



Introduction to probability distribution and petroleum project economics

*Pre-course reading for
PETRAD's courses on petroleum exploration & production*

**Dr Erik Jarlsby
2007**

*www.petrad.no
www.erikjarlsby.no*

I N T R O D U C T I O N

PETRAD conducts courses of 8 weeks on the "upstream" part of the oil and gas industry, that is exploration and production, as well as shorter courses under particular arrangements. Course participants are exposed to a range of technical, operational, commercial and regulatory issues. In the full courses and in some of the shorter courses, the learning process is supported by participants' contribution to group work in which a simulated process of resource assessment and development is carried out.

Economic analysis is a central feature of the upstream oil and gas industry, as indeed in many other industries. Being an extractive industry that is exposed to a range of formidable uncertainties, the upstream petroleum industry faces particular complexity in assessing the economics of its business opportunities. Our courses address a number of these challenges. Project economics and uncertainty are recurring themes both in course presentations and participants' case work.

This paper is mandatory pre-course reading for our course participants. Its purpose is to introduce participants to concepts and terminology which they will encounter during the courses and which are presumed prior knowledge. Some course participants are already thoroughly familiar with these concepts, and can make light use of this paper. For others, it will be unfamiliar ground, and these participants are encouraged to study this paper prior to commencing the course. In our experience, there are also numerous professionals in the industry who encounter these terms and concepts frequently without ever really being introduced to them. We hope that the paper will be useful for them also.

The paper deals with two major themes: Probability and project economics. They are related, because the understanding of uncertainty and probabilities is important in economic evaluation. The presentation here is brief, providing the participant only with the basics that are essential in the course context. A large literature exists covering both fields, and a few useful books for further reading are referenced at the end of the paper.

The paper includes calculation exercises which you are encouraged to perform on a Microsoft Excel or similar computer spreadsheet. An Excel file with proposed solutions to the exercises is also attached. At the end of the paper is a References section, which lists a selection of books that offer a broader guidance to the topics covered in this paper, and related topics.

Introduction to probability distribution and petroleum project economics

C O N T E N T S

Part 1: Probability distribution	page 1
1.1 Introducing probability and normal distribution	1
1.2 Cumulative probability distribution	3
1.3. Generating a view on expected outcome and uncertainty	4
1.4. Continuous and discrete distributions	6
1.5. Normal vs. lognormal distribution	7
1.6. Triangular distribution	9
1.7. Other probability distributions	10
Part 2: Project economics	12
2.1. Cash flows	12
2.2 The time value of money	14
2.3. Net Present Value and Internal Rate of Return	16
2.4. Dealing with inflation in project economics	17
2.5 Choice of discount rates and investment decisions	18
2.6 Complex cash flows: Beware of the dual IRRs	19
2.7. Criticism and new approaches to project valuation	21
2.8 Economics and strategy	22
Part 3: Exercises	24
E-1: Generating normal distribution in an Excel spreadsheet	24
E-2: The Standard Normal Distribution	29
E-3: Applying normal distribution to a manufacturing quality problem	29
E-4: Cash flow analysis for an investment in machinery	30
E-5: Cash flow discounting	30
E-6: Present values with a simple decision tree	31
Further reading	32

PART 1 : PROBABILITY DISTRIBUTION

This part will take you through a rather informal introduction to probability distribution, and some of the key terms and concepts used there. Probability analysis is part of the wider field of statistical analysis. The accompanying exercises in Part 3 are designed to assist and enhance the learning experience from this paper.

1.1 Introducing probability and normal distribution

In business analysis we often have to make assumptions about something that we do not know for certain. Consider the following questions:

- (1) By how many \$ per barrel will the price of Brent oil rise or fall over the next 30 days?
- (2) What will be the price of light, sweet crude oil in 2050, on average ?
- (3) Will the UK be a net importer of oil in 2020 ?
- (4) How much will it cost to drill a well ?
- (5) What is the total recoverable amount of oil in the Brent oil field ?

Any of these questions might be relevant for your business, depending on what your business is. We can't know the answers for certain, at least not yet. In business analysis, it is quite common to take a deterministic approach to such questions, that is, we figure out one reasonable answer to each question and base our analysis on that number only, ignoring the uncertainty. The alternative is to treat these parameters as stochastic variables, i.e. uncertainties.

Parameters such as those reflected in the 5 questions above, might be required inputs for a decision making process. If we use a sophisticated software tool for economic analysis, we are likely to find there a facility for treating such parameters stochastically. That means, we recognise the uncertainty and don't have just one assumed answer for each parameter.

We turn to the first question raised above: Assume we need to form a view of how much the price of Brent crude oil will change over the next 30 days. Our basic view is that the price will fall by \$2. We will revert in chapter 1.3 to how we might arrive at such a view. But we recognise that this is far from certain. The price may fall further, or it may instead rise.

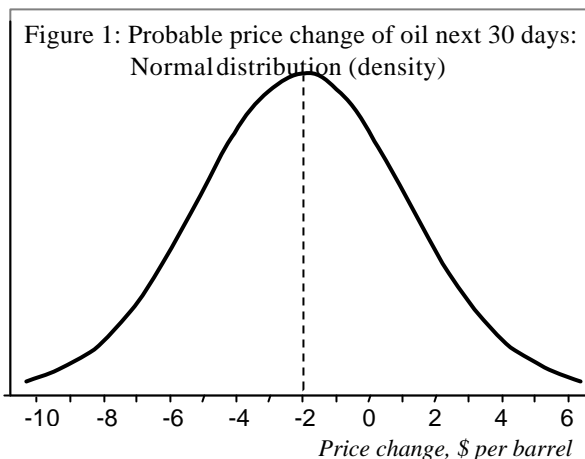
Figure 1 gives an example on how we might describe such an assessment of probable price changes. We have a bell-shaped curve which peaks at -2, which is our basic view of what will happen. Further out to the left, the curve shows lower values, indicating that such larger price drops are possible but distinctly less possible. To the right, there are also lower values indicating the lesser likelihood of price increases. We see that the curve on both sides never

Box 1: Deterministic – Stochastic

The terms *deterministic* and *stochastic* are usually applied to processes. A stochastic process involves uncertainty, whereas a deterministic process does not. Stochastic problem solving or modelling therefore involves variables that are treated as uncertain. Instead of being given just one assumed value, such variables are given a *probability distribution*, i.e. a mathematical description of the likelihood of different values that the variable can have.

quite reaches zero. This indicates that there is no definite limit to how much the price of oil can rise or fall, but changes beyond the shown range have very small probabilities.

A bell-shaped curve like Figure 1 is called a normal distribution. This is an important concept in statistics, and often used as basis for probability analysis. Many phenomena in life which are quantifiable and can vary, tend to be variable in ways that resemble a normal distribution.



To arrive at a normal distribution like the one in figure 1, we need to define two parameters: First, the mean, i.e. our view of what is most likely to happen, in this case a price change of -2 \$ per barrel.

Box 2: Mean; Standard Deviation

The *mean* is the average value of a population, a sample or a probability distribution, when all values are counted equally. Expressed mathematically:

$$\bar{x} = (\sum x) / n$$

where x is each observed or expected value, \bar{x} is the mean of all x , and n is the number of x .

The *standard deviation* is a measure of the tendency for the individual values to differ from the mean. A large standard deviation implies a large spread between values or a large uncertainty. Its formula (which you don't need to memorise), is:

$$\sqrt{\frac{\sum (x - \bar{x})^2}{n}}$$

Second, the standard deviation, which is a measure of how variable and uncertain this phenomenon is. In this case, a standard deviation of 3 \$ per barrel was used.

The standard deviation has the quality that in a normal distribution, about 68% of all occurrences fall within the range of the standard deviation on both sides of the mean.¹ In this case, that range is from -5 to +1 \$ per barrel, as possible changes from today's oil price.² This implies, there is a likelihood of about 0.68 (i.e. 68 per cent) that the oil price in 30 days will be within -5 \$ to +1 \$ per barrel as a change from today's price. There is therefore a probability of 0.32 that the price of oil will end up outside that range.³ The normal distribution is symmetrical around the

mean. The probability of 0.32 therefore has two equal parts: A probability of 0.16 on the negative side (price falling more than \$5) and also a probability of 0.16 on the positive side (price rising more than \$1). In summary, we are implying the following probabilities about the price of oil 30 days from now:

- Probability of 0.16 that the price will fall by at least \$5 per barrel
- Probability of 0.16 that the price will rise by at least \$1 per barrel
- Probability of 0.68 that the price will be within the range of -5 to +1 \$ from today's price.

¹The following applies to normal distributions generally: Approximately 68% of all occurrences are within +/- 1 standard deviation from the mean; 95% are within +/- 2 standard deviations and 99.5% are within +/- 3 standard deviations.

² We chose -2 as the mean and 3 as the standard deviation. This gives a range from -5 to +1, because -2-3=-5 and -2+3=1.

³ The combined probability of all possibilities is by definition 1 (or 100%). The number 0.32 referred in the text is therefore arrived at as $1 - 0.68 = 0.32$, where 0.68 was the probability that the price will end up within the described range.

1.2. Cumulative probability distribution

Many people find that a curve like Figure 1 illustrates a point well: We easily see the points that (1) a price change of -2\$ is considered likely; and (2) price changes both lower and higher than -2 are possible, but less likely the further away from the mean expectation of -2\$. But this graph is not easy to work with beyond that. Note that it has no vertical scale, for which additional assumptions would be necessary.

The bell-shaped curve in figure 1 is also called a density function. Figure 2 contains the same information as figure 1, but shown as an inverse cumulative distribution graph. On the vertical axis we now have a scale from 0 to 1, which indicates probabilities. A probability of 0 for an outcome means that it certainly will not occur, and a probability of 1 for an outcome means that it will certainly occur. The graph curve indicates, for each price change value found on the horizontal scale, the probability that the price in 30 days will be at least that much.

The vertical dotted line at the price change value of -2, which is our chosen mean, crosses the graph curve at the cumulative probability value 0.5. The implication of this is that there is a probability of 0.5 (50 per cent) that the price of oil in 30 days will be not lower than -2 \$ per barrel as a change from today's price.

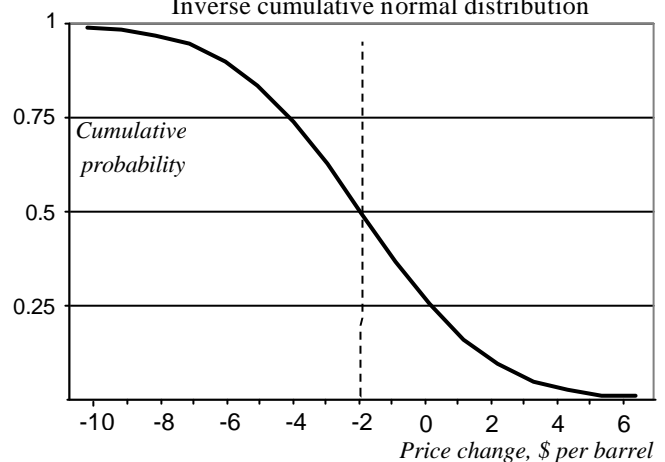
The graph in figure 2 shows an *inverse* cumulative distribution, because it starts with something that has a very high probability (that the price in 30 days will be not lower than \$10 below the current price). The cumulative probability curve falls as the horizontal scale shows higher figures for the oil price, because this lowers the probability of the price staying above those levels.

