Projecting US Oil Production

Posted by Stuart Staniford on January 12, 2006 - 7:59am
Topic: Supply/Production
Tags: hubbert linearization, hubbert peak, oil prices, peak oil, united states oil production [list all tags]

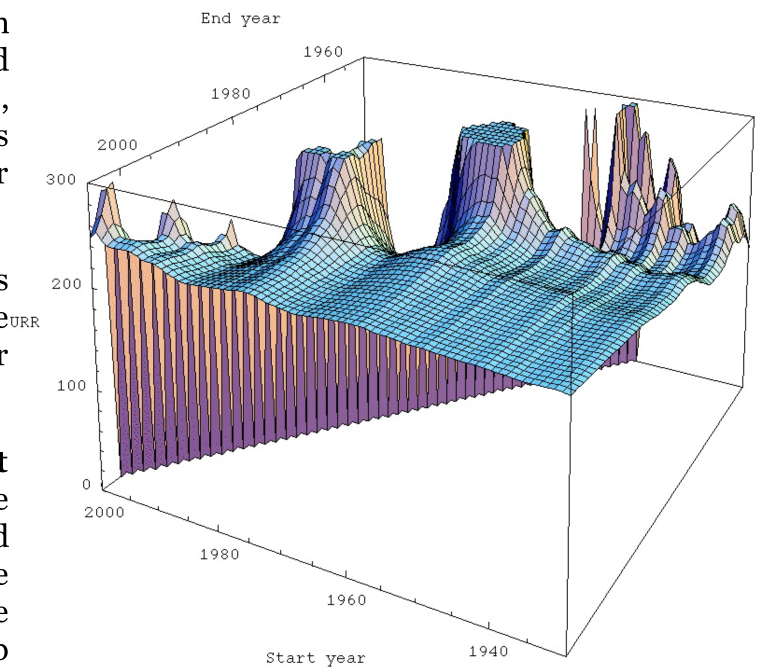

I began a discussion of US oil production in Four US Linearizations, and then continued in Predicting US Production with Gaussians, and Linearizing a Gaussian. This post wraps up my analysis of the US URR (at least for now).

The post that follows below the fold is rather technical, and filled with large images. Here's the executive summary for those who would like to skip the details.

Update [2006-1-13 3:22:37 by Stuart Staniford]: After spending more time thinking during the waking hours, I decided my approach had a flaw, which I've corrected. It changes my URR estimate very marginally from the original 218 ± 8gb to 219 ± 8gb. Details in another update inside the post. Apologies for any confusion.

- I estimate that the ultimately recoverable resource for US field crude production as measured by the EIA will be 219 ± 8gb. Since the EIA believes we've used 192gb so far, the remaining balance is 27 ± 8gb. These are intended as two-sigma error bars, and the figures exclude NGL (see EIA definitions).
- This conclusion come from fitting both Gaussian and Logistic curves to the data in two different ways each, doing extensive stability analysis, and making judgements about the level of agreement throughout the regions of stable prediction in all methods.
- I show that on the US data, Hubbert linearization is the most broadly stable prediction technique of the four I considered (despite the fact that the Gaussian actually fits the data better than the logistic).
- In particular, the linearization is the only one of the techniques considered that has any significant domain of reliability before the peak.
- However, the Gaussian is likely more accurate now that it is well constrained by a long history of data.
- The caveat to this extrapolation is, while these models seem to fit US production amazingly well, we still lack a deep understanding of why this is true. Therefore, there is some risk that the conditions which cause them to fit well might change in the future, thus breaking

the projections. I would be surprised if this happens in the case of the US, however.

