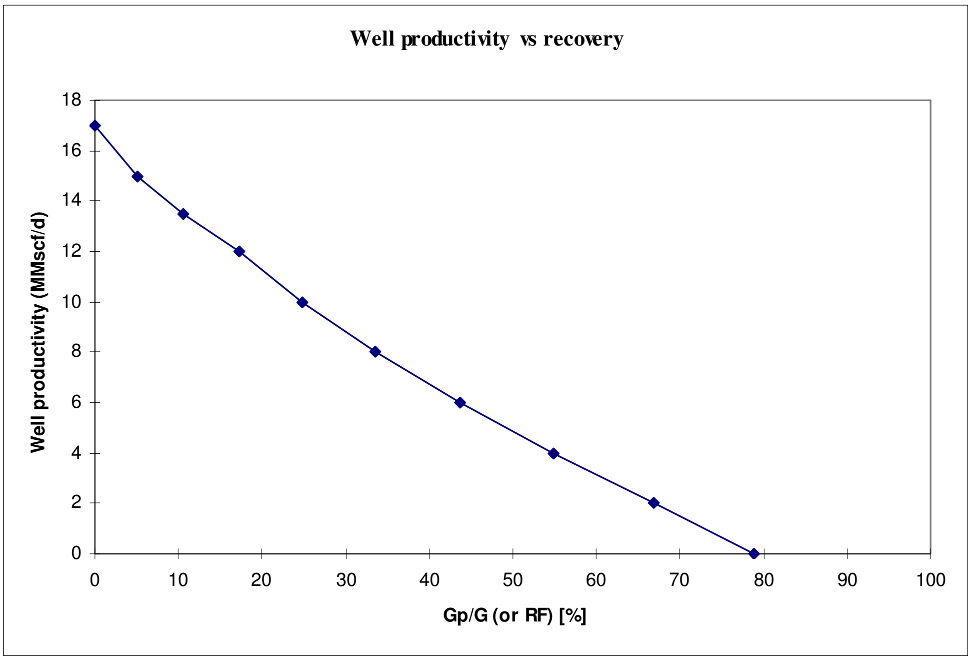

2.6 Production profile and well requirement

Production year	Wet gas rate (MMscf/d)	Gp (wet gas) (MMscf)	Gp/G (or RF)	Rate/well (MMscf/d)	No. Wells required
1	75	27.4	11.5	13	6
2	75	54.8	23.1	10	8
3	75	82.2	34.6	8	10
4	75	109.6	46.2	6	12
5	75	136.9	57.7	4	19
6	60	158.8	67.0	2.5	30
7	45	175.3	73.9	1	45
8	30	186.2	78.5	0	infinite

2.7 Interpretation of the above result

This profile implies that wells need to be drilled every year to maintain the required plateau. This is unreasonable for a small field developed as a subsea tieback, due the mob/demob costs of a drilling facility.

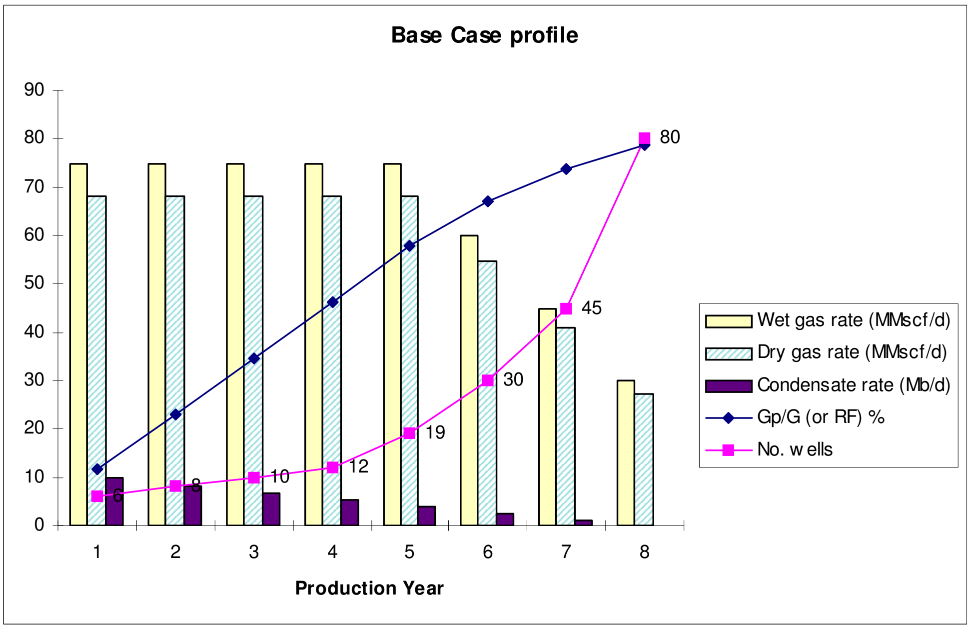

As an alternative the same method can be used to show a profile which results from drilling the first 5 wells and then allowing decline to occur. Only economics can determine the better choice.

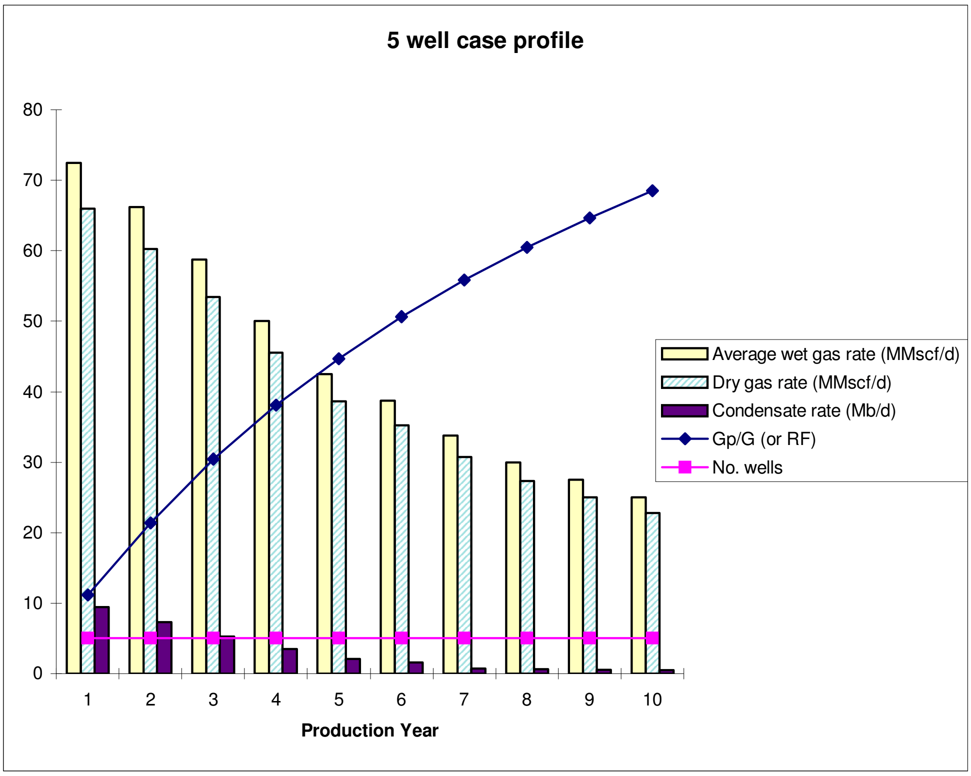

2.8 Reviewing the assumptions

Having created a simple model, it is useful to go back and examine the assumptions to check whether they are valid for the conclusions. For example, consider a model based on the assumption of a given cumulative oil - cumulative gross relationship. If such a model is applied to potential changes in the well configuration, there is a danger of error. The changes in the well configuration may modify the cumulative oil - cumulative gross relationship and so render the initial assumption invalid.

The first key assumptions in the gas field development model are those made in the material balance equation. It is assumed that the field behaves like a tank, with no pressure barriers or compartmentalisation, that pressures equalize very quickly across the whole field, except in the immediate vicinity of the wells and that the effects of any aquifers can be ignored.

These are often very reasonable assumptions for a gas field, because of the low viscosity of gas. However, in this case, the material balance assumptions may be questionable, because of the low reservoir permeabilities.

For the calculations of well inflow, it is assumed that each well will encounter similar permeabilities to the original well. This is a very important assumption. It should be discussed in detail with the geologist working on the field. It would also be worthwhile to calculate the field development profiles for a range of different average permeabilities.

It is assumed that productivity indices (PI) will be unaffected by the number of wells drilled. However, the drainage radius, r_e, features in the formula for PI. The more wells there are, the smaller the drainage radius for each well, but since the drainage radius only occurs in a logarithmic term ($\ln(r_e/r_w)$) the effects will be small and can probably be ignored.

Constructing mathematical models

The process of constructing mathematical models (either algebraic or spreadsheet-based) of a practical problem can be split into five stages -

1. Observing the problem in the real world.
2. Creating a simplified mathematical description.
3. Carrying out deductions/calculations.
4. Reaching conclusions and getting answers in the model.
5. Relating the model conclusions and answers back to the real world.

It is useful to consider, as an example of this process, the most successful such model - reservoir simulation.

There is not a great deal to be said about the first stage, other than to stress the importance of speaking to as many different people as possible, especially the field staff. It is they, after all, who are the closest to what is physically happening in the production process.

The second stage, the creation of a simplified mathematical description, involves choosing what approximations and assumptions to use. One should not be afraid of bold approximations. Reservoir simulators approximate the reservoir as a collection of interconnected tanks. Usually, this description works extremely badly at first, but then the tank properties are adjusted until the simulator starts behaving like the real field.

It is very important that the assumptions and approximations do not create false answers. There are three ways of ensuring that this does not happen -

1. Limit the assumptions to the physics of what is going on. In our example of a simulator, the key assumptions are the conservation of matter and Darcy's Law, for calculating the flow rates between the different tanks. For simpler models, it may be appropriate to make an assumption of a constant pressure drop between two points in the system, rather than a less physical assumption of a constant flow rate.
2. If one has to make assumptions with little physical basis - for example, an assumption of exponential decline - then one can still have full confidence in the conclusions if one can show that the conclusions are very insensitive to the assumption and its alternatives.
3. If neither of the two methods above can be used, then a way round the problem is to express the conclusions as a function of the assumptions. An example of this might be when one wants to apply Buckley-Leverett type calculations to a field that is expected to be mid-way between the conditions for segregated flow and diffuse flow. One could calculate the fractional flow curve for a variety of Corey exponents, between those measured on the rock curves and values of 1, appropriate for segregated flow.

When it comes to deductions and calculations, the strongest reasoning is formal mathematics. You are strongly encouraged to try to work through your problem as if it were a mathematical problem. This may be difficult, but the clarity and generality of mathematical results make it worthwhile.

If, for example, in a water injection scheme, flow rates were proportional to reservoir thicknesses, then the time to water break-through would be independent of thicknesses. If you created a spreadsheet and found this out, it would be difficult to tell, without further work, whether the independence of water break-through to reservoir thickness was an accident of the reservoir parameters you were using, or whether it was more general. In a mathematical argument, the thickness terms would have cancelled each other out, and it would be clear that this was always the case (with the assumption that flow-rate was proportional to reservoir thickness).

Having said that, it is important to be clear that the algebraic approach is not obligatory. In practice, most people are happier working with spreadsheet models. Almost any simple model can be constructed using a spreadsheet approach.

In reaching conclusions, it useful to be guided by the principle that, whatever is happening, it is simple. If the universe can be explained in a few equations, then it is likely that our oil fields will behave as simple systems, describable by a few key parameters. So, if you have not reached a simple answer, you should go back and see whether there are further deductions or calculations that could be carried out.

When relating the results back to the real world, it is worthwhile to restate the key assumptions, as we did in example 2. This reminds us of the limitations of any kind of modeling - that all we can ever demonstrate are statements of the form "If A, then B." Our conclusions will always be subject to the uncertainty as to whether our assumption A is true. Even if A can be directly measured, there is always the uncertainty as to whether our measurements are representative.

Problem 3 - Optimising the producer:injector ratio

In any field development involving injection, whether water, gas, WAG or whatever, it is important to get the correct balance between the number of production wells and the number of injection wells.

In a number of fields, the correct answer became apparent only after production had started, and it became necessary to make extensive (and expensive) modifications to the field development plan. In some of these fields, there was no way of knowing the correct balance in advance, because it depended on factors (such as the extent that faults were sealing) that could only be measured during actual production. However, in others, the modifications would have been avoidable if more attention had been given to calculating the optimal producer:injector ratio, and examining how sensitive it was to uncertainties in parameters such as average injectivity indices.

So, the problem is "What is the optimal producer:injector ratio in a water-flood?" A water-flood is chosen for simplicity - similar methods apply to other floods. Assume that there is no gas cap.

The answers that we will propose are based on simplifications that are particularly appropriate for a pattern flood, but are still somewhat applicable to flank injection.

It is suggested that, as a first step, you may wish to draw up an influence diagram.

3.1 Influence diagram for the factor affecting the producer:injector ratio

A possible influence diagram may run along the lines of the one illustrated below.

Influence diagram for the producer:injector ratio

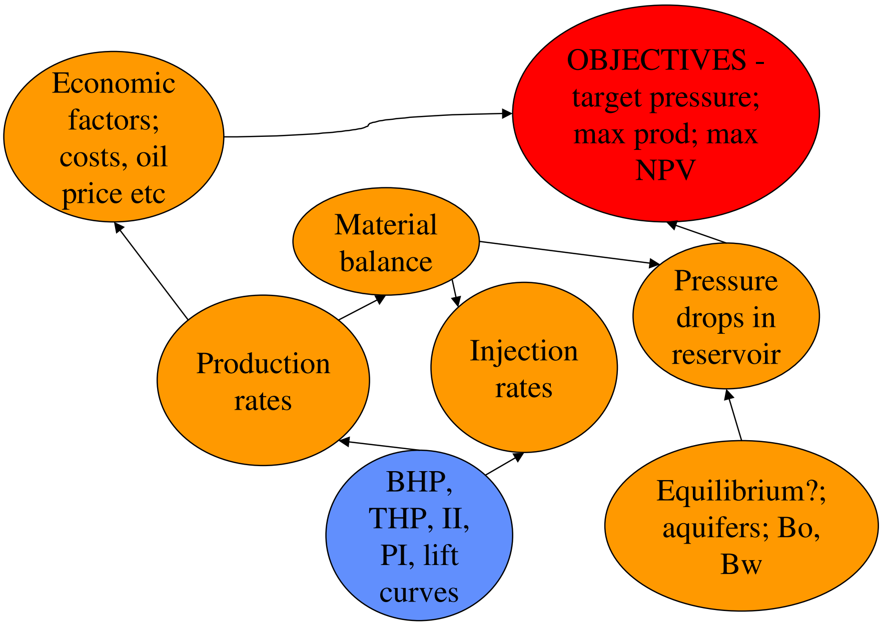

It can be seen that there are very many considerations, but we have to make the choice as to

which factors to include in the simplified mathematical model. That is your next task.

3.2 Simplified mathematical description

In almost all cases, it is best to start with the simplest possible model. Even if it does not exactly match your situation, the simplest model will give insight into the more complicated cases. Often, you will find that the answer for the simplest case can be modified to cover much more complicated cases.

Here, the key assumptions in building the model could be:-

1. When injection starts, the reservoir quickly reaches a stable state in which the simple *material balance* equation applies:- Volume injected = volume produced (in reservoir volumes).
2. Aquifer effects can be ignored.
3. Whatever the ratio, the injectors will be operating at a fixed BHP (as high as possible), and the producers will be operating at a fixed BHP (as low as possible). More complicated assumptions, involving lift curves, could be made - the overall calculation scheme would remain the same.
4. The average productivity and injectivity indices can be considered to be constant over the life of the field.
5. If the field NPV is at its highest possible value, then the ratio average well life-time cost : initial oil production rate is at its lowest possible value. (This can be deduced from other, more basic economic assumptions, but here we will just take it as an assumption).