Here is an example of how you can read the graph in figure 2. You might ask the question, what is the probability of the oil price rising, even though our best estimate is that it will fall by \$2 per barrel? Identify the value 0 on the horizontal scale. This represents today's price of oil (zero change). Use a ruler to make a vertical line up from the value 0. It hits the graph curve near the horizontal line indicating 0.25 as a cumulative probability. This implies a 0.25 chance that the oil price in 30 days will be higher than today's price. Correspondingly, there is a 0.75 chance that it will fall.

This probability of 0.25 for a price rise results from the two parameters which we chose: Our mean expectation of a \$2 fall in prices, and our standard deviation of \$3, as a measure how uncertain we consider the estimate to be. A bigger standard deviation would make the curves in figures 1 and 2 flatter and more extended, and would increase the probability of a price rise even with the same mean expectation of a \$2 drop.

EXERCISE: At this stage it is suggested you do exercise E-1, which you find in the exercise section at the back of this paper. You will need a computer with Microsoft Excel (or similar)

Figure 2: Probable price change of oil next 30 days:
Inverse cumulative normal distribution



spreadsheet software. The exercise will show you how to generate normal distributions with the spreadsheet, and make graphs like figures 1 and 2. You need some basic knowledge of using the spreadsheet, but the exercise will take you through the things you need to do.

1.3. Generating a view on expected value and uncertainty

How can we generate a view of the mean and the standard deviation in a case like this? The process is essentially judgemental, but should be supported by professional knowledge about the phenomenon in question. This knowledge may come from empirical (actual) observations of the past, or from theoretical models. Here is how you might go about it in the case of the 30-days oil price forecast, in three steps:

As a first step, we need to define precisely the variable which we analyse. In this case, we need to get clear about which oil price we are looking at. For Brent crude oil, there are essentially two types of prices which we can observe. There are the prices from futures trading, which for Brent crude oil traditionally has been on the International Petroleum Exchange (IPE) in London (www.theipe.com).⁴ Alternatively, you can choose the prices that are formed in direct trade between buyers and sellers of oil cargoes, as reported daily by agencies such as Platts (www.platts.com) or Petroleum Argus (www.argusonline.com). In this example we select the Dated Brent quotation by Platts. Since this quotation comes a range (for instance, \$45.20 - \$45.25, depending on what happened in the market during that particular day), we specifically select the high end of the range (or the low end or the average, as long as we are clear about it). The Dated Brent quotation gives us prices that are agreed today for deliveries of oil to be made within a couple of weeks. By selecting this quotation, we have also specified the quality of oil which will underlie our analysis (the Brent Blend, whose quality is well known to market participants) and the mode of delivery (FOB at the loading terminal on the Shetland Islands), both of which are important conditions affecting the price of oil.⁵

The second step is to form a mean expectation about the outcome of the defined variable, in this case Platts' Dated Brent quotation 30 days hence. If you follow the market closely, you may form your independent view dependent on your perception of short-term supply balances, evolving political tensions, the weather in important consuming areas, and other factors which you expect will influence the outcome. You may have access to market reports and the views of analysts whom you are inclined to believe. For oil prices, you can also look at forward quotations. Platts and Argus quote not only prices for prompt delivery (Dated cargoes), but also prices agreed now for delivery in the following few months. These prices for later deliveries are usually different from the dated prices.⁶ You can find the Platts quote for the month that covers your target date, and this gives you an idea of what the market

⁴ As of 2005, the New York Mercantile Exchange (www.nymex.com) is launching its own trade of Brent futures in competition with the IPE.

⁵ The observable market prices of oil are specific to each quality grade, delivery location and time of delivery. By evolved habit, the prices for various grades of oil in many parts of the world are agreed in terms of Brent crude, i.e. differentials to the observable market prices for Brent crude. For this and other reasons the world market prices for oil tend to be highly synchronised, so that movements in the price of Brent crude tends to occur simultaneously with similar price movements for other crudes all over the world. The role of Brent as a "marker" crude (i.e. a widely used indication of oil prices in general) may diminish in the future since production is declining in the oil fields feeding into the Brent pipeline and terminal system.

⁶ It is referred to as Backwardation when prices for later delivery are lower than for prompt delivery. The opposite is referred to as Contango.

expects generally for that period. But we should recognise that the Platts quote covers the full month as an average, whereas we are trying to form an expectation for the price in one particular day during that month. This is one reason why we may choose a different estimate than the one derived directly from Platts.

The third step, and the trickiest, is to form a view of the uncertainty. For now we will assume that the uncertainty can be expressed in the form of a normal distribution, but that will not always be the case. In later sections we will review other distributions.

Uncertainty is expressed as the standard deviation, in the form of a number. The standard deviation is not the kind of parameter that most people could estimate well off-hand, so we should definitely look for some facts to back up our assessment. We can, for instance, look to records of past events of a similar nature and analyse their variability. This would certainly be possible for the price of Brent crude oil, for which time series of all past quotations can be obtained. You can then select 30-day intervals covering the entire history of Brent price quotations, or a part of it which you find to be relevant. Here are two ways in which you can use such data:

- (1) You can determine the Dated Brent price at the beginning and end of each 30-days period, which gives you a set of price changes. From this price set you can calculate the mean (which will be close to zero) and the standard deviation. This standard deviation indicates how much the price of oil tends to change over 30-day periods.

For calculating the mean, you can use the spreadsheet function `AVERAGE(range)`, where *range* is the series of cells containing all your monthly changes. For the standard deviation, you can use the spreadsheet function `STDEVP(range)`. The P in the function name is appropriate if you work with the entire population of past cases, as was suggested here. In other cases, we might not analyse the full population but a sample, in which case the standard deviation should be calculated with the formula `STDEV(range)`.

- (2) You can make a similar calculation as

Box 3: Population and sample

A *population* consists of all elements of the relevant set. (If you are concerned with exploration wells drilled in the North Sea, then the population is all such wells ever drilled in the North Sea.) A *sample* is a selection of the population. (It would be a smaller number of wells that have been drilled in the North Sea.) For practical reasons we may study a sample in order to make conclusions about the entire set, provided we work with a sample that represents the population well. Preferably it should be obtained by *random sampling*, i.e. a process by which all members of the population had the same chance of becoming members of the sample.

Box 4: Common notations

It is common to use different mathematical symbols representing the mean and the standard deviation, depending on whether they apply to a population or a sample:

- m (Greek small “mju”) – Mean of a population
- S (Greek small “sigma”) – Standard deviation of a population
- \bar{x} (read “x bar”) – Mean of a sample, when observed values are denoted as x
- s - Standard deviation of a sample

above, but comparing the end-of-period dated Brent prices to the forward quotations at the beginning of the period. This gives you an indication of how the price change over 30 days has tended to be different from what market participants expected at the beginning of the period, and is conceptually a better proxy for the uncertainty in this case, than looking at changes to the dated price. The standard deviation from this calculation may be taken as a reasonable indication of the

uncertainty which you might want to attach to your own estimate of -\$2 per barrel.

In addition to using quantitative analysis of the past as a reference for selecting a standard deviation, you should also take your knowledge of the situation at hand into consideration. This may lead you to setting the standard deviation bigger or smaller than the analysis of historic data alone would suggest. In the case of the 30-days oil price forecast, you may consider the following issues:

- Is the oil price now higher or lower than it has been on average during the period for which you have the Brent quotations ? Higher oil prices are generally associated with bigger price fluctuations measured in \$ per barrel.

- Is price of oil perceived as being more or less volatile (and, by extension, uncertain) now than it has usually been ? This might be associated with particular events disturbing the market, political turmoil, etc.

1.4. Continuous and discrete distributions

So far we have worked on the assumption of continuous distributions of the variable. That is, the variable (such as the price of Brent crude in 30 days, or the total recoverable quantity of oil in a field) can take an infinite number of different possible values, in infinitely small steps along the scale.

The opposite of continuous distribution is discrete distribution. Here, the possible outcomes are restricted to certain specified values. An example of a discrete distribution is the result of rolling dice. The outcome can only be 1, 2, 3, 4, 5 or 6 - never anything between. The probability of obtaining each of these specific values is exactly 1/6 (assuming the die is fair).

Another example of a discrete distribution is the question (3) in the beginning of this text: Will the UK be a net importer of oil in 2020 ? As posed, this question has only two possible outcomes, Yes and No (which we can express numerically as 1 and 0). An informed analyst would surely not assign a 0.5 probability to the Yes outcome, but assign a higher or lower probability based on analysis of the oil production potential and domestic consumption projections for oil products. We might also pose the question differently: What will be the net imports of oil to the UK in 2020? This question can be answered with a continuous probability distribution.

In continuous distribution, probabilities are not assigned to specific values, only to ranges of values. If the distribution is truly continuous, then the possibility of the oil price falling by *exactly* \$2 per barrel has no probability. If pressed for an answer, a statistician might tell you that it has a probability of zero. That may frustrate you since you had chosen this as the most likely outcome. The reason for the frustration is then, that the outcome of -\$2 is just one of infinitely many outcomes, another one of which is -\$2.0000000000000001, which is not \$2. So the chance of this hitting exactly -\$2 is infinitely small (which statisticians treat as zero), even if this is where the probability distribution is at its densest.

In a continuous distribution therefore, no probability is assigned to the exact value of -\$2 per barrel. But a probability can be assigned to a range of values represented by the index value -\$2. That range may be, for instance, from -1.50 to -2.50. The Excel formula which we used


for exercise E1, does this. If you have done the exercise, you will see this probability as 0.13298... associated with the index value -2 in the density function.

There is a clear and important distinction in theory between continuous and discrete distributions, but less so in practice when dealing with available figures. Theoretically, when people agree to buy and sell oil, they might agree any price that can be expressed as a number on an invoice. In practice, outright prices agreed and reported for oil are usually in full dollarcents per barrel. When we finally get to know the price of Brent crude in 30 days as reported by Platts, it will be in the form of a figure expressed in full cents, because that is how Platts reports. This makes it part of a discrete distribution of possible prices, not continuous. However, a discrete distribution like this, with a dense set of index values (\$40.00, \$40.01, \$40.02,), can for most purposes be treated like a continuous distribution.

1.5. Normal vs. lognormal distribution

We turn now to the second question posed at the start of this text: What will be the price of light, sweet crude oil in 2050, on average? This is much further into the future than conventional market analysis will cover - never mind that conventional market analysis often gets it very wrong also when it comes to predicting short-term oil prices. There are some publications that try to describe very long term developments in global energy,⁷ but they will not give you any solid and precise price forecasts. There are obviously enormous uncertainties with regard to new technologies for making energy available, and other things that will affect the outcome.

