Mineral and Petroleum Economics

Overview

There are few topics more central to understanding oil and mining than the economics which underlie these industries.

This Topic Overview reviews the specific features of the economics of oil and mining, looks at "money in, money out" – how the economics of a gold mine or oil field, or any other commodity, stack up in terms of simple cash flows. A third section then layers on these real world complexities. How do operators arrive at end estimates of profitability? What measurements do they use, and how do they handle constantly changing operating and market conditions? How do governments do the same?

It includes the following chapters:

1. Key Economic Features of Mining and Petroleum Projects
2. Money In and Out: Project Cash Flows
3. Profitability

1. Key economic features of mining and petroleum projects

There are a number of features which make mining and oil and gas different from other economic activities. The following is a list of characteristics which influence the investment of trillions of dollars in the extractive industries.

Projects exploit a finite resource

Sub-soil resources are "one offs". There is a certain amount of oil, or gold, or copper, under the ground and once you dig it out, refine it and sell it, it is gone. It is, in other words, depletable. This makes it fundamentally different from, say, agriculture, where the same crops can be grown year after year, or even manufacturing where, even if it relies on these same raw materials as inputs, most of the business model is in the value added of industrial processes that are not depletable in the same way. Considered economically, that puts any resource owner in a position of producing out natural "capital", rather than earning revenues with replaceable inputs like labour. For a company, that then brings considerations of how to control the rate of production to maximize profit. The more you produce, the less you have left – unless you find new fields.

For governments it poses the challenge of how to manage what is called "intergenerational justice". Are you cashing in today at the expense of your children tomorrow? And then again, this important theoretical concept – that natural resources are depletable – needs to be tempered by what actually happens under market conditions. Geological depletion suggests an absolute amount of raw materials on the planet. But technology constantly improves, especially when driven by the incentive of higher prices and greater profits. Oil can be sucked out from places where it could not even be detected a generation ago. Miners who 30 years ago would have had to dig a deep vertical shaft to reach a few limited veins of gold these days can bring in diggers to remove – economically – tens of millions of tonnes of "overburden" and create a massive open mine, reaching far more of the resource.

Long, costly exploration periods

While the lore and popular image of extractive industries is steeped in booms and get-rich-quick movements, from the Klondike in 1849 to gold rushes across Africa today, the reality of industrial level production is somewhat different. Thorough exploration often takes years and costs millions, sometimes billions, of dollars before a prospect is declared commercial and real development begins. A single deep sea exploratory oil well can cost $100 million. And always with the risk of total failure. In economic terms, this creates what is sometimes called a "hurdle" before a decision to invest is taken. A prospect needs not just to contain copper, but *enough* copper, extractable at low *enough* cost, to be declared commercial.

Large up-front investments

Exploration costs pale in comparison to the money needed at the next stage. Developing a huge project takes billions of dollars, and the trend in both oil and mining in the last couple of decades has been towards more and more "megaprojects" – single concession areas which require more than a billion dollars up front investment. Moreover the structure of extractives is generally that that investment is needed "up front", at the start of a project, whereas production, revenues and profit lie years, sometimes even decades, down the track. For investors, time is money in a very real sense. The potential returns need to be that much greater the further in the future they lie, compared to investment outlay now. This relationship is factored into all investor metrics, as we shall see in the later sections.

Significant geological, political, environmental risks

Extractive industries are controversial because, when so many different things *can* go wrong, there will always be projects where they do go wrong. Analysts often talk in terms of "below ground" risk and "above ground" risk. Below ground risk means, even after all that exploration, there could be less resource, or it is more difficult to get at, than anticipated. In the past few years there have been several spectacular cases of such misjudgement. Rio Tinto had to write off billions of dollars, in a case which led to the resignation of its then CEO, when coal reserves in Mozambique turned out to be much lower grade than expected. The British oil company Genel wrote off half a billion dollars in a single oil field in northern Iraq in 2015 after it was forced to conclude it had overestimated the amount of oil and gone about drilling for it in the wrong way (reservoir management). Above ground risk is everything else, ranging from a coup d'etat – or an election – changing the government who signed a deal, and with it the political will to keep to the original terms, changes in law, lack of manpower or changes in global markets.

Sophisticated management and specialized technology

The constant evolution of technology may make more possible. But it also massively increases complexity. More things can go wrong. For example, gas can be brought to market globally if it is liquefied. But that means cooling it to minus 180 degrees and keeping it at that temperature on huge tankers running trans-continental voyages. Lower and lower grades of gold can be mined out – many operations now mine out ores with less than a gram of gold per tonne of ore – but only if all the many processes needed work well by themselves and interact with each other smoothly.

Prices (mostly) set on international markets; price volatility

Costs may be hard to predict accurately. Prices are impossible. The last few years have given ample demonstration of massive volatility in commodity prices, as what was touted as a supercycle boom in the years after 2000 culminated in a mid-2014 price crash of historic proportions. Moreover, experts do not agree on future forecasting, or even the reasons for price movements in the past.

High costs of abandonment

Large extractive projects not only cost a lot of money to set up. They cost another lot of money to leave in any kind of responsible way. "Abandonment", or decommissioning a mine or oil field, can take years and cost millions, and those costs have to be factored into investment decisions from the beginning.

Significant environmental impact & risks

The Deepwater Horizon blow out in the Gulf of Mexico in mid-2010 has so far cost BP, the operating company, at least $20 billion. The Paris climate change agreement of 2015 aims to factor in the impact of fossil fuels on global warming. But costs under business-as-usual scenarios are incalculable. In mining, too, tailings piles can subside, toxic chemicals used for processing can leak, dams can collapse, or the project can end up competing for water in the long term with the growing needs of communities in the region.

High community impact

The phrase "Resource Curse" is contentious. Experts disagree about what causes it, or even if it exists at all. But its mere presence points to one indisputable fact: the powerful impact extractive industries have. This impact can make itself felt at national level, as with the so-called "petrostates", and even more intensively at the level of local communities. Single projects can trigger, or at least play into, tensions within communities, between them and outsiders brought in to operate mines and fields. They can be a region's largest employer and most visible connection to the global economy. They can run a region's major infrastructure and use huge amounts of its other resources such as land and water. It is little wonder that companies spend a lot of time and money in "corporate social responsibility", trying to win support at all levels of the societies they operate in. And, conversely, the last few years have seen numerous disputes between companies and communities, in Zambia, Peru, Indonesia, Thailand, South Africa, Greece, Afghanistan, the United States, the United Kingdom, Nigeria, and many other countries.