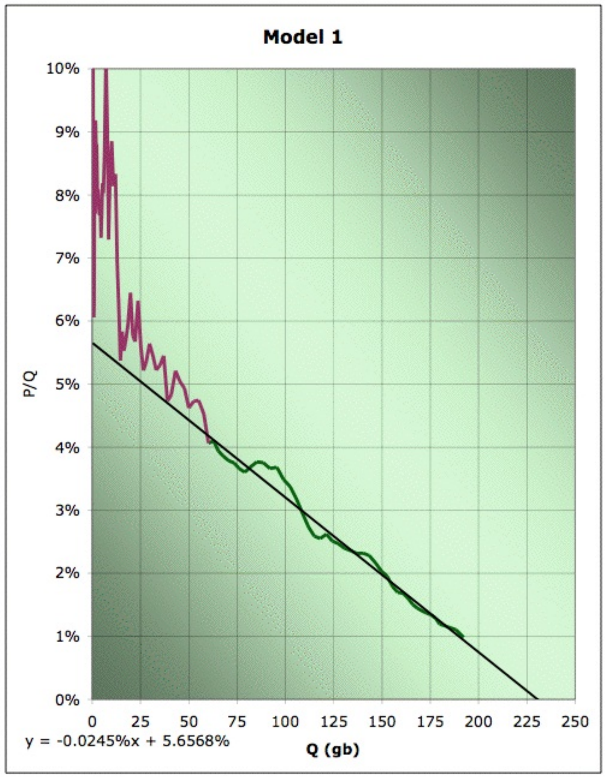

Recall that the other day I was making pictures like the one to the right, where I explore what happens when we try to extrapolate US production via a linearization, but we start to mess around with where we start and end the linear fit. The specific one shown is Hubbert linearization of EIA field crude production with linear fit from 1958-2005.

Long experience has taught us that the linearization generally does a bad job in the early part of the history, (though I didn't know why till know) so we don't usually start at the beginning. But then the question becomes how sensitive our answer is to where we start. Deffeyes picks 1958, but why, and is that a good place? Additionally, we'd like to know how sensitive it is to where we end fitting. In the past, we had less data and so had to stop the fit sooner. If our answer has been changing a lot in the recent past, that suggests we wouldn't have a lot of confidence that it wouldn't continue to be volatile in the future also.

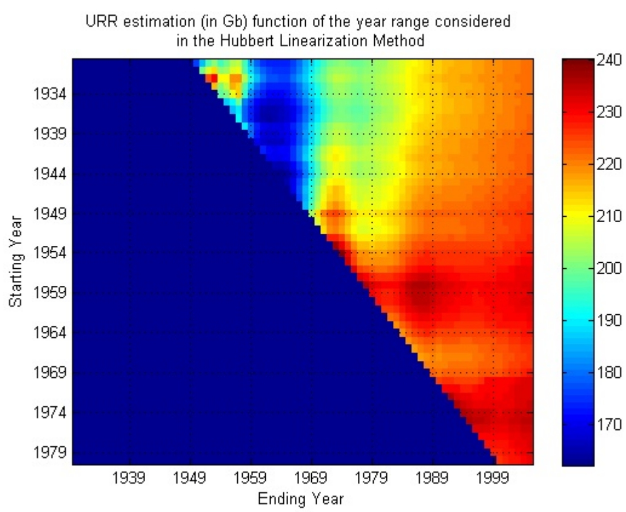

Next, **Khebab** made the very nice density plot to the right, where he explored a fairly large space of possible starting and ending points. The way to read his graph is as follows. The ending year is on the X-axis, and the starting year is on the Y-axis. For each point on the graph, he has done a linearization, and predicted the final US URR (ultimately recoverable resource). A color-code denotes the answer, and you can see it's been fairly closely around 220gb-230gb for quite a while (the right hand side of the picture), with only modest fluctuations. Encouraging.

I wanted to move the ball a little further down the field. I had two major goals. One was to come up with a quantitative error bar for the URR estimate. The second was to explore which of various prediction techniques does the best job. I noted in a piece on Predicting US Production with Gaussians, that the Gaussian actually fits the US data better, especially in the early stages. Indeed, that seems to be the reason why the linearization doesn't quite work at early times in the production history - the early tail is not matching a logistic well, but is matching a Gaussian well. So I wanted to explore more thoroughly projection with Gaussians too.

