Forecasting Oil Production using Data in the BP *Statistical Review of World Energy*

R. W. Bentley
Visiting Research Fellow, Dept. Cybernetics,
University of Reading, UK.

1. CLASS ASSIGNMENT FOR STUDENTS

This assignment involves the estimation of future oil production by country and globally.

The data source you will use is the BP *Statistical Review of World Energy*. With only these data it is *not possible* to make a useful study of future global oil supply. This is unfortunate, as many analysts have relied on these data for just this purpose.

However, by combining the data with additional knowledge, and some simplifying assumptions, useful predictions of future oil supply can be made. In particular, estimates can be made for the dates when specific countries will reach their maximum production peaks of *conventional* oil, after which their production of this class of oil will decline. With the same technique, the date of the global peak in conventional oil production can also be estimated. In addition, an estimate can be made of the date of the global production peak of *'all-oil'*, though admittedly here several caveats apply.

These notes explain some of the problems with the data and the implied methodologies in the BP *Statistical Review*, and also set out how useful oil predictions can be made. The following are suggested steps, but other approaches are also valid.

Step 1: ***Use data from the BP*** *Statistical Review of World Energy.*

Do this by downloading the long-run (since 1965) spreadsheet data in the file: 'Statistical Review 2014 workbook' (these data are in MS Excel '.xlsx' format, and the file is 1.6MB in size), from the site:

www.bp.com/statisticalreview, and look for 'downloads'

(or via the full address: *http://www.bp.com/en/global/corporate/about-bp/energy-economics/statistical-review-of-world-energy/statistical-review-downloads.html*)

The workbook download you need is obtained by clicking the middle 'download rectangle' in the top row, of the six 'download rectangles' given on this web page.

Step 2: ***Use the oil classification of the BP*** *Stats. Review.*

For simplicity, initially follow the oil classification used in the BP *Stats. Review* data, which lumps together as 'oil' all the main types of fossil oil produced today. [1]

Step 3: ***Do not use R/P ratios.***

Do not use reserves-to-production ratios ('R/P' ratios) to give an indication of future security of supply, despite the fact that this is done in the *Statistical Review* itself (pages 6 and 7), and by many analysts.

This is because a region (or the world) may have a quite secure-looking R/P ratio, but be close to – or already past – its physical maximum in oil production; its 'oil peak'.

Take the example of Norway. Here the 2013 proved reserves (as given by *Stats. Review*, page 6) are 8.7 thousand million barrels, (i.e., billion barrels, Gb).

Its production in 2013 (page 8) was 1837 thousand barrels a day. This has to be multiplied by 365 (the days/year) and divided by a million to give Norway's annual oil production in billion barrels; which is thus 0.67 Gb/yr.

Thus Norway's R/P ratio is 8.7 / 0.67 = 12.9 (page 6, final column); i.e., Norway's current proved reserves are sufficient to give the country 12.9 years' of supply at the current rate of production. This would seem to be a fairly comfortable number, especially as new oil fields are still being found, and also the country is improving extraction rates from its existing oil fields.

But the problem is that Norway is already *past its peak* of conventional oil production, and its production has been in terminal decline since the year 2001 (see the *Stats. Review* spreadsheet production data). The reason for Norway's oil peak is simple: not enough new oil fields of adequate size have been found in Norway's waters for many years.

(And note that back in 2001, the year of Norway's peak, there was little in the BP *Stats.* data to indicate that the peak was at hand. In that year, despite a small hiccough in 1998 and 1999, Norwegian oil production had been on a steadily rising trend for many years; proved reserves stood near their all-time high; and - as today - new fields were being discovered and recovery rates improving.)

And what about the R/P ratio of the world? BP *Stats. Review*, page 6, shows that today's global proved reserves today are sufficient to provide 53.3 years' of global production at the current rate. This reassuring 'over 50 years of secure production' is often quoted to prove that any potential global oil supply constraint must be very distant. And - as in the case of Norway - not only does the world have substantial *proved* reserves, but large volumes of oil are still being discovered annually around the world, and significant improvements being made in global oil recovery rates. [2]

But as we know from Norway's case, the question to ask is *not* 'How many years' of proved reserves remain?', but 'How close is the world to its physical peak in production?' You will answer this question for conventional oil in Step 8; and - if you have time, and wish to do so - carry out the calculation for 'all-liquids' in Step 12.