There are three general observations to be made about this list.

First, the fact that these are specific characteristics of the oil and mining industries is not supposed to mean that each one is unique to those industries. Many other economic sectors can be characterised by one or more of these attributes. Commercial real estate development, for example, can require comparable amounts of up-front investment. IT deals with similar technical complexity. Agriculture faces great volatility on world markets. It is the combination of all these elements together, and their intensity, which is unique to extractive industries. Like a fingerprint, as it were.

Second, some of these components are not directly financial, or at least quantifiable. But they all have financial implications for the management of extractives projects. Labour unrest, a change of government, or community tensions can all lead to stoppages, which are directly mapable as lost production and revenues. All of these variables are factored into the agreements which investors strike with governments, and integrated into contract terms and liabilities. One way or another, each of these factors ends up as a number, or part of a number.

Lastly, although at a media or policy level, there is much talk of the oil or mining industry, either globally or at a national level, in terms of economics it is the project which predominates and is the natural unit of analysis. A single "project" could encompass hundreds of wells, or seams, and combine many differently acquired concessions or license areas. But, generally speaking, it has its own business model, physical infrastructure, tax treatment, incorporation structures and finance mechanisms. Large multinational producers may operate scores of major projects in as many countries. But they are constantly seeking to match their forecast of global demand against a supply driven by the specifics, and resource bases, of all of their projects.

And each project is highly specific. The geology of one field may be significantly easier or harder than its neighbour. The same geology could spread across two or more jurisdictions with radically different "above ground" risk. Consider, for example, that the Eagleford shale formation of Texas extends hundreds of kilometres beyond the US border into Mexico. On one side of the border there are thousands of wells. On the other side, none.

The global portfolio of one company can be radically different to another, based on different competitive advantages, or indeed different views of the market. Shale gas in the United States was pioneered in the 2000s by then unheard of companies like Devvon Energy and Chesapeake Resources because the existing supermajors like ExxonMobil and Chevron initially saw less margin in it *for them*. Vertically integrated oil companies like BP and Total may be more likely to retain operations at all stages in the value chain, from the well head through to the petrol station, even through downturns in one or another stage in that chain, such as refining, than other companies which have never had that history or business model. Which is radically different to explorer companies, which routinely “flip” assets once they have made major discoveries, and stay resolutely out of major production themselves. Or Chinese companies owned or strongly influenced by the Chinese state, which will integrate factors such as security of access to resources into their thinking in a way a Western company never will. Newmont Mining can explain its decision (in mid-2016) to sell a mine in Indonesia to investors as a bid to concentrate purely on the production of gold, in this case divesting from major copper production, in the same market that its competitor Rio Tinto doubles down on copper production by moving ahead with an extended investment in Mongolia.

Each of these global portfolios is different. But each of them carries the same relationship to the project level. In supply terms, the portfolio is simply the sum of the parts - the projects. The project is the natural unit of analysis of economics.

2. Money in and out: project cash flows

The economics of extractives projects are complex, and have evolved around structures which are a combination of market economics, international best practise, and national law over a long period of time. But they can be broken down into three successive stages: cash flow, or the underlying “raw” economics of how much investment is needed to produce a resource selling for how much; “fiscal”, or what rules a host government uses to levy royalties, taxes and other revenue streams from the cash flows; and project finance, or how fiscalised revenue flows are affected by such features as incorporation structures, debt financing and intra-business group arrangements of finance and transfer pricing.

This section deals with cash flows, which itself breaks down into three categories: production, price, and costs.

Production

Once the quantity of the commodity has been estimated (see the previous section), a company needs to decide its production profile. Even at the highest level of proven reserves, this cannot be arbitrary. That is to say, a company will rarely be able to treat the resource as a liquid asset, capable of any configuration of production. There are two guiding factors which will determine how much of the resource is produced when: geology, and market conditions.

Many oil fields follow a pattern in which the first year or two of production are "ramping up", a "production plateau" is then reached which could last several years, and then there is often a long, and slow decline in production. And this is a pattern which is mainly determined by geology. When the oil field "comes online", pressure within the reservoir is at its greatest. It may take a year or two to complete the infrastructure of wells and pipelines which allow the company to take advantage of this pressure, hence the ramp up period. Once it has, the field will move to plateau production. Generally, though, the natural pressure in the reservoir starts to weaken quite quickly – as oil or gas gets produced out, the pressure is released and the remaining oil and gas flows through the reservoir more slowly, resulting in lower production levels.

Classic production profile figure

It is important to emphasise, though, that such a *production profile* is purely schematic. Each field's production history and future profile is, in the end, unique. But also, different classes of asset are likely to have different generic profiles. A field producing oil or gas through fracking techniques, for example, is likely to have a shorter lifespan than its *conventional* counterpart, with a shorter plateau and depletion rates, or the decline in the tail, which is steeper.

The company has options to influence this process, within limits. It can choose to produce all, or some part of the oil rising through natural pressure. It could decide to increase the pressure by various enhanced recovery techniques, such as injecting the field with gas, water or dioxide, but these all require extra capital investment. Such processes of *reservoir management* have a number of aims: to maximise production of the field over the long term, or to increase production when the market is booming and there is more profit to be made, or to find a sweet spot of maximum production against a given level of capital investment.

There are similar considerations in mining. A resource estimate, as we have seen, will generally include different quantities of reserves at different levels of confidence, and a big factor in determining the degrees of confidence is market price. It is common, therefore, for an operation to plan production in a way which seeks a path to the highest grade ores, but that will also need to be gauged against the specific geology of that mine. A gold mine, for example, will have several different future production scenarios, each mapped against a particular "cut-off grade" of the ore. More ambitious scenarios, involving lower grade ores, will only be brought into play if market prices can support them. But also, just as with oil and gas, a mine may have incremental development options: to progress from an open cast mine to an underground shaft, for example, or open up a new operation nearby, or to process tailings which may contain low levels of resource that have been left until now.

The big picture, then, is that it is important to see levels of production within the life of a project as a dynamic variable, capable of being influenced by many factors on an ongoing basis.

Price and valuation

Much of the material in this sub-section is dealt with in the separate paper on Oil and Mineral Trading. What is offered here is a potted summary.