We might come up with the view of \$50 per barrel as our best shot, but with such a large uncertainty that both \$5 and \$500 per barrel are also considered possible outcomes – and equally so. Our best-judgement probability distribution might look like Figure 3. This looks much like the normal distribution in figure 1, but the horizontal scale is of a different nature. The index values \$5 and \$500 are now placed in equal distance from the middle value of \$50. Arithmetically this does not hold, since the difference (\$500-\$50) is not the same as (\$50-\$5). But the distances are the same in the proportional sense: $\$500/\$50 = \$50/\5 . This makes it a logarithmic scale. The distribution that is normal (with the symmetric, bell-shaped form) on a logarithmic scale, is called a lognormal distribution.



Lognormal and similar distributions are often relevant in stochastic modelling for economic purposes. This is because many phenomena that we need to analyse, have a lower boundary but no upper boundary. This is certainly the case for such important parameters as the future price of oil, the amount of oil in a reservoir, and the number of drilling hours needed to reach the reservoir. None of these values will ever be negative, so zero is a lower boundary. But at the high end there is usually no absolute limit, only diminishing probabilities.

⁷ for instance Shell's booklet "Energy needs, Choices and Possibilities: Scenarios to 2050", at www.shell.com

The symmetry which we assumed for the regular, normal distribution, is not really there. The normal distribution assumes no lower or higher limit to possible values - only that the values far from the mean will have tiny probabilities. In reality, probability distributions for variables such as the future price of oil, should have an absolute lower limit, because the price will surely never be less than 0. In our previous case on the 30-days price forecast, we could ignore this effect without much error, because nearly all of the probability distribution was far away from the lower boundary of 0 (assuming oil prices as of 2005).

In the present case of considering the price of oil in 2050, we cannot ignore this issue, because we consider \$500 as a significant possibility, whereas an oil price of -\$400 is not a relevant possibility. At first sight, the lognormal distribution in figure 3 also appears to have no lower boundary - the curve will never reach zero no matter how far to the left or right we extend it. But the horizontal axis values also never reach zero if extended to the right, they

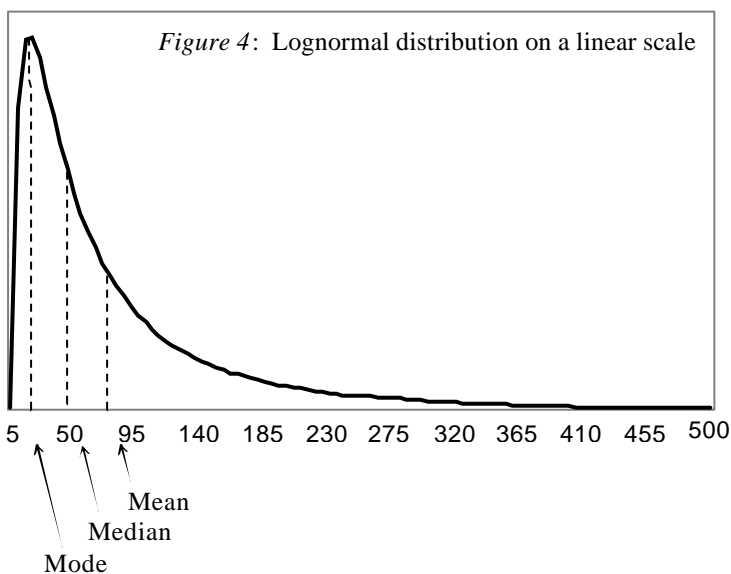
Box 5: Measures of central tendency

(See figure 4 for examples)

The *Mean* is the average of the population, i.e. average of all values in a discrete distribution.

The *Median* is the middle value, i.e. the population is equally large above the median and below the median. It is also called the 50-percentile.

The *Mode* is the value that has the densest distribution, i.e. the peak of the distribution curve.



just become ever smaller positive numbers. Therefore, the value 0 is a lower boundary for this lognormal distribution.

When shown on a linear scale, the lognormal distribution is strongly skewed to the left (Figure 4). This graphical presentation also reveals that the three common measures of central tendency, the Mode, Median and Mean, are all different in this distribution (at around 20, 50 and 80 respectively). In the normal distribution, these three values coincide (See Box 5 for explanation of these terms).

Lognormal distributions can be handled numerically by converting all values to logarithms, and then associating a mean and a standard deviation with the obtained values. Microsoft Excel has a function for cumulative lognormal distribution:

$$\text{LOGNORMDIST}(n; \text{mean.ln}; \text{sd.ln})$$

where

- n is the index value for which the cumulative probability is calculated
- mean.ln is the mean of $\ln(x)$
- sd.ln is the standard deviation of $\ln(x)$
- $\ln(x)$ is the natural logarithm function applied to the population values

Figures 3 and 4 were generated using the following values:


Mean of $\ln(x) = \ln(50)$
Standard deviation of $\ln(x) = \ln(3)$

1.6. Triangular distribution

Sometimes we shouldn't make things so complicated. Imagine the following conversation between a planning analyst and a drilling superintendent working on a hydrocarbon prospect somewhere offshore Africa:

Drilling superintendent: We've got plans ready for drilling that appraisal well.
Planning analyst: How much do you expect it will cost?
Drilling superintendent: 9.4 million US dollars.
Planning analyst: Are you sure ?
Drilling superintendent: No, but that is our budget estimate. Things can go against us, and then it will cost more.
Planning analyst: How much could it cost in the worst case?
Drilling superintendent: I'd say, certainly not more than 14 million. That's really the worst case.
Planning analyst: And the best case?
Drilling superintendent: 7 million. That's drilling at more than record speed with no hickups whatsoever.
Planning analyst: Great – can you give me the standard deviation of the probability distribution of the drilling costs expressed as natural logarithms and based on a lognormal distribution?
Drilling superintendent: No.

The analyst might have dropped his last question, up to which the two were getting on well. If the superintendent is an experienced professional with a good judgement of such things, he had already given the analyst a suitable basis for a probability distribution, as shown in figure 5. This is a triangular distribution defined by the three numbers stated by the drilling superintendent: Minimum (\$7 million), most likely (\$9,8 million) and maximum (\$14 million).



This methodology is simple and often appropriate, bearing in mind that real-life probability distributions are often based largely on subjective judgement. One should however be aware of certain statistical properties of this kind of probability distribution.

Even if there would be some chance that the cost may fall outside the assessed minimum-maximum range (as numerous drillers have found out the hard way), the suggestion here is to disregard that possibility. As shown in figure 5, the superintendent's statement about \$9.8 million, his budget estimate was taken to be the *mode*, i.e. the index value with the highest probability density. Note that this is not the same as the *mean* or the *median*, both of which

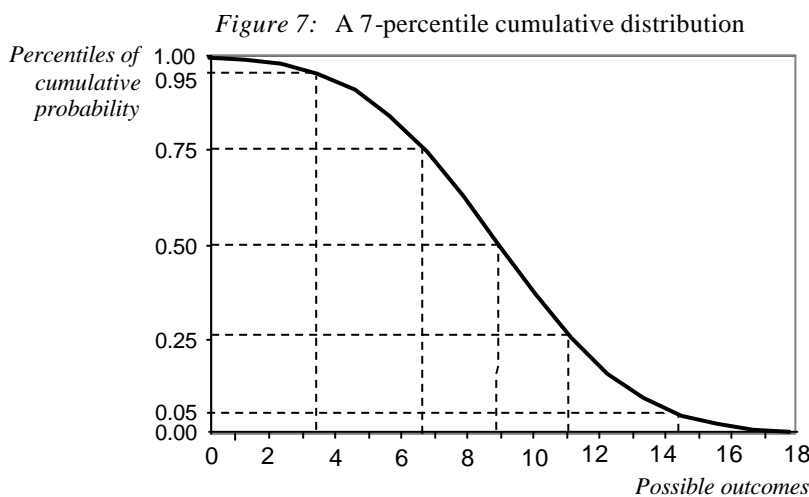
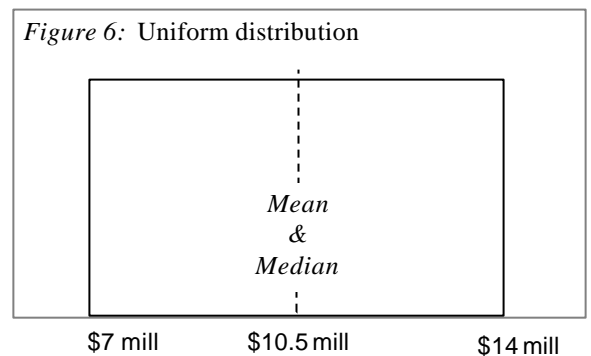
are higher than the mode in this case because of the way the distribution is skewed.⁸ By this distribution, the well cost will probably end up higher than the superintendent's estimate. This is not in contradiction to accepting his best estimate of what the cost will be.

This implies some interpretation about what the superintendent meant by what he said, and what he might have a basis for saying. To some extent, it depends on the culture and norms of the organisation with respect to cost control and operational responsibility. As a practically oriented line manager, the superintendent probably has his mind set on \$9.8 million as a reasonable number that he should be able to stay within, based on the information he has received about the geological conditions. But he also recognises the possibility of the costs ending up considerably higher, and to a lesser extent lower, than his estimate, for reasons of which he is not in full control.

1.7. Other probability distributions

The simplest probability distribution, or rather, lack of such, is the *constant*. Specifying a variable as a constant means just assigning one value to it, not allowing for any possibility that it may turn out differently.

A *uniform distribution* is also rather simple: It is defined by a minimum and maximum, and the probability density is the same across that entire range (Figure 6). We might have this if someone gives a cost estimate of "anywhere between \$7 million and \$14 million" - without any clue as to what may be more or less likely within that range.



Multiple percentiles distributions can be used if we have a direct basis for knowing the possible outcomes associated with probabilities. One example is the *7-percentile distribution*, which allows for specifying cumulative probabilities at the percentiles of 0%, 5%, 25%, 50%, 75%, 95% and 100%. The data for a multiple percentiles distribution may, for instance, come from the observation of a previous case to which we assume that

the case we are looking at will be similar. Figure 7 shows a cumulative distribution defined by 7 percentiles, where the outcome is considered certain to be in a range of 0 to 18.

⁸ The mean of a triangular distribution is calculated as the average of the three values, $m = (L + M + H) / 3$, where L is the minimum, M is the mode and H is the maximum. The median T_{50} lies between the mode and the mean. $T_{50} = H - \sqrt{0.5(H-L)(H-M)}$ if the mode is to the left of centre, otherwise $T_{50} = L + \sqrt{0.5(H-L)(M-L)}$

In literature references can be found to several other distribution forms which will not be reviewed here. Among them are *Beta* and *Weibull* distributions, whose shapes can resemble lognormal distributions but are defined differently; and *Exponential* distributions, whose density distributions are strictly falling or rising. There are also several distribution forms which apply specifically to discrete distributions, such as the *binomial* distribution and *Poisson* distribution.

PART 2 : PROJECT ECONOMICS

Should our organisation invest capital in a particular development project ? Such decisions in the oil and gas industry sometimes commit very large sums of money, and will essentially shape the future of the organisation involved. This part will introduce the techniques and considerations that are commonly applied for analysing potential projects in preparation for making such investment decisions, large or small.