There is one more key assumption to make, namely identifying the objective. What is your choice of objective?

3.3 Objectives

Three possible objectives are:-

i. achieving a target reservoir pressure;
ii. maximising production per well;
iii. maximising NPV.

Solutions will be given for all three and, later on, there will be a discussion at to which is the most reasonable objective.

The simplest case is that in which the objective is to achieve a target reservoir pressure. Please express injection rates and production rates as functions of the reservoir pressure and the number of injectors and producers. By equating injection and production (in reservoir volumes), derive an expression that gives the producer:injector ratio as a function of average reservoir pressure for a field that is still producing at 0% water-cut.

It is useful, at this point, to agree on some notation

Pprod = Flowing bottom hole pressure (FBHP) in the producers

Pinj = Flowing bottom hole pressure in the injectors

$\overline{P}$ = Average reservoir pressure

m = number of injectors

n = number of producers

r = producer:injector ratio = n/m

II = average injectivity index

PI = average productivity index

Bo = Oil formation volume factor (we assume that the water formation volume factor, Bw, is approximately 1)

3.4 Achieving a target reservoir pressure

In almost every water-flood, there will be a limit to Pprod, dictated by the need to keep the pressure in the near well-bore to above bubble-point. There may also be a higher limit, e.g. the limit imposed by the need to keep ESP inlet pressures above (or at least not much below) the bubble-point, to avoid pump reliability problems. There will also be limits to bottom-hole injection pressures – perhaps dictated by pump limits, or by a desire not to exceed the fracture propagation pressure. In most cases, it will worthwhile to operate the wells at these limits – otherwise, one could achieve higher production or get by with fewer injection wells.

In some fields, where there is a possibility of causing cross-flow, for example, it may be desirable to aim for a target average reservoir pressure, $\overline{P}$. The first step to take is to note that average reservoir pressure is stable when the injection rate (in reservoir volumes) equals the production rate (in reservoir volumes).

Now, if m is the number of injectors and n is the number of producers,

Reservoir injection rate = mII(Pinj- $\bar{P}$)

Reservoir production rate = nPI.Bo(Π-Pprod)

So,

mII(Pinj- $\bar{P}$) = nPI.Bo($\bar{P}$ -Pprod)

Hence, if stability is to be achieved and the wells are all to operate at capacity,

Producer:injector ratio = n/m = II(Pinj- $\bar{P}$) / [PI.Bo.($\bar{P}$ -Pprod)]

The next stage in this problem is to formulate strategies for calculating the producer:injector ratios needed to achieve the other two objectives - maximising production and maximising NPV (obviously, these two targets cannot be simultaneously achieved - we need to find one producer ratio for maximising production and another for maximising NPV). So your task is to answer the question "How can we calculate the producer:injector ratio that maximises production? How can we calculate the producer:injector ratio that maximises NPV." It is suggested that you focus on the job of creating an equation that has, as one of its terms, the appropriate ratios. Once we have such an equation, it is usually straightforward to solve it - either algebraically or in a spreadsheet.

3.5 Strategies for calculating the optimal ratios

In the previous section, where the objective was to achieve a target reservoir pressure, it was easy to find an equation - we knew the value of the target. For our current objectives, it is reasonably straightforward to express production as a function of the producer:injector ratio. However, we do not know what will be the value of production - only that it should be maximised.

What we do know is that, since it will be a maximum, i.e. "at the top of the hill", then its derivative with respect to the producer:injector ratio will be zero. This will give us the required equation.

In slightly more detail, what we will do is:-

a) Express the average reservoir pressure as a function of the producer:injector ratio.
b) Express the production rate as a function of the average reservoir pressure, hence as a function of the producer:injector ratio.
c) Calculate the derivative of the production rate with respect to the producer:injector ratio.
d) Set this derivative to zero.
e) Manipulate the equation from step (e) to give an expression for the producer:injector ratio.

When it comes to maximising NPV, it would be complicated to to derive an expression for NPV as a function of the producer:injector ratio (although not impossible - the theory outlined in Problem 4 provides one method to do so). For this problem, we will assume that maximising NPV is equivalent to minimising the ratio (call it ψ)
average well life-time cost : initial oil production rate
is at its lowest possible value. (This can be deduced from the economic

assumptions in Problem 4, but here we will just take it as an assumption).

The next task comes in two parts

a) Express the average reservoir pressure as a function of the producer:injector ratio.
b) Derive a way of expressing the number of producers in terms of the total number of wells and the producer:injector ratio. We will need such an expression later on.

(Hint - consider the equation $n = n.\frac{(n+m)}{(n+m)}$).

3.6 Expressing $\overline{P}$ as a function of the producer injector ratio

There is often flexibility about the choice of average reservoir pressure at which to operate the field. The field may be large in area, with relatively low permeability and thick shale barriers above and below. In such a case, increasing the field pressure to above the aquifer pressure will have few consequences in terms of oil pushed out into the aquifer, since the flow rates are small, and could be reduced further by arranging for there to be an outer ring of injectors. Hence, $\overline{P}$ can be chosen to give the optimal production rates.

From equating injection rates with production rates, we got the equation

mII(Pinj- $\overline{P}$) = nPI.Bo($\overline{P}$ -Pprod)

It can be seen that, in the medium term (long enought for pressures to reach equilibrium, but not long enought for there to be significant changes because of changing saturations and water-cut), that $\overline{P}$ is determined by the injection-production pattern, assuming aquifer influxes are negligible. Re-arranging the equation gives

(nPI.Bo + mII). $\overline{P}$ = mII.Pinj + n.PI.Bo.Pprod

hence

$\overline{P}$ = [mII.Pinj + n.PI.Bo.Pprod] / [n.PI.Bo + m.II] =

= [II.Pinj + r.PI.Bo.Pprod] / [r.PI.Bo + II]   where r = n/m

It is interesting to see that, for very low values of r (i.e. very few producers), this expression suggests that the reservoir pressure would tend towards Pinj - the injection BHP. For very high values of r (i.e. very few injectors), the reservoir pressure would tend towards Pprod - the producers' BHP. This is illustrated in the plot below, for a hypothetical field with

- injector BHP = 4000 psi
- producer BHP = 2000 psi
- average PI = 4 stb/day/psi
- average II = 10 stb/day/psi
- Oil formation volume factor, Bo, = 1.1 rb/stb.

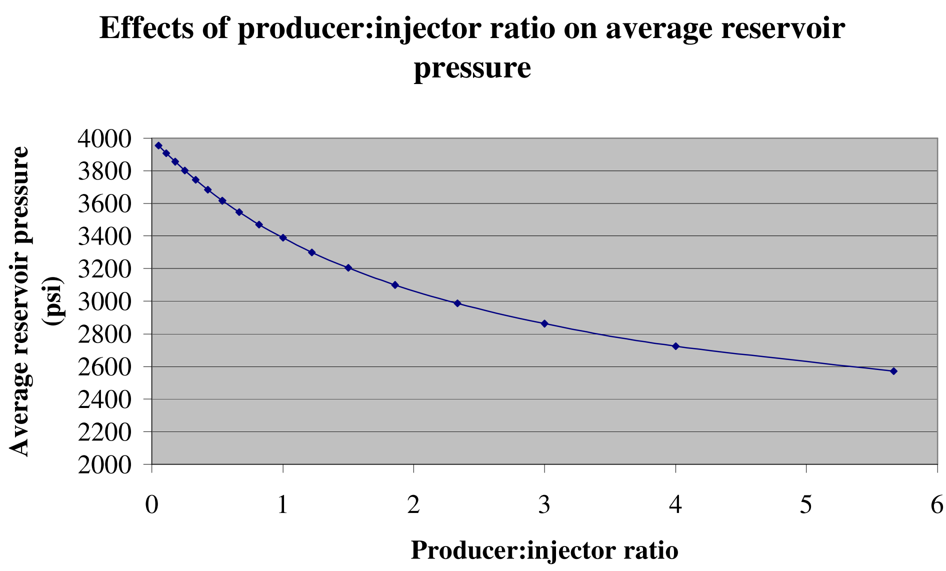

To answer the second part of the question, it suffices to note that

$$n = n.\frac{n+m}{n+m} = (n+m).\frac{n}{n+m} = (n+m).\frac{1}{1+{}^{m}\!/_{n}} = (n+m).\frac{1}{1+{}^{1}\!/_{r}}$$

These two results give us the ingredients needed for the next task, which is

- express the initial oil production rate as a function of the producer:injector ratio
- calculate the producer:injector ratio that gives the maximum initial oil production rate per well.

Attached overleaf is an outline of a possible answer, minus the actual mathematics. You may find this helpful. If you prefer, at this stage, to leave all the algebra to one side, you may instead address the problem numerically by, for example, writing a spreadsheet to carry out these calculations.

3.7 Calculating the producer:injector ratio that maximises production

The initial oil production per production well is simply equal to

(Drawdown) x (Productivity index)

so with n production wells,

Initial oil production = n x (Drawdown) x (Average productivity index0

= n.PI.($\overline{P}$ -Pprod)

Substituting the expression for $\overline{P}$ from section 3.6 gives

$$\text{Initial oil production} = \text{n.PI.}\left[\frac{\text{II.Pinj} + \text{r.PI.Bo.Pprod}}{\text{r.PI.Bo} + \text{II}} - Pprod\right] = n.PI.\frac{II.(Pinj - Pprod)}{r.PI.Bo + II}$$

Expressing the number of producers, n, in terms of (n+m) and r gives

Initial oil production =

$$\frac{(\text{n}+\text{m})}{1+\frac{1}{\text{r}}}.\frac{\text{PI.II.}(\text{Pinj}-\text{Pprod})}{\text{r.PI.Bo}+\text{II}} = \frac{(\text{n}+\text{m}).\text{PI.II.}(\text{Pinj}-\text{Pprod})}{\text{r.PI.Bo}+\text{PI.Bo}+\text{II}+\frac{\text{II}}{\text{r}}}$$

To find the maximum initial oil production (for a fixed total number of wells), one can calculate the derivative of the above expression (with respect to r) and set it to zero. Alternatively, one can simply note that, in the final expression for initial oil production, all the terms in the numerator (i.e. above the line) are constant (with respect to r). So the inital oil production is maximised iff the denominator (i.e. the terms below the line) are minimised.

Let us call the denominator α, where α = (r.PI.Bo + PI.Bo + II + II/r)

Then α is minimised when $\frac{d\alpha}{dr} = 0$

$$\text{Now} \qquad \frac{d\alpha}{dr} = \text{PI.Bo} - \frac{\text{II}}{\text{r}^2} \qquad \text{so when } \frac{d\alpha}{dr} = 0$$

$$\text{r}^2 = \frac{\text{II}}{\text{PI.Bo}}$$

$$\text{r} = \sqrt{\frac{\text{II}}{\text{PI.Bo}}}$$

We have shown that production is maximised, for a given number of wells when

Producer:injector ratio = √(II/PI.Bo)

This can be illustrated by the example from section 3.6. For this example,

Producer:injector ratio that maximises production = √(II/PI.Bo) = √(10/(4 x 1.1) = 1.5

It can be seen from the plot below that this matches the results obtained by using a spreadsheet to calculate the production obtained from a whole range of different producer:injector ratios.

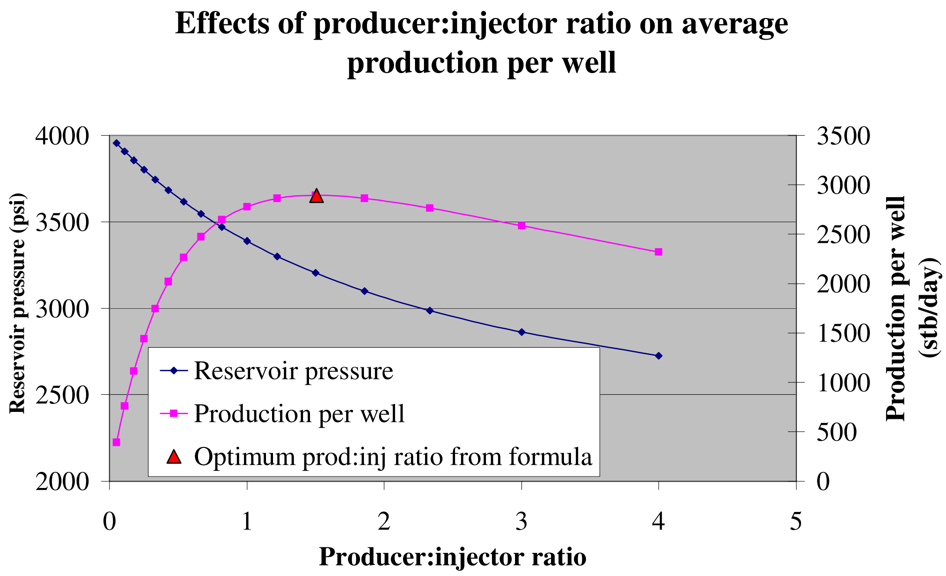

If injector wells and producer wells cost the same to construct and to operate, then this ratio will also be the economically optimal ratio. The next stage in the task is to find the economically optimal ratio in the more general case, in which the costs are not the same. Remember that we assume that the economically optimal ratio is the one that minimises the ratio:

average well life-time cost : initial oil production rate.

3.8 Calculating the economically optimal producer:injector ratio

We will show that the economically optimal ratio of producers to injectors is √(Xi.II/(Xp.PI.Bo), where Xi and Xp are the lifetime costs of an injector and a producer, respectively.

By assumption 5, NPV is minimised when the ratio (call it Ψ)

average lifetime well costs : gross production

is minimised.

Now, if, as before, m is the number of injectors, n is the number of producers and r is the producer:injector ratio is the proportion of injectors to total well stock.

and (1-r) is the proportion of producers to the total well stock, then

$$\psi = \frac{m.X_i + n.X_p}{q}$$

Expressing q as $\dfrac{(n+m).PI.II.(Pinj - Pprod)}{r.PI.Bo + PI.Bo + II + II/r}$ (see section 3.7) gives

$$\psi = \frac{(m.X_i + n.X_p).\left(r.PI.Bo + PI.Bo + II + II/r\right)}{(n+m).PI.II.(Pinj - Pprod)}$$

(A)

Since $\dfrac{\mathrm{m}}{\mathrm{n+m}} = \dfrac{\mathrm{m}/\mathrm{m}}{\mathrm{n}/\mathrm{m} + \mathrm{m}/\mathrm{m}} = \dfrac{1}{\mathrm{r}+1}$

and $\dfrac{\mathrm{n}}{\mathrm{n+m}} = \dfrac{\mathrm{n}/\mathrm{m}}{\mathrm{n}/\mathrm{m} + \mathrm{m}/\mathrm{m}} = \dfrac{\mathrm{r}}{\mathrm{r}+1}$

equation (A) can be simplified as follows

$$\psi = \frac{(m.X_i + n.X_p).\left(r.PI.Bo + PI.Bo + II + II/r\right)}{(n+m).PI.II.(Pinj - Pprod)} = \frac{\left(\frac{1}{r+1}.X_i + \frac{r}{r+1}.X_p\right).\left((r+1).PI.Bo + \left(1 + 1/r\right)II\right)}{PI.II.(Pinj - Pprod)}$$

$$= \frac{\left(X_i + r.X_p\right)\left(PI.Bo + \frac{II}{r}\right)}{PI.II.(Pinj - Pprod)}$$ by noting that $(1+1/r) = (r+1)/r$ and then cancelling $(r+1)$ terms

Taking the partial derivative of ψ with respect to r

$$\frac{\partial\psi}{\partial r} = \frac{X_p.\left(PI.Bo + \frac{II}{r}\right) - \left(X_i + X_p\right).\frac{II}{r^2}}{PI.II.(Pinj - Pprod)} = \frac{X_p.PI.Bo - X_i.\frac{II}{r^2}}{PI.II.(Pinj - Pprod)}$$

NPV will be at a maximum whenever the cost:initial production ratio, ψ, is at a minimum, which is when the partial derivative ∂ψ/∂r is zero.