The single biggest concept to grasp in terms of commodity values is that it is not as simple as price, for two reasons. First, relatively few commodities are sold on a like-for-like basis, like potatoes, say, or wheat, but instead valued indirectly against "benchmark" grades. Second, significant proportions of commodities are never sold on an open market at all. They are traded between business groups in related party transactions, or as part of "term contracts" lasting many years. As an indicator of the scale of this non-market trading, the OECD has estimated that over half all international trade that takes place every day is between related parties, not "arms length". And the proportion might be on the rise. The US government has estimated that in 2014 some 42% of oil imports were between related parties, compared to just 23% in 2002.

This means that in terms of economics, we are really talking about "valuation", not just price.

Benchmarks, spot markets and liquidity

Many commodities that are sold in open market are sold against benchmark prices, rather than directly. Of the thousands of grades of crude oil that exist, for example, most are sold against one of less than a dozen benchmark grades, such as Brent, West Texas Intermediate (WTI), Dubai Light, Urals Heavy Blend, and so on. This means they will not be quoted as a direct price themselves, but instead as "*Brent plus or minus*", so to speak. One grade could be sold at Brent plus 3 dollars, another one at Brent minus 2 dollars. As long as this *differential* holds, the price of the crude will go up and down in sync with Brent itself.

Benchmarks evolved in different markets to serve as reference prices against which other commodities can be priced. Each is likely to be specific to the markets on which it is used as a benchmark. So in the case of crude oil, Brent functions as a benchmark in the UK and Europe, WTI in the United States, Dubai Light in the Middle East, and so on. With iron ore, the London Metal Exchange offers one benchmark price, but because Australia and Brazil are major producers, they offer additional benchmark prices.

Differentials exist because of differences in the underlying quality of the commodity, and market related factors, like how close it is physically to the market, and what the specific market demand is at any given moment for that grade of commodity. Crude oil, for example, is not a useable end product. It has to be refined to make petrol, kerosene, tar, petrochemicals and a host of other products. The range of end products producible from the crude oil going in is called "*the slate*", and differs widely from one grade of crude to the next. A barrel of *light* and *sweet* crude may generate a lot more high value end products like petrol and diesel, than another barrel of crude which is *heavy* and *sour* (containing high amounts of sulphur).

And different commodities also vary widely by how much of a spot market exists.

Oil is a "thick" market – there are thousands of buyers and sellers, and benchmark prices are relatively well established. So is gold, where the end product of one Troy ounce of gold is directly comparable around the world. But natural gas, by comparison, has no single global market price because it is so difficult to transport. At any one time, prices for a million British Thermal Units (BTU), the main unit of sale, could vary by as much as 600% between North America, Europe and east Asia. And there are other commodities, such as uranium or mineral sands, where there are only a small number of producers, such that even if quantities are sold directly in port markets, they are relatively few and far between, and mean a spot price cannot be reliably established.

It is commonly assumed that open, competitive markets are the norm for international trade. But not only is this not entirely true today, as we have seen, because of the prevalence of intra-group transactions, spot markets are a relatively recent invention altogether. The oil industry was historically dominated by the “Seven Sisters”, the British, French and American supermajors who were vertically integrated in industry terms. As long as the controlled upstream production, most oil was being transferred from their own upstream arms to refining, petrochemicals and retail distribution arms within the same groups, and relatively little was being traded on an open market. It was only when economic nationalism created much more active state involvement among newly independent producing countries, in the 1960s and 1970s, which included direct ownership of a portion of the oil, that spot markets developed, and trading houses such as Phibro, Glencore and Vitol became actors with significant impact on global markets.

Contractual terms relating to price and valuation

Since the price fetched is such a critical question, extractives contracts often contain previsions about how valuation is to be handled. There are several critical features:

Point of valuation: Transporting a commodity from where it is produced to an end market is often expensive. The location at which a commodity is sold along the value chain therefore influences the price. The further from market, the lower the price.

Commodity basket: A commodity can be value relative to a single benchmark, or a “basket” of several. In such cases, the basket is likely to be weighted, with each component of it accounting for a percentage of the basis of comparison.

Currency and date: Valuation formulae will also specify a given date of sale of any benchmark and currency, to guard against day to day fluctuations in the commodities and foreign exchange markets. Any of the transactions in the value chain which relate to revenue flows that will eventually be taxed in local currencies are likely to specify which rates of conversion were used.

Market volatility

The 21st century BCE has already seen massive volatility in market prices. From 2000 up until 2014, oil and minerals saw a boom in both prices which led some analysts to talk in terms of a supercycle, and others to suggest that commodities were now under such pressure from rising demand, itself driven by population growth and a rise in prosperity, that the higher price levels now achieved represented new equilibria – a long-term step change in market conditions. And then, beginning in 2012-13 with gold but accelerating in 2014 with oil, prices crashed. This is simply the latest round of volatility, which is a constant feature of commodity markets as long as they have existed.

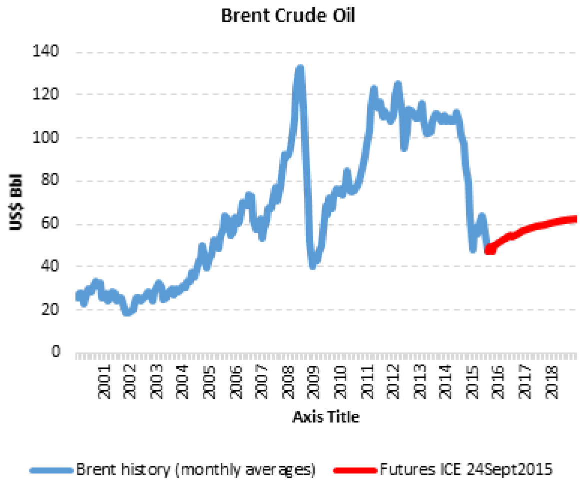

One important thing to understand is that the relationship of price to demand is *nonlinear* – that is to say, prices rise and fall by percentages which are much greater than the rise and fall in underlying demand and supply. The first decade of the 21st century, for example, saw the emergence of China as a "swing consumer" of many commodities – a large economy with such high economic growth that it could drive pricing in world markets like iron ore, aluminium, tungsten, and coal. In the case of oil and gas, the emergence of a new segment of the industry in the United States based on fracking was also considered a key factor in the crash of 2014, and a response driven particularly by Saudi Arabia to try and drive the new shale producers out by flooding the market. But in both cases shifts of demand and supply were more modest than the price swings. Shale production in the US started in 2005 and focused on gas. In terms of oil, total production added may have been, cumulatively, four to five million barrels a day. Certainly significant for the United States, but since global consumption of oil has been running at about 90 million barrels a day through this period, it represents a rise in supply of perhaps 6%. The other factor often cited was in the slowing down of Asian economic growth, which led to a drop in demand. But this drop, in terms of global demand, might have represented perhaps 1-2%. The cumulative fluctuation of both demand and supply, then, was less than ten percent, and yet prices were halved, from about $100 per barrel at the start of 2014 to below $50 per barrel by the end of 2015.