Step 4: *Changes in proved reserves are a very poor indicator of oil discovery.*

Do not use changes in proved reserves data as being indicative of the history of oil discovery in a region. However, this also is often done; e.g. in past issues of the BP *Stats. Review*, or in the *Oil & Gas Journal*, or by chairmen of oil companies (who ought to know better), declaring reassuringly that: 'Global proved oil reserves have increased yet-again this year.'

The problems here are a little complex, but briefly:

i. Proved reserves (so-called '1P') data for a specific oil field, or for a region, have historically been very conservative numbers, and usually significantly lower than the more likely 'proved-plus-probable' ('2P') reserves estimates. Over time, as more knowledge about fields accumulates - and in the case of large fields, as these are drilled more extensively - the oil in the proved reserves estimates grows toward the larger 2P values. Thus growth of proved reserves in a region from one year to the next *may mean* that more oil has been discovered, but may often *also mean* simply that the 1P values are now closer to the more realistic 2P ones. [3]

ii. An important anomaly arises with OPEC proved reserves data. Here, against logic, for many years the 1P (proved reserves) have been *larger* than the 2P (proved-plus-probable) values. This has often been due to manoeuvring by OPEC countries to obtain better oil production quotas under OPEC rules. For this reason OPEC proved reserves data must be handled with a great deal of caution.

iii. The proved reserves data in the BP *Stats*. reflect different categories of oil (not just conventional oil), so - as the data for Canada and Venezuela show - large increases in a country's proved reserves can result from the inclusion of new classes of oil. There is nothing intrinsically wrong with this approach, but if one is trying to find out when *conventional* oil production in a region will peak, these other classes of oil must be stripped out.

Step 5: *What to do.*

If we cannot use R/P ratios (Step 3), nor rely on *changes* in proved reserves (Step 4), what do we do?

The answer is to use the 'mid-point peaking' rule-of-thumb. This says that production of *conventional* oil in a region typically peaks (i.e., reaches a maximum and then declines) – due primarily to physical processes – when about half of the region's original amount of conventional recoverable oil has been produced.

The total original amount of conventional oil in a region (i.e., before production started) is called its conventional oil 'ultimately recoverable resource' ('URR').

Once production in a region has been underway for some period, this conventional oil URR is thus composed of:

- *the* total conventional oil produced to-date from the region,
- *plus*: the region's current conventional oil 2P reserves (i.e., the oil that has been discovered, but not yet produced),

- *plus*: the conventional oil in fields not yet discovered, i.e., the region's yet-to-find.

Thus: URR of a region = oil produced to-date + current 2P reserves + yet-to-find.

And thus, in turn, the 'mid-point' rule says that production of conventional oil in a region typically reaches its physical maximum ('peak') when about 50% of this URR has been produced.

But how can this rule be applied to estimate the date of peak conventional oil production in a region (such as a specific country) when we have only the BP *Stats*. data to hand?

This can be done; but there are several problems, and we will deal with them in turn.

(i). Categories of oil.

As we know, BP *Stats*. data account for 'all-oil', not just conventional oil for which the 'mid-point' rule holds. But the oil produced in these other categories of oil is – except for Canada and Venezuela – not much, so initially just use the BP *Stats*. data without adjustment.

(ii). Finding total oil production to-date if production started before 1965.

Total oil production to-date in some regions, such as Norway or the UK, can be found easily enough by simply summing the data in the BP *Stats*. spreadsheet you have downloaded (recalling of course to multiply the production data given in '000 bbl/day by 365 and dividing by a million to give total production to-date in Gb).

But what to do if production in the country started before 1965, the start-date in BP *Stats*.? This applies to quite a number of countries, including the major producers such as the US, Russia and Middle East countries.

In this case use three simple approximations:

(a). Guess (or look up on the web) the start date of oil production in the country concerned.

(b). Then assume oil production grew from zero in the start year in a simple straight-line fashion to the volume per day given by the 1965. Make a simple calculation of the total oil volume that would have been produced up to end-1964. (It is a triangle, use: half base * height).

(c). Multiply this total by 0.7 to roughly allow for the fact that production typically increases exponentially, rather than linearly.

Take the case of the U.S. Though production here started in the 1860s, assume little was produced to 1900, and take this as the US start date. Since US production was 3.29 Gb/yr. in 1965, cumulative 'straight-line' production from 1900 to 1964 would be: 0.5 * 64 yrs * 3.29 Gb/yr. = 105.3 Gb; and roughly cumulative 'exponential' production of about (0.7 * 105.3 =) 73.7 Gb. This is a pretty good approximation, at least for the US.