Now

$$\frac{\partial \psi}{\partial r} = 0$$

$$\Rightarrow X_p.PI.Bo - X_i.\frac{II}{r^2} = 0$$

$$\Rightarrow \frac{1}{r^2} = \frac{X_p.PI.Bo}{X_i.II}$$

$$\Rightarrow r = \sqrt{\frac{X_i.II}{X_p.PI.Bo}}$$

i.e. The NPV is maximised when the producer:injector ratio is √(Xi.II/(Xp.PI.Bo)).

Problem 4 - What is the optimal number of wells to drill on a field?

This is the problem that gave rise to the idea of this course. The question of how many wells to drill (or, equivalently, of the optimal well spacing) is a very basic one. The conventional answer is usually either along the lines "It is decided on a field-by-field basis by the reservoir engineer, in conjunction with the economist, geologist and facilities" or even by rules of thumb such as "Assume 500 metre well spacing for this type of field." Using the techniques that we have described in this course quickly gives an answer in the form of a simple formula for the optimal number of wells to drill.

The stages in addressing the problem are

- building a qualitative description of what happens as well numbers vary
- calculating the effects of well numbers on the production profile
- simplifying economics
- finding the number of wells that maximises NPV
- modifying the method to take account of value:investment hurdles.

The first stage in the task of solving the problem can be phrased as -

"Describe in general terms what happens when you vary the number of planned wells in a field development." You may find it helpful to draw up an influence diagram.

4.1 The effects of varying the number of wells drilled

Before going into details, it is useful to create a qualitative description of the mechanism governing the choice of number of wells. Such a qualitative description can be guide to building a quantitative model, and also provides a quality control or reality check on the quantitative model -

when we ask ourselves "Does the quantitative model match the qualitative description?"

The effects of changing well numbers are illustrated in the influence diagram below.

Influence diagram - effects of changing well numbers

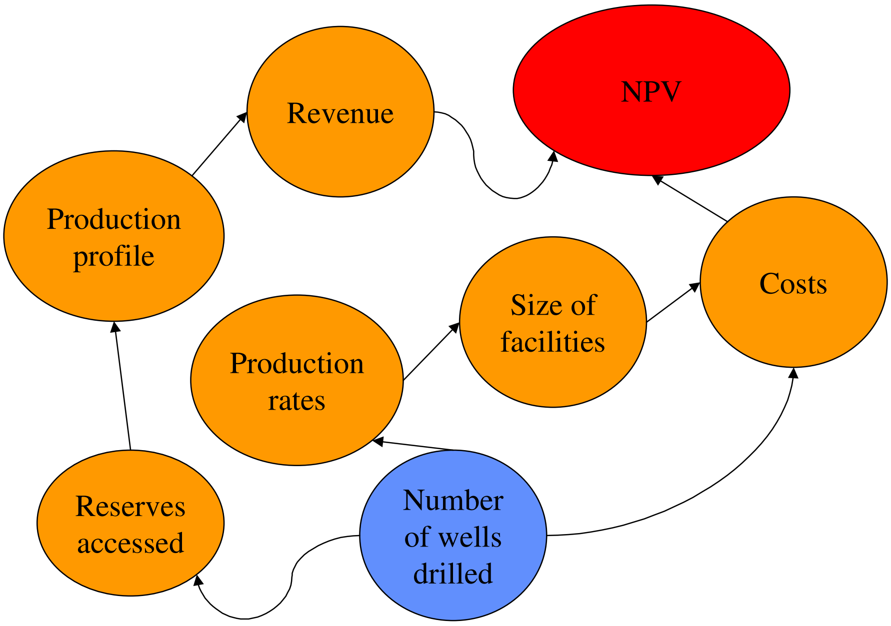

The first effect of increasing the number of wells in a field development plan is to increase costs. Assuming that the facilities are scaled to the number of wells, both opex and capex are likely to be approximately a linear function of the number of wells in the plan. In consequence, the present value of the total costs will also be approximately a linear function of the number of wells, of the form a + b.N, where

- a is the fixed expenditure necessary irrespective of the number of wells - e.g. much of pipeline construction costs or the baseline facilities costs
- N is the number of wells drilled
- b represents the costs that vary with the number of wells - this obviously includes the direct costs of drilling the wells, but also includes the extra costs incurred in facilities etc.

On the benefits side, increasing the number of wells speeds up the field production, so you get your oil faster. This obviously increases the present value of the production. Increasing the number of wells also usually increases reserves and ultimate recovery (although there are circumstances where the opposite can happen, such as in a water-flood where increases in throughput rate can damage the displacement efficiency).

The benefits of speeding up production and the increases in reserves are both subject to diminishing returns and will always be less than the value of having all the available oil for sale now. Using this value as an asymptote, and the fact that the benefit of no wells is zero, we can construct a benefit vs well numbers relationship that will, qualitatively, look like the relationship in the plot below.

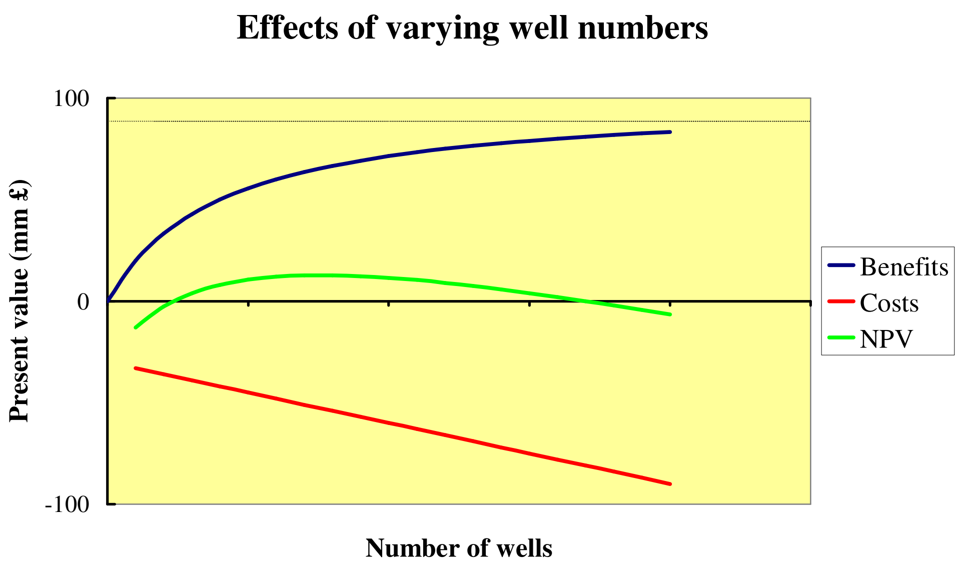

Adding together the present value of the benefits and the present value of the costs (a negative quantity) gives the NPV. It can be seen that the straight line of costs and the diminishing returns on the benefits side give rise to an NPV function that rises to a maximum and then declines.

In the next stages of this example, we will quantify the relationships above and look at ways of explicitly finding the number of wells that maximises NPV. In reality, your objectives are likely to be to maximise NPV subject to a value:investment hurdle. We will discuss later on how NPV maximisation can be modified in a simple way to take account of any value:investment hurdles.

The next stage in the problem is to look at the benefits side in more detail. Please address the question "How could you quantify the effects on the production profile of changes in well numbers?" The aim, at this stage, is still to increase our understanding, so we can make some simplifications, to prevent the problem becoming too difficult. It is suggested that you make the assumptions:-

- reserves are independent of well numbers
- initial well rates are independent of well numbers
- the field shows exponential decline
- all wells start production at the same time
- abandonment effects can be ignored.

4.2 Quantifying the effects of well numbers on the production profile

We have made an assumption about the shape of the production profile. We know, for a given number of wells, what is the initial production rate. Finally, the area under the production profile is known - this is equal to the reserves. These three pieces of information are sufficient to calculate the decline rate and so construct the production profile.

Mathematically,

If q is the average initial well rate and there are N wells, then the initial field production rate is N.q.

By the assumption that there is exponential decline, with decline rate, a, (this varies with the number of wells)

Oil production rate = $N.q.e^{-at}$

Since reserves, R, are equal to the total production over the life of the field

$$R = \int_0^\infty N.q.e^{-a.t}dt = \left(-\frac{N.q}{a}.e^{-at}\right]_0^\infty = N.q/a$$

so a = N.q/R

This method is not limited to exponential decline. A similar approach of making an assumption about the shape of the production profile can be taken for constant production rates, or for a plateau period followed by decline.

It can be argued that reserves are equal to the production from time = 0 to abandonment time, and that the integration should run between these two times. An easy way round this problem is to use for R not the economic reserves, but the technical reserves - the oil that would be recovered if production went on for a very long time[2].

It is clear that it would be desirable to combine the production profile with economics to give the present value of the production, as a function of well numbers. A simple way of doing this is to feed each profile into an economics spreadsheet. However, remembering that the objective in this problem is to gain understanding, it may be more informative to try to simply the economics into a simple equation or two and combine these directly with the production profile.

So, the next stage is to answer the questions "How can you simplify project economics? What equations would you expect (with suitable co-efficients) to produce approximately the same results as the economic models currently in use?"

Hint - you may want to estimate typical value for the effects on project NPV of increases in oil production, CAPEX and OPEX at various different times in the project, i.e. filling in the table below.

[2] To be even more precise, R should be the area under the extrapolation of exponential decline to time = ∞. This is to avoid the effects of fields switching to hyperbolic production late in field life, a phenomenon that can give rise to significant extra production, but which is almost totally irrelevant to field development planning, because of time and discounting effects.

EFFECT ON NPV OF CHANGES	**Timing**	
Change	Year 1	Year 5
+ 1 stb oil		
+ 1 MM stb oil		
+ $1 CAPEX		
+ $1 MM CAPEX		
+ $1 OPEX		
+ $1 MM OPEX		

4.3 Simplifying economics

The first point to note is that most economic models behave approximately linearly. If spending an extra £1,000 opex in year 1 reduces NPV by £350, then, usually an extra £1,000,000 in opex in year 1 will reduce NPV by £350,000.

Changes in production and capex are also usually linear in their effects on NPV. Changes in capex generally cause changes over several years (because of the tax treatment of depreciation) but each of these changes is usually linear to the change in capex, so the combined effect is still linear.

The second point is that discounting applies equally to both revenue and costs. Combining discounting and linearity suggests that it may be possible to simplify economics to

$$NPV = \sum_{Year\ i=1}^{i=n} \frac{L.P_i - b.C_i - c.E_i}{(1+d)^i}$$

where

d = discount rate

Pi = production in year i

Ci = capex in year i

Ei = opex in year i

and L, b, c are the appropriate co-efficients (their values vary according to the tax and royalty system and the oil price etc).

The co-efficient "L" can be called the "net oil price". It is a measure of how much NPV increases for each extra barrel of oil sold. Similarly, "b" can be called the "net cost of capex" and c the "net cost of opex."

It is important to note that we are *not* assuming that the NPV is linear with respect to the oil price - often this is not the case, especially with high royalties. Instead, we are assuming that, for a given oil price, NPV is linear with respect to production.

Money spent on capex can be offset against tax, but the tax benefits are usually spread-out over several years. Hence, the value of these tax benefits depends on the discount rate in use. Consequently, the net cost of capex depends on the discount rate in use.

Your next task is to describe how to calculate the number of wells which maximises NPV. You may describe the process either for an analytic model or for a spreadsheet model.

4.4 Finding the number of wells which maximises NPV

Analytically

The main steps might be

1. Change the production profile from section 4.2 into a present value (PV) of revenue by multiplying the production at time t, P_t, by L / $(1+d)^t$ (i.e. net oil price x discount factor) and summing up (integrating). The integration is straight-forward, since both discounting and exponential decline can be combined into a single exponential function. PV revenue ends up by being a function of the number of wells, since production rates depend on the number of wells..
2. Similarly, calculate a present value of costs. Subtract PV costs from PV revenue to give NPV as a function of numbers of wells.
3. Calculate the derivative of NPV with respect to number of wells. When NPV is at its maximum, the derivative will be zero. Hence, setting the derivative to zero gives an equation that can be solved to give the number of wells that maximises NPV.
4. A further step that can be carried out is to calculate the maximum NPV, by putting the NPV-maximising number of wells into the NPV formula.

Doing this yields the following equations

$$\text{Number of wells that maximises NPV} = \frac{R.d}{q} . \left(\sqrt{L - \frac{\alpha.E}{q}} - \sqrt{\frac{C.d + E - \alpha.E}{q}} \right) \Big/ \sqrt{\frac{C.d + E - \alpha.E}{q}}$$

$$\text{NPV with this number of wells} = R.\left(\sqrt{L - \frac{\alpha.E}{q}} - \sqrt{\frac{C.d + E - \alpha.E}{q}} \right)^2 - D$$

where $\alpha = (1+d)^{-Tab}$ and T_{ab} (abandonment time) = $\dfrac{R}{q.N}.\ln\left(\dfrac{L.q}{E}\right)$

and

q – initial oil production per well per year, averaged over all wells, including injectors

N – total number of wells, including injectors

R – technical reserves i.e. the amount of oil that could be recovered if the field were run for a very long time

L – net revenue per tonne of oil (i.e. after all taxes and royalties, including profit tax)

d – discount rate

C – net capital cost per well

D – net capital costs not related to numbers of wells, e.g. roads and pipelines

E – net opex per well

A full proof of these equations is given in Appendix 1.

In Excel

It is straight-forward to construct production profiles and costs from input that includes the number of wells. Then it is simply a question of calculating the economics, either from the simplified economic model we have proposed or by creating direct links to the full economic model.

To find the optimal number of wells, the easiest method is to use "Solver" (click on "Tools" and then "Solver...") and ask it to maximise NPV by varying the number of wells.

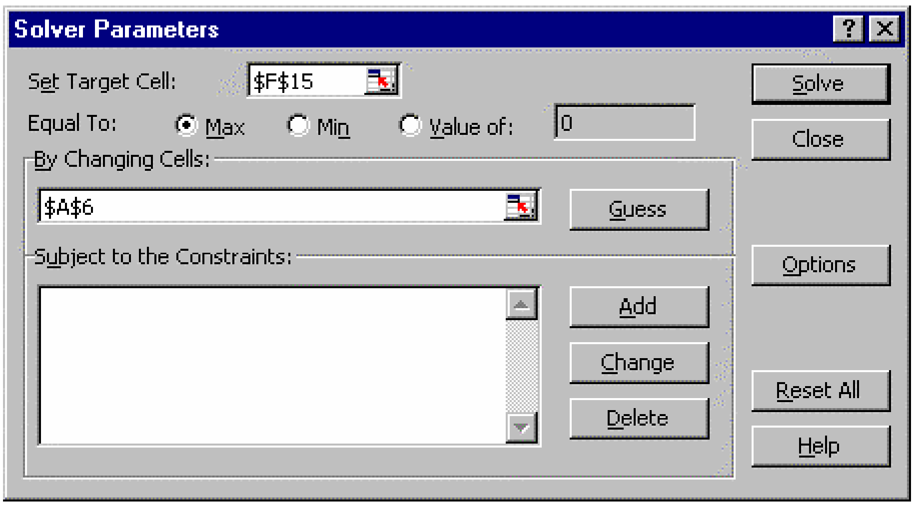

The formulas as they stand show the number of wells required to maximise NPV. In the real world, oil companies are faced with the problem of constraints limiting what is practical. At this point, you may like to pause briefly to consider what are the major constraints faced by typical oil companies.

One of the most important constraints is that companies usually have only a fixed amount of capital available, or at least limits on the amount of capital employed (gearing limits etc). The final stage of the problem then is "How do we find the optimal number of wells, given that the total amount of capital available to the oil company is limited and must be put to best use across a whole portfolio of possible investment project?"

4.5 Incorporating value:investment hurdles

At the company level, the problem is to maximise NPV, subject to the limits on the amount of capital employed. Assuming that there are sufficient number of projects available and no other constraints intrude (such as limits on management capacity), the method to do this is to rank all the projects in order of decreasing NPV:Capex ratio, and then choose the projects in this order until one has reached the limit on the amount of capital employed.