Financialisation and speculation

One feature of markets in the last couple of decades is their increasing *financialisation* – the ratio of investment in financial instruments that are tied to underlying commodity prices in some way, but are not sales of physical oil, or gold, or iron ore. Significant use of such instruments first evolved in the1980s, when major consumers of physical commodities such as airlines needing large amounts of fuel, or manufacturers needing significant quantities of, say, steel, began to buy futures and options in commodities, to allow financial planning and protect themselves against market volatility. If an airline, for example, knew that it would need a million gallons of jet fuel in six months time, it could pay for an option to buy it at a fixed price. If the market price stayed below that then they would not exercise the option and it would expire. But if prices soared, they would be protected against those rises. Producers too might want to insure themselves against price drops, in what became known as "hedging".

By the early 1990s there was more "paper oil" being traded than physical oil. But the relation between the two was altered dramatically in the next 20 years, until by 2010 trade in "paper oil" was an estimated 20 to 30 times higher than the physical commodity itself. Even within physical markets a barrel of oil might be traded ten times between being loaded on a ship in, say, Nigeria, and offloaded at Rotterdam. Markets also created commodity-tracking financial instruments which rose or fell directly in response to the price of benchmarks, or a basket of different commodities, without necessarily being directly invested in an underlying consignment of the commodity. Such financialisation has led to a debate about the extent to which speculation has fuelled volatility on global markets. Experts are divided. Some say the speculation has played an extreme and negative role, while others maintain that speculation, on balance, is positive, and that markets continue to be governed by "fundamentals", the supply and demand curves for the commodities themselves.

Costs

Costs in extractive industries tend to fall into two categories: capital expenditure (or "capex") which is required to develop and maintain an industrial project, and operating expenditure ("opex") to run it day to day. These are significant because they have different treatment in the tax codes of most countries – in other words the same amount of overall investment in two different projects could lead to different rates of profitability, depending on the mix of categories of spending in the total.

Capital expenditure

The large up front expenditures usually needed to develop a mine or an oilfield, one of the distinguishing characteristics of the industries mentioned in the introduction, are capex – building the mine itself, sinking wells, building access roads, pipelines, on-site processing facilities and so on. These can often run into billions of dollars. The Gorgon liquefied natural gas project in Australia, for example, was completed in 2015 at a cost of over $40 billion. Another project under consideration, the Libra oil field off the Brazilian coast, could involve up to $90 billion of capex.

Mining projects are usually not quite as large. Nevertheless, in 2016 Rio Tinto went ahead with an expansion to the Oyu Tolgoi copper and gold mine in Mongolia scheduled at $6 billion, and the Simandou iron ore mine in Guinea, whose estimated capex could exceed $20 billion, since it involves building a new deepwater port.

Capex includes spending on exploration, initial development, and all major infrastructural spend. It also includes abandonment costs, although special provision is often made for these since a major part of these are incurred after the mine has stopped producing, and cannot therefore be paid for out of current revenues.

The critical thing about capital expenditure is that once it is done it cannot be undone. This sometimes gives rise to what economists have called the "time inconsistency" problem. Companies are keen to sign agreements that last for the lifetime of a project before risking large amounts of capital in developing a project. At that time, with no production or revenues in sight, governments are also keen to make the investment happen.

But once the project goes ahead, and in particular the capex has been sunk, that dynamic can change. The government may see considerable production, and revenues, materialising but without major tax benefits – because the project is in a "cost recovery" stage where most revenues are going towards repaying the large initial capex. This can lead to a government changing its view about what is reasonable, and demanding changes in the contract, knowing that the investor cannot now walk away from the project without losing all its capex.

Although cos recovery happens early in the life of a project, capex is *depreciated* in accounting terms. That is to say, its impact on cost recovery or taxation is staggered over a few years, sliced into tranches that go into the mix every year, rather than being "expensed", or immediately deductible.

Operating expenditure

Operating expenditures include salaries, day to day running costs of the mine or oilfield and its associated infrastructure such as access roads or pipelines, some categories of transportation, and overheads such as management and administration. By definition, operating expenditure must come later in the life cycle of a project since they can only begin once production starts, which in megaprojects is often five or more years after the start of the project. If capex is often the leading cost factor in determining whether a decision will be made to invest, opex is often key in deciding whether a mine or oilfield continues in operation or not under changing market conditions. And once there is production, operating costs are immediately expensable from turnover.

Differences between oil and mining cost structures

Although a typical project has a lifecycle of the kind described, with high and early capex followed by opex, there are some structural differences in these schematic views between the oil and mining industries. While "conventional" oil loads all capex up front to the start of a project, mining has a category of capex called "sustaining capital" which remains a factor throughout the life of a project. Sustaining capital is used for the maintenance of vehicle fleets, and for replacing any part of the wear and tear caused on equipment caused by mining on an industrial scale. Significant sustaining capital costs will changes the cost and risk profile of a whole project, since capex is now needed continuously.

In hydrocarbons, unconventional projects such as fracking or tar sands have a different cost profile too. Tar sands is essentially a form of mining and so entails what are effectively sustaining capital costs. In fracking there is a need for continuous capex, not so much for maintenance purposes as due to the fact that each well produces less oil or than a conventional field and has a much shorter life. A shale project therefore has more wells, which are drilled continuously over most of the life of a project, in contrast to conventional oil where most drilling work is done at the start.

The rise of costs, and the rise of the mega-project

Industrial scale extractives projects have always been relatively large. But they have got bigger and bigger in recent decades. In 2015 Ernst and Young published a review of some 365 “megaprojects” – those requiring over a billion dollars of investment – and concluded that this scale of project had become a new norm in the oil industry as the end of “easy oil” forced companies to look for new resources in more and more challenging places with more and more sophisticated technology: the Arctic, the deep sea, coal bed methane, liquefied natural gas (LNG).

