Economics of Oil, Gas and Energy

Week 5: Risk and uncertainty in oil and gas production

Week 5 introduction

In previous weeks we have worked our way down the value chain from overall energy demand in an economy to refined products demand and from there through refining to crude supply.

This week we dive to the very bottom of the supply chain for oil (and gas)—the value-adding segment responsible for taking unknown quantities of oil in unknown locations deep beneath the earth's surface and finding and extracting it to create crude supply. This necessarily brings us face to face with the subject of risk.

However, before launching into this topic, it will be useful for you to have a picture of where we will be headed in the coming weeks. To this end, below is a figure reproduced from the Week 8 Weekly Notes.

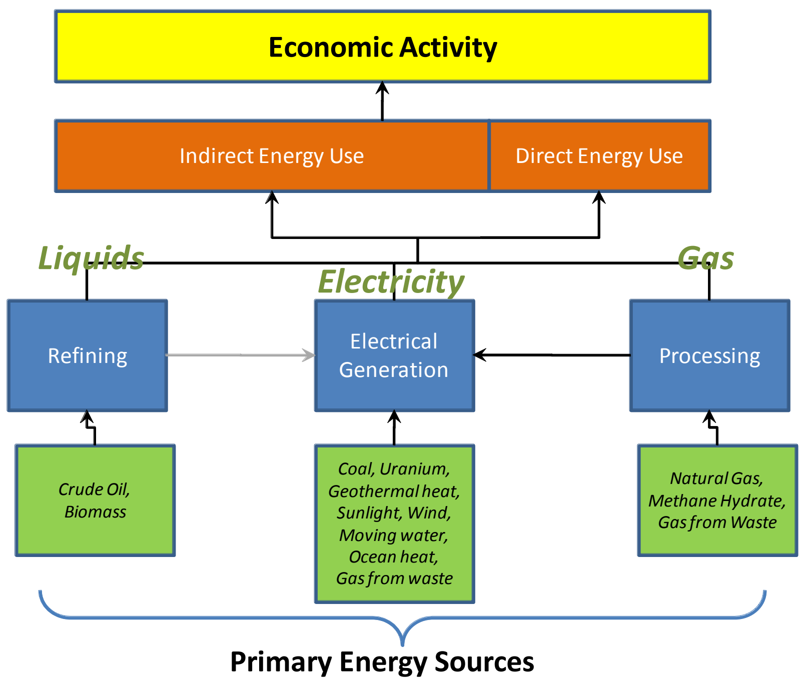

Figure 5.1. The energy supply "web"

This figure will be more understandable to you by the time we reach Week 8, but you can see here that the energy supply chain might more properly be called a "supply web," owing to interactions that occur at the level of conversion.

You see here that there are really three component "supply chains" that provide us useable energy in the three forms we require them—liquids, electricity, and gas. Each rests on a different combination of primary energy sources (showing in green). Each has a "conversion" component (showing in blue).

Thus far we have focused attention mostly on the liquids chain. We began at the top of the energy supply web in Week 2, exploring the link between energy use and economic activity and distinguishing indirect from direct energy use. In Week 3 we examined the refining segment of the liquids supply chain and in Week 4 we tied this to crude supply in global markets.

In future weeks we will explore the electricity supply chain, the gas supply chain, and will introduce the conventional and unconventional energy sources that compete with oil and gas and supply electrical generation, highlighting the prospects for new supply sources and new technologies in each chain.

This week is devoted to examining the bottom of both the liquids and gas supply chains with respect to the conventional energy sources—conventional oil and conventional natural gas.

Basic geology concepts

Textbook reading (Downey: Chapter 6, 'Exploration and Production' pp. 88–94)

To understand the risks involved in finding and producing oil and gas, you first need to have a basic understanding of a few geological concepts.

Chapter 6 of *Oil 101* gives you an excellent picture of exploration and production ("E&P") and contains almost everything you will need to know about it for this week. But here are three concepts to bear in mind as you read that may help you extract the essentials for working ahead:

First, **oil and gas are primarily found in sedimentary rock**, that is, rock that has been formed over geological time by the continual erosion of surface rock by weathering. The sediments formed by erosion bury previously-eroded sediments in layers that can become very deep over time as new layers are added. Surface conditions at the time a layer is laid down determine the composition of the sediment in that layer. If the sediment is formed, say, in a river delta, it will likely contain the remains of marine organisms. In a later geological period, tectonic activity may cause the river delta to uplift and no longer be a river delta but a desert area devoid of such organisms where sand from wind erosion accumulates on top of it—burying it, over a long time, in sand that in turn when buried becomes sandstone. This creates distinct sedimentary layers that accumulate to great depths in some areas of the earth's crust, some of which contain the detritus of marine organisms ("productive zones," in oil geologist-speak) and some of which do not.

Second, as described in *Oil 101*, **the resulting heat and pressures from overlying rock differ as to depth and have different effects on the buried organisms**. Typically, in the deepest zones, only gas can form, not oil. Above this, one can find pools of exclusively oil, exclusively gas, or pools containing both oil and gas (called "associated" gas).[1] In such cases, the producer is faced with a "joint products" problem somewhat like the one we saw in connection with refining, and one thing he/she sometimes must decide is whether it is more profitable (given prevailing and predicted oil and gas prices) to produce and market the associated gas or re-inject it into the reservoir to enhance the recovery of perhaps more profitable oil. Further, in making a drilling decision, the explorationist must take the potential gas/oil ratio into account in evaluating what drilling success will deliver by way of profit. This week we shall not be dealing with this problem; rather we shall assume the explorationist will discover either exclusively gas or exclusively oil.

Third, **a single well may produce from more than one zone at different depths, or sometimes, from different zones sequentially**. In real-world analyses this must be considered and modelled. In our modelling of a drilling decision this week we will not be considering this, but you should be aware of it.

What do geology and geophysics tell us?

Textbook reading (Downey: Chapter 6, 'Exploration and Production' pp. 98–101)

Geology and geophysics ("G&G" in oil industry jargon) has, as its goal, reducing drilling risk.

The main tool of the geophysicist is seismic testing, wherein sound waves are used to image sedimentary layers and other rock formations.

The critical thing to bear in mind about this for your work this week is that seismic testing **cannot see oil or gas**; it can only see rock. There is no technology available today that can actually look into the earth and see where the oil or gas is. The significance of this will become clearer to you when we later discuss the concept of value of information.

Geologists and geophysicists use seismic surveys to develop theories about the geological history of the region they are looking at. Often there are competing theories for explaining the seismic record. Geologists additionally rely on their experience and knowledge of the geology of the area. They can use this to help identify various sedimentary layers and rock formations ("structures") that may

[1] For completeness, it should be said that above the layer where temperatures are insufficient for oil to form, one sometimes finds "shallow gas".

contain oil. This does not tell them that a sedimentary layer known to be productive in another location will actually be oil-bearing at this location, but it can give them greater confidence that potentially productive zones may be accessible. Different geologists may have different theories for how the formations visible in the seismic came to be as they are; in doing this, they will often hypothesise descriptions of the geological dynamics over long eons of time. The prevailing view that has the greater internal consistency and cogency is one that tends to gain the most credibility. But the explorationist in charge will often find him- or herself faced with assigning different likelihoods to more than one plausible scenario.

Clues can also be had from technical methods other than seismic, such as remote imaging, magnetic imaging and others.

But, bottom line, geology and geophysics can give the explorationist only imperfect information.

The four geologic determinants of drilling success

Textbook reading (Downey: Chapter 6, ‘Exploration and Production’ pp. 94–101)

When oil is formed, differential pressures acting in the formation usually cause it to migrate, so it is not enough for the geologist to simply identify a particular formation likely to have given rise to the creation of oil. The presence of this “source rock” is *sine qua non* for drilling a successful well, but other factors must be taken into consideration, such as the pathway the oil may have followed in its migration away from the source. Such a pathway must be constituted by rock that is both *porous* (has spaces between its granules) and *permeable* (has channels connecting the pores so the oil can flow).

When geologists see a particular structure that they think might contain oil (or gas), they must consider four main things:

1. Whether there is *source rock* in the area they are looking;
2. Whether there is a plausible pathway whereby the oil could have found its way to the structure; i.e., whether *migration* has occurred;
3. Whether the structure they are looking at has *reservoir rock* that can accommodate oil in the present day (porous and permeable); and,
4. Whether the structure they are looking at has the characteristics to *trap* the oil and prevent it from having migrated to the surface in ancient times and disappeared.