For the projects chosen, let H (the "hurdle" rate) be the lower limit to the NPV:Capex ratio i.e. the NPV:Capex ratio for the last project chosen.To determine, for a new project, the optimal number of wells, it is useful to consider the project as consisting of

(Minimalist development option) + (Series of increments to the mininalist development option)

(Note that the split is purely conceptual; all the increments start at time = 0).

All the increments can be analysed as if they were separate projects. They pass the screening criteria if their NPV:Capex ratio is greater than H

i.e. $\delta NPV / \delta Capex \geq H$

So, while the NPV/Capex ratio for additional wells is greater than H, the wells are worth adding to the development scheme. The optimal total number of wells is reached when the limit is reached, so the criteria for determining the total number of wells is

$\delta NPV / \delta Capex = H$

This can be converted into a more workable criterion as follows

$\delta NPV / \delta Capex = H$

$\Leftrightarrow \delta(NPV\text{-}H.Capex) / \delta Capex = 0$

$\Leftrightarrow [\delta(NPV\text{-}H.Capex) / \delta N] \times [\delta N / \delta Capex] = 0$ where N is the number of wells

$\Leftrightarrow \delta(NPV\text{-}H.Capex) / \delta N = 0$ since $\delta N / \delta Capex \neq 0$ (it never costs an infinite amount of money to drill a new well).

Hence, the new problem (find the number of wells that maximises corporate NPV subject to a limit on capital employed) can be converted into the old (find the number of wells that maximises NPV) by using an artifical NPV, defined to be

NPV′ = NPV – H.Capex = NPV – H.(C.N + D)

and artificial C′ and D′ defined to be

C′ = (1 + H) x C

D′ = (1 + H) x D

Then

$$\text{Optimal number of wells, } N = \frac{R.d}{q}.\left(\sqrt{L - \frac{\alpha.E}{q}} - \sqrt{\frac{C'.d + E - \alpha.E}{q}}\right) \Bigg/ \sqrt{\frac{C'.d + E - \alpha.E}{q}}$$

$$\text{NPV with this number of wells } = NPV' + H.(C.N + D) = R.\left(\sqrt{L - \frac{\alpha.E}{q}} - \sqrt{\frac{C'.d + E - \alpha.E}{q}}\right)^2 + H.C.N - D$$

(the –H.D and +H.D terms cancel each other out).

This method of including VIR hurdles in the form of an extra cost in capital can also be applied in many different problems.

All that remains to be done is to run through an example illustrating how the method works. You may choose to apply the formulas directly, or to use "Screener", an Excel Add-in with the necessary set of functions. You may wish to use parameters from a field you are familiar with, or you can use the following set, based roughly on the Britannia field.

Reserves - 3 tcf gas

Initial well rates - 10 bcf per year

Capex per well - $10,000,000

Opex per well per year - $1,000,000

Fixed capex - $500,000,000

Net gas price - $1,000 per mm scf

Capex factor - 0.8

Opex factor - 0.6

VIR hurdle - 30%

Discount rate - 10%

4.6 Results from gas field example

Using the parameters from Section 4.5 (for gas field roughly analagous to Britannia) gives

Optimal number of wells (taking into account VIR hurdle) - 47

Number of wells that maximises NPV - 58

The relationship between well numbers and NPV for this field is illustrated in the plot below.

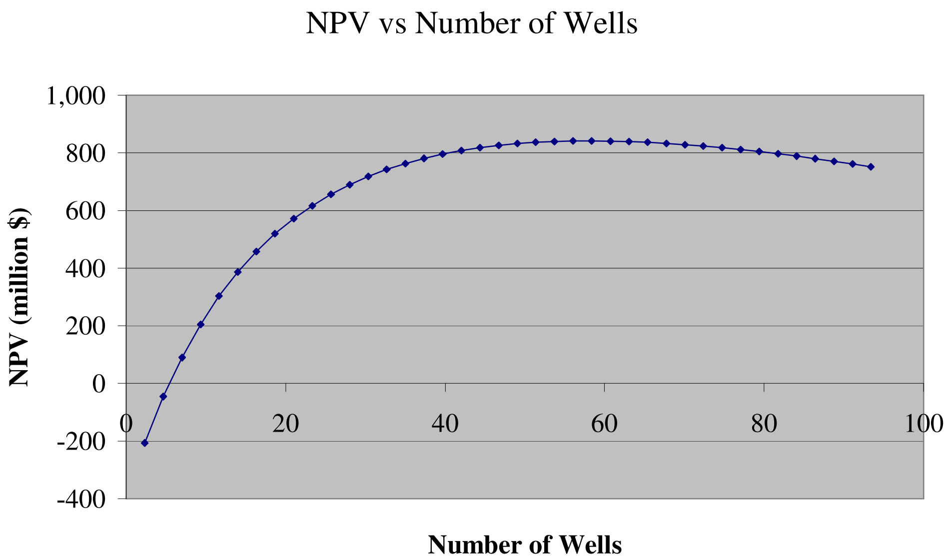

Problem 5 - Project optimisation under a fixed production limit

We have described the process of creating simple models and have gone through a number of examples illustrating this process. However, you may be wondering how realistic it is to attempt such an approach when confronted with a real problem and a limited time frame.

It is true that all the examples have benefited from a gradual simplification, as a result of further work on the problems - in the preparation of this course, for example. When the examples were unsolved problems, they appeared to be rather difficult. Then came a breakthrough and a solution was found. The initial formulation of the solution was often rather long and involved. But, in time, short cuts and simplifications were spotted.

When you start on a real problem, it will appear difficult - probably much more difficult than any of our examples. But you should persevere for the following reasons

- Finding a solution may require time, but you will be surprised at how much ground you can cover in even a week. Such an investment of time should be compared with the six months plus required by detailed reservoir simulation modeling.
- Even if, after a week, you have not found a solution, you will have gained enormously in terms of understanding the problem. It should be remembered that mathematics is the language of engineering and all our training is given in mathematical terms, yet, once we are on the job, we do not spend much time on the same mathematics that seemed so essential to our professors. In this context, any additional time doing mathematics is likely to be valuable.
- Like training for a sport or learning to use a piece of software, these methods become much easier with practice.

- There are a number of tricks that can help finding solutions, some of which are mentioned below. We will illustrate some of them in this example.

Many of you, especially from research backgrounds, already take a similarly mathematical approach, as a guide to and a quality control on the reservoir simulation work, for example. For you, our message would be

- Document your work thoroughly.
- Try to extend these approaches from the purely technical to include economics - it gives enormous scope for optimising development plans.

It is worthwhile to make a final point about the time required for creating simple models. Although one might expect to make significant progress in a week, it is very difficult to estimate what that progress might be. In consequence, the planning and scheduling of such work should be much looser than the planning of a reservoir simulation study. In detailed scheduling of the work, there is a lot to be said for allowing the individual engineer to follow his instincts. If an area of work seems very interesting, it may be because one has grasped, on a subconscious level, that it is the key area.

There are a number of common tricks or strategies that can help us in solving our problems. These include -

1. Solve a bigger problem or a smaller problem. In this case, one might start with the bigger problem of "What should the field unit do with its portfolio of satellites and projects, given the gas supply contract?" The answer is that, since revenue is fixed, it should minimise the present value of its costs profile. So the company should choose whatever combination of satellite start-up dates and development options that minimises, at the company level, PV costs or, equivalently, that minimises UTC.

2. Make conjectures. It is not the end of the world if they are proved wrong. (What is important, though, is that you try to establish whether they are right or wrong). In this case, you might guess that it would be sensible, at the level of your individual field, to aim to minimise UTC.

3. Construct examples. These can be used for testing conjectures and also for gaining a feel for what is and what is not important. The examples might be typical cases, but it often more instructive to consider extreme cases. These test the boundaries of the problem. Moreover, extreme cases can often be constructed so as to be more easily solved than the general case. For our problem, useful examples to consider would be, firstly, the case where there are a large number of identical satellites; secondly, the case in which the field unit has a single good satellite and a large number of mediocre projects which neither make nor lose any money.

4. Draw a diagram if at all possible.

5. Try to characterise the problem. What sort of solution might you expect to find? In this case, it is easy to show that changes to well numbers for one satellite affect the choice of development plan for all the other satellites. For a given mathematical model, the task of finding the optimal solution is inherently a multi-field problem. It may be possible to find a rule (similar to "Minimise UTC for the field") that can be applied on a field-by-field basis. However, such a rule would only give an approximation to the optimal solution (except, possibly, if it were applied in an iterative fashion). Hence, we know that there is no point looking for an exact solution for our problem.

6. Try to transform your problem into a problem with a known solution. For the specific case in which the field unit has a single good satellite and a large number of mediocre projects which neither lose nor make money, our problem can be transformed into the problem of maximising NPV for a single field.

7. Examine your basic assumptions. If there does not appear to be a mathematically simple solution, examine your basic assumptions. In reality, most problems have simple (not necessarily easy) solutions - it is very difficult, for example, to find any physics that cannot be described in a couple of equations. However, the wrong basic assumptions may make it impossible to find the simple solution. Even well-established assumptions may prove to be questionable - for example, the assumption (used in Shell, but not in academic work on rock physics) that the net effective rock stress = Overburden - pore pressure.

8. Dimensional analysis. This is particularly useful as a check on the accuracy of a long derivation of a formula.

9. Plough through the algebra. If even if you do not know where you are going, you are likely to find some interesting results. The formal manipulation of mathematical symbols is the mechanisation of a logical thinking process. This mechanisation helps us in the same way that machines help us in the physical world.

10. Consider what is the minimum you need to know in order to make a decision. For example, in deciding whether to go ahead with an exploration well, it may be sufficient to establish whether the probability of success (POS) is greater than 25%. It may be unnecessary to estimate the actual value of the POS.

The problem that we have chosen to illustrate some of the tricks is as follows: -

You are a team leader for a team working on a recent discovery of a satellite to a large gas field. The field and its satellites are being produced under a long-term gas supply contract, in which the gas production total is effectively fixed. You have been told to come up with a field development plan.

You have asked what should be the economic objective of the field development plan and have been told "Don't worry too much about the details of the gas supply contract. Concentrate on putting together a sound technical plan. It will then be passed on to the field business management, where the plan will be ranked against other projects - other satellites and changes such as platform gas compression upgrades."

Despite this, you decide that it is important to do at least some optimisation of the development plan. The overall problem is "What should be the economic objective of the field development plan?" In solving this problem, your first task is to decide whether, in general, it is correct to aim to minimise UTC for your satellite.

It is suggested that you examine the example in which the field unit has a single good satellite and a large number of mediocre projects which neither lose nor make money.

5.1 Effects of minimising project UTC

It can be shown that minimising project UTC is not always a good idea. When UTC on a good project is minimised, there is a danger that this brings forward much higher expenditure on other, less good projects. This may harm the company economics.

Consider the example in which the company has a single good project and a large number of mediocre projects, which neither lose nor make money. In these circumstances, since

Overall NPV = NPV of project 1 + NPV of project 2 + ...

the overall NPV is simply the NPV of the single good project. Consequently, the optimal solution for the company as a whole is to maximise the NPV of the single good project.

Let us compare this with the possible effects of minimising project UTC. If the good project has very low fixed costs, in the sense of chapter 4, (so all costs are proportional to the number of wells drilled), then

Project UTC = (PV of project costs) / (PV of project production)

= (PV of well costs) / (PV of well production)

Hence, the development that minimises UTC is the one that maximises (PV of well production) i.e. the one that minimises well decline i.e. as few as possible wells are drilled.

With very few wells, then, although project UTC is low, so is project NPV and consequently so is the company NPV. This establishes that minimising project UTC is not a good general rule.

The final task the problem is for you to hazard a guess at what might be a better general rule. As a hint, you might like to spend a little time considering the following line of thought. The benefits to the field unit of producing gas from your satellite are not those of selling the gas (because production from your field implies foregoing production elsewhere), but those of avoiding the expenditure necessary to produce from elsewhere. The costs to the field unit of producing gas from your field are clear. If you have both costs and benefits, you may be in a position to do some fairly conventional economic calculations.

5.2 A general rule for optimising field development plans under a production limit

When optimising field development plans under a fixed production limit, a good objective is to maximise a "shadow" NPV for each project - the NPV calculated using a shadow gas or oil price equal to the UTC of production from alternative projects.

As mentioned above, benefits to the field unit of producing gas from your satellite are not those of selling the gas (because production from your satellite implies foregoing production elsewhere), but those of avoiding the expenditure necessary for other projects. Let us assume that this expenditure is some constant, L, times the amount of gas produced.

This is a reasonable approximation, since the changes made to our project will only have a second order effect on other projects. It is likely that the expenditure per unit of gas necessary to produce from other projects will rise in time (as the better projects are done first), but such effects will be no larger than the normal effects of oil and gas price uncertainty.

Clearly, the constant, L, can be identified with an "average" UTC of production from alternative projects. These alternative projects are generally those lower down the ranking than our project. After all, if we decide to cut down on the planned production from our project, the extra production needed will come from bringing forward production from projects lower down the ranking. The extra production will not come from better projects, because these are already being carried out in preference to our project.

Such a definition of the benefits of production from our project ensures that the overall field unit NPV changes in line with any changes in the calculated benefits. Since field unit NPV obviously changes in line with any changes to costs of production on our project, field unit NPV can be expressed as

Field unit NPV = A + PV(Benefits from our project) - PV(Costs from our project)

where A is constant whatever we do to our project.

Hence, it is clear that field unit NPV is maximised whenever we maximise the shadow NPV defined as

Shadow NPV = PV(benefits) - PV(Costs) = PV(L x gas production) - PV(Costs)

This shadow NPV is clearly similar to a conventional NPV, except that, instead of a gas price, it uses L, the UTC of production from alternative projects.

A mathematical proof of this line of argument is given in Appendix 2 (although it is expressed in terms of oil production under an OPEC quota, rather than gas production under a fixed production contract). It starts from the same key assumption that costs per unit of alternative production is constant. The proof shows that, if this assumption is correct, then it is optimal to maximise shadow NPV. However, it should be borne in mind that the assumption will never be fully correct, so the validity of maximising shadow NPV depends on how close the assumption is to reality.

It can also be shown that this method gives the correct answers for the two examples mentioned in 5.1 - firstly, the case where there are a large number of identical projects; secondly, the case in which the field unit has a single good project and a large number of mediocre projects.

PDO (Petroleum Development Oman) is a Shell company that has been operating, since 1998, under a fixed production limit. It is believed that they use a two stage procedure in field development planning that is entirely equivalent to maximising shadow NPV.

The first stage in the PDO procedure, for optimising planning for a given field, is to produce an approximate ranking in terms of UTC. This enables them to identify the UTC of alternative production and the costs saved per barrel of production from the given field. The second stage is to try to maximise the costs saved by production from the given field.

Now

Costs saved from given field = $PV(production_{this\ field})$ x Costs saved per barrel = $PV(production_{this\ field})$ x $(UTC_{alt}-UTC_{this\ field})$

but since $UTC_{this\ field} = PV(Costs_{\ this\ field}) / PV(production_{\ this\ field})$

Costs saved = $PV(production_{\ this\ field})$ x UTC_{alt} - $PV(Costs_{\ this\ field})$ = Shadow NPV

i.e. The two procedures are identical.

We consider that the shadow NPV formulation is slightly preferable to the PDO procedure. With the NPV formulation, it is clear that one can use existing methods derived for maximising conventional NPV. With the PDO procedure, this is less obvious.

Using simple models for business decisions

You may have some doubts as to whether it is appropriate to use simple models for making important business decisions that may involve many millions of pounds or dollars. To answer these questions, we will examine the uses of simple models.

There are three main ways in which simple models can be used

1. To provide starting points for more detailed work.
2. As independent checks on the results of more detailed work.
3. Directly by themselves for decision making.