The emphasis here is on the financial aspects of the evaluation. That is, we will take as a starting point certain assumptions about the costs and incomes which are expected from the project, and analyse the financial implications of those costs and incomes. The technical and geological knowledge which is required in order to generate good assumptions about such costs and incomes, will not be provided here, but are topics of the PETRAD courses.

Chapters 2.1 through 2.4 in this part cover the basic techniques and concepts of economics for the purpose of investment decisions. Course participants should be familiar with these concepts to facilitate learning in the courses. The remaining chapters discuss issues that are important for a good understanding of project economics, but are not necessarily critical to participation in the course.

2.1. Cash flows

Project economics is usually done as cash flow analysis. It works with *cash flows*, which are payments. From the point of view of a firm, a cash flow is something that you can normally identify as a movement on the firm's bank statement in retrospect.

Someone may come up with the following proposal: Our organisation should spend \$90,000 to buy a machine, which will enable us to increase sales revenues by \$30,000 per year. In order to analyse the project economics of this proposal, you also need to know at least the following:

- (1) Is \$90,000 the total expenditure of acquiring the machine, or will there be added initial costs for freight, installation, etc.?
- (2) How much will it cost per year to operate and maintain the machine ?
- (3) For how long can we expect the machine to be useful before it becomes either worn out or obsolete?
- (4) Will it have any residual value at the end of that period ?
- (5) Are the revenues of \$30,000 per year truly incremental, or will there be side effects on other parts of the firm's sales ?

Figure 8 shows the cash flows of this project based on the information available. Figure 9 shows the same as a histogram, each column representing a year.


For the purpose of illustration the following answers to the above questions were assumed: (1) Additional expenditures for freight, installation etc will be \$10,000, so that the total expenditure of acquiring the machine will be \$100,000. (2) Annual costs of extra manpower, maintenance, electricity etc. will be \$10,000. (3) 10 years. (4) No residual value, as we assume the machine can be sold for scraps at a value that just covers the cost of removal. (5) No such side effects assumed.

Figure 8: Cash Flow Calculation (Taxes not considered)

Unit: US dollars

Year	Investment	Revenues	Operating	Salvage	Net cash flow
0	-100 000				-100 000
1		30 000	-10 000	0	20 000
2		30 000	-10 000	0	20 000
3		30 000	-10 000	0	20 000
4		30 000	-10 000	0	20 000
5		30 000	-10 000	0	20 000
6		30 000	-10 000	0	20 000
7		30 000	-10 000	0	20 000
8		30 000	-10 000	0	20 000
9		30 000	-10 000	0	20 000
10		30 000	-10 000	0	20 000
11		0	0	0	0

Figure 9: Histogram, net cash flow from machine project



Note the following features of the cash flow table in figure 8:

- Each line represents a year. The start year is identified here as year 0. (If you wish to show the actual years (2006 etc), it is advisable to do so in an extra column.)
- Outgoing cashflows are shown as negatives, ingoing

as positives.

- Net Cash Flow for each year is the net sum of payments (inflows minus outflows).
- When displaying a table, it is always advisable to specify units clearly (US Dollars in this case)

The project may have an impact on the company's taxes, which is not shown here. The Net cash flow shown is therefore a pre-tax cash flow. With knowledge of the applicable tax system it would be possible to expand the table with more columns to show taxes and the after-tax net cash flow.

Note that *profit*, and certain other terms found in a firm's income statement, are not cash flows. An income statement contains certain components which are closely linked to cash flows (such as sales revenues and labour costs) and other components which are not cash flows (such as depreciation). The firm's balance sheet also does not show cash flows.⁹

A cash flow for the purpose of calculating project economics can also be a saved expenditure. We might modify the example above so that the machine would not generate any extra

⁹ It is however possible to use information from the firm's income statement, balance sheet and some supplementary information in order to calculate the firm's total cash flow.

revenue, but it would make the organisation avoid certain costs which otherwise would be incurred, of \$30,000 per year. The cash flow would then be precisely the same as shown in figure 8.

It is important to identify cash flows not only by how large they will be, but also by their timing, i.e. when will they occur. In project economics it is common to work on an annualised basis, i.e. we identify cash flows as a sum for each year. In some cases it can be relevant to work with shorter periods in project economics, which will not be done in this paper.¹⁰ For cash flow analysis done on the basis of a table such as Figure 8, the presumption is that all payments occur at similar times in all years, which we can take to be either the middle of all years or spread evenly throughout the year. This

would be consistent with a case where the investment would be paid for on July 1st in Year 0, and revenues would start flowing evenly in the following January.

2.2. The time value of money

A key premise in project economics is that the value of money is different depending on when it will be paid. A certain sum of money is considered less valuable if it will be paid in the distant future than if it will be paid sooner.

Our parents might tell us how much money they had to live on as young adults, and the amounts may seem to us as ridiculously low. This is not only due to rising wealth, but also to the inflation that has affected the value of money in most parts of the world. In project economics however, we consider the time value of money not primarily in terms of inflation, but for two other reasons: Interest and risk.

In the example above of the machinery, we would invest this year in order to cash in net \$20,000 next year and for each following year to the 10th year. But if we expect to get \$20,000 both in year 1 and year 10, then the last \$20,000 are less valuable to us than the first \$20,000. There are two reasons for this:

Box 6: Cost, expenditure and expense

Accountants make an important distinction between *expenditures* and *expenses*. An expenditure is essentially a payment out, and is therefore a cash flow. An expense is a consumption of value. The term *cost* in normal parlance may mean either of the two.

Assume we invest now in a machine for which we pay \$100,000. We will start using the machine in January next year and it will become obsolete or worn out in 10 years. We have an expenditure of \$100,000 this year, from paying for the machine. But this is not an expense during this year, because we have just exchanged one asset (money) for another (machine). If we keep the machine for its useful life, we will eventually account for the \$100,000 as expense. This happens over the course of time when the machine is useful, and is called *depreciation*. Depreciation is an accounting procedure for allocating the cost of assets such as machinery over their time of use, and is not a cash flow.

In the stated example, the *cash flow* from buying the machine is -\$100,000 in the current year. This would be the relevant input for cash flow analysis.

¹⁰ Firms do *liquidity budgeting* of cash flows on a monthly, weekly or even daily basis for a different purpose than project economics, namely to optimise short-term cash management and ensure that funds are at all times available to meet payment obligations.

- (1) Interest: the first \$20,000 can be invested safely in a bank account to earn interest, and therefore it will have grown to a larger amount by year 10.¹¹ Alternatively, we can consider the way that the early \$20,000 can be used to repay company loans and thus save interest expenses.
- (2) Risk: What do we really know about the future in 10 years? We may be fairly confident about being able to cash in the \$20,000 next year, but the long term future would normally be considered less certain. Most firms are risk averse. That means, they would tend to dislike the higher uncertainty associated with the \$20,000 in year 10, even if the uncertainty works both ways (i.e. the net cash in year 10 could equally well be higher as lower than \$20,000.)

The discount rate encompasses both interest and risk as compensation for the time value of money. A discount rate is expressed in per cent per year, for instance 10 % per year, in the same format an interest rate. Discount rates are however generally higher than the prevailing interest rates, since they also include a compensation for risk. Chapter 2.5 gives further considerations on what an appropriate discount factor may be. In project economics, discount rates are used to reduce the value of future payments. Computationally they therefore work in the reverse of interest rates.

Figure 10 shows a calculation where a discount rate of 10 % p.a. is used to discount (reduce) the value of cash flows from the machinery investment example. We again have one line per year as indexed by the first column. The second column contains the same net cash flows as in figure 8. The third column, “Discount factor”, contains conversions of the 10% discount rate to annual factors based on the following formula:

$$d_n = (1 + i)^{-n}$$

where n is each year, with 0 representing the base year to which values are discounted;
 d_n is the discount factor applicable for year n ;
 i is the discount rate (in this case 10%, expressed as 0.10).

If you do this in Microsoft Excel, the formula for the factor for year 1 (placed in cell G30) is:

$$=(1+\$E\$20)^{(-A30)}$$

where cell E20 contains the discount rate (10% in this case)
 cell A30 contains the year, numbered as 0, 1, 2, with 0 being the base year.

This Excel formula for year 1 yields the value 0.909. It can be copied vertically to the other years, including year 0.

The last column in Figure 10 shows the discounted cash flows, which are obtained by multiplying each annual cash flow with the annual discount factor from the two preceding columns. For year 1, this indicates that the \$20,000 which we expect to obtain in that year, have a present value of \$18,182. For year 10, the \$20,000 which we expect to obtain, have a present value of \$7,711. The questions answered here, are: What is the value today (i.e. in

¹¹ For instance at 5 % annual interest, \$20,000 will grow to \$ 31,026 in 9 years.

year 0) of the expected \$20,000 in year 1 and year 10 ? We put the same value on the expectation of having \$20,000 in year 10, as on having \$7,711 now.

The table in Figure 10 also shows that a particular sum of money (\$20,000) decreases in present value the further into the future we expect to get it.

2.3. Net Present Value and Internal Rate of Return

The Net Present Value of a project is the sum of its discounted cash flows. It expresses the value of all payments resulting from the project, referred to a common point in time (the “present”, or at the middle of year 0).

At the bottom of Figure 10 are two sums: One for the net cash flow, and one for the discounted cash flows. The sum of cash flows for all years, \$100,000, indicates that we expect the project to generate revenues in excess of expenditures in the amount of \$100,000 over its lifetime. An important consideration is however that a large bulk of expenditures have to be made initially, whereas the revenues will come later. This makes the project distinctly less attractive than indicated only by the \$100,000 surplus.

Figure 10: Cash flow discounting based on a 10% discount rate

US\$ Year	Net cash flow	Discount factor	Discounted cash flows
0	-100 000	1.000	-100 000
1	20 000	0.909	18 182
2	20 000	0.826	16 529
3	20 000	0.751	15 026
4	20 000	0.683	13 660
5	20 000	0.621	12 418
6	20 000	0.564	11 289
7	20 000	0.513	10 263
8	20 000	0.467	9 330
9	20 000	0.424	8 482
10	20 000	0.386	7 711
11	0	0.350	0
Sum	100 000		22 891

Net present value

Box 7: Discount rate, Net present value and Internal rate of return

A *discount rate* is a parameter for assigning different value to amounts depending on when they will occur. A discount rate of i , stated as a percentage, implies that a payment is considered i % less valuable for each year into the future that it will take place.

The *net present value (NPV)* of a project or cash flow series is the sum of all payments, discounted to a base point in time, applying a chosen discount rate.

The *internal rate of return (IRR)* of a project or cash flow series is a discount rate that yields a net present value of 0.

The other sum, \$22,981 as the sum of discounted cash flows, is a much better indication of the attractiveness of this project. It answers the question: How much is the expectation of these future payments, including what we have to invest up front, worth today? This is the pre-tax Net Present Value (NPV) of the project. In calculating this, we have converted all future expected payments to their present equivalent amounts, and added it all up.¹²

One interpretation of the net present value is that the investment will create value for the company in the amount of the NPV.