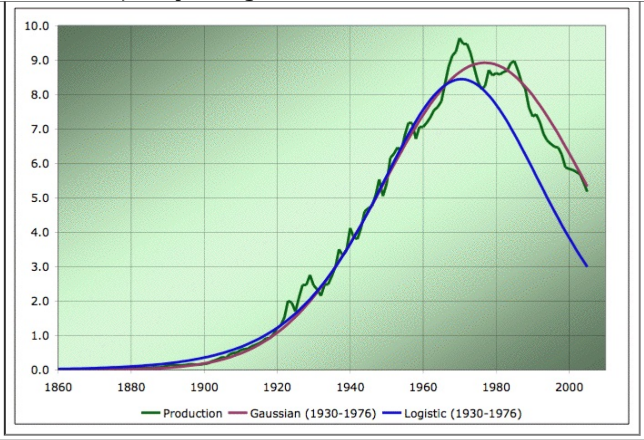

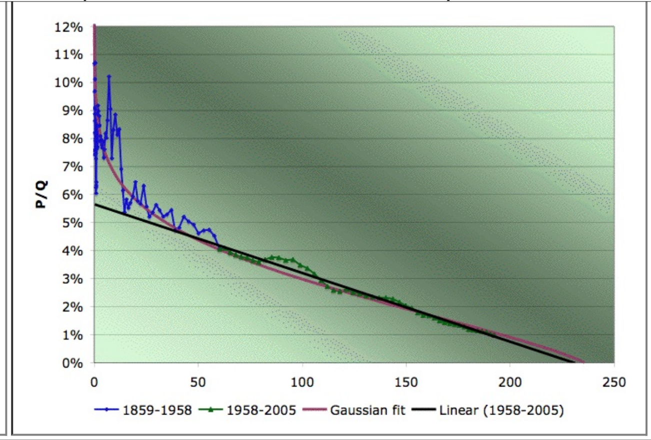

Thus in my analysis I explore predicting URR via four different fit/extrapolate techniques. They are:

1. Making a linearization of the data, and then extrapolating the straight line fit out to the X-axis to get the URR.
2. Directly fitting a Hubbert peak (first derivative of the logistic) to the production data in the P/t (production versus time) domain.
3. Fitting a quadratic to the log of the production data as a way to estimate the Gaussian parameters.
4. Directly fitting a Gaussian to the production data in the P/t domain.

The first and third involve linear fits, while the second and fourth required non-linear iterative fits. However, modern computing equipment and sofware being what it is, the difference is barely noticeable any more. For all of these techniques, I repeated the fit at a sizeable range of starting and ending years. The following plots are the result.

These are analagous to Khebab's plot above. Notice that in the far back there is a region where the end year is before the start year, which doesn't make any sense. So I just set the answer to always be zero there. In the foreground of the plots, there is a more-or-less flat horizontal area which I refer to as the *zone of stable prediction*. Typically it involves having the end time fairly recent, and the start time reasonably early (but the exact nature of the stable region is dependent on technique). As you move around in that area, the answer doesn't change too much. However, as you move back into the plot, things go haywire. The curtain across the middle is the area where the start and end time are a very small number of years apart. Clearly, when that's true, the fit is unlikely to work well as it becomes very susceptible to the noise in the data.