These four geologic determinants of drilling success (source, migration, reservoir, and trap) are all uncertainties and must each be considered to assess the likelihood of drilling success.[2]

[2] Note that Downey (*Oil 101*) combines determinants 1 and 2 into a single determinant. Some geologists think of these two combined this way; most do not.

The Venn diagram in *Oil 101* (page 98, Figure 6-8) gives you a conceptual picture of how the probabilities of each of these elements combine together to give a usually low probability of drilling success for any one exploration well. For those of you familiar with probability theory, you will appreciate why the combined probability is typically low. Specifically, if we designate the four component probabilities as p_S (probability the source rock exists), p_M (probability migration to the structure has occurred), p_R (probability that reservoir rock exists), and p_T (probability of a trap), the probability that all four will simultaneously occur is

$$P = p_S \, p_M \, p_R \, p_T$$

Even if the probability that each of these is true is 50%, the joint probability will be only 6.25%.

In actuality, it is not quite as simple as this since the above equation assumes the probabilities are independent of each other (whereas, for example, the presence of a porous/permeable migration path may increase the likelihood there exists reservoir rock, making independence not a valid assumption). But, the conclusion is valid: the probability of success for an exploration well (and particularly a "wildcat" well, about which more below) is generally very low because the dependence is weak and because the likelihoods of each of the components of success are generally low and must all simultaneously prove true.

Exploration is a risky business.

Wildcat drilling versus development drilling

Textbook reading (Downey: Chapter 6, 'Exploration and Production' pp. 101–124 and Margonelli: Chapter 4, 'Drilling Rig')

From the point of view of a decision maker, there are two fundamentally different types of drilling decisions. The first has to do with exploration drilling in a new area or new field that has not been drilled before. This carries very high risk, as you will appreciate from the discussion in the last section (exploration drilling is extremely expensive and the probability of success is low). This kind of drilling is called "wildcat" drilling. While it is high risk, it can have high reward. Finding a new field or new productive zone in a previously unexplored area can generate huge potential returns to the company with drilling rights to the area.

Once a field has been discovered and drilling has proved successful, drillers can rely on information gleaned from the previously-drilled well(s). Those wells will typically have identified productive zones through well logging and core sampling and this creates confidence that these zones will be productive in nearby areas. Successful drilling also tends to confirm the assumptions of some geologists and refute those of others.[3] The more that is understood about the "play," the lower the drilling risk. While it is still exploration drilling, it is not "wildcat" drilling.

The least risky drilling is development drilling, drilling for purposes of producing more of the oil faster. Well logging, pressure testing and coring from previous wells combined with newly-informed geology and geophysics will give strong clues about the areal extent of the productive zone and reduce the risk. Further, the production behaviour of past wells can give reservoir engineers knowledge of reservoir characteristics that makes development well placement more efficient.

So when in later sections we talk about the probability of drilling success, bear in mind that this probability can be vastly different across different drilling types.

The economics of oil and gas supply

Finding oil and gas deep within the earth and moving them to their respective markets requires several steps along their respective supply chains, each segment carrying its own economics and requiring its own set of decisions.

The supply chain for the oil part of the liquids supply chain looks like the following:

Figure 5.2. Liquids supply chain detail: oil supply chain

The supply chain for the natural gas part of the gas supply chain looks like the following:

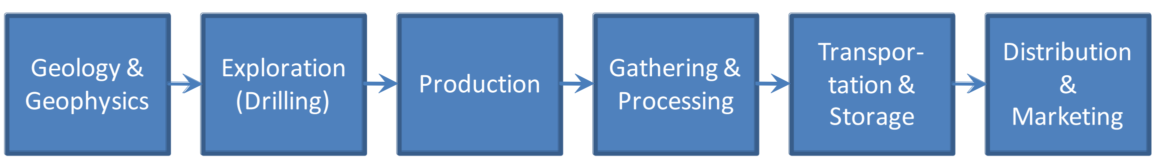

Figure 5.3. Gas supply chain detail: natural gas supply chain

In the oil supply chain, crude oil must first be discovered and produced. Then it must be moved to the refining segment via pipeline, tanker ship, or occasionally, tanker truck. In Week 3 we examined the refining segment. Then, the refined products must be distributed to the marketing segment again via pipeline, tanker ship, or most frequently, tanker truck. Different products move through different channels to market. For example, jet fuel is frequently transported from the

[3] And not just successful drilling, but "dry holes." Dry holes, while expensive, often convey positive (and therefore valuable) information about the geology.

refinery to local airports via pipeline since the volumes are sufficiently large. But gasoline and diesel are nearly always transported to petrol stations via tanker trucks operated by “jobbers.” Sometimes refined products move via pipeline to ports for export, are carried via smaller tanker ships to the receiving port, and from there are distributed via tanker truck.

In the natural gas supply chain, the gas must first be discovered and produced just as with crude oil. The next segment in the supply chain is gathering and processing (described in Week 6) followed by a transportation and storage segment (usually transportation is via pipeline but with a growing movement toward moving liquefied gas (LNG) via tanker internationally—described also in Week 6). Since it is a gaseous fluid, part of the transportation segment involves storage facilities (again described in Week 6). Distribution to industrial, commercial, and residential customers is nearly always via pipeline, although there are exceptions.

The economics of oil production is driven largely by oil price, which, as we saw in Week 4, is effectively set by OPEC. For an individual oil producer, his/her actual realised price is set by the “netback” the producer receives from the particular refinery supplied—the price the refiner sets at the “refinery gate” less the cost to transport the crude to that location. The price the refiner sets at the “refinery gate” will be determined in turn by the global crude price adjusted for crude transportation costs (and adjusted for crude quality differentials as described in Week 3).

The other main components of oil production economics are capital cost and operating costs (and of course production volumes, which are determined from recoverable reserves and decline rates as described in Week 4). In general capital (and to a lesser degree operating) costs have increased substantially over the last two decades as easier-to-access oil deposits have been largely exploited and the majority of remaining exploration and production opportunities have required moving to more hostile locations, such as Arctic regions and deeper offshore. Very expensive exploration and production technologies are needed for exploitation in these more hostile environments. You no doubt gained an appreciation of the sophisticated equipment and facilities needed, and the correspondingly great expense involved in such endeavours, from your reading of *Oil 101*, Chapter 6, pp. 101–124 in association with the previous section. On top of this, these costs are virtually never known for certain until the exploration wells have been drilled and the production facilities have been put in place and have been operating for some time.

And naturally, for the explorationist, a key to exploration economics is the likelihood of drilling success and the likely recoverable reserves discovered given drilling success. This bulk of this week’s discussion focuses on how to evaluate exploration economics given these uncertainties. But, as you will see, much also depends on the (mostly capital) costs of drilling.

For natural gas, the story is almost identical for the geology & geophysics, exploration drilling, and production segments. As you will see in Week 6, the price of gas is set in part by the global crude price and in part by refined product price differentials.

To summarise the key risks involved at the base of the oil and gas supply chain, we have: the geological risks of drilling (drilling success and recoverable reserves discovered); the uncertain capital costs of exploration drilling, the uncertain capital and operating costs of production; and, of arguably greatest magnitude, the uncertain price of oil over the (sometimes decades) over which production takes place.

Not only are the exploration and production segments of the oil and gas supply chain highly risky, they represent a substantial burden on the economy. Globally, oil and gas companies spend hundreds of billions of dollars annually for exploration and production.[4] This represents a significant drain on the limited capital resources of the global economy, and this burden can only be expected to increase.

Global exploration opportunities

Historically, oil and gas discoveries in the late 1800s and early 1900s were mostly confined to the United States and Russia, although in the 1890s Royal Dutch Shell discovered oil in Indonesia. Production in the Middle East commenced with the discovery by what is now British Petroleum in Iran (then Persia) in 1908. In 1938, oil was discovered in Saudi Arabia, an event that was to shape the character of global oil markets to the present day and undoubtedly far beyond, as elucidated in Week 4. Since that event, oil has been discovered in sedimentary basins throughout most of the world (see *Oil 101*, p. 95, Figure 6-5 for a map of the world's sedimentary basins).

By comparison to today, these oil deposits were easy to find. Particularly in the Middle East, comparatively little effort was required to expand the world's known reserves substantially. That picture has changed.