They also found that as projects got bigger, so did the likelihood that they would overrun on costs, and time to production. The oil industry spent about $100 billion a year on development in 2000, but this had risen to over $600 billion a year by 2010. The Ernst and Young report suggested this was likely to rise again to an average of over $900 billion a year in the coming years. To some extent development costs are driven by prices. When there is a boom market, producers have more incentive to develop new resources, and so costs naturally rise in those areas where there is limited supply, such as renting drill rigs and personnel. But the extent to which costs have risen seems to go beyond simple reaction to changes in demand in the market.

The study estimated that these megaprojects were more than 50% likely to entail significant cost overruns, more than 70% likely to incur time overruns, and that the cost overruns would average 57%.

It is the same story in mining. Another Ernst and Young study in the same year showed mining megaprojects likely to average cost overruns of 62%.

The phenomenon is not limited to the extractives sectors. There is a growing body of literature among economists about overruns in sectors such as public infrastructure, utilities, airports and Chinese construction. So far the common factor seems to be that since megaprojects take a long time, there is more chance unaccounted for externalities can impact execution, such as higher than anticipated general inflation rate, for example.

But the fact that the megaproject is accounting for more and more development and production in extractives needs to be seriously factored in by governments. All estimates of revenue flows depend on forward looking estimates, with the company often the sole source. If those estimates, influenced by the rise of the megaproject, are systemically optimistic, then estimates of how much money when need to be revised down.

Economic Rent

At the heart of extractives projects is the economic concept of “rent”. Literally hundreds of volumes have been written about economic rents in a literature which goes back to David Ricardo, one of the earliest economists, writing in the early 19th century.

At its simplest, though, rent could perhaps be described as the compound effect of three core principles: first, that producers in a given market have to sell at the same price but can have radically different cost bases, which are “blocky”, or discontinuous, rather than “smooth”, or continuous; second, that a large portion of those costs may have weak correlation to price, meaning that superprofits – or superlosses – could be created by changing market conditions without any significant change in the operation of the project itself; third, that such changes in the market can be *non-linear* – a relatively small change in demand or supply can drive a large change in price, and a still larger change in profitability.

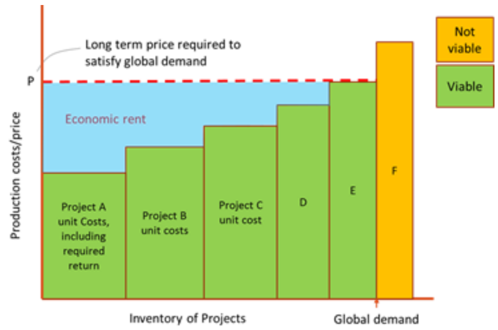

In the chart on the right, the cost bases of different projects producing a mythical commodity are represented on the horizontal axis, left to right. The amount it costs each project to produce is how high each green block extends up the vertical axis. So project A, on the left, has costs only half of project F, on the right hand side. The line P running horizontally is price, and the blue area in between the top of each block therefore represents profits – *which are therefore different for each project against the same market price*. This kind of chart as a whole is known as a *cost curve*.

If this seems abstract, a real world example can clarify. Imagine that the cost curve is for oil, project A on the left is Saudi Arabia, and project F on the right is drilling in the Arctic, or the Canadian tar sands. Even if the details are a closely guarded secret, we know that Saudi production costs remain the lowest in the world – with operating costs of perhaps somewhere between $5 and $10 per barrel. While a barrel of oil from the Canadian tar sands might cost $70 per barrel to produce. This means that at the price point P, Saudi Arabia can still make large profits out of its production whereas tar sand producers must either take losses – because their cost base (project F, in orange) is now higher than the market price – or stop production. If project F were not tar sands but shale production in the USA, the chart would reflect what many analysts have said was Saudi thinking in continuing large scale production ("flooding the market") while shale production grew, in the hope of forcing prices (the line P) down the chart to a point where the Saudi industry could still survive, while US shale producers could not.

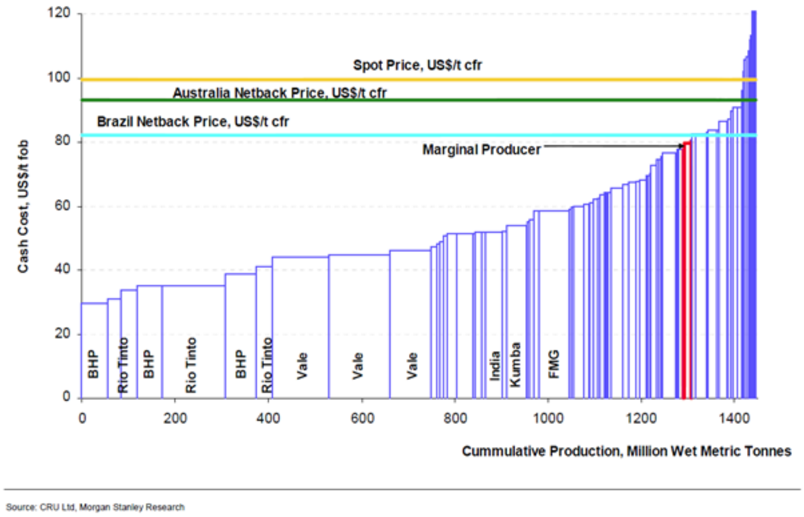

Real cost curves of actual commodity prices look messier than this neat example, but the same principle is in play. Take this cost curve of the iron ore industry.

These real projects show estimated cost curves for major companies such as BHP, Rio Tinto and Vale, in 2014. Again, the projects on the left show low cost projects which are earning higher profits than the projects on the right of the chart, against the benchmark prices represented by the three horizontal lines across the top – the then current spot price, and "netbacks", or prices set at the export point, in two of the major producers Australia and Brazil.