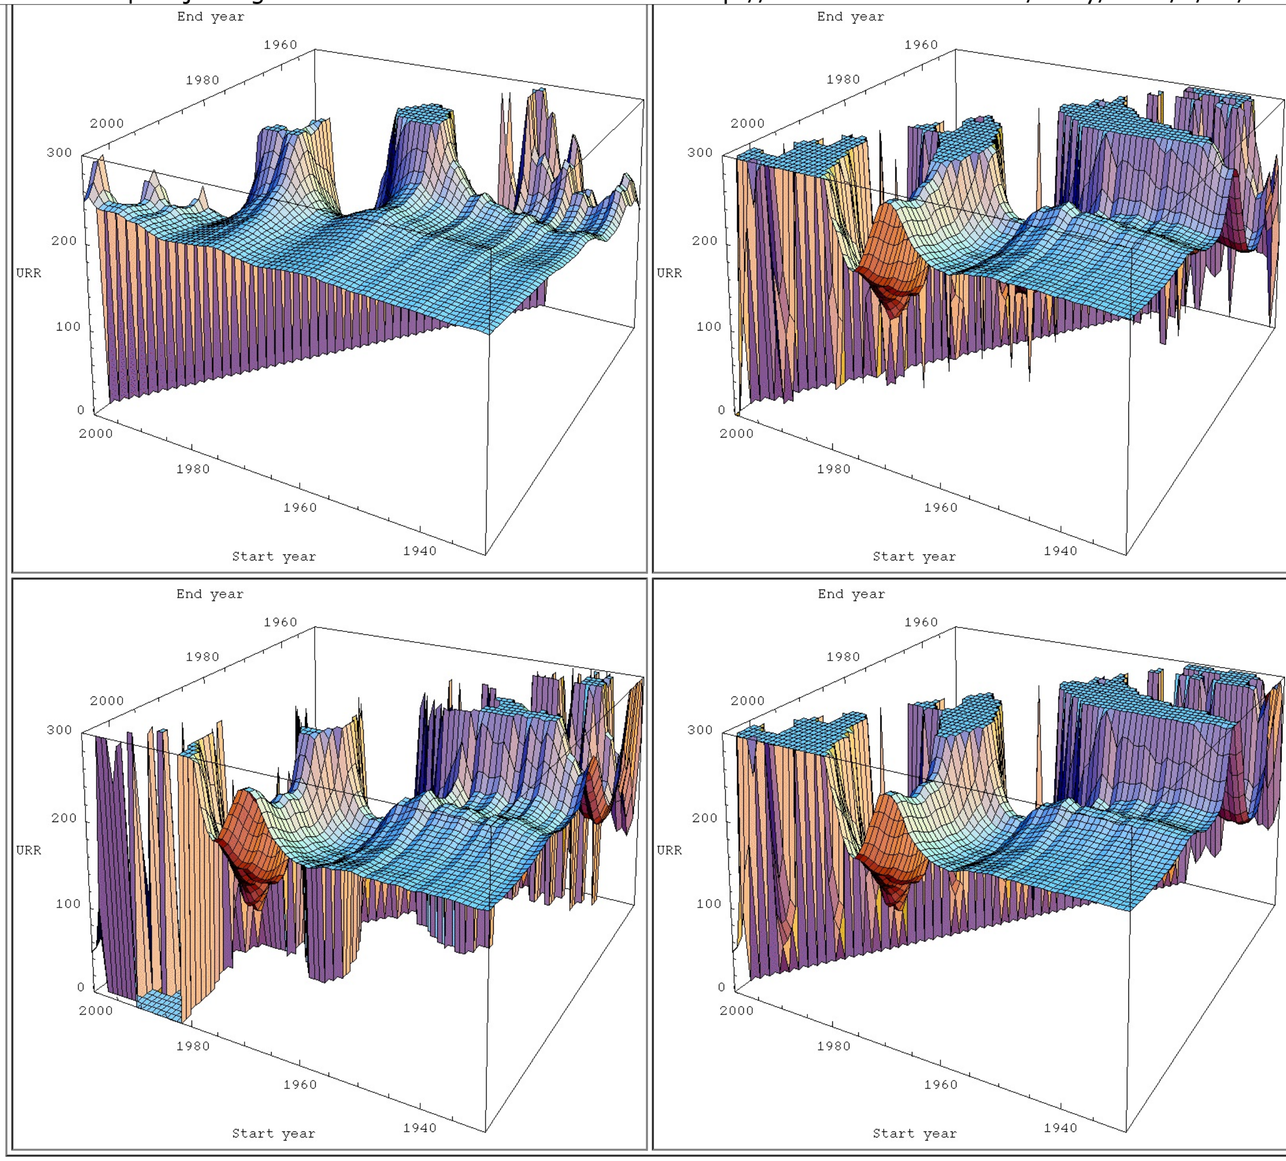

Stability assessments of estimates for ultimately recoverable US oil production (URR). In each case, URR is plotted against the start and end of the range of years used for fitting. Click to enlarge each figure. Four prediction techniques are used: top left is Hubbert linearization. Top right is direct fitting of the logistic curve in the production versus time domain. Lower left is a Gaussian model based on fitting a quadratic to the log of production. Lower right is direct fitting of a Gaussian in the production versus time domain. The underlying data is from the EIA estimate of field crude production.