Adding on cumulative production from 1965 to to-date gives total US cumulative oil production to-date as 237.0 Gb.

(iii). Can we use proved reserves data, if these data are so poor?

This is an important question, and the crux of this assignment. In the past *proved* reserves ('1P') data have been of no use at all for oil forecasting purposes because such data were very conservative numbers. But in recent years, for most countries, these proved reserves have grown to be at least reasonably close to the more accurate 'proved-plus-probable' ('2P') reserves data held in the oil industry datasets. So initially in this assignment assume that all current proved reserves data as given in the BP *Stats. Review* are good enough to use in the equation for URR given above.

(iv). How to estimate a country's yet-to-find?

The 'mid-point' peaking rule requires that an estimate to be made of the region's oil not yet discovered (its 'yet-to-find').

For analysts with access to proper oil industry proved-plus-probable (2P) data, this is not as difficult as it might appear, and there are many ways of making such an estimate, such as looking at the past oil 2P discovery history in the region, and extrapolating this forward.

But because the BP *Stats.* data do not give the oil discovery history in terms of 2P data, you must make some assumption about the yet-to-find if you are to use the 'mid-point' peaking rule to determine the date when peak production in a region is likely to occur (or has occurred if this peak was in the past, as is the case for conventional oil for many countries).

Here external experience is of help. While the amount of conventional oil yet-to-find is quite variable between counties, examination the historical 2P oil discovery data has told analysts that the total amount of conventional oil yet-to-find globally is probably not large. So initially in your calculations assume the yet-to-find in each country you examine adds just 15% to the total of: (cumulative production to-date plus reserves).

Step 6: *Two examples of applying the 'Mid-Point' rule.*

So now you are in a position to use the BP *Stats*. data to apply this 'mid-point' peak rule to a specific country to see when to expect its peak in production of conventional oil:

First pick the UK.

Look at the spreadsheet. A very small amount of oil production in the UK came from on-shore fields before 1965, but oil production in any volume (from off-shore) only started in 1976. So add up the total UK oil production from 1965 to to-date in the spreadsheet.

Remember that the oil production is in thousands of barrels per day, so has to be multiplied by 365 (i.e., days/yr.), then divided by a million to give billion barrels in a year. Summing the data gives UK total oil production to-date (ignoring the very small amount *before* 1965) as: 27.8 billion barrels (i.e., 27.8 Gb).

Next you need the BP *Stats*. data for reserves, where the proved reserves can be read from the BP *Stats. Review* page 6. These are: 3.0 thousand million barrels (i.e., 3.0 Gb).

So UK total oil discovered to-date = total oil produced to-date + reserves = 27.8 + 3.0 = 30.8 Gb

Adding the crude approximation of 15% for yet-to-find gives the UK's URR for conventional oil as 35.4 Gb, and where the 'mid-point' rule expects the production peak to occur at half of this, i.e., at 17.7 Gb.

Going back to the spreadsheet, examination of the production data shows that a total of 17.7 Gb of oil had been produced in the UK by 1998. Since the actual date of the UK production peak was 1999, as examination (or better, plotting) of the spreadsheet data of UK production will show, this indicates the usefulness of this technique, as least in this example.[4]

Take one more example, that of the US.

As you have already seen, US cumulative production to-date is 237.0 Gb. Adding on proved reserves of 44.2 Gb gives the US total discovered to-date of 281.2 Gb; and adding on an additional 15% for yet-to-find gives the US URR as 323 Gb. The US peak would be expected when half of this, 161.5 Gb, had been produced. From the spreadsheet, we can see that this quantity of cumulative production was reached in 1987.

In fact the US peak was quite a bit earlier, in 1971. The explanation for this discrepancy is several-fold: experience shows that many regions in fact see their peak of production some way *before* half of their URR is reached; in the US the largest field (Prudhoe Bay, in Alaska) was only discovered in 1968, so it did not impact the date of peak, but did significantly slow the post-peak decline; the current US proved reserves have grown in recent times due to exploitation of non-conventional oil (here, shale 'light-tight' oil); and 15% is a generous quantity to allow for the yet-to-find of conventional oil in a country so well explored.