In the first two, simple models are used in addition to conventional models. One would to have more confidence in the results of a project that used both approaches than in a project that relied entirely on reservoir simulation. The important question then is "Is this extra confidence worth the effort? Is it worthwhile to construct simple models in addition to conventional models?"

Our belief is that it is not only worthwhile, but essential, for three reasons. Firstly, conventional models are so complicated and difficult to audit that even elaborate review and challenge procedures are no guarantee that there are not errors in the models. The best check on simulator predictions is to reproduce the same results by different calculations - simple models. Secondly, conventional methods, while suitable for evaluating development proposals, are very helpful in designing developments - they are like a navigation system that can tell you where you are, but not how to get to your destination. In contrast, simple models can be very helpful for design work. Thirdly, simple models can give good initial results quickly. This can speed up projects considerably, saving time and money.

The third way of using simple models - by themselves - is not recommended, but it is no worse than using reservoir simulators by themselves. Our experience has been that when there is a big disagreement between simulator results and the results of a simple model, it is more often the simulator that is in error.

Simple models often require more thought than does a reservoir simulation model and so may be regarded as "more difficult." But these difficulties can be overcome with practice. It is worth it, because the prize is better field development plans.

Problem 6 - Optimising water injection on Catbert Alpha

The final problem of the course is posed in a less structured way than the earlier problems. The reason for this is that real problems do not come neatly structured. It is hoped that by this stage you will be able to start providing your own structure for solving the problem.

The Catbert Alpha platform provides water injection from a pair of high-pressure pumps (with a common discharge manifold) along four routes

- directly into Catbert Block I platform wells
- through a flowline to the Catbert Undersea Manifold Centre (UMC) and then into the Catbert Block I UMC wells
- directly into Catbert Block II wells
- through booster pumps and along a flowline to the Dogbert drilling centre and then into the Dogbert wells.

With Dogbert showing better than expected production characteristics, it was found that the water-injection capacity on the Catbert Alpha platform was less than desired. The bottle-neck was the twin high-pressure pumps. These were limited to a maximum pressure of 218 atm. At this pressure, the pumps could deliver 14,000 m^3/day of water.

The more recent development, Dogbert, was at a much lower water-cut (10%) than the other fields (typically 55%). Moreover, because of different tax regimes, production from Dogbert and Block II gave twice as much net revenue per barrel of oil produced than did Catbert Block I.

It was argued that since the net revenue per barrel of water injected was four times higher for Dogbert than for Block I, net revenue could be increased by switching some injection from Block I to Dogbert. Hence, injection into Dogbert should be maximised, which requires all its injection wells to be open and the discharge pressure of the HP pumps to be maximised. To keep the discharge pressure up, some Block I injection wells should be closed in as required.

You are asked to address the problem "Is this argument correct? How might you go about deciding the optimal injection policy?" To answer these questions, you may well require further information, which can be obtained from the workshop leaders.

Hint - think through what happens if you start with all the wells open and then procede to shut-in some of the Block I injectors. Also, consider the operating envelope for the pumps, illustrated below.

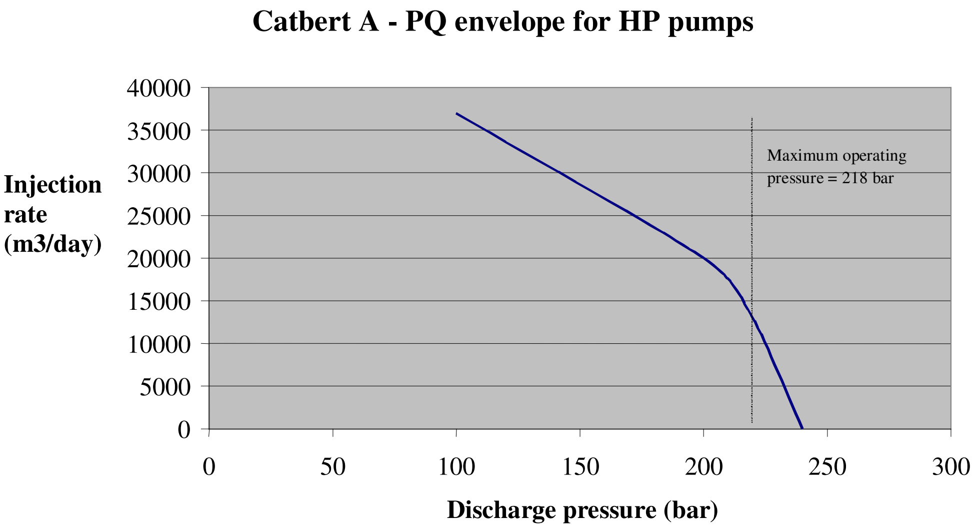

6.1 Water injection optimisation

The argument for prioritising Dogbert is incorrect. As presented, it relies on the implicit assumption that injection water can be simply switched around - that to get an extra 1 m^3/day into Dogbert, it suffices to cut Block I injection by 1 m^3/day. In reality, it turned out to be necessary to cut Block I injection by approximately 30 m^3/day. Even if Dogbert injection is four times more valuable than Block I injection, it is clearly uneconomic to prioritise Dogbert.

To start, it is useful to think through the process of shutting-in or choking back Block I injectors. If all the wells in the system are open and can, between them, take more water than the pumps can deliver at maximum pressure, then the pressure at the pump discharge will fall quickly until equilibrium is reached. At equilibrium, the pumps can deliver as much as the wells can take. The equilibrium pressure will be below the maximum pump pressure. It is likely that the pumps can deliver more at a lower discharge pressure than they can at the maximum pressure. To verify this, we can examine the pump PQ (pressure - flowrate) relationship, which shows how the maximum flowrate possible from the pumps varies with the discharge pressure.

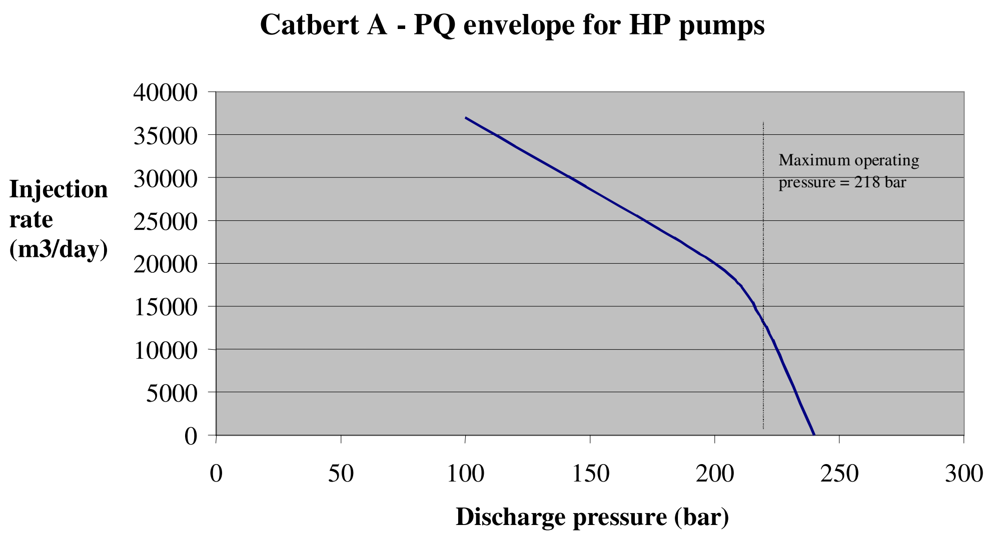

It is clear from the pump PQ relationship that if you start with all the wells open and then gradually shut Block I injectors, you will increase the HP discharge pressure and reduce the overall flow-rate from the pump. If, then, the overall flow-rate is being reduced, it is clearly not correct to assume that

1 m^3/day less into Block I = 1 m^3/day more into Dogbert.

To make an estimate of what is the correct ratio of (injection lost to Block I): (injection gained by Dogbert), one can follow a procedure similar to calculating well production rates by finding the intersection of the well inflow performance curve and the tubing lift curve when both are plotted against BHP.

One needs to know what injection rate each reservoir would take as a function of the pump discharge pressure, as in the plot below.

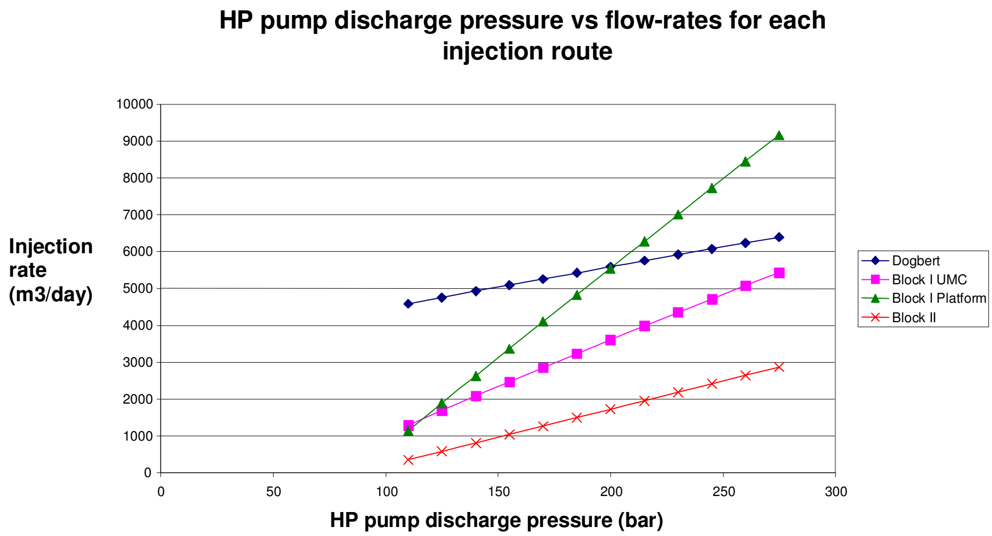

When the combined flow-rate into all the open injection routes (i.e. what the fields can take) is plotted together with the HP pumps' P-Q curve (i.e. what the pumps can deliver), the intersection of the two curves gives the flow-rate that would occur for that combination of open injection routes.

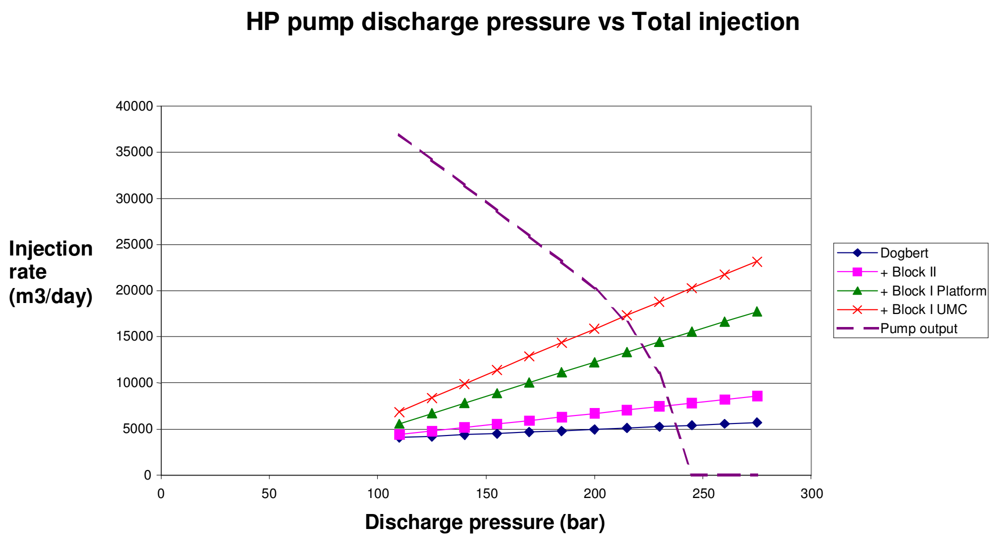

From the gradient of the Dogbert injection rate vs pump discharge pressure curve, one can calculate how much the discharge pressure has to be increased to inject an extra 1 m^3/day into Dogbert - 0.1 bar, as it turns out.

This increase of 0.1 bar would increase injection into Block II as well - by 1.5 m^3/day (again, this can be calculated from gradient of the Block II PQ curve). It would decrease injection from the pump by 30 m^3/day (from the gradient of the pump PQ curve).

Since

Change in injection from the pump = change in Dogbert injection + change in Block II injection + change in Block I injection

it follows that

Change in Block I injection = Change in injection from the pump - change in Dogbert injection - change in Block II injection = -30 -1 -1.5 = -32.5 m^3/day.

Even though injection into Dogbert and Block II is more valuable than injection into Block I, the loss of injection into Block I is so large that one would lose money if one shut in Block I wells to prioritise Dogbert. The actual calculation of the reservoir PQ relationships can be done either with runs of the reservoir simulators or by repeated application of the well inflow arguments - finding the equilibrium flow-rates.

Forecasting infill well performance in a mature field

The following example addresses the issue of forecasting well performance of infill wells in a mature field, where a large amount of real production history is available. The example demonstrates that real fields can show certain simple patterns of behaviour, even in complicated matters such as individual well performance.

It is typical that during the drilling history of a field, the most prolific wells are drilled first, on the basis that those targets are most readily identified from seismic data. This is followed by the drilling of progressively less attractive targets, as the prime locations become more difficult to identify. In other words a law of diminishing return, or a "creaming curve" (in explorationists terms) is established.

The following diagram illustrates the creaming curve for exploration successes in the UKCS.

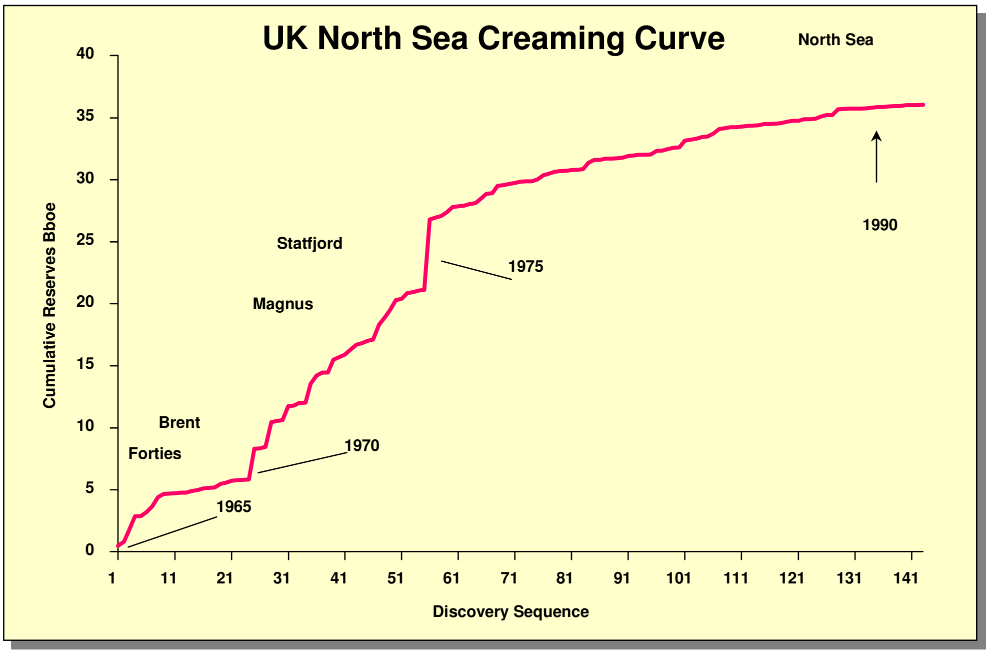

In an infill drilling situation, an equivalent of the creaming curve can be used to predict the performance of sequential infill drilling targets, and this can be based on the well performance to date.

The following example will illustrate the application of this tool in the prediction of infill well performance in a large water-flooded oil field, with more than 20 years of production history.

The following diagram shows the recovery per well for all wells drilled since the beginning of production in the field. A clear trend in the decline per well is apparent. Each well is represented by its expected UR, and redrills for mechanical purposes are added to the original well.