Now, the most interesting thing to me is that the Hubbert linearization technique (top left) produces the largest zone of stability - there's a noticeably larger smooth area towards the front of the plot, and as it goes haywire towards the middle, it goes less haywire less quickly. The other three techniques are all of roughly similar quality to one another.

This is particularly interesting given that the Gaussian actually models the data better. However, there's a big difference between making pretty fits of past data and actually being able to predict new data robustly. In general, making pretty fits to past data is best served by having a model with a lot of parameters in. That means the curve has more degrees of freedom with which to wiggle itself around to the nicest shape lying along the data. However, lots of parameters means that there are likely to be more different ways that the model can get close to the data, and that allows for greater uncertainty in what the parameters actually are, which allows them to get more wrong. Models with too many parameters can suffer from *overfitting* in which the regression chooses a model that is essentially optimized to the particular noise in the data, and is losing touch with the true dynamics that would allow it to successfully extrapolate outside of the range where the data is.

A simple model (ie fewer parameters), even if it doesn't actually fit the data as well, may do better just because the few parameters it has are better constrained. The linearization trick has the merit of removing one parameter from the situation (the date of the peak), which then means we are in a better position to estimate the others (at least better when we have only a marginally adequate part of the data history). At least, that's my best guess as to what's going on.

The next four plots are essentially the same thing for the same four techniques, except drawn as contour plots rather than rendered as three dimensional surfaces. This allows us to see what the zones of stability are like a little more quantitatively. In each case, blue corresponds to a URR of zero (or undefined), red corresponds to 300gb or more, and green is the 150gb point. The contours are 3gb apart.