But despite this discrepancy between the predicted date of peak of 1987 and the actual peak in 1971, again the value of the 'mid-point' technique is clear: it gives a useful indication of the production peak - and hence of the subsequent decline - of conventional oil in a region, and avoids being misled by the size of proved reserves, R/P ratios, or production trends etc.

Step 7: *Determine the date of Conventional Oil peak in other countries of interest.*

Now apply the method just outlined to your choice of any <u>three</u> other countries of interest (such as, for example, Russia, Canada, Saudi Arabia, Iraq, Iran, Mexico, China, India, Indonesia, Nigeria; or your own country if it is an oil producer and listed in the BP *Stats.* data); to see when to expect their peaks in conventional oil production.

If a country you have chosen is *not yet past* its peak (unlike the UK and US examples given above), you must project that country's oil production forward (making a reasonable assumption on how fast this production will grow each year) until the 'mid-point' rule says that the country's conventional oil production peak will have been reached. This applies, for example, to Middle East OPEC countries, Nigeria, and a number of other countries. (And to carry out this calculation, make the over-simplifying assumption that the country's proved reserves do not change over this projected future period.)

List the key data you use in each country's calculation, as given above in Step 6 for the UK and US.

Step 8: *The Global conventional oil peak*

Apply the same 'mid-point' method to roughly estimate the date of peak of *global* conventional oil production.

Step 9: *Compare your results to Globalshift's (and Campbell's if possible).*

Compare your results for *conventional* oil production with the graphs of past and predicted future *'all-oil'* production on the Globalshift website: *www.globalshift.co.uk* (Go to the home page, then select a region of the world from the banner near the top, and then a country from those in the subsequent list.)

(And if you have access to *Campbell's Atlas of Oil & Gas Depletion* (Springer), compare your results to the data for '*Regular Conventional*' oil.)

Step 10: *Advice to Governments*

On the basis of your findings in Step 7, state briefly how you would advise the governments of the three countries you selected.

Now look at the global picture. Once global conventional oil production approaches its peak, and can no longer increase sufficiently to meet rising demand, the additional oil required must come from the non-conventional sources of oil. Often these are more expensive to produce than conventional oil, in part because of the additional energy required for their production. And in terms of the world price for oil, in an open market it is the price of the marginal barrel (i.e., the most expensive of these 'extra' barrels of oil required to meet demand) that sets the overall price of oil.

Using this information, and your finding from Step 8 of the likely timing of the global peak in conventional oil production, state briefly how you might advise an international body such as the UN.

2. CONCLUSIONS

In this assignment you have investigated a simple way to estimate the date of the peak of production of *conventional* oil in regions (such as specific countries), and in the world as a whole, using the BP *Stats. Review* data combined with external knowledge and assumptions.

And also as a result you have found out how not to be misled by R/P ratios, current trends in production, apparent increases in proved reserves, or high 'ultimates' as apparent proof that there can be no supply-side constraints to global oil supply for many years to come.

3. POSSIBLE ADDITIONAL WORK

If desired, and if time is available.

Step 11: *Special handling of OPEC countries' reserves; also Canada & Venezuela oil types.*

The above analysis can be extended and refined; for example:

a). Recall the remark above that proved '1P' reserves as reported in the BP *Stats. Review* in some OPEC countries are probably *larger* than the more likely '2P' reserves. So reduce proved reserves in these countries by some amount (such as by their cumulative production during the total time their 1P reserves remained static), and re-calculate their expected dates of conventional oil peak.

b). Treat Canada and Venezuela differently because of their tar sand, and Orinoco oils, respectively; making guesses (or assumptions, based on web data) for the production of these oils and subtract off to give estimated production of conventional only; and for reserves likewise subtract off the tar sands and Orinoco Belt proved reserves shown at the bottom of page 6. Then use these revised data in a 'mid-point' calculation to arrive at the two countries' dates of expected peaks in their production of *conventional* oil.

Hence modify your earlier calculation of the expected date of the global peak to more closely reflect the date of global peak as it applies to conventional oil only.

Step 12: *Forecasting the Global production of 'All-Liquids'*

Move beyond the peak of *conventional* oil, and see what is likely globally from the production of non-conventional oils. Recall that these include 'light-tight' oil (shale oil); tar sands and Orinoco oils; oil from shale rock ('kerogen oil'); oil produced by processing coal, and gas - but remembering in the latter case the approaching production peak of *conventional* gas); and oil from biofuels. Use the two reserves figures you have just used (for tar sands and Orinoco oils), plus other data from the web.