Today, while major exploration efforts continue in areas known to have proven fruitful in the past, these efforts are showing diminishing returns. North Sea oil was discovered in the late 1960s, and production reached a peak around 1999 and has been steadily declining since. This oil was much more expensive to find and produce than oil found previously, owing to the cost of platforms designed for offshore production in challenging North Sea conditions. In general, offshore drilling and production are significantly more expensive than onshore.

---

[4] See, for instance, GlobalData's Global Oil and Gas Capital Expenditure Outlook – 2010, available at http://www.investorideas.com/Research/PDFs/CAPEX_PR.pdf.

Mexico has gone through a similar cycle, becoming a significant oil exporter by the mid-1970s, but production peaked and has since continued to decline. Similarly, production in the United States peaked in 1970 and has declined ever since. While Russian oil production fell significantly in the 1990s, it increased significantly in the 2000s, making Russia an oil producer of comparable size to Saudi Arabia. Nonetheless, most oil analysts believe that increasing global oil production through new major discoveries is becoming less likely over time. A recent exception has been a major offshore discovery by Brazil.

On the surface, the massive oil production potential of Saudi Arabia makes it by far the most significant opportunity for exploration activity. However, the Kingdom has long determined to restrict exploration activity, worrying that increased production capability could potentially force it to produce at levels that would be far above optimal in terms of sustaining global oil prices at high levels.

All this means that new exploration opportunities will likely be restricted to developing marginal fields in existing areas, with finding costs per barrel steadily increasing since fewer barrels will be found on average per unit of exploration capital spent. The only other possibility is moving into more remote and hostile areas, which again promises increased finding costs. There may be significant remaining potential in very deep offshore areas, but deep-sea drilling and production is extremely expensive and current technology places limits on the attainable water depth, although technology advances are steadily increasing the depths at which drilling and production can be attained.

For natural gas, the story is similar, although the geographic distribution of new exploration opportunities is different. Substantial increases in gas reserves have been seen recently in the Middle East, primarily among OPEC countries and in the Asia-Pacific region. Russia is by far the largest gas producer, and Russian reserves have been holding steady. Central Asia shows promise as a new and productive exploration place for natural gas.

As with oil, new gas reserves will generally come along with higher finding costs. However, new technology shows promise in the United States for recovering so-called “tight” gas, which some believe may represent substantial new reserves.

Overall, exploration opportunities for both oil and gas will come with increasing challenges, both finding cost-wise and technical feasibility-wise.

Assessing uncertainties

We now return to the root of both the oil and natural gas supply chains, geology and geophysics and exploration drilling. These are the segments arguably fraught with the greatest risk.

In the following section we shall be introducing you to decision and risk analysis, but central to this discipline is the quantification of uncertainties.

For a decision like a drilling decision, there are two types of uncertainties that need to be quantified: *discrete*, and specifically *binary* uncertainties (uncertainties with only two possible outcomes, like a well will be successful or not), and *continuous*, or *range-based* uncertainties (like the range of possible future oil prices).

Decision analysts adopt a *Bayesian* perspective in treating uncertainty.[5] Because so many decisions relate to new and unique situations, it is not generally possible to rely on historical data in the way classical statisticians do (a company, in assessing the likelihood that a contemplated acquisition will be successful, cannot do an experiment where they make the acquisition a dozen times and see how many times it is successful). Instead, Bayesians treat uncertainty as *subjective* and rely on assessment of probabilities from experts to obtain quantitative measures. To be sure, these experts can rely on historical data and statistical analyses in forming their judgement, but such data are frequently of limited value to, say, a geologist assessing the likelihood source rock exists in a particular area.

The assessment of uncertainty has advanced to something of an art form. The task is complicated by the fact that psychologists have identified certain "cognitive biases" that appear to operate in individuals (across all nations and cultures, incidentally). By way of example, it can be shown[6] that when asked to report a 10%-90% range of uncertainty individuals tend to report their ranges as narrower than they should be to accurately reflect their own, true internal beliefs about the degree of uncertainty they have. This is largely due to a translation problem: people are not very skilled at taking their beliefs and translating them into numbers, into percentages.

To account for this, decision analysts have developed tools that make it easier for the expert to do this translation. A particularly useful tool is the probability wheel. For those of you interested, Box 1 gives a description of the probability wheel.

Range-based uncertainties are typically assessed by picking three points off the continuous distribution residing in the head of the expert—the so-called P10, P50, P90 points. The P10 is the low value, technically the 10^{th} percentile of the expert's distribution (the number—the quantity of reserves, say—that is low enough such that the expert believes there is only a 10 percent chance the *actual* reserves will turn out to be lower than this quantity). Similarly, the P50 number is the 50^{th} percentile of the expert's distribution (technically, the median). Likewise, the P90 is the 90^{th} percentile (the number the expert believes is high enough such that there is a 90 percent chance the actual value is less than that quantity).

[5] So-named for Thomas Bayes, an eighteenth-century mathematician and clergyman.

[6] Tversky, A. & Kahneman, D. (1974) 'Judgment under uncertainty: heuristics and biases', *Science*, 185, pp.1124-1131.

For a binary uncertainty (event happens/does not happen), the uncertainty is assessed as a single percentage—the likelihood the event will happen (the percentage on the other "leg" of the uncertain event can then be imputed as 100% minus the assessed number).

The probability wheel is used for both range-based and binary uncertainties.

In the drilling decision described below we use three uncertainties: drilling success (binary uncertainty), recoverable reserves (range-based uncertainty), and future oil price (range-based uncertainty).

Box 1: The probability wheel

The probability wheel is a tool developed by Dr. Carl Spetzler, co-founder of Strategic Decisions Group headquartered in Palo Alto, California. The tool presents people with a series of hypothetical "bets," from which subjects can extract numerical assessments of confidence. Once the bets have been completed, people agree that the numbers they arrive at reflect all the nuances they want to consider, including knowledge, intuition and even feelings about any given possible event. The event can be, in particular, the event that an oil well successfully discovers economic quantities of crude.

The probability wheel is similar to a pie chart containing two adjustable colours. It can show all red, or it can show mostly red with a slim pie-slice of blue or any other proportion of red to blue. On the back is a scale (hidden from the subject to prevent them "locking" on some specific number) that tells what percentage of the wheel is showing red at any given time. Attached to the centre of the wheel is a pointer, which can be spun much like the pointer in a board game.

The figure below depicts the process of using the probability wheel to extract a quantitative assessment from the subject. The branches emanating from the yellow square depict a choice we ask the subject to make.

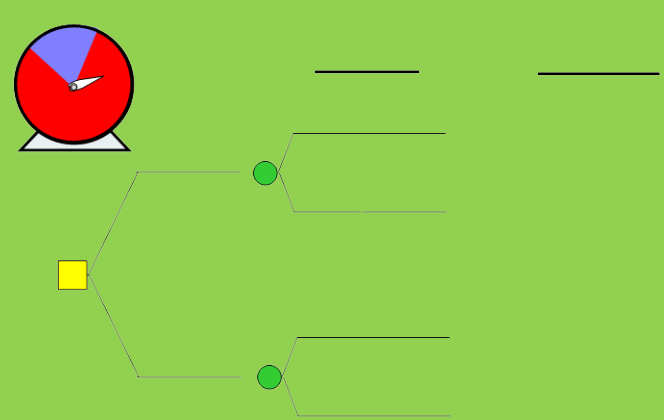

Figure 5.4. Using the probability wheel to assess the likelihood of drilling success

Suppose the subject is considering: how confident am I that this particular well will be successful? To help the subject quantify her answer, she plays a game that involves making a large (though fictitious) bet. The player can choose to place her bet in one of two ways: she can put her money on the probability wheel or on a “clairvoyant”. The clairvoyant, for purposes of the game, is a hypothetical person who can see into the situation and know with absolute certainty what the reality is. In our case, we could postulate to the subject that our “clairvoyant” knows with certainty whether a well drilled in some specific location will be successful or not.

The rewards are as follows: if the player chooses to bet on the wheel and the pointer ends up on red, she will win $100 million. If the pointer ends up on blue, she will win nothing. Alternatively, if the player chooses to bet on the clairvoyant, and the clairvoyant reveals that the well will indeed be successful, the player wins $100 million. If the clairvoyant reveals that the well will be a dry hole, the player wins nothing. The question for the player, at each stage in the game, is whether to bet on the wheel or the clairvoyant so as to maximise her chance of winning. Note that the reward structure is irrelevant, so long as the reward structure is the same for both alternatives. The reward could be “you get to live” and the cost could be “you are shot dead.” The objective is to focus attention on the reward structure as if it were real.