¹² The Microsoft Excel spreadsheet software provides a more direct way of calculating the net present value, in the form of the function: NPV (*discount_rate*; *cashflow_values*).

In a spreadsheet containing the data in figure 9 with the discount rate in cell E20, the net cash flow for year 0 in cell F29 and the other years following on below F29, the formula is: F29+NPV(\$E\$20;F30:F49)

The value for year 0, from cell F29, is kept out of the NPV formula and added separately. The reason for this is that the NPV will discount back to one period earlier than the first value in the range.

We might say the machine investment project makes our company \$22,981 more valuable once the project is embarked upon, bearing in mind that the value of the company also includes the expected future earnings from assets that it has acquired.

The net present value of a series of cash flows depends strongly on the chosen discount rate. If we change it from 10% to 20% in our example, the net present value changes from (positive) \$ 22,981 to (negative) 16,151. A high discount rate implies that little value is assigned to payments in the more distant future. Investment projects generally have lower or negative net present values if high discount rates are applied, because the future revenues are then not valued sufficiently to make up for the initial outlay.

Decision makers also like to see the profitability of projects expressed in a form that can be compared to interest rates. This would be the internal rate of return.

We found that the NPV of our machine investment project was positive at a discount rate of 10%, and negative at 20%. In between 10% and 20%, there is a discount rate that makes the NPV of this project exactly 0. This is the internal rate of return (IRR). In this case it works out to 15.1%. It indicates an expectation that the project will generate similar returns as putting money in the bank and obtaining 15.1% effective interest.

Besides NPV and IRR, certain other measures of economic performance are often used in project assessment. The most common are described in Box 8. These indicators would normally be used as supplementary information, more seldom as stand-alone criteria for investment decisions.

EXERCISE: At this point it is suggested you do exercise E-4, which is described in Part 3. The exercise amounts to constructing in an Excel spreadsheet the example that has been described in chapters 2.1 – 2.3, and manipulating the input parameters to see how the resulting NPV and IRR are affected. Also do exercises E-5 and E-6. The exercises should also give you a feel for the importance of timing in project economics: A delay in revenues can have a severe impact on the economic merits of a project.

Box 8: Supplementary indicators of project economics

Profitability index: A ratio obtained by dividing the net present value by the original investment.

Payback: The time taken from start of investment to recovering the original investment as cumulative cash flow.

Maximum exposure: The largest cumulative negative cash flow early in the project. (Often, but not always, the same as the original investment amount)

2.4 Dealing with inflation in project economics

The 1970's and 1980's was a time of high inflation for many of the world's economies. Inflation has since fallen back, and in the early 2000's it was usually less than 4% in the major economies. This has tended to simplify things in project economics. However, as long as there is some inflation, we need to be clear about how we account for it in project economics, especially since these calculations often extend many years into the future.

A distinction is made between nominal monetary terms and constant monetary terms as modes of project economic calculations. A calculation in nominal monetary terms means that all monetary values in the calculation are to be taken literally, regardless of when they arise. A calculation in constant (or "real") monetary terms displays monetary values by

constant purchasing power. The constant purchasing power must be referred to a year, which should be the same year as used for discounting purposes. In a calculation in constant monetary terms, if we assume 3% annual inflation and show a value of \$100 in year 2, then this represents \$106.09 ($=\$100 \times 1.03 \times 1.03$) in actually expected money.

There is a need for consistency in applying either nominal or constant monetary terms. In a constant-terms calculation, the discount rate must also be understood in constant monetary terms, i.e. a percentage over and above the assumed inflation. There is a potential complication in calculating taxes, because depreciation, which has an impact on taxes, usually must be calculated in nominal terms. Any assumed inflation should be linked to the currency in which the calculation is done, i.e. in a calculation made in US Dollars, it is appropriate to use inflation as assumed in the USA (usually the consumer price inflation), even if the project takes place in a different part of the world.

The cash flow calculation in figure 8 has the appearance of being in constant monetary terms, since revenues and costs do not increase in time. In a low-inflation environment such as the early 2000's, these could be realistic cash flows even in a nominal calculation. Many traded goods see little or no price increases, due to competitive pressures and productivity gains. Futures prices for crude oil are often backwardated, i.e. oil for delivery in the far future is traded at lower prices than oil for earlier delivery, which is the opposite of what one would expect under inflation. It can also be reasonable to assume no increase in labour costs under the assumption that salary increases will be offset by productivity gains.

Sometimes project economics involve two or more different currencies, and assumptions about exchange rates must be incorporated in the economics calculation. The inflation assumed for currency them can be quite different. At least in the long term, one would then expect exchange rates to adjust broadly to differences in inflation, so that the currency that exhibits high inflation, will normally be on a falling trend against the low-inflation currency. In the short and medium term, this relationship may be distorted by market developments and sovereign monetary policies, such as fixed-rate regimes.

2.5 Choice of discount rates and investment decisions

Project economics are usually done for the purpose of preparing for investment decisions. Many firms use the NPV calculation as a criterion for deciding on investments. A discount rate is then applied which reflects the firm's required return on such investments. An investment is considered profitable if this discount rate results in a positive NPV. The profit-optimising firm should select a portfolio of investments which maximises positive NPV overall.¹³

The choice of discount rate(s) to be used in project economics, and, more generally, what financial criteria to apply for investment decisions, is often a matter of corporate policy. As a member of a large organisation, you might be instructed by headquarters in this regard.

Financial theory provides guidance in how firms can determine the required returns to be used as selection criterion for investments. It starts from the premise that firms obtain capital from financial markets where investors demand returns based on what they can obtain from

¹³ Using NPVs as the criterion for selecting investments sometimes results in different decisions than using IRRs. The NPV approach is generally recommended in theoretic literature.

other projects with similar risk. Investors demand compensation for taking risk. A distinction is made between so-called “systematic” risk and individual project risk. Investors can largely offset individual project risk by spreading their funds on many different assets with uncorrelated risks, and therefore should not need to demand any significant risk premium for that type of risk.. For a full explanation of this issue, refer to a good textbook on financial management and capital budgeting.¹⁴

A company might apply a higher discount rate to oil project A than to oil project B, because project A is perceived as more risky. This is not necessarily consistent with financial theory. It depends on what possibility the company and its owners has for diversifying risk, i.e. taking a low stake in A while also investing in other projects whose risks are uncorrelated with the riskiness of A. If this possibility exists, financial theory suggests that the same discount rate should be applied to project A as to oil projects in general.

A high risk premium for a project is warranted if the project is risky and requires that a company commits a big part of its resources, so that a negative outcome of the project would hurt the company badly. The high risk premium can be incorporated either by requiring a large positive NPV using a standard discount rate, or by using a larger than normal discount rate - the latter with some reservations which will be reviewed in the next chapter.

It is sometimes argued that public entities - State agencies, public utility firms etc - can and should apply lower discount rates than usually done by private-sector investors, including firms operating internationally. There is some merit in this on the basis that the comparison to anticipated returns in equity markets is less relevant for such entities, who pursue other goals than profit maximisation in the purely financial sense. Also, public entities may well perceive risk differently from foreign investors when investing in long-term assets within their own country. In some countries the required return rate for public investment in infrastructure and utility services has been the subject of much politicised debate, resulting in public entities applying explicitly lower discount rates than normally would be expected from private investors.

What is generally important, is that economic calculations for a project are based on estimates which are well understood in terms of their associated uncertainty. Financial theory generally presumes that economics are calculated on the basis of mean probabilities, which are not necessarily the same as line managers will state as their best estimates¹⁵. Advanced decision-support software such as GeoX¹⁶ provides for probability distributions also of financial outcomes, which can help companies understanding the implications of the financial risks associated with projects.

2.6 Complex cash flows: Beware of the dual IRRs

Consider this case: An offshore oil field has been operational for several years, and is still expected to produce 100 million remaining barrels over the next 15 years. The production

¹⁴ Commonly applied concepts in this regard are known as CAPM (Capital Assets Pricing Model) and WACC (Weighted Average Cost of Capital)

¹⁵ See for instance the discussion of triangular distributions in chapter 1.6, where the mean differs from the mode of the probability distribution.

¹⁶ The GeoX software is provided by Geoknowledge (www.geoknowledge.com), and is used for participants' case work in PETRAD courses.

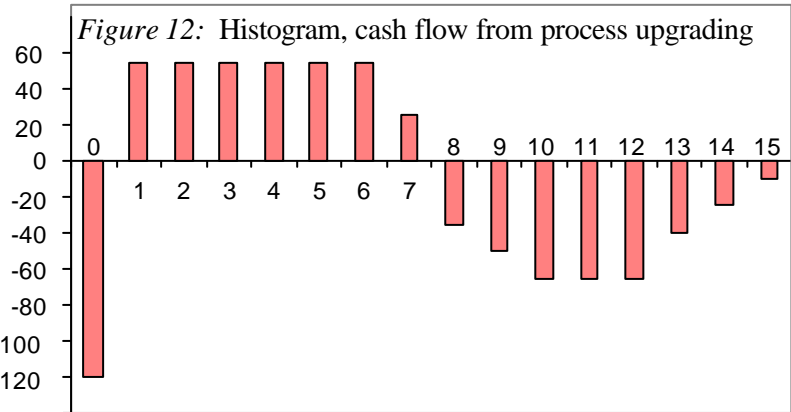
rate is presently constrained by processing and water injection capacities. A proposal is made to spend \$ 120 million to upgrade capacities, resulting from which production would increase by 20% right away. It would however not change the total amount of oil ultimately produced from the field, but it would reduce the time taken to deplete the field from 15 to 12 remaining years.

Figure 11 shows a summary of cash flows that may result from this. The cash flow of the process upgrading project, shown in the last column, is derived as the difference between two cash flows: The total cash flow from running the field as it is for the remaining 15 years, and the total cash flow from the field if the investment is made. (Cash flows unaffected by the investment are not shown.)

Figure 11: Incremental cash flow for upgrading oilfield processing \$ million

Year	Net cash flow without process upgrading	Net cash flow with process upgrading	Project cash flow: Effect of upgrading
0		-120	-120
1	280	335	55
2	280	335	55
3	280	335	55
4	280	335	55
5	280	335	55
6	280	335	55
7	250	275	25
8	220	185	-35
9	175	125	-50
10	130	65	-65
11	100	35	-65
12	70	5	-65
13	40	0	-40
14	25	0	-25
15	10	0	-10
16	0	0	0
Sum	2 700	2 580	-120
NPV @ 15%	1 353	1 372	18

Running this in MS Excel obtains an IRR of 9% from the last column. This strikes as odd since a positive NPV of \$ 18 million is shown when discounting at 15%: In normal investments, we get negative NPV's when discounting at higher rates than the IRR, and positive NPV's at lower discount rates. Here, the reverse is the case.



The anomaly is due to the pattern of cash flows resulting from this project: Cash flow changes direction from negative to positive, and then back to negative. After the initial outlay (-120), we get several years of positive cash from the increased oil production. But ultimate oil

recovery was going to be unaffected, so there will be some years towards the end where oil production will be less as a result of the project, than it would have been without it. Hence the negative cash flows in later years.