Couple these findings with the global peak you have above for conventional oil, to allow a picture to be built up of the possible future global supply of 'all-liquids' (i.e., all types of oil, including biofuels and oil produced chemically from coal & gas), to see what sort of oil future the world faces. For forecasting purposes, assume for simplicity that global conventional oil production in a region, *once it is past its conventional* *oil production peak*, declines at 3% a year (where this corresponds to a decline of ~5% p.a. from the fields in the region that themselves are past their production peak, offset by an increase of ~2% p.a. from newer fields in the region that are coming on-stream, including - in time - fields still yet-to-find).

There are, of course, significant caveats to such an 'all-liquids' forecast. These include:

(i). Higher prices can help find more oil, and also make more of the non-conventional oils (including probably even 'kerogen' oil) economic to produce.

(ii). A high oil cost is ultimately self-limiting, as it leads to oil 'demand destruction', and damages overall economic activity;

(iii). Production of significant quantities of non-conventional oil may be restricted under forthcoming climate change agreements: if we burn all the oil we know about, including all of the non-conventionals, the global 2 ºC limit is far exceeded.

Step 13: *Other approaches for predicting the date of peak production of conventional oil in a region using the BP Stats. Review data.*

In addition to the 'mid-point' rule used above, there are two other simple approaches that can be used to estimate the date of peak conventional oil production in a region.

(i). PFC Energy's '60%' rule.

The 'PFC Energy' rule-of-thumb is that production of conventional oil in a region typically reaches its peak when about 60% of the total conventional oil that has been discovered in the region to-date has been produced.

That is, for a region, for conventional oil:

Total oil discovered to-date = oil produced to-date + current 2P reserves

with the 'PFC Energy' rule saying that production of conventional oil in a region typically reaches its peak for primarily physical reasons when about 60% of this 'total oil discovered to-date' has been produced.

This rule is thus simpler than the 'mid-point' rule, in that no estimate is required of the region's yet-to find. But the rule does rely on the reserves data being proved-plus-probable (2P) reserves, not proved reserves, and so can be rather unreliable when using BP *Stats.* data.

(ii). Hubbert Linearisation

An even simpler rule assumes that cumulative production of oil in a region follows a 'logistic' curve (the derivative of which, for oil, is known as the 'Hubbert curve'). In this case a clever linearisation of this curve (by plotting production data in a region against suitable axes, see below) gives not only the quantity of oil that will have been produced by the peak data (and hence allows the peak date to be determined - either historically if in the past; or by extrapolation if in the future), but also the region's URR.

The approach, for a given region, is to plot on a graph's ordinate (vertical) axis for each year: the region's (annual production divided by its cumulative production); and on the abscissa (horizontal axis) for the same year the region's (cumulative production).

Such a graph indicates the region's URR by the value indicated by extrapolating the (notionally) straight line produced to the abscissa; with the region's production peak occurring when cumulative production reaches half this value. (This technique works because a logistic curve is symmetrical, and peak in such a curve occurs at exactly half of URR).

This approach is simpler than either the 'mid-point' approach or the PFC '60%' rule, as it requires only the region's annual production data, which are usually fairly well known. But because it assumes production is following a logistic curve fairly accurately, it is often only reliable later in a region's life, once the peak date is near or past.

If you have the interest and time, by all means try either or both of these techniques on one or more of the countries that you have already investigated for date of peak under the 'mid-point' rule, and see how each of the three the techniques compares for ease of use and accuracy.

Step 14: *Net energy*

Of considerable interest going forward is how the generally lower energy-return-on-energy-invested (EROEI) ratios of nearly all new sources of energy (such as non-conventional oil and gas, as well as the renewables) will impact mankind's ability to access energy. Some analysis suggests that the impact will be large.

To investigate the oil aspect of this question, take your global 'all-liquids' forecast and modify the production quantities of the various categories of oil you assume (such as conventional oil, 'light-tight', tar sands, biofuels etc.) by the EROEI ratios listed below (converting EROEI numbers to 'net-yield' numbers, and making estimates where needed) to change your forecast from 'total barrels of liquids' to a forecast of 'useful net-energy' barrels of liquids; and hence compare the two forecasts and determine the implications for mankind.

Table: Estimates of Energy-return-on-energy-invested (EROEI, or EROI) ratios for different types of oil.