At the beginning of the betting the subject will not be able to state her true confidence level as a percentage, so she first considers some extreme bets, setting the wheel at say 100 percent red and 0 percent blue. Then she asks herself: would I rather bet my money that the spinner will end up on red, or that the clairvoyant will reveal that the well has been successful.

A quick glance at the wheel makes it clear to this player that the odds are overwhelming for the spinner to end on red. If, by contrast, she believes (but at this point cannot explicitly articulate) that there is only about a 60 percent chance that the well will be successful, she will quickly decide she has a better chance of winning her $100 million by betting on the wheel.

Next, she adjusts the wheel to the opposite extreme: 0 percent red and 100 percent blue. This time the player perceives the assurance that the spinner will end up in blue; she opts for the event (producing well), which now seems a safer bet. From this point, a series of increasingly fine adjustments is made, much like the adjustments that take place in an ophthalmologist’s offices. The patient being fitted for eyeglasses is asked to read an eye chart through lenses of differing strengths—stronger, then less strong. The subject continues to adjust the amount of red on the wheel, each time asking her to choose whether she prefers to bet her money on the wheel or the event of drilling success. Each time she finds herself choosing the wheel; she decreases the amount of red and asks again. Each time she chooses the event; she increases the amount of red and asks again. When the amount of red showing leaves her just as willing to bet on the wheel as on the event—when, in other words she is indifferent about her choice of bets—then her degree of confidence has been quantified. It takes about twenty minutes to teach people the probability wheel technique.

At the end, what we have is an assessed percentage probability for the event on which the expert has declared they would be willing to bet their money (or their life, if necessary). Importantly, at no point has the subject been asked to state a number.

Similar methods are used for assessing the P10, P50, P90 points of a range-based uncertainty.

Decision and risk analysis: the expected value concept

This week we will be evaluating—and you will be analysing—a drilling decision. To evaluate decisions, many companies (especially oil and gas companies) use a methodology called by many of them decision and risk analysis (DRA). The academic community calls this discipline simply decision analysis (DA), but we shall from this point forward use the DRA descriptor.

To understand the core concept of DRA consider the figure below:

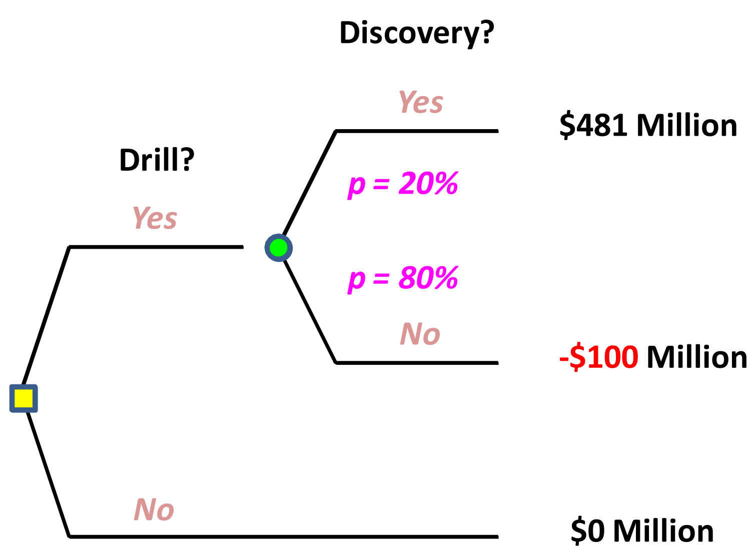

Figure 5.5. The drilling decision

Here we see a decision cast in DRA terms. In DRA, the convention is that a square or rectangle represents a decision. A circle or ellipse represents an uncertainty. The *decision* has two alternatives: drill or don't drill. If the drill alternative is chosen, there is an *uncertainty* as to whether the result will be a discovery. As shown here, there is a 20% probability that the result will be a discovery. To the right of the decision and the uncertainty is the third component of any decision—the *value* component. Here we have depicted the outcome values, for simplicity, as *certain* values. For what we need to describe at the moment, think about these values this way: if the driller drills the well and it is a discovery, the drilling company will receive a tax-free check for $481 million the moment the discovery is confirmed; if the driller drills the well and it is a dry hole, the company will have to write a check for $100 million (the cost of the well). If the driller decides not to drill, the value of choosing this decision alternative is zero (“nothing ventured, nothing gained”).

How do we evaluate which is the best alternative for the driller? You can see that if the driller decides to drill, the odds heavily favour losing $100 million (in fact the odds in favour of losing are 4 to 1). But note, too, that if the well is a discovery, the upside is much larger than what the driller loses if it is a dry hole. A simplistic analysis would say that the driller should go with the odds—the odds against making a discovery are 4 to 1, so the odds considered by themselves say don’t drill. But what if the check the driller receives for a discovery were not $481 million but, say, $300 billion? A moment’s thought will tell you this makes the drill alternative highly attractive, despite the odds against success.

The decision analyst will look at this is by appealing to the concept of “expected value.” The idea behind *expected value* is to take the outcome values and weigh them by their likelihood of happening. Mathematically, if we call the value of outcome #1 v_1 and the value of outcome #2 v_2 (outcomes #1 and #2 can be any two outcomes of a binary event) while the respective probabilities of these outcomes occurring are p_1 and p_2, the expected value is defined as:

$$EV = p_1 v_1 + p_2 v_2$$

For the example in Figure 5.5, the expected value of choosing the drill alternative is

$$EV = 0.2 \cdot \$481 + 0.8 \cdot \left(-\$100\right) = \$16.3$$

That is, the expected value of drilling is **$16.3 million**, whereas the expected value of not drilling is $0. On this basis, the analysis says the driller should proceed with the drill decision since it is the alternative with the highest monetary value.

Note that the expected value is a summary measure that contains within it the consideration of downside risk: it weighs the good outcome against the bad outcome and gives a value that in this case favours the good outcome (considering the likelihood of its occurrence).

Now, you may feel uncomfortable with the apparently mechanical and formulaic nature of this way to look at a decision. For instance, if you think of applying it to a personal decision you may face involving a financial investment, you may conclude that you would look at it a little differently. Most people do. If the investment opportunity is very large relative to your financial resources, for instance, you will likely want to pay more attention to the magnitude of the loss you could bear, and especially if the odds favour it happening. The downside risk may simply be too great for you.

For personal decisions, decision analysts accommodate such considerations by applying what is called utility theory, which is a way to incorporate into the analysis allowance for different individuals' "risk attitude," or "risk tolerance." This is a topic that is beyond the scope of this module, but it may assuage your concerns to know that DRA as a methodology can account for these considerations.[7]

For corporate decisions, the use of expected value as a measure of decision merit takes us deep into the subject of the theory of finance, but suffice it to say for our purposes here that the modern theory of finance argues in its favour for publicly-held companies who account for their shareholders' preferences by assuming these shareholders are diversified.[8] That said, managers of even publicly held companies will typically want to look explicitly at the downside risk, a subject we turn to next.

Decision and risk analysis: the concept of downside risk

A full-fledged decision analysis of a major decision will consider many uncertainties (and often, many different component decisions). This means the universe of possible outcomes considered is large and the outcomes therefore numerous, unlike the situation depicted in Figure 5.5.

With numerous outcomes comes a way to usefully depict downside risk. For each decision alternative, we can construct an "S-curve", more properly called a cumulative probability distribution (in your business statistics courses, you may have seen this referred to as an "Ogive"). In general, if there are n decision alternatives under consideration, there will be n S-curves, one for each alternative. An example is shown below:

[7] Tersely stated, decision analysts encode, using a highly structured set of questions, the decision maker's "utility curve." This curve maps (usually monetary) outcome values onto utility values that reflect the decision maker's tolerance for (or equivalently, aversion to) risk. The curve is used to translate the raw (usually monetary) outcomes into utility values. Rather than calculate the expected value as in the equations above, an identical calculation is used to calculate the expected utility. When the expected utility is reconverted, using the curve, to (usually monetary) terms, the resulting value is the "risk-adjusted" value (sometimes called the "certain equivalent"). In general, the risk-adjusted value will be less than the expected value. If the decision maker is highly risk-averse, the risk-adjusted value can be negative. In the drill decision, if the risk-adjusted value is negative, the driller should choose not to drill (even though the expected value is positive), given his/her attitude towards risk.