Contrary to ordinary investment projects, this project does not generate extra revenue. Its possible merit is in accelerating a given total revenue, so that it is obtained earlier than would be the case without the project.

It is possible to calculate two IRRs from this cash flow series: One at 9% and one at 34%. Neither is very helpful in determining the merits of the project. In particular, beware of any IRR (the 9% in this case) that has the property that the NPV at a discount rate slightly lower than the IRR is negative. Such an IRR will increase if the project deteriorates, and does not indicate project returns in the ordinary sense.

Also in a project like this, it makes little sense to set the discount rate high in order to compensate for risk. The NPV in this case turns out higher at 20% discounting than at 15% and 10% (but lower again at 25%).

Ordinary NPV calculation gives an indication of whether this project has merit or not, but one should be careful about not using a too high discount rate, which would put the project in a too favourable light. In addition to running NPV at the firm's ordinary discount rate, it is suggested to also run NPV at a discount rate corresponding to the firm's cost of borrowing, which is lower than the ordinary discount rate.¹⁷ If this yields a negative NPV (which it may well do, even if the other NPV is positive), then the project does not have clear merit unless it is considered to reduce the firm's overall exposure to financial risk.

2.7. Criticism and new approaches to project valuation

The reversed cash flow case in the previous section points to some general problems with the discounted cash flow approach to project valuation. David Laughton, in a series of articles in the *Energy Journal*, vol.19 nr 1 (1998), highlighted some perceived shortcomings of the traditional approach to project economics with special relevance for petroleum. Conventional cash flow discounting as a basis for investment decisions was believed to frequently result in the following errors:

- (1) Underestimating the relative value of long-term reserves;
- (2) Bias towards building too much capacity, leading to inefficiencies when production declines;
- (3) Undervaluing future costs, leading to a preference for low-quality equipment with low initial costs but higher future maintenance costs;
- (4) Inability to evaluate the unique risk profiles represented by similar projects under different fiscal regimes;
- (5) Not valuing flexibility.

The three first points can be related to the level of discount rates (required returns) selected. Using too high discount rates for NPV-based investment decisions can certainly lead to those problems. Building a risk factor into the discount rate implies that risk increases exponentially with time. This may be a fair representation of certain types of risk, but not all.

¹⁷ Make sure there is consistency in the currency applied.

One simplified remedy for this problem would be to apply a lower discount rate than conventionally used, and then applying a requirement that the NPV thus calculated for a project, must be positive above a certain ratio to the initial capital commitment (the profitability index, see Box 8). This would tend to favour investments with long-term benefits over those with short-term benefits, compared to the conventional approach. Laughton (1998) proposes a more elaborate approach which involves decomposing the cash flow into individual components to which different discount rates may be applied.

Point (5) above relates to flexibility as the ability of firms to respond to uncertainty. Part 1 of this paper discussed probability distributions as a way of describing uncertainty related to the variables that determine the profitability of a project. Not considered there was the reality that well-managed firms can often find ways of adjusting their projects and operations in response to different outcomes of those variables.

One example can be related to the oilfield referred to in figure 11. In the base case (no investment), the field is assumed to keep producing for another 15 years, and certain assumptions were made about the oil price and the production rates. These are two uncertain variables, for which probability distributions can be generated. If one or both of them turn out less favourable than projected in the latter years, it will hurt the firm's cash flow during those years. But the firm can surely have the option to close the field down earlier than after the 15th year, should such circumstances arise, which would probably dampen the impact of the reduced revenues by eliminating operations costs for some period. The possibility of closing down at a different time than originally foreseen, is a real option for the oil firm, and has value to it as a way of responding to uncertainty.

Box 9: Real options

Real options is a term used for the recognised ability of firms to enhance value through flexibility, i.e. the ability to act in different ways in the future so as to adjust to outcomes that are initially uncertain.

Real options has been a prominent feature of economic and management literature especially in the 1990's. There are techniques developed for assigning values to real options, and to describe ways in which they can be used in response to unfolding events. The most common such technique is the decision tree, which features nodes of uncertainty and decisions. Exercise E-6 in part 3 is a simple application of this approach.

The implication of the real options approach is that firms should sometimes be willing to pay an extra cost in order to acquire flexibility, as such flexibility can become quite valuable in the future. Astute business people have of course always made such decisions based on more intuitive perceptions, but the real options literature offers some more formalised approaches.

2.8. Economics and strategy

Net present values and the related concepts are a guide to making decisions on investments, with the implication that investments with a positive NPV have merit for the organisation. Those who have an interest in the discipline of business strategy, will argue that such quantitative measures of profitability cannot be the only criterion for how the organisation should deploy its resources: Investments must above all be consistent with the organisation's strategy.

Competitive strategy is concerned with how firms can develop the ability to survive and prosper in competitive environments. The development of strategy for a firm is usually a deliberate process lead by its senior management and Board, and resulting in formal decisions. In the strategic perspective, decisions are viewed as a pattern which must be coherent. Individual decisions are then evaluated not only by their NPVs, but above all by their contribution to moving the firm in the intended strategic direction. The discipline of strategy is mainly qualitative rather than quantitative in orientation, dealing in terms such as core competence, strategic intent, entry barrier and sustainable competitive advantage - terms for which there are no formulas in Excel.

Decision making in organisations have a strategic dimension in addition to the quantitative dimension dealt with in the preceding chapters. Sometimes investment opportunities should be foregone in spite of showing attractive economics, because they would detract from the firm's strategic direction. Sometimes investments should be made in spite of showing unattractive economics when viewed in isolation, but contributing essentially to building strategic resources. And sometimes, projects are misguidedly labelled "strategic" as a poor excuse for pursuing someone's unprofitable pet idea.

Strategies are however not successful if failing to meet the organisation's financial objectives. Applied on that level, NPV and related techniques are relevant criteria also for strategic decision making.

PART 3 : EXERCISES

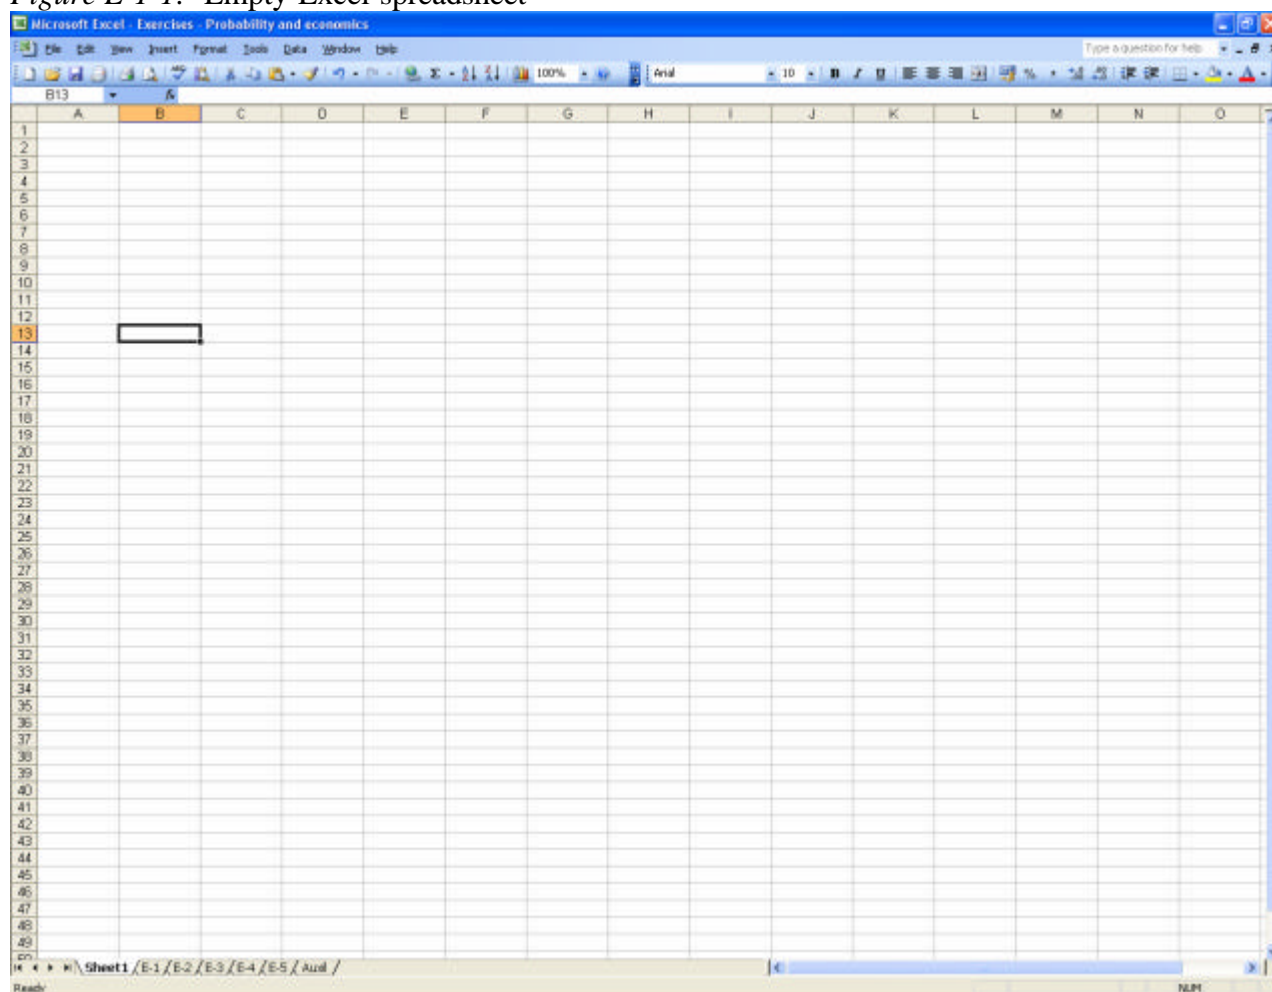
Solutions to the exercises are found in an accompanying Excel spreadsheet named “Exercises - Probability and economics.xls”. See your communication from PETRAD on how to get the spreadsheet, or contact the PETRAD secretariat.

Exercise E-1: Generating normal distribution in an Excel spreadsheet

This exercise will guide you through the process of generating a normal distribution in a Microsoft Excel spreadsheet, which will be expressed as probability numbers and graphs similar to figures 1 and 2 in the text. Basic previous skill in working with Excel is assumed. It should be possible to do the exercise also if you work with a different spreadsheet programme than Excel, but the layout and terminology may differ from what will be shown here.

Step 1: Open a new workbook and get an empty spreadsheet looking approximately like this (Figure E-1-1). The precise layout may vary somewhat, especially in the upper part, depending on how your software has been configured.

Figure E-1-1: Empty Excel spreadsheet



Note that in Excel, cells are referred to by columns and rows as indicated in the grey frames. The active cell as shown here is B13. You move this around by the four arrows on your keyboard.


Step 2: Write text into cells as follows:

- Cell A1 Exercise E-1: Probability distributions for oil price change in 30 days
- Cell A3 INPUTS
- Cell C4 \$ per barrel
- Cell A5 Today's price
- Cell A6 Mean expected change
- Cell A7 Standard deviation
- Cell A10 CALCULATED VALUES
- Cell A11 Density function.
- Cell C11 Cumulative distribution

For tidy formatting, you can turn cells A1, A3 and A10 to bold font by clicking the B on the menu above.