Interplay between local and global cost curves

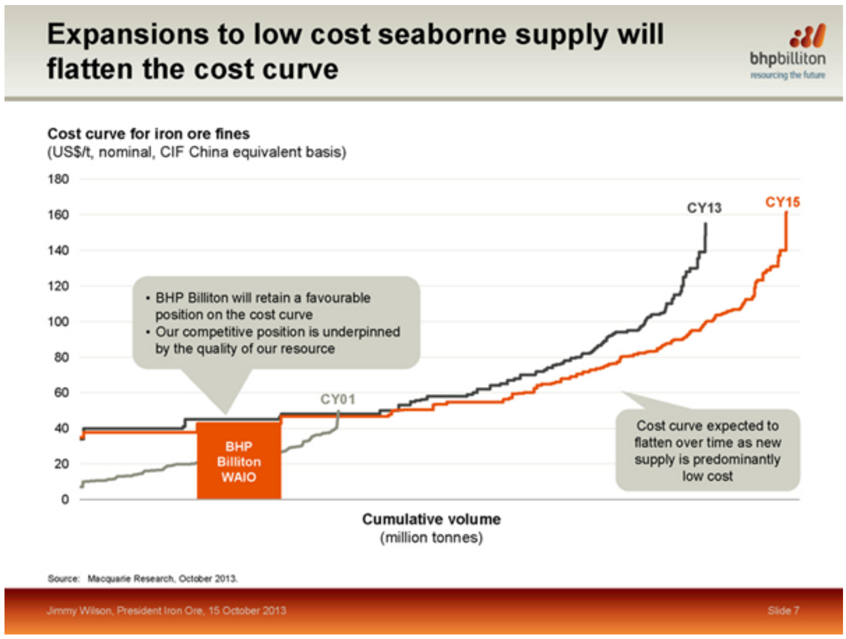

A company will always be looking for a project with a low cost base – on the left of the curve. There are even cases in commodities where a single big project could affect the whole of a global supply curve – the amount of a commodity available at a given price – as in the example below from BHP Billiton, where the company states to its investors that its new project, WAIO in Western Australia, will "flatten" the cost curve of iron ore globally. In other words, this one project will bring on enough new production of iron ore, at a low enough cost, that it will bring down prices in world markets. That might be bad news for BHP Billiton's competitors. But since the company's own production will have low costs, it will continue to make profits, even with lower prices, and its position will be strengthened vis a vis other producers.

The minimum necessary return on capital

Part of the theory of rent which is not readily intuitive to non-economists is that the cost base for each project includes what is defined as a minimum necessary return on capital. Going back to the theoretical cost curve above, then, this explains why project E would continue to produce. Its cost curve is right up against the price line, so that if it did not include any return on capital, it would be making no profit at all, and there would be no incentive to continue production.

This is intrinsic to the concept of economic rent. The blue area above the cost bases of each project are therefore not just profits, but profits above and beyond the minimum return to capital – in some cases they could be described as superprofits.

While the theory is relatively straightforward, its application is complex. First comes the question of how such a minimum return would be defined. Ricardo first posited this theory in terms of agricultural production in 19th century Scotland, when, it might be reasonable to suppose, land and food prices might be relatively stable. But fast forward to the 21st century, and these global industries, operating in a liberalised trade environment in financial markets that are digitised, and it can readily be seen that capital is so fluid across borders and economic sectors that such a "minimum return" is much harder to define even in theory. In practice, many financial models of extractives apply a simple rate of 8 to 10 percent per year, thereby defining "rent" (the blue area) as lying in profits above this.

The second practical difficulty with applying the theory of rent is that it assumes cost bases are easily knowable. But companies do not routinely publish costs of individual projects. These data are subject very much to the vagaries of compliance reporting and investor confidence seen in the section above on reserves and resources. Some companies publish some of the time. In both of the real cost curves included already, the cost bases are estimated, or imputed, by investment analysts or financial publishers. Such estimates can be informed and skilful, but there is no getting away from the fact that they are in most cases estimates, not facts. In the case of companies, such as BHP Billiton, they at least have firm knowledge of cost bases in *their own* projects to serve as a basis for estimating the cost bases of their competitors. Governments, however, may not have any actual operations and so rely totally on external databases, which are themselves estimates more often than not.

Who captures the rent?

The theory of rent sets the ground for a specific view of what taxation policy should be in the extractive industries, quite apart from other economic sectors. Since rent is defined as profits above and beyond what is needed for an investor to invest, in theory very high taxation of such rents does not deter investment. Since also, sub-soil resources in the vast majority of the world are publicly owned – that is to say, by the state, or at least managed by the state on behalf of the people – it follows that one objective of fiscal policy around extractives is to capture as much of the rent as possible.

It is important to understand that the specifics of extractives lead to these conclusions apart from any political or economic doctrine about the role of the state in the economy in general. The International Monetary Fund, for example, generally seen as an advocate of free market economics and global trade, consistently advocates that the state should try to capture as much of the economic rent of the oil and gas industries as possible.

The question, given the methodological issues raised above, is *how*. A separate paper in this series deals with the individual mechanisms of different kinds of royalty, taxes and so on which governments can deploy. There are a wide variety, each with strengths and limitations, acting on their own or as a suite, often called "the fiscal regime". It is worth noting in passing, though, that some of these mechanisms are explicitly targeted at capturing rent, such as a tax specific to the petroleum sector (in Nigeria, Norway, the United Kingdom and elsewhere), or a Resource Rent Tax (first introduced in Australia in the 1980s). Others are implemented in a way which implicitly targets rents, such as royalties which are staggered according to market price in the Ivory Coast. To relate it back to the cost curve, higher prices means the horizontal line of price goes higher, and the blue area increases. So a higher rate of royalty or tax corresponds to greater rents being earned.

3. Profitability

As noted, the stakes and risks in extractives projects are high. What guides investors in choosing projects are two factors: first, the underlying profitability of the project, both at the cash flow stage, and after taxation and the fiscal regime have been taken into account; second, how revenue flows stack up over time – on the principle that time is money.

The Investor Discount Rate

The way time is factored into financial analysis of a project is through the *discount rate*. Investors make projections of how much profit they will make in the future, and then analyse the times during the life of the project at which they expect these flows. Then they apply a discount rate to say that the further out in the future profits are, the more they should be *discounted* by comparison with the investment that needs to precede it to make it happen. The discount rate is expressed annually. For example, if a discount rate was 10%, an investor would calculate that for an investment of a million dollars, they would need to make a profit of over $1.1 billion by the end of a year before they started to generate returns that met their minimum return on capital. If it took two years for revenues to arrive, the minimum necessary return would be $1.21 billion – the same 10% applied in each of two successive years, with the added wrinkle that it is compounded, so that in the second year, the 10% would be applied not to the original billion dollars but to $1.1 billion, since the discount rate implies that is the total that should have been generated by then.

One way of looking at a discount rate is as a self-applied interest rate, applied to an investment that one is directly involved in. A normal investor might look at investment options, try and calculate future profits, and then subtract, or discount, general interest rates in the banking system. The thinking behind this is that any sum of money could earn such an interest rate safely in a bank, and that other investments therefore need to give higher returns. The "real" return of these other projects is therefore the extent to which they outperform the general interest rate.