[8] Privately held companies, whose owners are generally not widely diversified in their investment portfolios, will want to incorporate utility considerations.

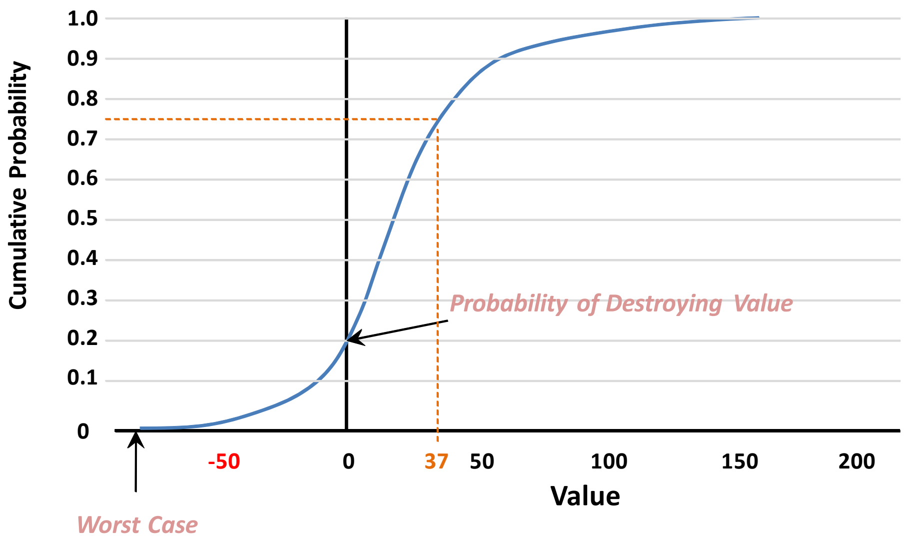

Figure 5.6. Example of an S-Curve

Most of you will be more familiar with a probability curve that is properly called a "density function", often a bell-shaped curve that shows the value of an outcome on the horizontal axis and its probability of occurrence on the vertical axis. By contrast, the cumulative curve is formed by taking such a density function and integrating it.

Putting the curve in this cumulative form allows us to extract risk-related information that is not readily apparent to us from a simple density function curve. The way to read this curve is this: pick some probability on the vertical axis on the left, say 0.75 (meaning 75%); draw a horizontal line from this point to the right until it intersects the S-curve; then drop a vertical line from this intersection point until it intersects the horizontal axis. For this particular curve, you can see that the intersection point has a value of 37 (units can be Euros, dollars, pounds…). The interpretation of what you have done is this: for the decision alternative represented by this S-curve, there is a 75% chance that choosing this alternative will result in a value of *less than or equal* to 37. Alternatively stated, there is a 25% chance the alternative will deliver *greater than* 37.

Similarly, you can see for this S-curve that there is a 90% chance that the alternative will deliver *less than or equal* 60 units of value. (You can't see this exact number in this example, but it is around 60.) An important point on the curve is the point it crosses the zero value on the horizontal axis. You can see that this point corresponds to a cumulative probability value of 0.2. This says that if one were to choose this decision alternative, there is a 20% chance it would deliver a value *less than or equal* zero. Stated another way, there is a 20% chance this alternative will deliver negative value. If the value measure used is NPV, we can say there is a 20% probability of *destroying shareholder value* with this alternative.

Thus, the S-curve gives us a handy way to measure the risk of choosing a particular alternative. In addition to looking at the probability of destroying shareholder value, managers will typically want to look at the worst case outcome, showing on the left in Figure 5.6. There are more sophisticated interpretations of risk available from the S-curve we will not describe except to say that you can imagine that if we shaded the area under the curve on the left side of the zero value line it conveys information about the risk of this alternative.

Naturally, the S-curve gives us a useful way to compare different alternatives' risks.

A drilling decision

This week you will be analysing a drilling decision. For pedagogy purposes, the decision is simplified in several significant ways compared to a real drilling decision, but this simplified version is sufficient to give you a fairly complete understanding of how to do such an analysis.

The decision you will address is illustrated below:

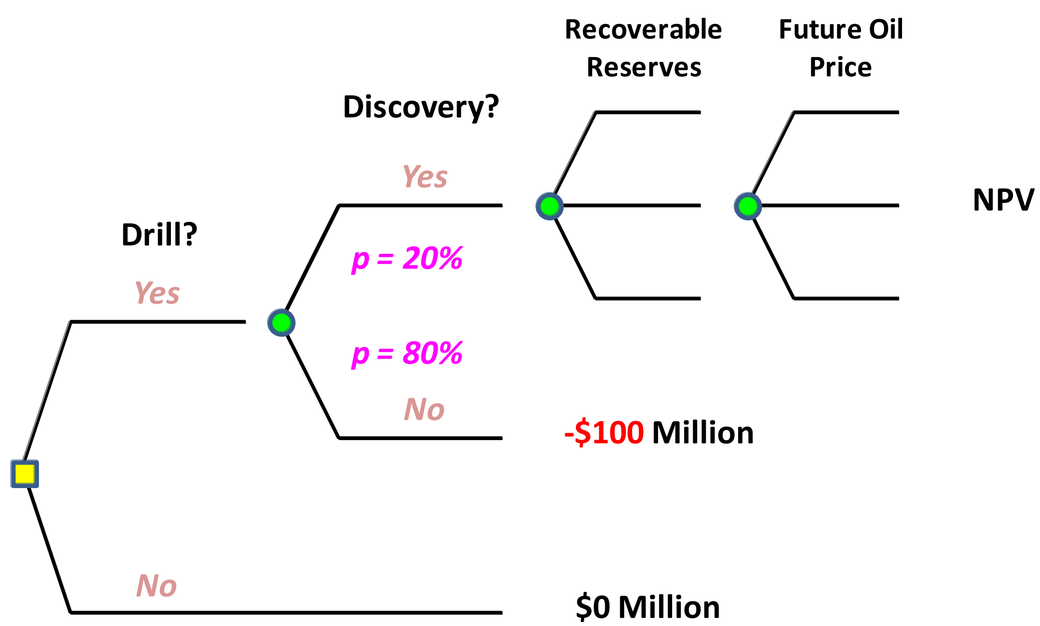

Figure 5.7. The (more complete) drilling decision

This figure looks much like Figure 5.5, but two uncertainties have been added in this case. First, this new tree structure contemplates that the quantity of actual reserves recoverable given a successful discovery are unknown. The driller does not know, prior to deciding to drill, whether a discovery will find a large or small quantity of reserves. Second, when those reserves are produced, it is not known what price the oil can be sold for, since production will typically occur over decades and the future oil price—as you saw from your work in Week 4—is highly uncertain.

You see to the right the designator "NPV." In connection with Figure 5.5, we made the simplifying assumption that the driller would be handed a check the moment the discovery is confirmed. But in reality, most oil companies drill for their own account—drill to produce and sell the discovered oil or gas to the market. This means revenues will be flowing to them over time (and expenses will be incurred over time). From Week 3, you will remember being introduced to the concept of net present value (NPV), and how it is used to value cash flows spread out over time. We need to do the same thing here to adequately represent a (relatively) realistic drilling decision.

Note that this is a "conceptual" decision tree. It is merely representative of the actual tree that must be used for analysis. Specifically, to make the image compact and manageable the two uncertainties on the right are each showing as having three branches but these three-branch "nodes" are shown side-by-side. This part of the tree should properly be shown as a tree with nine branches because there are nine possible outcomes: one can have an outcome involving high recoverable reserves along with low oil price, high recoverable reserves along with high oil price, low recoverable reserves with mid-level oil price, etc. We need to introduce you to this compact representation at this point because in later sections the trees would otherwise become too large and unmanageable. But realise that the actual tree structure for this part of the tree looks as follows:

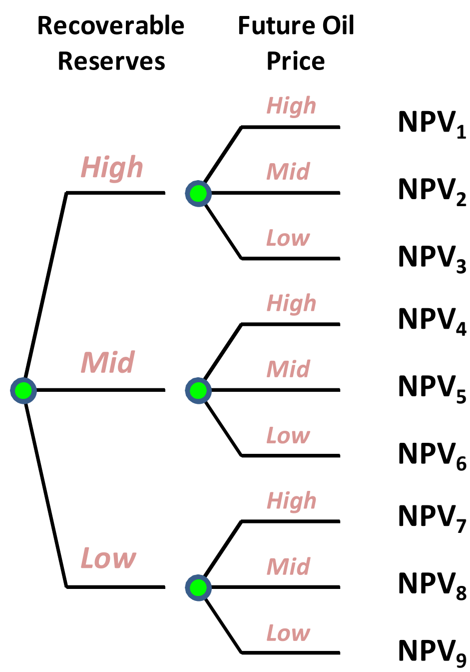

Figure 5.8. The full depiction of the two range-based uncertainties

Each of the nine outcomes will result in a different (in general) NPV outcome—will deliver a different NPV to the driller.