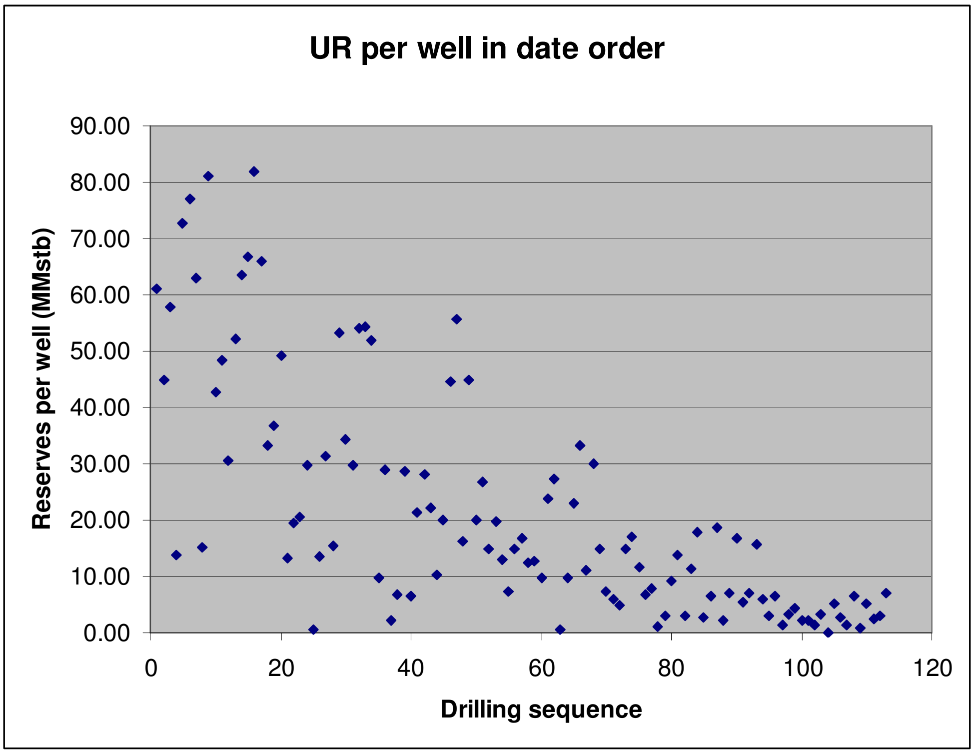

The diagram below shows the recovery per well for all producers drilled into the field since between 1992 and 2001. In order to ensure that the results do not penalise the wells drilled last, each well's ultimate recovery has been predicted from decline curve analysis on each well. In other words the recovery per well is the forecast UR over its lifetime, not cumulative production to 2001.

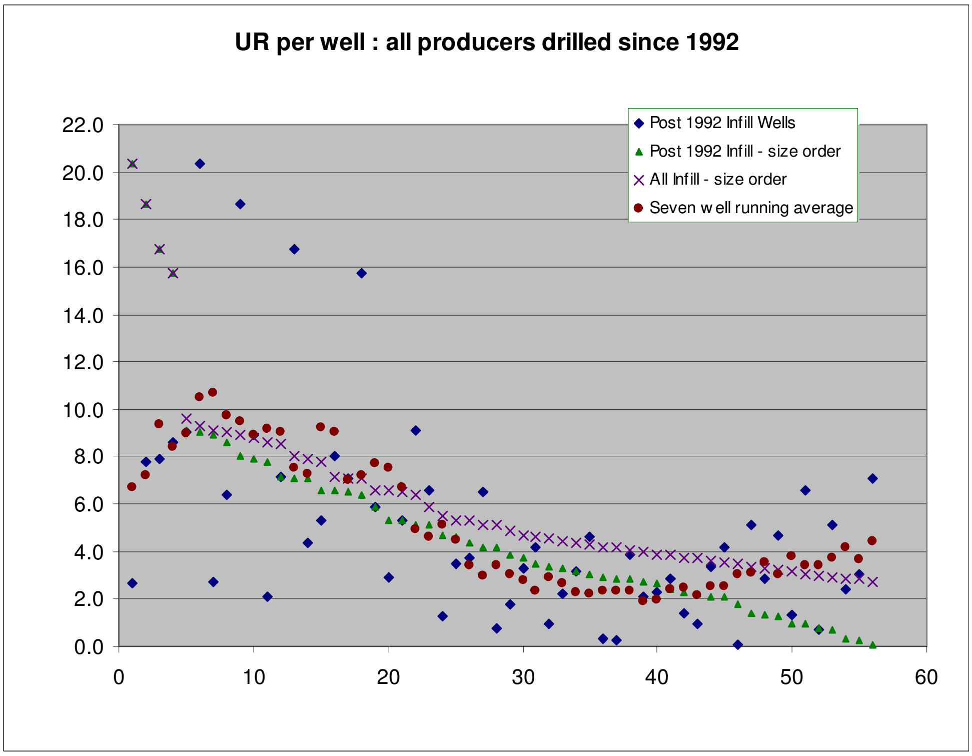

The above plot indicates a general trend of diminishing returns, through which a trend-line may be fitted as a tool for predicting future infill well performance. The problem is one of deciding which trend to adopt.

Clearly when the data is organised in terms of size order, it will trend downwards. In terms of the performance of the development team, the data in drilling sequence is more informative. In the above example, the historic data follows a reasonable trend up to infill well number 40. However, subsequent wells 40-60 show a significant improvement in the infill well performance, which requires an explanation. In this case it is the application of 4-D seismic interpretation methods, which improves the definition of attractive infill targets, and results in enhanced performance compared to the recent history.

The above data set includes all producers drilled since 1992. Not all of these were strictly speaking infill wells, as some were replacements for mechanical failures in existing producers, some were sidetracks of disappointing development wells, and some were late development wells drilled into parts of the field where development had been purposely delayed. A round of data screening identified the true infill wells, and only these are included in the plot below.

.

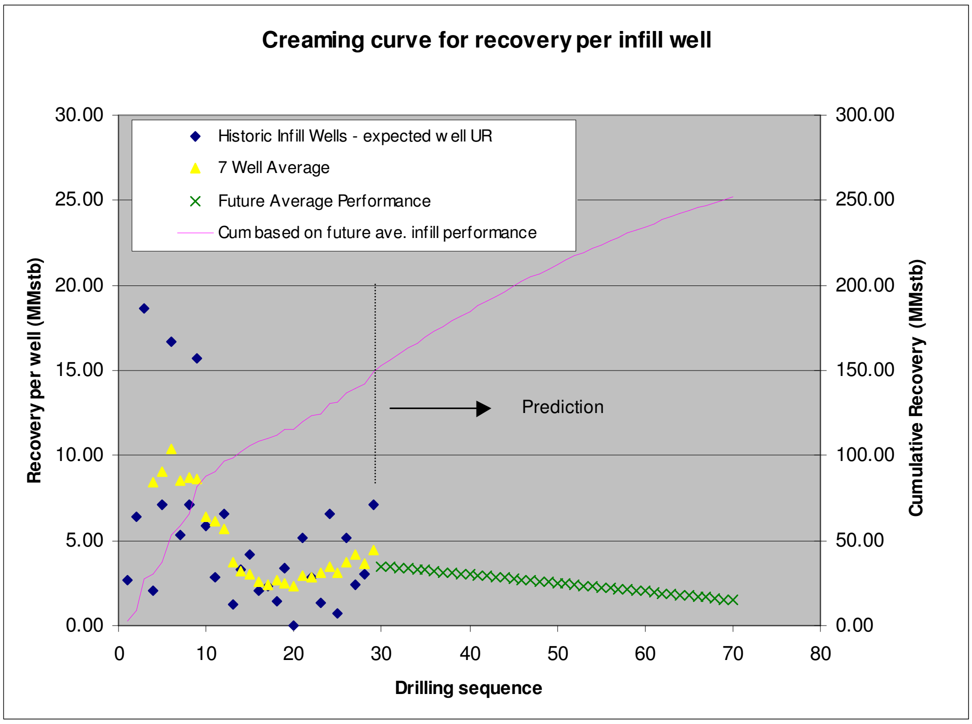

The above plot shows the cleaned up data set, comprising only true infill producers drilled since 1992. The creaming curve is used to predict the future infill well performance, providing an estimated remaining reserve of 100 MMstb from infill drilling. The well performance ranges from 3.5 MMstb for the first infill well, down to 1.5 MMstb for the fortieth infill location. This performance can be used to create an infill production profile.

The full field reservoir simulation model for this field provides an alternative way of predicting remaining reserves. With 40 infill targets included in the simulation model, an infill reserves figure of 125 MMstb is predicted, ie 25% above that generated by the analytical approach.

The following graph shows what the performance of the infill wells would need to be to meet this prediction.

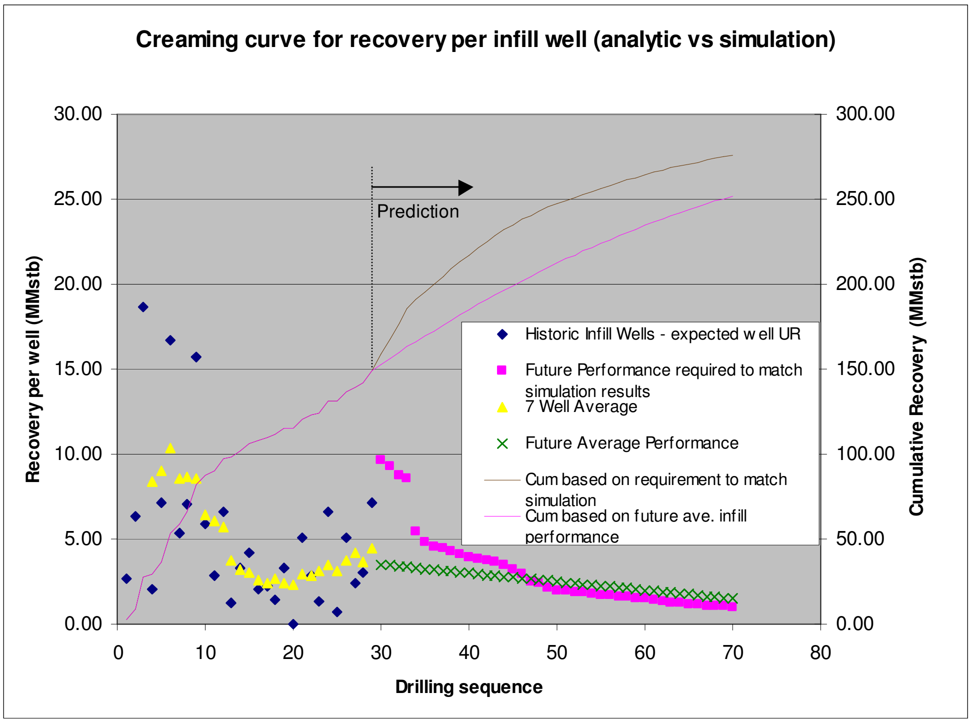

It is clear that the performance required to meet the simulation prediction will need to significantly outstrip the historic performance, and therefore casts some doubt as to the confidence which should be attached to the simulation forecast.

One of the reasons why the simulation model may be over-estimating reserves from infill activity is that it is probably modelling the reservoir as more homogeneous than reality. This is of course a common concern with the reservoir description.

In summary, an analytical approach has been used to sense-check a simulation result, demonstrating that some caution should be used in using the simulation result for assessing the benefit of the long-term infill programme. The two methods may be considered as a "top-down" approach (the simulation model) and a "bottoms-up" approach by extrapolating the historic well performance to date.

Non-exponential decline

As has been discussed, it is often useful to manipulate or create approximate production profiles. To be realistic, these profiles should honour the total amount of recoverable reserves and their shapes should match approximately whatever is typical for the type of reservoir considered. For fields showing exponential decline, it is straight-forward to construct a reasonable decline curve, using the relationship

Expected decline rate = (Production rate) / (Remaining technically recoverable reserves)

For fields showing hyperbolic decline, it is much more difficult to construct a reasonable relationship. The reason for this is that, in going from exponential to hyperbolic decline, the implicit strong relationship between profiles and recoverable reserves has been lost.

An alternative extension of exponential decline is given below, called the "C-curve." Derived during work on the Alba field (heavy oil, with underlying water, in a high permeability turbidite sand), this is based on extending the cumulative production equation, rather than the production rates equation, and yields the following relationship

$$Q_{oil} = R.\left[1 - \sqrt[b]{\frac{a}{a - 1 + e^{\left(a.b.Q_{liquid}/R\right)}}}\right]$$

where

Q_{oil} = cummulative oil production

R = technically recoverable reserves

Q_{liquid} = cummulative liquid production

a = constant (0.213 for Alba Extreme South)

b = constant (5.62 for Alba Extreme South)

Initial assumptions

The justification for exponential decline is usually quoted as the empirical observation that

$$\frac{dq_{oil}}{dt} \Big/ q_{oil} = -a$$

where

q_{oil} = oil production rate

t = time

a = constant.

This is then extended to hyperbolic decline in the form

$$\frac{dq_{oil}}{dt} \Big/ q_{oil} = -a.q_{oil}^{b}$$

However, examination of the physics of simple systems that exhibit exponential decline (such as the decay of radioactive particles, or production of a gas field under pure depletion) suggests that the fundamental driver is that the decay/production rate is proportional to the remaining population/reserves. In our case, this would be

$$\frac{dQ_{oil}}{dt} \Big/ \left(R - Q_{oil}\right) = -a$$

where

Q_{oil} = cummulative oil production

R = ultimate recovery

The "C-curve" method is to extend this relationship to the more general form

$$\frac{dQ_{oil}}{dt} \Big/ (R - Q_{oil}) = -(a + \beta (R - Q_{oil})^b)$$

or, in a dimensionaless form

$$\frac{dr}{dx} \Big/ r = -(a + \beta r^b)$$

where r = 1 – Q_{oil} / R

and x is any measure of field aging – very Q_{liquid} / R

The plot below illustrates how this relationship is sufficiently general to give a good match to typical individual well declines in Alba, as predicted in reservoir simulation.

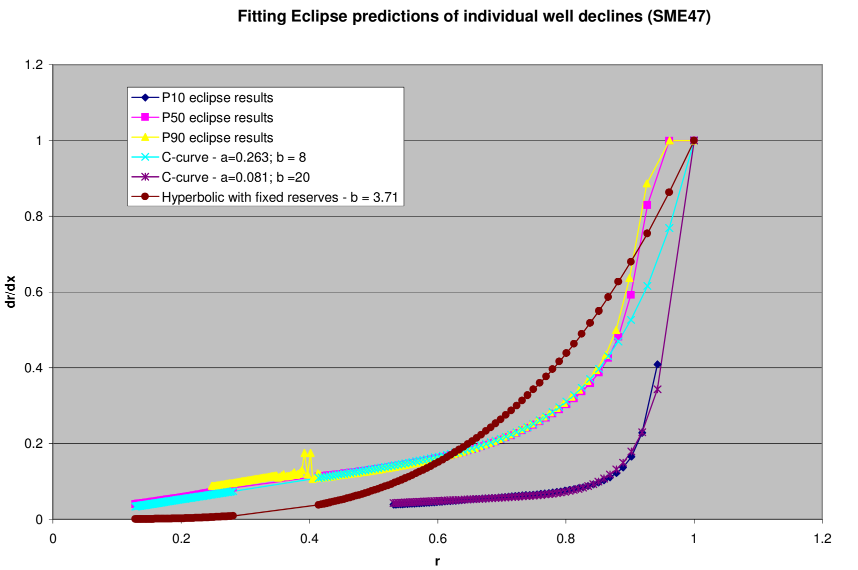

The equivalent relationship in hyperbolic decline is

$$\frac{dr}{dx} \Big/ r = -\beta r^b$$

It can be seen from the plot that hyperbolic decline does not give a good fit, at least not with the same figure for ultimate reserves. On Alba, it was found that hyperbolic decline gave reasonable fits only with high exponents (i.e. close to harmonic decline) and enormously high reserves figures (e.g. 1100 million bbls for Alba Extreme South, compared to a STOIIP of 320 million bbls).