That general interest rate would also form a part of a discount rate, since the investor would either be investing their own capital, which they could equally well deposit in a bank, or borrowing it, in which they case they have to pay interest on the loan. But it will also include all other factors which add up to what the opportunity cost is of investing in this project as opposed to another one. These could include: what are other mining or oil projects available to an investor – large multinational groups run internal competition processes between new projects bidding for management attention and shareholder capital; there could be expansion options within existing projects, or even investment not in increased production but in processes and technologies that would cut costs, and thereby increase profit margins on the same turnover. Conversely, what are the risks of one project relative to another? The same gold mine in terms of project economics might look very different in terms of “above ground” risk depending on whether it was located in a politically stable and relatively prosperous country or, by a poor and politically turbulent country. In other words what is the “country risk” of a given project?

All of these factors can be packaged into an indicator known as the Weighted Average Cost of Capital (WACC). Technically speaking, this analyses what an investment is costing in terms both of debt taken on to finance it, and the cost of any core capital a company is investing. In theory those calculations in turn should have factored in things like country risk, since they should affect both the cost of borrowing and the cost of capital.

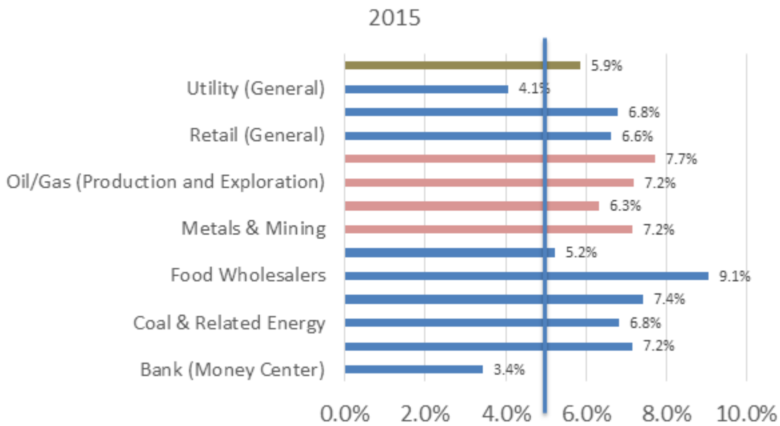

WACCs are compiled across industries and the chart below, compiled by Aswath Damodaran, a financial analyst and professor at New York University, shows how different sectors compare.

In reality there are many other wrinkles, such as whether these rates are compiled in terms of “real” interest rates (stripping out inflation) or “nominal” interest rates, and whether data are equally available from all parts of an industry, particularly oil and mining where so many projects take place in remote parts of the world. The WACC rates here are indicative of the way company discount rates are formed, and should not be taken to be definitive, and companies often keep the discount rates they are using in particular projects, and the way they were compiled, a closely guarded secret. Ultimately, discount rates, like rent, are a concept which it is relatively easy to conceptualise but often extremely complex to formulate and execute.

One consequence of factoring in not only much profit, but when that profit arrives in the life cycle of a project, is that analysis may yield results that are counterintuitive. For instance, one project may have lower overall returns in cash terms than another, but be a more attractive investment, either because the profits are calculated to arrive sooner, or because, for one reason or another the company is applying a lower discount rate.

Another key factor to bear in mind here is the longevity of extractives projects. Since a discount rate, like an interest rate, is expressed annually, and *compounded*, the effect of short-term costs can be dramatic in projects which measure decades. The Libra offshore project referenced above requires $90 billion of investment over the next few years. The size of the oil field is considered to be enormous, and the project could well last into the middle or even the latter half of the 21st century. But how much profit would the project need to be earning in the 2040s, or 2050s, to justify tens of billions of dollars of investment now? It is another area in which impact is non-linear. It is therefore no surprise that companies are extremely sensitive to costs incurred up front, and that fiscal tools such as signature bonuses, which occur right at the start of a project, are unpopular.

Indicators of Profitability

The two core indicators used by investors to measure profitability, the Net Present Value (NPV) and the Internal Rate of Return (IRR) both relate to the discount rate. We examine each of them in turn.

Net Present Value

The Net Present Value is exactly what it says: a calculation about an investment, or series of investment in the future, which yield revenue streams, also in the future, expressed as a value of money right now today.

Perhaps the easiest way to think about it is as working like interest rates do, but in reverse. Say an investor puts $100 into a bank account where it earns 10% interest a year. After one year, the investor would have $110. After two years, the investor would have $121. Notice $121 not $120. This is because the interest *compounds* – interest earned one year gets added into the base on which interest is calculated the next. So in the first year, 10% of $100 is $10. But that interest from the first year is added into the base for the second year, the 10% interest rate is applied on $110 – not the original $100. This makes interest of $11, not $10, in the second year and so we arrive at $121 returns after two years.

Given these facts, we can say that $100 today will be worth $121 in two years time if invested in this way.

All the NPV does is to reverse the direction of this calculation. Instead of using an interest rate to calculate the *future* value of a current investment, the investor uses a discount rate to calculate the *present* value of future investments. If in this case the discount rate was 10%, in other words the investor needed a 10% return on capital each year, then the same conditions which lead to a value of $121 in two years time lead to a Net Present Value of exactly 0 (once the original $100 investment has also been subtracted). That $121 withdrawn two years from now, in other words, is the same value as $100 now – because time is money.

Obviously everything here does then depend on the discount rate used. With the same initial investment and subsequent returns, the NPV turns positive with a discount rate of anything less than 10%. For example, the same return of $121 in two years would yield an NPV of $9.75 at a discount rate of 5%. If, on the other hand, the discount rate was set higher, at 12%, (because the investment was perceived to be risky), then the NPV would be -$3.75. And, it is important to note, that negative NPV has been achieved with the same results in *nominal* terms: the investor is still estimating investing $100 now and receiving $121 in two years time.

No investor will invest against a negative NPV, as it implies they could earn more doing something else with the same amount of money.

Another key feature of the NPV is that it provides apples-for-apples across geography, time, and commodity. A company could calculate NPV against an iron ore prospect in West Africa to last 35 years, and a gold mine in Indonesia lasting only eight years and the results would be directly comparable.