To do the analysis, we need to attach probabilities to each of the branches. Figure 5.7 shows the probability of discovery we used before: 20%. To attach probabilities to the branches of the range-based uncertainties we rely on the P10, P50, and P90 assessments we obtain from experts.

There is a subtlety here. The P10, P50, P90 assessments are in effect points taken from a *continuous* distribution residing in the expert's head (and extracted by the decision analyst via assessment). What is needed here is instead a representation that is *discrete*, since we are representing what is in fact a continuous distribution by a three-branch node. To avoid delving into the rather involved mathematics of this question, suspend your disbelief and take it as (an entirely non-obvious) given that we can take the expert's P10, P50, P90 points and validly represent his/her continuous distribution by a three-branch node whose high, mid, and low values are the 10/50/90 values, with probabilities assigned to the branches of 25%, 50%, 25%, respectively.

So, in Figure 5.8 think of replacing the words "high" with a probability of 25% (or 0.25), the words "mid" with the probability 0.50, and the words "low" with the probability 0.25. Note that by doing so we are in this case making an assumption of *independence* between these two uncertainties. This reflects the idea that, for instance, knowledge of actual reserves the driller discovers tells us nothing about the global price of oil. In a later section, you will learn how to deal with the situation where there *is* dependence between uncertainties.

One solves this new tree similarly to the tree of Figure 5.5. We calculate the expected value of each alternative. In this case, this is somewhat more involved, but basically one begins with the nodes on the right-hand side of the tree, calculates the expected NPV from the endpoint NPVs associated with that node and that node's probabilities, replaces the node with this calculated expected value, and uses this expected value as the endpoint value when calculating the expected value of the next node to the left of it. This is called "rolling back" the tree. In your Hand-in Assignment this week you will not have to do this—the interactive model for this week does it for you—but you should know that this is how it is done.

The NPV associated with each different endpoint is a separate calculation that uses the recoverable reserves and oil price contained in the pathway through the tree that attaches to the endpoint. For example, NPV_6 in Figure 5.8 is calculated using the mid value of recoverable reserves and the low value of oil price. For a given level of recoverable reserves, production from the well over time is calculated using the relationship among reserves, initial production, and decline rates you saw in Week 4. This is all encoded in the interactive model if you want further detail on how this works.

A final note on a particular simplification adopted here. The future oil price is treated as if it were constant for all time—there are high, mid, and low values for it as you have seen, but each is treated as fixed. Methods are available to consider oil price time dynamics, but are more involved than we need to use here.[9]

Decision and risk analysis: the value of information concept

The analysis described above is sufficient to give the driller guidance on whether or not to drill the well (and you will be examining this question under different input assumptions in this week's assignment), but we can take the analysis a step further to understand how the driller's decision might change if he/she had better information to go on.

We begin simply. We ask ourselves the question, "how would the driller's decision be affected if he/she knew for certain the recoverable reserves before drilling?"

To put this in a DRA framework, it is useful to picture that the driller hypothetically has available to him/her a "clairvoyant," someone who has access to *perfect information*. In this case we postulate that the clairvoyant can see into the ground and determine with perfect accuracy the quantity of recoverable reserves.

Conceptually, what we do is change the ordering of the nodes showing in Figure 5.7. In DRA, the ordering of uncertainty nodes relative to decision nodes indicates the timing of uncertainty resolution with respect to decisions. Specifically, if an uncertainty node precedes a decision node, this structure says the uncertainty will be completely resolved (i.e., will become a certainty) prior to the time the decision will be made. We postulate that the clairvoyant will completely resolve the uncertainty about reserves. So the ordering of the nodes changes as indicated below:

[9] Basically, these methods involve partitioning the time horizon into periods and then obtaining assessments for each period conditional on the oil price outcomes in previous periods. Often these prices are further conditioned on underlying drivers.

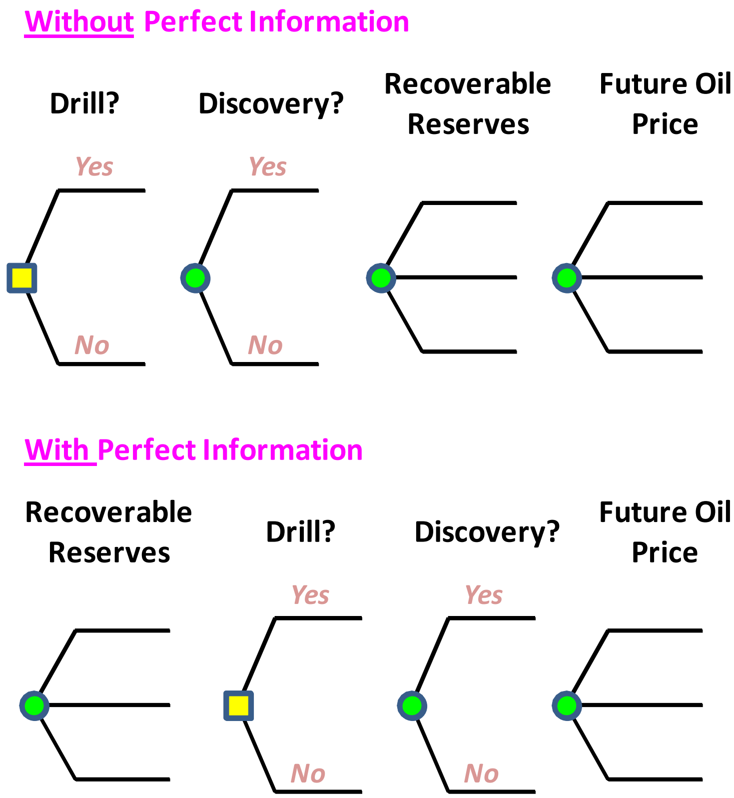

Figure 5.9. The decision with and without perfect information

This is only a conceptual tree. To do the actual calculation we must re-order the nodes in the actual tree. Mechanically, this involves taking the Recoverable Reserves node in Figure 5.7, pulling it out in front of the decision node, and cloning the decision node to attach to each of the branches of the Recoverable Reserves node. This is shown below:

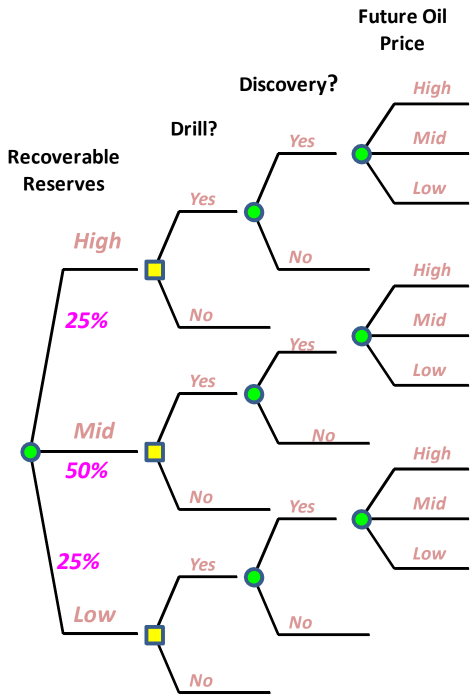

Figure 5.10. The re-ordered tree

We calculate the expected value given this new structure. Then, the *value of perfect information* is calculated as: the expected value *with* perfect information (the value we calculate from this new tree) less the expected value *without* perfect information (the value we calculated from the tree of Figure 5.5).