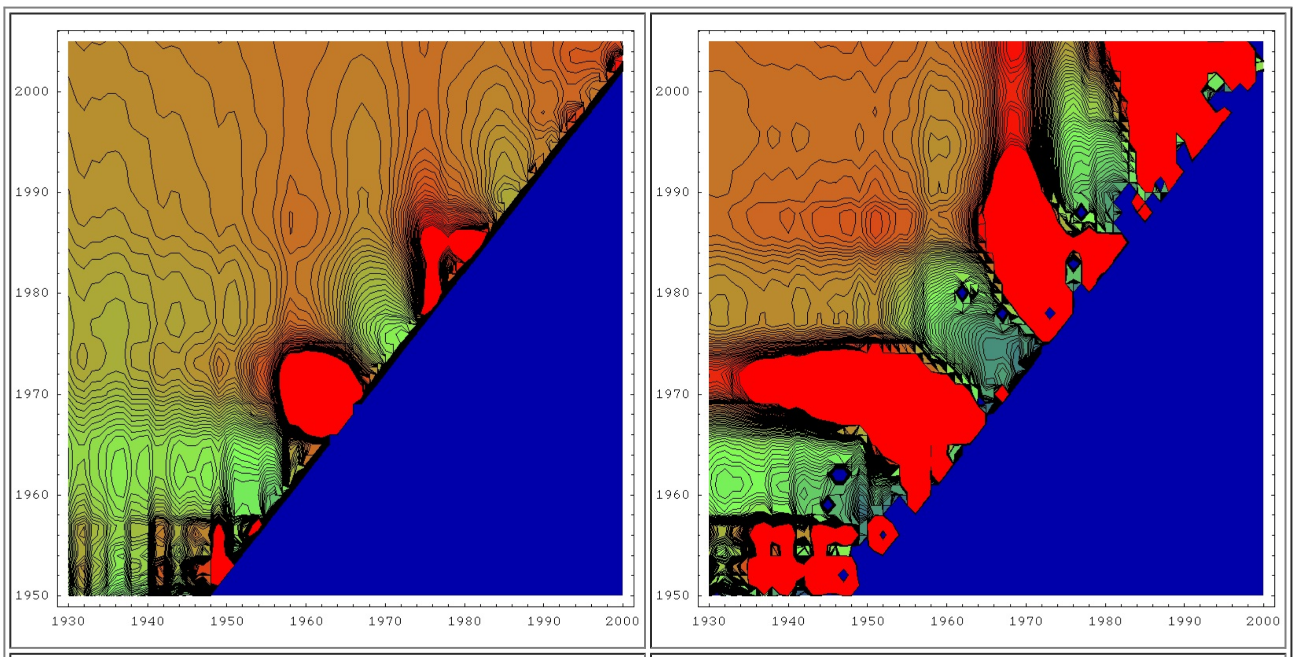

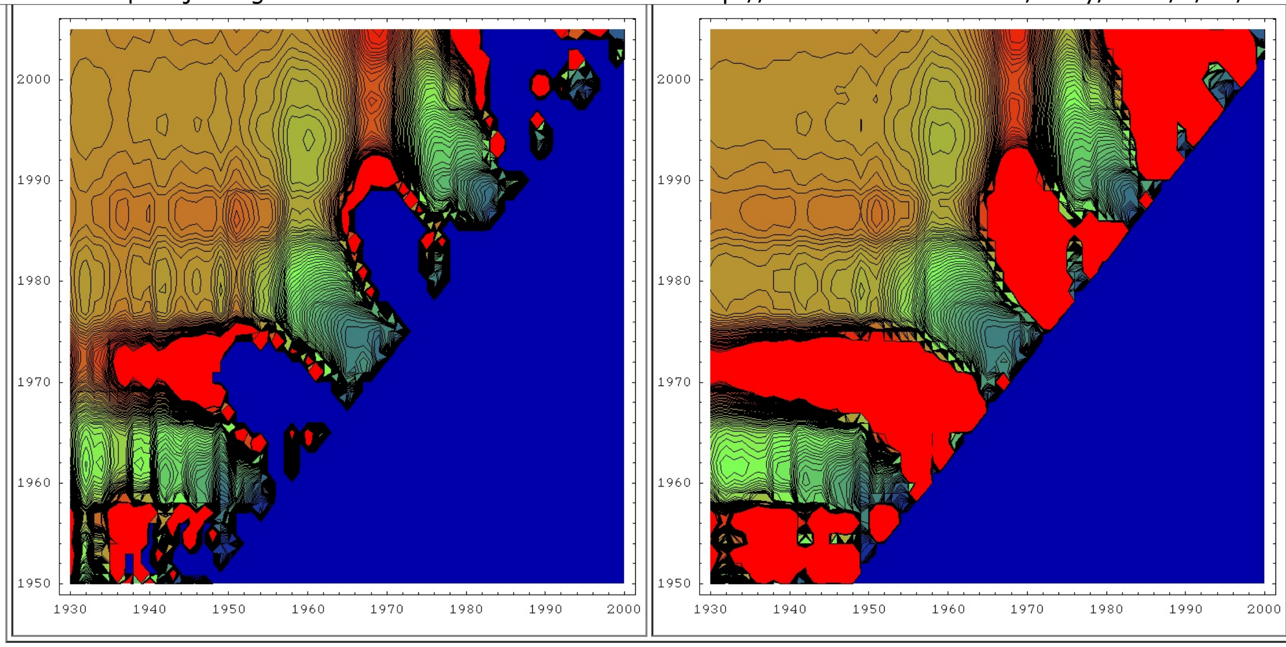

Stability assessments of estimates for ultimately recoverable US oil production (URR). In each case, URR is plotted against the start (x-axis) and end (y-axis) of the range of years used for fitting. The zone of stability is generally in the upper left of each plot. Click to enlarge each figure. Blue is URR=0, Red is URR>=300gb, and contours are 3gb apart. Four prediction techniques are used: top left is Hubbert linearization. Top right is direct fitting of the logistic curve in the production versus time domain. Lower left is a Gaussian model based on fitting a quadratic to the log of production. Lower right is direct fitting of a Gaussian in the production versus time domain. The underlying data is from the EIA estimate of field crude production.