From a government's point of view, it is also very important to focus the time scope of an NPV calculation, and to distinguish between *life cycle* on the one hand, and *point forward* project economics on the other. This is because price volatility and the long lifetimes of extractives projects could combine to give radically different NPV results from the same project. When commodity prices crashed in 2014, there were cases of companies presenting negative NPVs to governments to persuade them to relax terms of royalties, or taxation. But if such NPVs were determined over the life cycle of a project, including capex that had already been sunk, then they would not be the relevant metric. Instead, the government should be looking at the point forward NPV – is this mine, or oil field, going to yield an operating profit from this point forward? The life cycle metrics were estimated by the company before it took its final investment decision. If that decision has been taken, and the investment is sunk, it is only the point forward economics, expressed in the NPV, which will determine whether continued production is viable or not.

The last point about the NPV to note is that, precisely because it tracks flows of money against time, the same absolute cash flows over a whole project lifecycle could yield different NPV values depending on the life stage of the project. Suppose, for example, the $121 was returned to the investor over the course of two years but in a different time sequence. Instead of receiving nothing for 24 months and then $121 at the end, the investor got $21 after 12 months, and then the return of the original $100 investment at 24 months. Although the figures are the same in absolute terms, the NPV is now positive (slightly, at $1.58) because the investor got more payback sooner.

Internal Rate of Return

The IRR is the twin sister of the Net Present Value. Instead of being an absolute sum of money, it is expressed as a percentage in the same way an interest rate would be. So in the example above, the IRR is 10%, just as the interest rate is. It is the fact that the IRR is *the same as* the discount rate which makes the NPV 0 in this case. If the IRR is higher than the discount rate, the NPV will always be positive, and if it is less than the discount rate the NPV will be negative. So the IRR and NPV are really different expressions of the same calculations. As a rule of thumb, investors are often looking for an IRR of 15% or more, after all costs including taxes. Though there is no standardised rate.

Pre- and post-fiscal profitability

Investors will analyse profitability both before and after anticipated tax liabilities. In cases where terms for taxation can effectively be negotiated (such as an individually negotiated contract, for example), the comparison might lead them to specific suggestions on royalty rates, cost recovery and depreciation procedures, or a host of other mechanisms that combine to form the "fiscal regime". Without diving too much into detail, *which are dealt with in the separate paper on Petroleum and Mining fiscal regimes*, it is important to understand that these different terms interact with each other in a way that makes it impossible to consider each in isolation. For example, the higher the royalty exercises on gross sales, the lower will be a company's profits, and therefore any revenue streams to government which are based on profits such as corporate income tax, a special sector tax, or resource rent tax. This is why it is impossible to truly understand the economics of any given project without a detailed financial model.

Government Returns

Governments of course need to run their own analyses of project economics like companies do.

Conventional financial analysis often assumes the same discount rate to be applied to a project and then run against revenue flows to both the investor and the government alike, with only minor divergences For instance, since in many projects the government makes no direct financial investment, an Internal Rate of Return cannot be calculated.

One key question is: are or should the considerations be exactly the same for a government as they are for a company? Since a discount rate is supposed to integrate fully opportunity cost, using the same discount rate as a company would imply the government faces the same range of risks and opportunities. If this is not the case, should a government construct its own discount rate? And if so, how?

"Government take"

The most broadly used metric to calculate returns to government is what is called "the government take". This takes the overall profits generated over the lifetime of a project and calculates what percentage of them go to the government, and what percentage to the investor. The International Monetary Fund deploys a specific implementation of government take in its modeling for governments known as the Average Effective Tax Rate (AETR).

Government take provides comparable numbers between one project and another. A bigger question is, are the projects themselves comparable? For instance, can we predict that because Project A has a higher government take than Project B, it is a "better" project for a government, and was "well negotiated"?

And the answer is: not necessarily. As one simple example, government takes tend to be considerably higher, generically, within the petroleum sector than they are in mining, in a way that even IMF economists have stated they cannot fully account for. Then there is the fact that price and cost volatility have constant, unpredictable and nonlinear impacts on government take, with the further complication that two otherwise identical projects could be affected differently in their government take depending on what the terms in the fiscal regime are. Third, since profitability and risk can vary so much within the same industry, it is natural that projects could have different government takes. In the Middle East, where large oil reservoirs are proved and production costs remain low, the level of risk for investors is much lower, leading to governments being able to impose mechanisms that capture a higher proportion of the economic rent, and create a higher government take. In Iraq, for example, the government's take in the massive oil projects it has signed with multinationals since 2009 average over 90% - higher than in most other countries. But is this because the Baghdad government has negotiated and run these contracts better? Or is it because they started from a much stronger negotiating position, because of the scale of their reserves and low production costs?

The Impact of Project Financing

One factor which has been underestimated by governments has been the impact of project financing and incorporation structures on revenue flows, and in particular tax liabilities. Because, as noted above, even well leveraged companies will borrow financing rather than invest core capital when they can, and because financing costs are often tax deductible, and many other tax revenues depend on which corporate entity where is engaged, governments can often get a nasty shock as projects develop. Time and again in the last few years, officials have found that a project which looks from the outside as though it should now have reached profitability, and should therefore yield corporate income tax or some other profit-related taxation is not doing so because of these structures. Positive cash flows do not translate into accounting profits because the operating entity itself may be loss-making, because of complex project finance and transfer pricing arrangements with other entities. So the project as a whole may be profitable at one and the same time as the taxable entities within it are not.

Integral then to a government understanding what its actual revenue flows look like from an extractives project is understanding the exact corporate structures combining to invest and operate it. In this sense, "beneficial ownership", covered in a separate paper, must also count as necessary information to reach a fully informed evaluation.

Materiality of different factors

With so much uncertainty, and so many potential variables affected by so many factors, many of which may lie beyond the control of anyone directly involved in a project, managing complexity becomes a key consideration for many governments. In the context of project economics, then, it is important to understand that different factors of uncertainty can have a significantly different material impact on revenue flows.

Future price is the single most material factor, and of course impossible to predict. But impossible to predict does not mean impossible to manage – with financial modelling governments, like companies, can understand how a range of different price points will interact with profitability and taxes. In terms of cost categories, development costs tend to be very material, because they tend to be both large and "frontloaded" – preceding revenue generation, for example. Other categories like exploration costs might be less material, while operating costs become gradually more material to predictions in the life of a project to the extent that they start to represent a higher and higher proportion of total anticipated *remaining* costs.

Not all revenue streams are equal. Public attention, or a local EITI process, might focus significantly on land surface rentals, or commitments to training costs, even when these represent tiny proportions of projected revenues. Accountants can subject cost recovery structures to detailed analysis of capital depreciation rules even if they have marginal impact on overall government take.

Coming to an understanding of how to weight various factors project by project can make overall management of project economics easier for governments.