The key idea behind this approach is that it is the dr/dx – r relationship that matters for creating life-of-well or life-of-field production profiles. The exact form of the relationship chosen does not matter much, providing it is sufficiently general to fit the shape of decline as observed in reality or as predicted in Eclipse. The C-curve relationship was chosen so as to be easily solvable to yield formulae that can be easily used and manipulated.

Solving for cummulative oil

Starting from the initial equation

$$\frac{dr}{dx} \Big/ r = -(a + \beta r^b) \qquad \text{(A)}$$

the variables can be split as follows

$$\frac{1}{r\left(a + \beta r^b\right)} dr = -dx \qquad \text{(B)}$$

Integrating both sides gives

$$\frac{1}{\mathrm{ab}} \ln\left(\frac{r^b}{a + \beta r^b}\right) = -\left(x + \alpha\right) \qquad \text{(C)}$$

where α is a constant.

[Proof of integration of left hand side –

$$\frac{d}{dr}\left(\frac{1}{\mathrm{ab}} \ln\left(\frac{r^b}{a + \beta r^b}\right)\right) = \frac{d}{dr}\left(\frac{1}{\mathrm{ab}} \ln\left(\frac{1}{ar^{-b} + \beta}\right)\right)$$

$$= \frac{1}{ab}\left(ar^{-b} + \beta\right) \cdot \frac{-1}{\left(ar^{-b} + \beta\right)^2} \cdot \left(-abr^{-n-1}\right) = \frac{1}{r\left(a + \beta r^n\right)}$$

.

]

Solving equation (C) for r gives

$$r = \sqrt[b]{\frac{a}{e^{ab(x+\alpha)} - \beta}}$$

The usual boundary conditions include

a) production starts with dry oil i.e. dr/dx = 1 when r =1

b) At the start of production (i.e. when r =1) x = 0

These boundary conditions allow us to express the α and β in terms of other variables, as follows

Condition (a) implies (from equation (A))

-1 = -(a + β)

i.e.

$\beta = 1 - a$

Applying condition (b) to equation (B) gives

$$\frac{1}{ab}\ln\left(\frac{1}{a+(1-a)1^{b}}\right) = -(0+\alpha)$$

i.e.

$\alpha = 0$

Applying these values of α and β to equation (C) gives

$$r = \sqrt[b]{\frac{a}{e^{abx}+a-1}}$$

Changing from r and x to R, Q_{liquid} and Q_{oil} gives

$$Q_{oil} = R.\left[1-\sqrt[b]{\frac{a}{a-1+e^{\left(a.b.Q_{liquid}/R\right)}}}\right]$$

Interestingly, it was found that the r, x formulation could also be used in rather different way. If x was taken to be the number of wells drilled on the field, and r to be the fraction of movable oil remaining at the end of field life after the x wells had been produced to very high water-cuts, then

$$\frac{dr}{dx} = -\beta r^{2}$$

was found to give a very good match to a set of simulator runs with varying numbers of wells, as is illustrated in the two plots below.

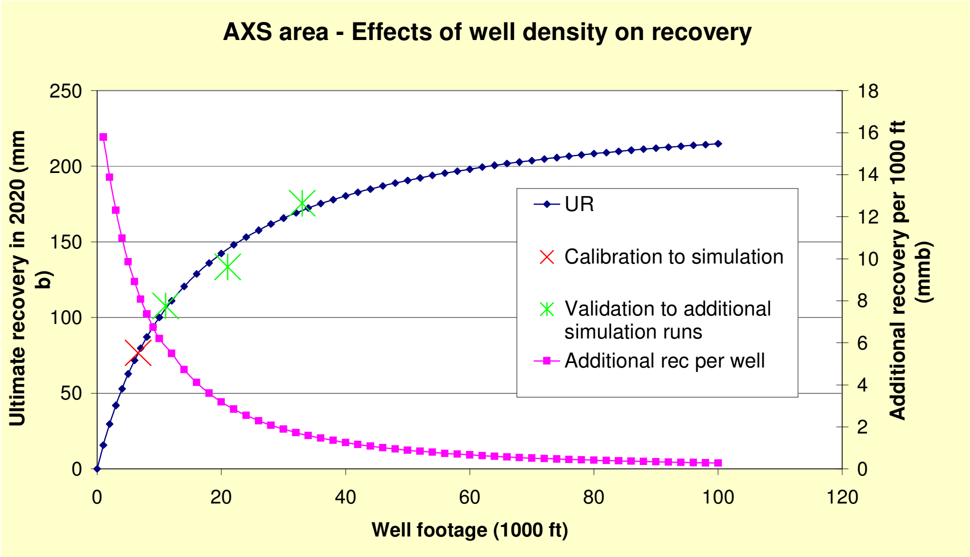

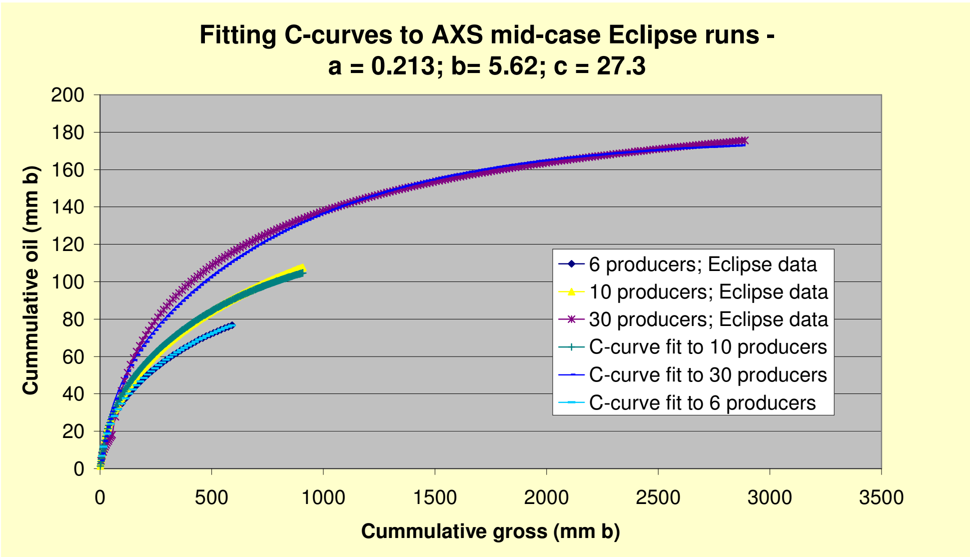

Appendix 1 – Mathematical proofs of the formulas in Problem 4

Assumptions and Definitions

Let us make the following assumptions to describe an oil-field development

1) Let N be the number of wells drilled. All the wells start up at time t=0.
2) If the field were run for an infinite time, the total production would be R (the technically recoverable reserves), independent of the number of wells.
3) The initial production rate per well is q, independent of the number of wells i.e. it is not affected by well spacing.

4) The field oil production rate follows exponential decline i.e.
   Field oil production rate = initial rate x e^{-at} = $N.q.e^{-at}$
5) The net oil price is a constant, L, after all taxes and deductions.
6) The net capital costs can be expressed as D + C.N; all capital expenditure happens at time t=0.
7) Net opex can be expressed as E.N per unit time. (The year is probably the most appropriate unit of time, but any unit can be used, providing it is the same for both E and q).

Theorem 1 – Expressing NPV as a function of well numbers

The NPV of a field run until abandonment can be expressed as

$$NPV = \frac{R.L}{1+\left(R.\ln(1+d)\big/ q.N\right)}.\left(1-\alpha.e^{-q.N.Tab/R}\right) - \frac{N.E}{\ln(1+d)}.(1-\alpha)-(C.N+D)$$

where $\alpha = (1+d)^{-Tab}$ and T_{ab} (abandonment time)= $\frac{R}{q.N}.\ln\left(\frac{L.q}{E}\right)$

and

q – initial oil production per well per year, averaged over all wells, including injectors

N – total number of wells, including injectors

R – technical reserves i.e. the amount of oil that could be recovered if the field were run for a very long time

L – net revenue per unit of oil (i.e. after all taxes and royalties, including profit tax)

d – discount rate

C – net capital cost per well

D – net capital costs not related to numbers of wells, e.g. roads and pipelines

E – net opex per well

Proof

NPV can be broken down into the component parts of the cash-flow

NPV = NPV(revenue stream) + NPV(Opex) + NPV(Capex)

where the NPV(Opex) and NPV(Capex) are, of course, negative.

Start by calculating the NPV of the revenue stream:-

By the assumption that there is exponential decline

Oil production rate = $N.q.e^{-at}$ = $N.q.e^{-(q.N/R)t}$

[Since, by the definition of technical reserves

$$R = \int_0^\infty N.q.e^{-a.t}dt = \left(-\frac{N.q}{a}.e^{-at}\right]_0^\infty = N.q/a$$

then a = N.q/R]

Oil revenue per unit time = (production rate) x (net oil price) = $N.L.q.e^{-q.N.t/R}$

By the definition of NPV

$$\text{NPV of revenue stream } = \int_0^{Tab} \frac{\text{Revenue per unit time}}{(1+d)^t} dt = \int_0^{Tab} N.L.q.e^{-(q.N/R+\ln(1+d)).t} dt$$

$$(\text{Since } (1+d)^t = e^{\ln((1+d)^{-t})} = e^{-t.\ln(1+d)})$$

$$= \left. \frac{N.L.q.e^{-(q.N/R+\ln(1+d))t}}{-(\frac{q.N}{R}+\ln(1+d))} \right|_0^{Tab} = \frac{N.L.q}{\left(\frac{q.N}{R}+\ln(1+d)\right)}.\left(1-e^{-(q.N/R+\ln(1+d)).Tab}\right)$$

$$= \frac{R.L}{\left(1+\frac{R.\ln(1+d)}{q.N}\right)}.\left(1-e^{-\ln(1+d).Tab}.e^{-q.N.Tab/R}\right) = \frac{R.L}{\left(1+\frac{R.\ln(1+d)}{q.N}\right)}.\left(1-(1+d)^{-Tab}.e^{-q.N.Tab/R}\right)$$

$$= \frac{R.L}{\left(1+\frac{R.\ln(1+d)}{q.N}\right)}.\left(1-\alpha.e^{-q.N.Tab/R}\right)$$

Looking at the Opex cash-flow

Opex per unit time = -N.E

$$\text{NPV of opex } = \int_0^{Tab} \frac{-N.E}{(1+d)^t} dt = -\int_0^{Tab} N.E.e^{-\ln(1+d).t} dt = \left. \frac{N.E.e^{-\ln(1+d).t}}{\ln(1+d)} \right|_0^{Tab} = -\frac{N.E}{\ln(1+d)}.\left(1-e^{-\ln(1+d).Tab}\right)$$

$$= -\frac{N.E}{\ln(1+d)}.\left(1-(1-d)^{-Tab}\right) = -\frac{N.E}{\ln(1+d)}.(1-\alpha)$$

Looking at Capex, since all the capital expenditure is assumed to occur at time t=0,

NPV of Capex = - (C.N + D)

By adding NPV(Revenue), NPV(Opex) and NPV(Capex), one obtains the desired formula.

Q.E.D

Theorem 2 – Number of wells giving the highest NPV

<u>Part I</u> – The number of wells giving the highest NPV, N_{opt}, can be calculated (iteratively) from the expressions

$$Nopt = \frac{R}{q}.\ln(1+d).\left[\sqrt{\frac{L.q-\alpha.E}{C.\ln(1+d)+E-\alpha.E}}x\left(\sqrt{\varepsilon^2+1}-\varepsilon\right)-1\right]$$

$$\text{where } \varepsilon = \frac{\alpha.E.\ln\left(\frac{L.q}{E}\right)}{\sqrt{4(L.q-\alpha.E).(C.\ln(1+d)+E-\alpha.E)}}$$

and $\alpha = (1+d)^{-Tab}$

$$and\ Tab = \frac{R}{q.N}.\ln\left(\frac{L.q}{E}\right)$$

Part II – If

a) L.q ≥ 2.[C.ln(1+d) + E] (i.e. the well is reasonably profitable at first) and
b) 0.2 ≥ α (equates, for d = 8%, to a field life ≥ 20 years)

then the approximation of setting ε = 0 gives approximate value for the number of wells, N_{approx} within the following bounds

$$1.31\ N_{opt} \geq N_{approx} \geq N_{opt}$$

Proof of Part I

Consider NPV as a function of well numbers and abandonment time. For a normally behaved function

$$\text{NPV is a maximum} \Leftrightarrow \frac{\partial \text{NPV}}{\partial \text{N}} = \frac{\partial \text{NPV}}{\partial \text{Tab}} = 0$$

Finding Tab such that $\frac{\partial \text{NPV}}{\partial \text{Tab}} = 0$ can be done by taking the partial derivatives or, more simply, by seeing that this occurs when revenue has dropped until it equals operating expenditure.

i.e. $N.L.q.e^{-q.N.Tab/R} = N.E$

so $-q.N.T_{ab}/R = \ln(E/(L.q))$

so $Tab = R/(q.N)\ .\ \ln(L.q/E)$

(N.B Also, for Tab optimal, $e^{-q.N.Tab/R}$ = E/(L.q))

For well numbers, using Theorem 1, and writing out the expression in full (not using α, for example)

$$\frac{\partial NPV}{\partial N} = \frac{\partial}{\partial N}\left(N\left[\frac{Lq}{\frac{qN}{R} + \ln(1+d)} . \left(1 - (1+d)^{-Tab}.e^{-(qNTab/R)}\right) - \frac{E(1-(1+d)^{-Tab})}{\ln(1+d)} - C \right] - D \right)$$

$$= \frac{Lq}{\frac{qN}{R} + \ln(1+d)} . \left(1 - (1+d)^{-Tab}.e^{-(qNTab/R)}\right) - \frac{E(1-(1+d)^{-Tab})}{\ln(1+d)} - C$$

$$- \frac{N.Lq}{\left(\frac{qN}{R} + \ln(1+d)\right)^2} . \left(1 - (1+d)^{-Tab}.e^{-(qNTab/R)}\right) \frac{q}{R}$$

$$+ \frac{N.Lq}{\frac{qN}{R} + \ln(1+d)} . \left(1 - (1+d)^{-Tab}.e^{-(qNTab/R)}\right)\left(-\frac{q.Tab}{R}\right)$$

Setting this partial derivative equal to zero, multiplying both sides of the resultant equation by $(qN/R + \ln(1+d))^2$, then using α as shorthand for $(1+d)^{-Tab}$, and the relationships that apply when T_{ab} is optimal, $e^{-q.N.Tab/R}$ = E/(L.q) etc, we get the following equation. N.B. This equation only applies when T_{ab} is optimal.

$$\left(\frac{q.N}{R} + \ln(1+d)\right).L.q.\left(1 - \alpha.\frac{E}{L.q}\right) - \left(\frac{q.N}{R} + \ln(1+d)\right)^2 . \left(\frac{E.(1-\alpha)}{\ln(1+d)} + C\right) - \frac{qN}{R}.L.q.\left(1 - \alpha\frac{E}{L.q}\right)$$

$$+ \ln\left(\frac{L.q}{E}\right).L.q.\left(\frac{q.N}{R} + \ln(1+d)\right).\alpha.\frac{E}{L.q} = 0$$

Cancelling terms and multiplying both sides by –1 gives

$$\left(C + \frac{E.(1-\alpha)}{\ln(1+d)}\right).\left(\frac{q.N}{R} + \ln(1+d)\right)^2 + \alpha.E.\ln\left(\frac{Lq}{E}\right)\left(\frac{q.N}{R} + \ln(1+d)\right) - \ln(1+d).(L.q - \alpha.E) = 0$$

This can be considered to a quadratic equation, with the “x” term being (q.N/R + ln(1+d)) and the other terms being as follows:-

a = C + E.(1-α)/(ln(1+d)

b = α.E.ln(L.q/E)

c = -ln(1+d).(L.q - α.E)

We will proceed here in Part I of this proof to solve the quadratic equation. In Part II, we will show that the “b” term has little effect, and can be ignored. (It is interesting to consider how the “b” term arose. Consider the NPV of the field. If the number of wells increases, then abandonment is brought forward, so the term representing the present value of the oil lost at abandonment is increased. This decreases the overall NPV, but as can be imagined, such effects are small. This will be proved later, in Part II).