So the first thing to become clear again is that the linearization (top left) has the largest region of approximate stability. Most importantly, it's the only technique that is approximately reliable before one actually hits peak (in the mid seventies). J.H. Laherrère in his paper The Hubbert Curve: It's Strengths and Weaknesses makes several Gaussian extrapolations from early on that fail, but here we can see more systematically that the linearization works best for early predictions.

However, another thing becomes clear too - it's stability region is not as flat as the smaller stability regions of the other techniques - in particular the Gaussian techniques. I think what's going on here is that once there really is enough history to constrain the parameters well, the Gaussian technique starts to do better because it is actually a better model of the data. As we know, in the linearization, the early data is always drifting upwards from the straight line, and this tends to distort the estimate unless we make the start date late, and then we cannot take advantage of as much data to fix the parameters as the Gaussian technique can. By using more data, the Gaussian can effectively bridge across the (rather lumpy) noise.

To come up with the URR and error estimates, I took a triangular region at the top left of each picture and grabbed all the different URR estimates out of them. I then got averages and standard deviations of those estimates. I used a larger triangle for the linearization than the other three techniques. Those numbers come out at:

Case	Triangle side (yrs)	Mean URR (gb)	Edge URR (gb)	Sigma (gb)
Linearization	40	215	225	10
Direct Logistic	30	232	236	5

Quadratic Gaussian	30	218	220	4
Direct Gaussian	30	217	218	4

You have to look at these in the context of the pictures above. The linearization URR trend is still going up, whereas the others are flat/wandering as they approach 2005. So I think the linearization is headed up towards the direct logistic estimate of 230gb or so. So the question becomes do we believe that answer or the Gaussian answer. I prefer the Gaussian at this stage, since it's well constrained with this much of the curve in view, and it seems to do a significantly better job of fitting overall, and especially in the early tail. Presumably, there is some central limit type reason for this (though I wish we knew exactly how that worked), and if so, we'd expect the late tail to be Gaussian also.

The main difference in the late tail is going to be as follows. The logistic curve has a decline rate that asymptotically approaches K. The Gaussian has a decline rate that increases at a fixed constant rate per year forever. So in the late tail, the Gaussian starts to decline a lot faster than the logistic. I believe this is why the Gaussian URR estimates are a little lower than the logistic ones.

Given all this, I take as my estimate the Gaussian estimates. What's in the table is the standard deviation, but what I quote above is a two-sigma error bar: 218 ± 8gb. Note, I am intentionally keeping the standard deviation here as the error bar rather than reducing it according to the number of observations since the noise here looks very lumpy rather than iid random. Therefore, I'm not assuming any potential for it to cancel (to be conservative).

My estimate can be contrasted with that of Deffeyes (228gb - based on linearization in Beyond Oil), and Bartlett of 222gb using Gaussians. Also, Khebab quotes a figure of 222gb and has some very interesting discussion but doesn't quote a single error bar. I don't quite agree with his technique there because he's effectively assuming random uncorrelated noise, and the real noise doesn't look like that - it's lumpy and nasty.

Update [2006-1-13 3:22:37 by Stuart Staniford]: I decided on reflection that there's a problem with estimating the URR by averaging over the triangle - it means that the most recent years in the production profile are underweighted in the overall average. Thus we fail to take account of the most recent data properly, which should inform us most. So I added an extra column to the table for the *Edge URR*, which is just averaged over the leading (most recent) edge of the triangle. I still use the fluctuations in the full triangle for my error bar estimate, however. Otherwise the reasoning is unchanged. So my estimate is now 219 ± 8gb.

Finally, one of the interesting discoveries I made in writing this post was that in making the plot to the right, I actually got very lucky. That is a pretty decent extrapolation that is sitting in a saddle between several areas of quite poor predictions. So sensitivity analysis is always a good idea.