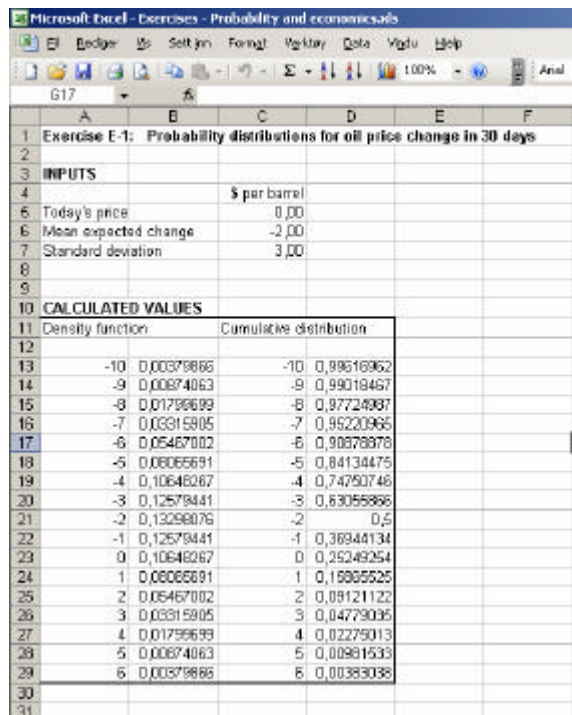
Then input the following numbers in cells:

- Cell C5 0
- Cell C6 -2
- Cell C7 3

If necessary, use these menu buttons  to get a suitable number of decimals. If your cell contents do not respond to clicking these buttons, then you have probably entered blanks etc with the numbers so that Excel treats them as texts not numbers. Try again, entering numbers only in those cells.

Your spreadsheet should now look like the top half of figure E-1-2 (down to row 11):

Figure E-1-2: Spreadsheet with basic inputs



	A	B	C	D	E	F
1	Exercise E-1: Probability distributions for oil price change in 30 days					
2						
3	INPUTS					
4			\$ per barrel			
5	Today's price		0.00			
6	Mean expected change		-2.00			
7	Standard deviation		3.00			
8						
9						
10	CALCULATED VALUES					
11	Density function		Cumulative distribution			
12						
13	-10	0.00379866		-10	0.98516962	
14	-9	0.00674063		-9	0.99019467	
15	-8	0.01796699		-8	0.97724967	
16	-7	0.03915905		-7	0.95220965	
17	-6	0.05467002		-6	0.90579978	
18	-5	0.06066691		-5	0.84134475	
19	-4	0.10646267		-4	0.74790746	
20	-3	0.12579441		-3	0.63055896	
21	-2	0.13296076		-2	0.5	
22	-1	0.12579441		-1	0.36944134	
23	0	0.10646267		0	0.25249254	
24	1	0.06066691		1	0.18985525	
25	2	0.05467002		2	0.09121122	
26	3	0.03915905		3	0.04779095	
27	4	0.01796699		4	0.02279013	
28	5	0.00674063		5	0.00981533	
29	6	0.00379866		6	0.00383038	
30						
31						

Step 3: In cell A13, input the value -10. Continue downward with -9, -8, always increasing the value by 1. In cell A23 this gives you the value 0, and you continue to cell A29 where you get the value 6.

Repeat this in cells C13 to C29. You can save effort by copying.

Move the cell marker to cell B13. We now enter the formula for normal distribution as follows:

$$=NORMDIST(A13;C6;C7;0)$$

When pressing Enter, the cell should show the value 0.00379866. This is the probability that the price change will be within the 1\$ range around -10 \$, given the mean expectation in cell C6 and the standard deviation in cell C7.

Note that our basic inputs for the mean expected change and the standard deviation are in cells C6 and C7. In the formula we just entered, the references to these cells are anchored with \$-signs so as not to change when we copy into other cells.

Copy this formula into the cells below down to B29.

If not familiar with copying cells, try this:

Move the cell marker to B13.

Hold the Shift key [\uparrow] while pressing C.

Move the cell marker down to B14.

Hold the Shift key [\uparrow] while using the down-arrow [\downarrow] to highlight the range of cells down to B29.

Hold the Shift key [\uparrow] while pressing V.

You should now have a spreadsheet looking like figure E-1-2, except that we don't have the values in the D column yet.

Note that the highest value of these calculated probabilities is found next to the value -2, which was our choice as the expected mean. The calculated probability of about 0.13 is not the probability of the outcome being exactly -2, but being within the 1\$ range that is closest to -2, i.e. from 1.50 to 2.50.

Step 4: Move the cell marker to D13. Now enter the same formula as in B13, but with two changes:

- insert 1- before NORMDIST
- The last figure should be 1 not 0

$$=1-NORMDIST(A13;C6;C7;1)$$

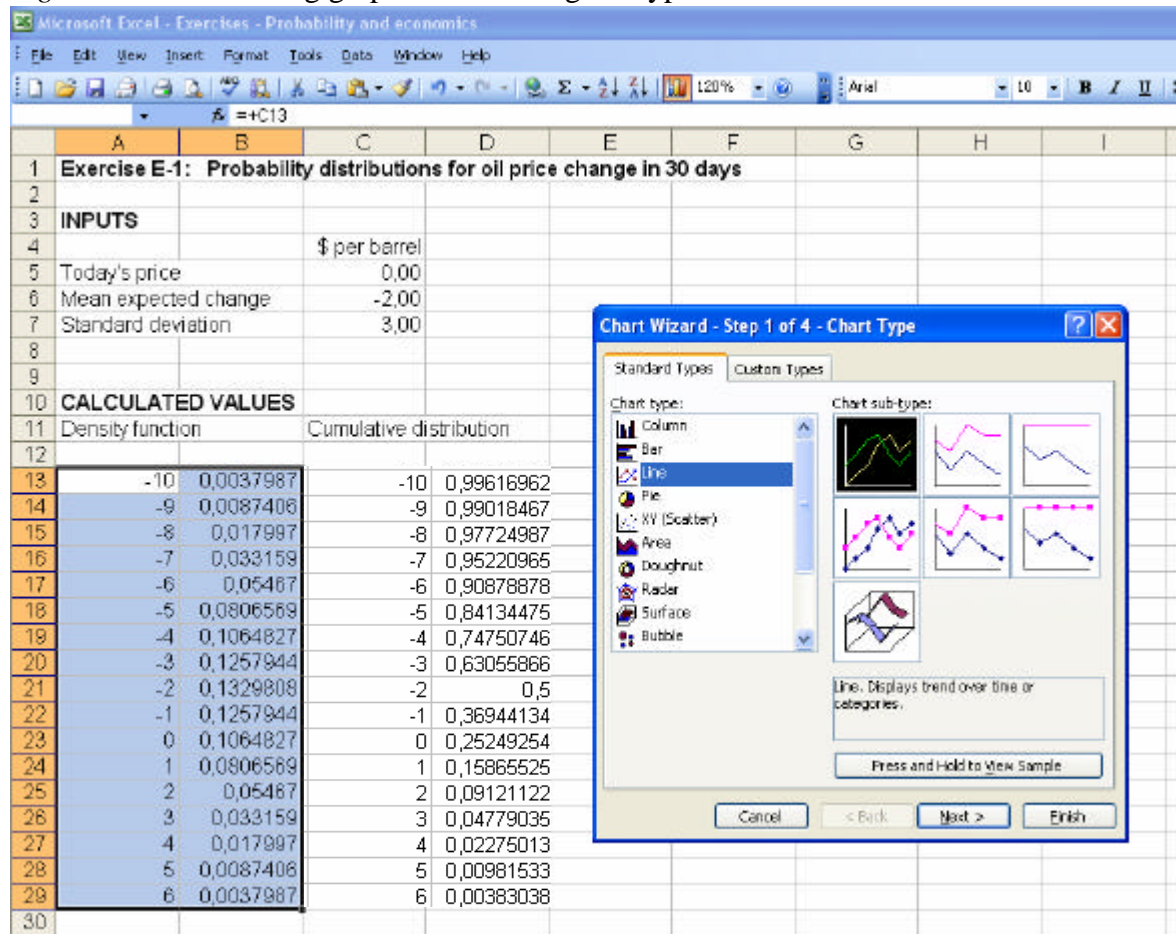
Copy this to each cell below, down to D29. Your spreadsheet should now look like figure E-1-2. The numbers in the D column are the cumulative function.

Next to the price change value -2, we now find the calculated value 0.5 . This indicates the probability that the price change will be -2 or lower (i.e. more negative than -2). The figure in cell D13 indicates that a price change of -10 or even further down has a probability of only 0.00383038. In the lower part we have figures getting close to 1 (but never reaching 1), indicating high probabilities that the price change will be at or less than the number next to it in the C column.

The last figure in the formula, which we changed from 0 to 1, is a switch to determine whether the formula will yield the density function (as now in column B) or the cumulative distribution (as in column D).

Step 5: We have now the basis for making graphs similar to those in figures 1 and 2. Use the Shift key and arrow keys to highlight the cell area A13:B29. Then click the Graph tool on the menu above, and select the line diagram and first version as shown in figure E-1-3.

Figure E-1-3: Selecting graph data and diagram type



Having obtained this screen, click Next. On the next image, you may have to deal with something which is not so elegant in Excel. Click the Series flag on top of the dialog box. You probably get the dialog box shown in figure E-1-4. When Series1 is highlighted on the left, and provided the graph shown on top is dominated by a straight diagonal, you need to copy the content of the middle dialog box into the bottom dialog box (indicated with the red arrow in figure E-1-4). Having done that, delete Series1 on the left (by clicking the Delete button when Series1 is highlighted).

This, Gates willing, should change the dialog box into what is shown in figure E-1-5. The graph in the upper part now looks like a bell curve. You can press OK, and place the resulting diagram in the spreadsheet.

Step 6: Your diagram now has the basic outline of Figure 1, but may need a bit of visual editing to get an appearance that you are happy with. You get going on that by right-clicking on the graph. This is where our detailed instruction stops, and we'll let you figure out for yourself. You repeat the process with the cumulative distribution in columns C and D. Our spreadsheet ended up looking like figure E-1-6.