So, ignoring the negative solution, the solution of the quadratic equation is:

$$x = \frac{\sqrt{b^2 - 4a.c} - b}{2a}$$

Re - arranging this gives

$$x = \sqrt{\frac{-c}{a}}.\left[\sqrt{1 + \left(\frac{b}{\sqrt{-4a.c}}\right)^2} - \frac{b}{\sqrt{-4a.c}}\right]$$

$$\textit{Defining } \varepsilon \text{ by } \varepsilon = \frac{b}{\sqrt{-4a.c}} = \frac{\alpha.E.\ln\left(\frac{L.q}{E}\right)}{\sqrt{4.(C.\ln(1+d) + E - \alpha.E).(L.q - \alpha.E)}}$$

and expanding x, c and a gives

$$\frac{q.N}{R} + \ln(1+d) = \sqrt{\frac{\ln(1+d).(L.q - \alpha.E)}{C + \frac{E.(1-\alpha)}{\ln(1+d)}}}.\left(\sqrt{\varepsilon^2 + 1} - \varepsilon\right)$$

$$= \ln(1+d)\sqrt{\frac{(L.q - \alpha.E)}{C.\ln(1+d) + E - \alpha.E}}.\left(\sqrt{\varepsilon^2 + 1} - \varepsilon\right)$$

Re - arranging this equation gives

$$\mathrm{N} = \frac{R}{q}.\ln(1+d).\left[\sqrt{\frac{(L.q - \alpha.E)}{C.\ln(1+d) + E - \alpha.E}}.\left(\sqrt{\varepsilon^2 + 1} - \varepsilon\right) - 1\right]$$

Proof of Theorem 2, Part II

As a first step, it is useful to note that, for $\varepsilon \geq 0$ (which is the case, providing $L.q \geq E$ – i.e. Year 1 net revenue for a well is greater than the opex for the well)

$1 \geq (\sqrt{(\varepsilon 2+1)} - \varepsilon) \geq 1-\varepsilon$ hence, Napprox ≥ Nopt

$$[\text{ Proof - } \left(\sqrt{\varepsilon^2 + 1} - \varepsilon\right)^2 = \varepsilon^2 + 1 - 2\varepsilon\sqrt{\varepsilon^2 + 1} + \varepsilon^2 = 1 + 2\varepsilon.\left(\varepsilon - \sqrt{\varepsilon^2 + 1}\right) \leq 1$$

since $\sqrt{\varepsilon^2 + 1} \geq \varepsilon$]

Examining ε^2 and dividing both the denominator and quotient by $(L.q)^2$ gives

$$\varepsilon^2 = \frac{\left(\ln\left(L.q/E\right)\right)^2 \alpha^2 \left(E/L.q\right)^2}{4\left[\dfrac{C.\ln(1+d)+(1-\alpha).E}{L.q}\right]\left(1-\dfrac{\alpha E}{L.q}\right)}$$

Since $\dfrac{\mathrm{E}}{\mathrm{L.q}} \leq \dfrac{C.\ln(1+d)+E}{L.q} \leq 0.5$ (by assumption 1) and $\alpha \leq 0.2$ (by assumption 2)

$$\left(1-\frac{\alpha E}{L.q}\right) \geq 1-0.2x0.5 = 0.9$$

Also

$$\frac{C.\ln(1+d)+(1-\alpha).E}{L.q} \geq (1-\alpha).\frac{E}{L.q} \geq 0.8\frac{E}{L.q}$$

Hence,

$$\varepsilon^2 \leq \frac{\left(\ln\left(L.q/E\right)\right)^2 0.2^2 \left(E/L.q\right)^2}{4x0.8x\left(E/L.q\right)x0.9} = 0.014\left(\ln\left(L.q/E\right)\right)^2.\frac{E}{L.q}$$

It can be easily shown that for 1 ≥ E/(L.q) ≥ 0, the maximum value of $(\ln(L.q/E))^2$.E/(L.q) is achieved when E/(L.q) = $1/(e^2)$ = 0.1353, which gives $(\ln(L.q/E))^2$.E/(L.q) = 0.541.

[Proof - Differentiate x.$(\ln(x))^2$ and set to zero].

Hence ε ≤ 0.014 x 0.541 = 0.00757

ε ≤ 0.087

$(\sqrt{(1+\varepsilon^2)} - \varepsilon) \geq (1-\varepsilon) \geq (1-0.087) = 0.93$

Before moving on to look at a lower limit for N_{opt}/N_{approx}, it is useful to establish a couple of small lemmas.

Lemma A

For u, v, w such that u ≥ v > 0 and v > w ≥ 0,

$$\frac{u-w}{v-w} \geq \frac{u}{v}$$

$$\Pr oof\ \frac{u-w}{v-w} - \frac{u}{v} = \frac{uv - vw - uv + uw}{v(v-w)} = \frac{(u-v).w}{v.(v-w)} \geq 0$$

Lemma B

For w constant and greater than zero, the function f (y) = [(1-w).y –1] / (y-1) is strictly increasing (i.e. $y_1 < y_2 \Rightarrow f(y_1) < f(y_2)$).

Proof Expressing f (y) as 1 – wy / (y-1) gives

$$\frac{df}{dy} = \frac{-w}{y-1} + \frac{w.y}{(y-1)} = \frac{w}{(y-1)} > 0$$

Moving back to the main proof, let us establish a lower bound on N_{opt}/N_{approx}

$$\frac{Nopt}{Napprox}=\frac{\sqrt{\frac{Lq-\alpha.E}{C.\ln(1+d)+E-\alpha.E}}.\left(\sqrt{1+\varepsilon^2}-\varepsilon\right)-1}{\sqrt{\frac{Lq-\alpha.E}{C.\ln(1+d)+E-\alpha.E}}-1}\geq\frac{0.93\sqrt{\frac{Lq-\alpha.E}{C.\ln(1+d)+E-\alpha.E}}.-1}{\sqrt{\frac{Lq-\alpha.E}{C.\ln(1+d)+E-\alpha.E}}-1}$$

By Lemma A and assumption (i)

$$\sqrt{\frac{Lq-\alpha.E}{C.\ln(1+d)+E-\alpha.E}}\geq\sqrt{\frac{Lq}{C.\ln(1+d)+E}}\geq\sqrt{2}$$

By Lemma B

$$\frac{0.93\sqrt{\frac{Lq-\alpha.E}{C.\ln(1+d)+E-\alpha.E}}.-1}{\sqrt{\frac{Lq-\alpha.E}{C.\ln(1+d)+E-\alpha.E}}-1}\geq\frac{0.93\sqrt{2}-1}{\sqrt{2}-1}=0.76$$

Hence $N_{opt}/N_{approx}\geq 0.76$

or equivalently $N_{approx}\leq 1.31N_{opt}$

Combining this result with the result established at the beginning of the Part II of the proof gives

$1.31\ N_{opt}\geq N_{approx}\geq N_{opt}$

Q.E.D.

Note – In order to express the expression for the approximate optimal number of wells in the form first quoted, i.e.

$$\text{Number of wells that maximises NPV} = \frac{\text{R.d}}{\text{q}}.\left(\sqrt{\text{L} - \frac{\alpha.\text{E}}{\text{q}}}-\sqrt{\frac{C.d+E-\alpha.E}{q}}\right)\Bigg/\sqrt{\frac{C.d+E-\alpha.E}{q}}$$

it suffices to rearrange the expression slightly and to note that for normal values of d (the discount rate), ln(1+d) ≈ d (e.g. ln(1+0.08) = 0.077)

Theorem 3 – NPV for a development with the approximately optimal number of wells

The NPV of a development with the approximately optimal number of wells, as defined in Theorem 2, is

$$NPV=R.\left(\sqrt{L-\frac{\alpha.E}{q}}-\sqrt{\frac{C.\ln(1+d)+(1-\alpha).E}{q}}\right)^2-D$$

where α = (1+d)-Tab

and Tab = R / (N.q) . ln(L.q/E)

<u>Proof</u>

Combining the results from Theorems 1 and 2

$$NPV = \frac{R.\ln(1+d)}{q}\left(\sqrt{\frac{L.q-\alpha.E}{C.\ln(1+d)+(1-\alpha).E}}-1\right)x$$

$$\left[\frac{L.q}{\frac{q.R.\ln(1+d)}{R.q}.\left(\sqrt{\frac{L.q-\alpha.E}{C.\ln(1+d)+(1-\alpha).E}}-1\right)+\ln(1+d)}.\left(1-\frac{\alpha.E}{L.q}\right)-\frac{E.(1-\alpha)}{\ln(1+d)}-C\right]$$

$$-D$$

$$= R.\left(\frac{\sqrt{a}}{\sqrt{b}}-1\right).\left[\frac{L}{\left(\sqrt{a}/\sqrt{b}\right)}.\left(1-\frac{\alpha.E}{L.q}\right)-\frac{E.(1-\alpha)}{q}-\frac{C.\ln(1+d)}{q}\right]-D$$

(where a = L.q - α.E

and b = C.ln(1+d) +(1-α).E)

$$= R.\left(\frac{\sqrt{a}}{\sqrt{b}}-1\right).\left[\frac{a.\sqrt{b}}{q.\sqrt{a}}-\frac{b}{q}\right]-D = R.\left(\sqrt{\frac{a}{q}}-\sqrt{\frac{b}{q}}\right).\left[\sqrt{\frac{a}{q}}-\sqrt{\frac{b}{q}}\right]-D$$

$$= R.\left(\sqrt{L-\frac{\alpha.E}{q}}-\sqrt{\frac{C.\ln(1+d)+(1-\alpha).E}{q}}\right)^2-D$$

Q.E.D.

Appendix 2 - Mathematical treatment of shadow NPV

Consider a company operating under a fixed production limit. Consider an individual, newly discovered field whose field development plan is to be optimised.

Using the notation

q for production rates

- Q for production totals
- c for expenditure rates
- C for cumulative expenditure
- subscript T for the company totals across all the fields

and dividing the company fields into three groups

1. fields started in production before our field
2. our field, whose development plan we are trying to optimise
3. fields that will be started in production after our field

let us assume the following

a) The aim of the company is to maximise the NPV of its cashflow.

b) The company cashflow is the sum of the cashflows from the individual fields
c) The company production rate, $q_T(t)$, cannot exceed a certain limit (note - t = time). There is an ample supply of profitable oil-fields, so the company will always aim to be at this limit, so q_T can be considered to be fixed.
d) The company production rate, q_T, can be expressed in terms of the production rates from the three groups of fields as follows:
$q_T = q_1(t) + q_2(t) + q_3(t)$ for all t.
e) If production is brought forward or moved back, then cumulative costs can be similarly brought forward or moved back in time. Formally, the cumulative costs associated with each group of fields can be expressed as functions of the cumulative production from that group
$C_1(Q_1)$, $C_2(Q_2)$, $C_3(Q_3)$
Remark - In reality, for each group i, there would be considerable expenditure before any production occurs i.e. while $Q_i = 0$. However, we can model this by assigning a little "pseudo" production to the initial expenditure. What we are aiming to capture is the way that expenditure is brought forward or put back with production.
f) Let us consider changes (parameterised by a real number x) to the field development plan for our field. We assume that changes to the field development plan affect our field's production profile and expenditure i.e. Q_2 and C_2 can be considered to be functions of x and t,
$Q_2(x,t)$ and $C_2(x,t)$

g) Changes to our field development plan do not affect production from fields that started production before our's i.e. x does not affect $Q_1(t)$ and $C_1(t)$.
h) Changes to our field development plan do not affect the production designs for subsequent fields, but our changes may affect timings for these fields i.e. the function $C_3(Q_3)$ is unchanged, but Q_3 should be considered to be a function of x and t. In fact, from assumption (d), it is straight-forward to conclude
$Q_3(x,t) = q_T.t - Q_1(t) - Q_2(x,t)$

i) C_3 is approximately a linear function of Q_3. i.e. dC_3/dQ_3 is a constant, which we will call the "marginal UTC of group 3", $UTCM_3$.
j) All our functions are continuous and differentiable etc.

Let us define the "shadow NPV", NPV_{SH}, for our field to be

$$NPV_{SH} = \int_{t=0}^{t=\infty} \frac{UTCM_3.q_2(t) - c_2(t)}{(1+d)^t} dt$$

i.e. the "NPV" calculated using $UTCM_3$ instead of the actual oil price.

The result we want to prove is the following -

Shadow NPV theorem

Our field development plan is optimal with respect to x (the parameter describing the choice of developments) if and only if the plan maximises shadow NPV.

Proof

Since the company aim is to maximise company NPV (assumption 1) and the oil production rate is fixed, our development plan is optimal iff it minimises the present value of the total costs profile.

The total costs profile can be split into three parts, corresponding to the three groups of fields. Let us consider the costs for the third group - the fields coming into production after our field - and examine how these costs are affected by our development plan.

$$\frac{\partial}{\partial x}(\text{PV costs of Group 3}) = \frac{\partial}{\partial x}\int \frac{\frac{\partial}{\partial t} C_3(Q_3(x,t))}{(1+d)^t}\,dt$$

$$= \frac{\partial}{\partial x}\int \frac{\frac{dC_3}{dQ_3}(Q_3(x,t)) \times \frac{\partial}{\partial t} Q_3(x,t)}{(1+d)^t}\,dt$$

Since $\frac{\partial}{\partial t} Q_3(x,t) = q_3(x,t) = q_T - q_1(t) - q_2(x,t)$

and $\frac{dC_3}{dQ_3}$ is a constant, UTCM_3

$$\frac{\partial}{\partial x}(\text{PV costs of Group 3}) = \frac{\partial}{\partial x}\int \frac{UTCM_3.(q_T - q_1(t) - q_2(x,t))}{(1+d)^t}\,dt$$

Since q_T and q_1 are unaffected by changes in x,

$$\frac{\partial}{\partial x}(\text{PV costs of Group 3}) = \frac{\partial}{\partial x}\int \frac{-UTCM_3.q_2(x,t)}{(1+d)^t}\,dt$$

Since the PV costs of Group 1 is independent of x,

$$\frac{\partial}{\partial x}(PV \text{ of combined costs}) = \frac{\partial}{\partial x}(PV \text{ costs of Group } 1) + \frac{\partial}{\partial x}(PV \text{ costs of Group } 2)$$

$$+ \frac{\partial}{\partial x}(PV \text{ of costs of Group } 3)$$

$$= 0 + \frac{\partial}{\partial x}(PV \text{ costs of Group } 2) - \frac{\partial}{\partial x}\int \frac{UTCM_3.q_2(x,t)}{(1+d)^t}dt$$

$$= -\frac{\partial}{\partial x}(\text{Shadow NPV})$$

Hence, with respect to our choice of field development plan for our particular field, maximising the shadow NPV for our field is equivalent to minimising PV of combined costs for the company as a whole.

Q.E.D

Just as with less formal reasoning, it can be useful to look back over the proof to identify what are the most important assumptions. Mathematical proofs are usually well suited to such an examination, since they make all the steps in an argument explicit and require us to be clear about precisely what our assumptions are. A good proof, moreover, should have been worked over and refined so that the main structure of the mathematical argument is clear, with uninteresting algebraic manipulation kept in the back-ground.

Here, it is clear that the key step in the proof is changing from a company-wide optimisation problem to a field-specific optimisation policy i.e. from (PV costs of Group 3) to ($UTCM_3 \times$ production of Group 2) in the $\partial/\partial x$ terms.

This step relies on assumption (i), which states
"C_3 is approximately a linear function of Q_3. i.e. dC_3/dQ_3 is a constant, $UTCM_3$."

This can be verified in a simple way by plotting planned cumulative expenditure verses planned cumulative production for the group of fields that will come into production after our field.