In the interactive model you will see a tree that looks exactly like this. A truncated version of the tree there (for Base Case assumptions) is shown below:

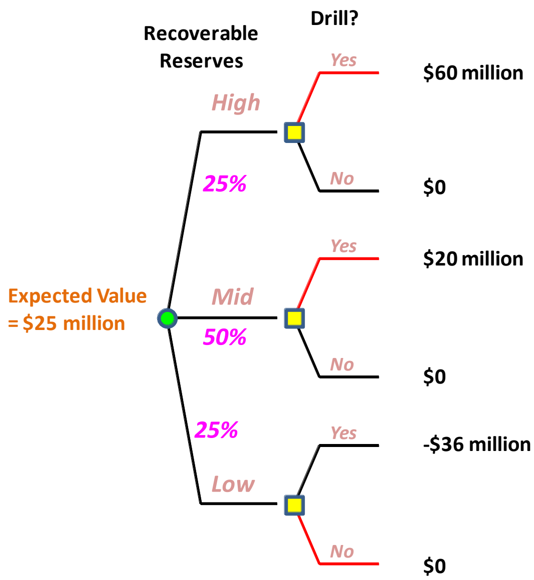

Figure 5.11. The *calculated* re-ordered tree

In calculating the expected value for a *with* perfect information tree like this, you only consider the outcome values associated with the decision chosen in each case (the respective decision choices are showing as red highlighted branches). Accordingly, the three outcome values to be multiplied by the 25%, 50%, 25% probabilities on the Recoverable Reserves node are $60 million, $20 million, and, not -$36 million, but $0.

You see that the value *with* perfect information is therefore $25 million and the value *without* perfect information is the $16 million we saw in connection with Figure 5.5 (these values are rounded). In Figure 5.5, we assumed the driller would know what the recoverable reserves were only *after* drilling. Accordingly, the value *of* perfect information in this case is the difference between the two, or $9 million.

What does this tell us? It only tells us that if he/she had access to a clairvoyant who could tell us for certain what the actual recoverable reserves were, it would be worth $9 million to the driller. Of course, such a clairvoyant is not available to the driller, but the calculation does give us an *upper bound* for what information of this nature is worth. So, for example, if a geology and geophysics (G&G) company approaches the driller and offers to do a seismic survey for $12 million, the driller can easily determine that the price is too high to be of value.

Figure 5.11 contains within it a very important concept, the *value of information principle*. Note that some of the decision branches in Figure 5.11 are coloured red. The red branches show the best decision alternative given different reports from the clairvoyant. You can see that if the clairvoyant reports the actual recoverable reserves to be in either the high or mid range, the best alternative is to drill; if the clairvoyant reports the reserves to be at the low end, the best alternative is to *not* drill (since drilling would result in an expected value of -$36 million). This allows us to state the principle:

VALUE OF INFORMATION PRINCIPLE: Information *only* has value if it has the potential to change a decision.

Figure 5.11 shows the decision changing depending on what information the clairvoyant reports. So there is value of information. The value of information principle is something you can think of applying generally in your careers. It is frequently the case that companies will pay large sums of money for information that does not change any decision. The value of information principle would say that this is a waste of money.

There is a point that needs clarification as it is often a source of confusion. Notice that the Recoverable Reserves node in Figure 5.11 has the probabilities 25%/50%/25% attached to the branches. A question that may have occurred to you is, “if the information we are getting is *perfect*, why on earth would we attach probabilities to different reserves levels?" The way to think about this is the following: The question we are asking here is whether or not the driller should consult the clairvoyant. Prior to determining whether this is a value-creating action, we do not know what the clairvoyant will report. The clairvoyant could report that the reserves are high, are mid-range, or low. So we need to attach probabilities to the possibility that any of these will occur. Since we do not know, prior to consulting the clairvoyant, whether this individual will report high, mid-range, or low, we must take our best shot at assessing these possibilities. The only probabilities we have available to us are the probabilities we originally assessed as to the likelihood that the *actual* values will turn out high, mid-range, or low—our so-called *prior* probabilities. Accordingly, we assign these probabilities to each of the possible clairvoyant reports.

The interactive model you will be using this week does all these calculations for you. It also colours the decision branches just as is shown in Figure 5.11. These can change if you change the input assumptions.

This is the value of *perfect* information calculation. In reality, what the driller will want to know is the value of obtaining *imperfect* information by means of conducting G&G. This is the final topic we address this week.

The value of imperfect information: geology and geophysics

The information the driller can obtain from G&G is imperfect. The very best G&G available on the planet delivers information shot through with uncertainty. But this does not mean it is without value.

The DRA methodology deals with this by explicitly considering the degree to which the information obtained may be unreliable. It does this by adding an additional uncertainty. For our drilling case, this is illustrated in the figure below:

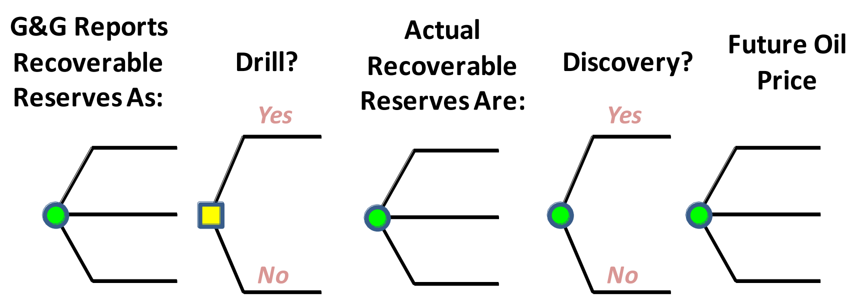

Figure 5.12. The tree with imperfect information

This conceptual tree shows that the actual uncertainty that will be fully resolved prior to the time the driller needs to take his/her decision comes from the G&G group reporting their best guess as to the recoverable reserves. Actual recoverable reserves will of course not be resolved until after the drill decision.

Prior to deciding whether to engage the G&G group, the driller does not know what they will report; but he/she will know their results (if the decision is made to consult them) prior to making the drill decision.

Setting up this tree requires making some additional probability assessments from experts (in reality, these can either be the driller's own experts assessing the reliability of a G&G company offering their services or the driller's own experts assessing their own proposed G&G program). But the probabilities to be assessed are illustrated below (ignore for the moment that the ordering of nodes in Figure 5.13 is different from the ordering in Figure 5.12—this is explained shortly):

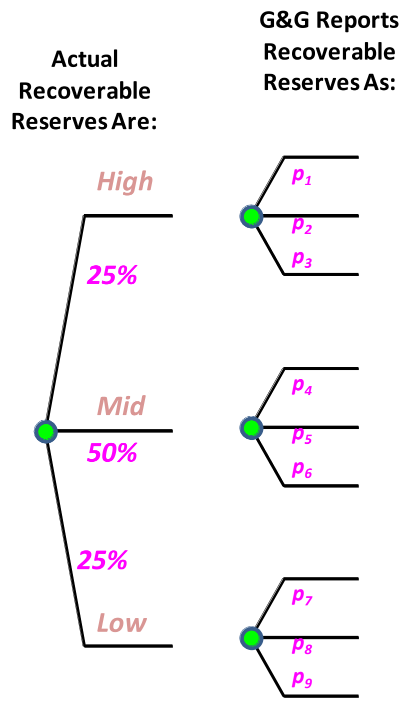

Figure 5.13. The probability assessment for evaluating imperfect information

What the G&G expert is asked to do is answer the following questions:

1. If I told you that the actual recoverable reserves were (for certain) *high*, what is the probability that the G&G study will report them as:
    a. High?
    b. Mid-range?
    c. Low?
2. If I told you that the actual recoverable reserves were (for certain) *mid range*, what is the probability that the G&G study will report them as:
    a. High?
    b. Mid-range?

   c. Low?
3. If I told you that the actual recoverable reserves were (for certain) *low*, what is the probability that the G&G study will report them as:
   a. High?
   b. Mid-range?
   c. Low?

With these nine probabilities, we can proceed with the analysis. Geologists and geophysicists will generally be able to answer these questions when presented this way. In fact, such assessments frequently lead to these individuals gaining greater insight into the problem.

Note that these two uncertainties are not independent. The probability p_1 will in general be different from the probabilities p_4 and p_7, even though they occupy the same "high" branch on the recoverable reserves node. This is how dependencies between uncertainties are treated in DRA.

But there is a mechanical problem with this approach. You will note that the ordering of the nodes in this assessment is different from the ordering shown as required in Figure 5.12. The probabilities must be assessed this way because the expert cannot reasonably assess them in the order we need them to solve the tree. For instance, if we were to ask the expert to give us his/her probability that actual reserves will prove to be high *given* that the G&G says they will be, the expert's assessment will be confounded by trying to think about, not just the reliability of the G&G, but also the likelihood that the reserves actually *are* high.