Figure E-1-4: Arranging the series

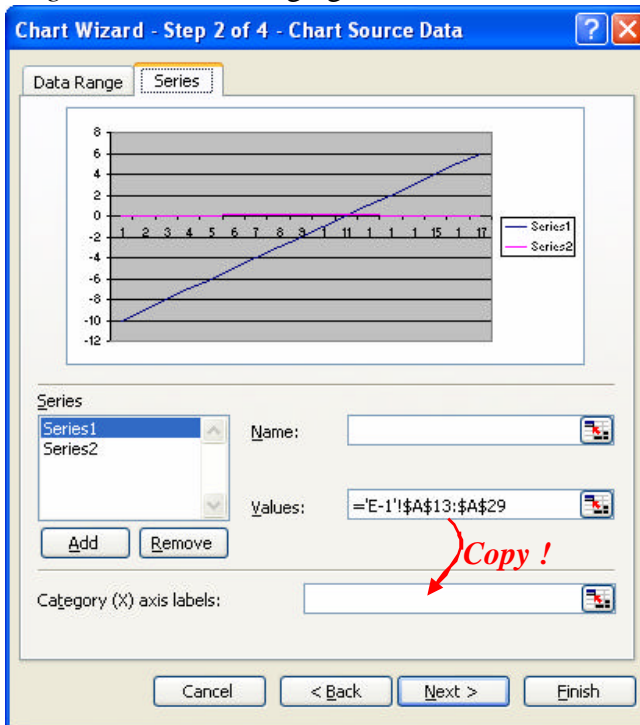


Figure E-1-5: Having arranged the series

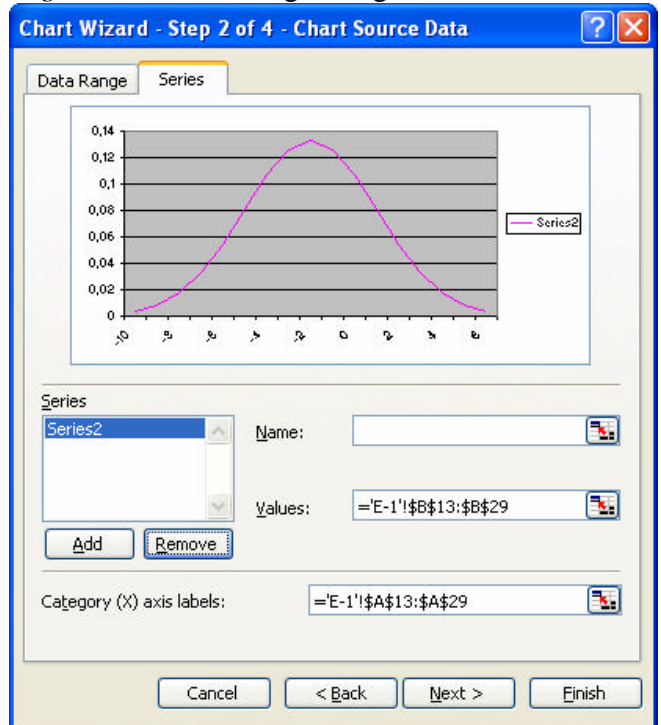

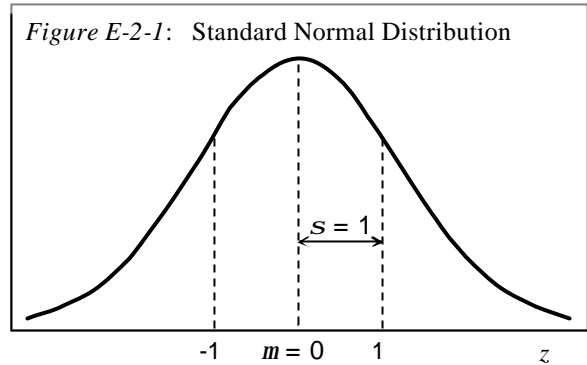


Figure E-1-6: Spreadsheet with density and inverse cumulative graphs



Exercise E-2: The Standard Normal Distribution

The Standard Normal Distribution is the distribution function which is normal, has a mean of 0 and has a standard deviation of 1. As in all distribution density functions, the total area under the curve $\equiv 1$. Find the range probabilities (=areas under the curve) in the standard normal distribution limited by the following index values z :



- Between 0 and 1.2
- Between 0 and 1.49
- Between -0.48 and 0
- Between -1.37 and 1.33
- Above 1.33

Solution hint: In Excel, use the function NORMSDIST(z), which gives you the cumulative probability of z in the standard normal distribution.

Exercise E-3: Applying normal distribution to a manufacturing quality problem (from Sincich: *Business statistics by example*)

Pacemakers are used to control the heartbeat of cardiac patients, with over 120,000 of the devices implanted each year. A single pacemaker is made up of several biomedical components that must be of a high quality for the pacemaker to work. It is vitally important for manufacturers of pacemakers to use parts that meet specifications. One particular plastic part, called a connector module, mounts on the top of the pacemaker. Connector modules are required to have a length between .304 inch and .322 inch to work properly. Any module with length outside these limits is “out-of specs”. One supplier of connector modules was found to have been shipping out-of-spec parts to the manufacturer for the past 12 months.

- The lengths of the connector modules produced by the supplier were found to follow an approximate normal distribution with mean $m = .3015$ inch and standard deviation $s = .0016$. Use this information to find the probability that the supplier produces an out-of-spec part.
- Once the problem was detected, the supplier’s inspection crew began to employ an automated data-collection system designed to improve product quality. After 2 months, the process was producing connector modules with mean $m = .3146$ inch and standard deviation $s = .0030$ inch. Find the probability that an out-of-spec part will be produced. Compare your answer to sub-question a.

Exercise E-4: Cash flow analysis for an investment in machinery

Please use an Excel or similar spreadsheet to construct the cash flow analysis referred in chapters 2.1 to 2.3. The base assumptions are:

Initial investment:	Price of machine	\$90 000
	Freight & installation	\$10 000
Annual operating expenditures		\$10 000
Annual revenues :		
Gross revenue generated from machine		\$30 000
Adjustments to revenue		\$0
Useful life of machine:		10 years
Net salvage value after useful life:		\$0
Discount rate:		10 % per year

Hints for constructing a good and tidy spreadsheet for cash flow analysis :

- Keep assumptions and calculation separate wherever practical. Change of one, visible input variable updates the full calculation.
- Use clear column & line headings and specify units, so as to make clear what is contained in every cell of the spreadsheet.
- Use format elements (lines, colours, font sizes etc) moderately to make a clearly structured spreadsheet

Use the Excel functions $NPV(\text{discount_rate}; \text{cashflow_range})$ and $IRR(\text{cashflow_range})$ to obtain project NPV and IRR. Beware that NPV discounts one year back from the first year included in the range.

Manipulate the values for investment, revenues, operating expenditures, useful life and discount rate to check for impacts on NPV and IRR.

Exercise E-5: Cash flow discounting

(From *Introduction to Petroleum Economics*, by Terje Sørenes and Olav Fjellså, PETRAD)

An investment project will give the following estimated cash flow in current (nominal) values: (-100, 20, 40, 80). The expected inflation is 2% per year.

- Deflate the cash flow (calculate the cash flow in real values (constant monetary value) referred to the first year.
- Calculate the project's net present value with a 10% discount rate in constant monetary value.
- Is the investment profitable? What is your decision?
- Calculate the project's net present value with the following discount rates: 0%, 5%, 15%, 20%. Illustrate the net present value profile as a graph of NPVs at different discount rates.
- What is the project's internal rate of return? Estimate this from the graph, as well as calculating it in a spreadsheet if you have access to one.
- Now suppose the cost of the project were to become 5% higher. What impact would this have on your IRR and your NPV at 10%?

- g) Go back to the original cost of -100, but now suppose that the earnings from the project are delayed by one year. Cash flows are now (-100, 20, 40, 80). What does this do to your IRR and NPV?

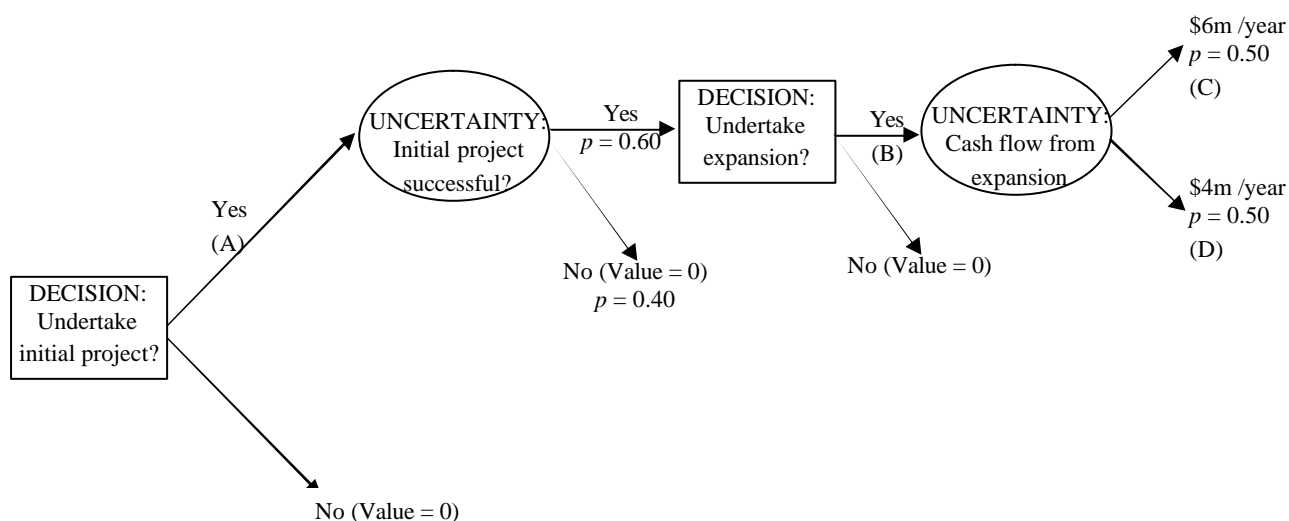
Exercise E-6: Present values with a simple decision tree

Queezy Drug Company is considering a new drug, which would be sold over the counter without a prescription. To develop the drug and to market it on a regional basis will cost \$12 million over the next 2 years (\$6 million in years 1 and 2 each). Expected cash inflows associated with the project for years 3 through 8 are \$1 million, \$2 million, \$4 million, \$4 million, \$3 million and \$1 million, respectively. These values represent the mean of a probability distribution.

If undertaking the initial project, the company will have the option to invest an additional \$8 million to secure the national market. This would be done only if the initial project is deemed successful on financial and medical criteria, for which there is a 0.60 probability. If undertaken, the additional project would increase cash flows in years 6 through 10 by \$6 million or \$4 million, each with a probability of 0.50.

The company's required rate of return is 14%.

- a) What is the net present value of the initial project? Is the project acceptable on this basis?
- b) Below is the problem structured as a decision tree. Calculate the expected values of points (A), (B), (C) and (D) while specifying the conditions required to realise these values. The values should be calculated as present values referred to year 0. *Hint*: Work from right to left in the figure.
- c) What is the worth of the project if we take account of the option to expand? Is the project acceptable on this basis?



F U R T H E R R E A D I N G :

The following is a selection of books which may be useful for further study of topics related to this paper.

Books on statistics and probability:

Kvanli, A.H., Guynes, C.S. and Pavur, R.J.: Introduction to Business Statistics. West Publishing Company.

Sincich, T. Business Statistics by example. Prentice Hall International Editions.

Books on investment economics and valuation:

Damodaran, A.: Investment valuation. John Wiley & Sons, Inc.

van Horne, J.C.: Financial Management and Policy. Prentice Hall International Editions.

McKinsey & Company, Inc: Valuation. Measuring and managing the value of companies. John Wiley & Sons, Inc