	Approx. EROEI range
Conventional oil: 1930 / 1970 / today	30 / 40 / 14
Shale oil ('light-tight' oil)	8?
Tar sands / Orinoco oil	1.5 - 8
Oil from kerogen ('oil shale')	5?
GTLs, CTLs	10?
Biodiesel, gasohol	~3

(Sources: Various; incl. C. Hall & J. Day, *American Scientist*, 97, 230-237, 2009.)

Note: If you search on the web for EROEI data, be aware of whether or not the 'internal' energy of the source material is included in the calculation of EROEI. For example, in determining the EROEI for 'kerogen' oil, a great deal of the heat need for retorting the kerogen can come from burning the kerogen itself. Including this energy leads to a low EROEI value. If only the energy purchased by the plant (the 'external' energy) is included, a higher EROEI value results.

4. ENDNOTES

1. Thus the main oil data in the BP *Stats. Review* include the following categories of oil:

- Conventional oil (i.e., normal 'flow-able' oil located in distinct oil fields, and which makes up the majority of all oil produced globally);
- Light-tight oil (also called 'shale oil': i.e., light oil locked in the rock pores, and which must be extracted by hydraulic fracturing - 'fracking' - of the rock);
- Oil extracted and upgraded from oil ('tar') sands (today, mostly in Canada);

- Orinoco oil (the heavy oil from the Orinoco Belt in Venezuela);
- Oil produced from oil shale (by retorting the oil pre-cursor, kerogen, in the shale rock; either on the surface following mining of the rock, or by *in-situ* methods). Very little of this type of oil is currently produced, though it is likely to become important in future.
- Condensate (oil that condenses naturally from gas streams, from either an oil well's gas cap, or a separate gas field, when this gas reaches the surface).
- NGLs (the oil-like 'natural gas liquids' that are produced by physical processes from gas in gas fields, and which constitute about 10% of total global oil supply. These are composed mostly of ethane to pentane, being either liquid at normal temperatures and pressures, or can be turned into a liquid by moderate pressure).

Note that three types of oil are excluded from the main BP *Stats. Review* oil data. These are: oil produced by chemical processes from coal or natural gas (coal- and gas-to-liquids, 'CTLs' & 'GTLs', respectively), and oil from biomass (biofuels). Production of these last three classes of oil is currently fairly small, and they are excluded from the main BP *Stats.* oil data because they are produced from very different sources, and by very different processes, than 'normal' oil.

Note also that while the BP *Stats.* includes oil from the Canadian oil sands and from the Venezuelan Orinoco Belt in the main oil data for reserves and production, for reserves these two classes of oil are included in the data, but also broken out separately at the bottom of the table. This is very useful when analysing the likely dates of the Canadian and Venezuelan production peaks of *conventional* oil, and also for estimating the possible future contributions from these two classes of oil.

2. Note that there is plenty of scope - at least in theory - for knowledge and technology to increase the recovery of oil from existing fields. Production of *conventional* oil typically recovers – on a global average – only about 40% of oil in fields, leaving behind 60% that cannot be produced by normal means. Thus the original *oil-in-place* in fields is - on average - usually more than twice the amount of oil considered *recoverable*. However, in your analysis deal only with the oil that is considered recoverable, as given by the BP *Stats.* reserves data; do not consider the oil in fields that is currently considered difficult or impossible to recover.

3. In collaboration with two students from a course possibly rather like yours, an academic paper was published in 2007 on the need to use such '2P' data:

Bentley, R.W., Mannan, S.A., and Wheeler, S.J. (2007). *Assessing the date of the global oil peak: The need to use 2P reserves*. Energy Policy, vol. 35, pp 6364–6382, Elsevier.

This paper drew on one of the oil industry's main datasets of '2P' data, that of IHS Energy Ltd. This was the company's 'PEPS' Exploration & Production dataset, purchased by the Oil Depletion Analysis Centre, London. The data that were published in the paper were with the permission of IHS Energy.

4. This precision in estimating the date of the UK peak from using a rough rule-of-thumb is something of a lucky fluke. The explanation is partly that BP *Stats.* 1P reserves for the UK underestimates the 2P reserves by a factor of two (an unusually large ratio compared to many other countries); coupled with the fact that the UK peaked later than expected due to its 'double-hump' production profile, itself partly an outcome of safety work carried out across all oil fields following the Piper-Alpha tragedy.