Fortunately, probability theory gives us a way to take the expert's assessment, given this ordering, and reverse the ordering to match what we need for the analysis. For those of you who are interested, this methodology is explained in Box 2. You will probably be pleased to know that the interactive model you will be working with this week does all this for you.

Geologists and geophysicists have long fought frustrating battles within their companies to make the case for the value they add. Especially in today's metric-dominated corporate world, this frequently is a challenge: G&G, the saying goes, has never produced a drop of oil; only drilling produces oil. "While we can certainly quantify the (often large) dollars you spend, how do we quantify the dollars you deliver us?" …is an argument commonly heard by the G&G community.

The value of imperfect information approach gives geologists and geophysicists ammunition for responding to this question. G&G delivers, not oil, but information. But the value of that information can be quantified.

Box 2: "Flipping" the tree

The procedure for reversing the ordering of two dependent nodes is based on something called Bayes' Theorem.

This can be illustrated mathematically (kind of inscrutable), or by using a simple tree. The simple tree used here has only two branches on each of the nodes, but it will give you the idea.

Let's say the uncertainty we are trying to resolve is whether or not it will rain tomorrow. We have available to us a rain detector, but this rain detector is not completely reliable. Suppose we have available to us an expert on rain detectors, who has assessed for us the reliability of this rain detector. Suppose also we have an expert meteorologist. Combining their expertise we determine the following:

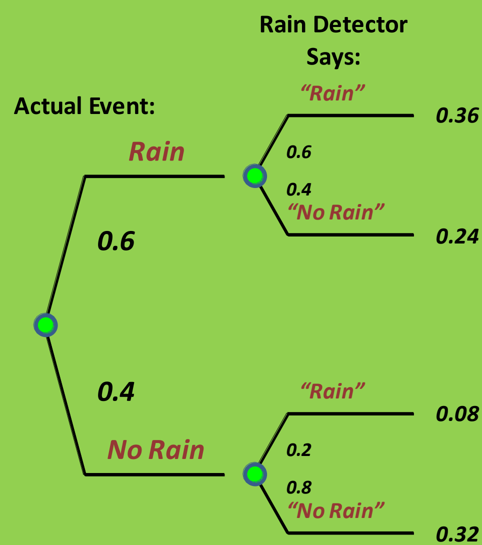

Figure 5.14. The original probability assessment

The meteorologist says the likelihood it will rain is 60%. The rain detector expert says that, if the rain detector calls for rain, there is a 60% chance it will prove correct. If the rain detector predicts no rain, the rain detector expert says there is an 80% chance it will be correct. (Note that in this example, the rain detector can deliver both "false positives" and "false negatives.") The numbers on the right show the "joint" probabilities of each of the outcomes. For example, the joint probability that it will both rain and the rain detector will predict rain is 0.6 x 0.6 = 0.36.

Now we introduce the tree in the "flipped" form in which we need it:

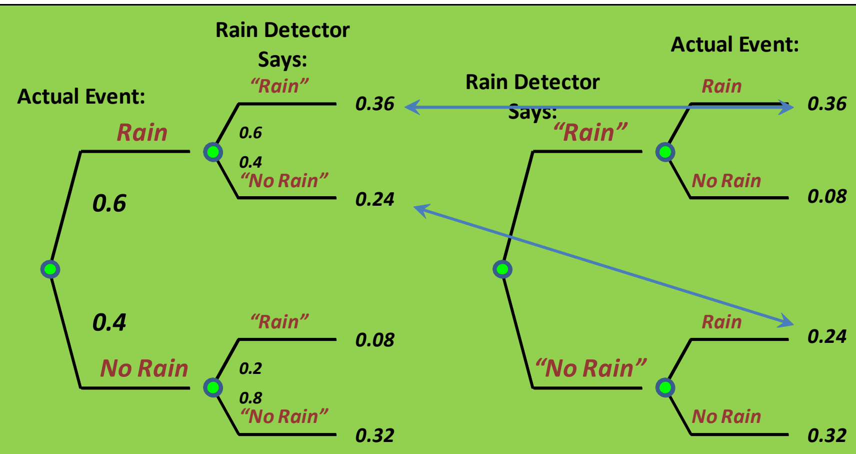

Figure 5.15. "Flipping" the tree, step 1

The first thing we do is match the joint probabilities in the "flipped" tree to those of the original tree. For example, the combined event that it actually rains but the rain detector calls for no rain has a probability of 24% (left tree); this joint probability we now associate with the same combination of events in the right tree that is the second outcome from the bottom.

Next we calculate for the right tree the overall probability that the rain detector will predict rain. This is simply the sum of the probabilities showing on the right for the "rain" event—if the rain detector calls for rain, it will either rain or not, but the probability it will do so is the probability it will actually rain plus the probability it will not. This calculation is shown below:

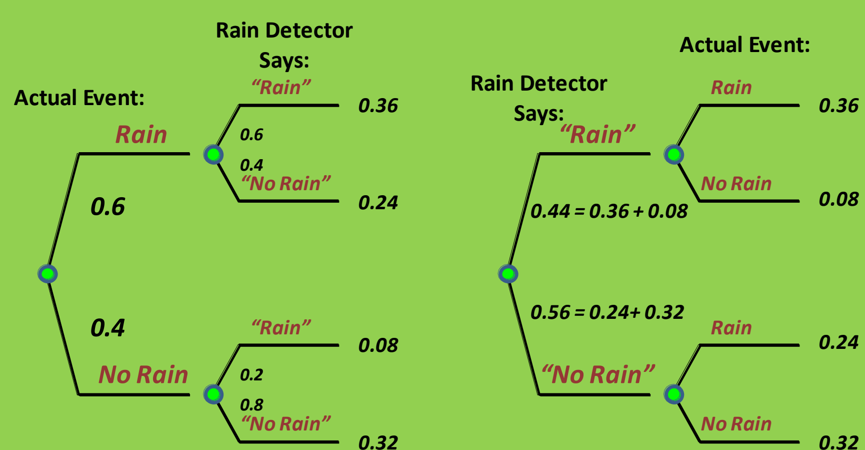

Figure 5.16. "Flipping" the tree, step 2

The final step is to figure out the correct probabilities for the Actual Event node. Since we know that any joint probability is obtained by multiplying them together, we can calculate these by division. For instance, for the top branch we know that the joint value of 0.36 must be the multiplication of 0.44 and whatever probability belongs on the top branch of the Actual Event. Accordingly, the correct probability is 0.36/0.44 = 0.818. Doing this for all four branches of the actual event delivers the final “flipped” tree shown below:

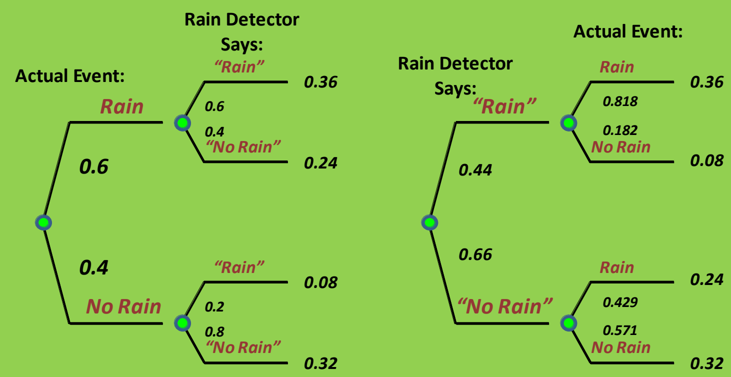

Figure 5.17. “Flipping” the tree, step 3

The probabilities showing on the branches of this final tree are *completely* consistent with the probabilities showing in Figure 5.14. If the experts believe the probability assessments they gave in Figure 5.14, they *must* believe these probabilities, if they are rational and consistent.

Lest you be concerned that you will be called upon to do such calculations, take comfort in the fact that the interactive model you will be working with does this for you. If you are interested in details, look in cells AV41:BM 68 of the “Value of Imperfect Info” tab.

Week 5 summary

This week you explored basic geology concepts and the role of geophysics and especially seismic in providing imperfect information to inform drilling decisions.

You also investigated the principles and tools of decision and risk analysis, including the use of the probability wheel to assess uncertainties, the expected value concept, the concept of downside risk, and the important concept of value of information, including the value of imperfect information.

With this background in hand, the interactive model for this week allows you to undertake a risk analysis of a drilling decision.