The Oil Drum
DISCUSSIONS ABOUT ENERGY AND OUR FUTURE

Well, could we linearize this, then?

Posted by Stuart Staniford on July 17, 2006 - 1:06am
Topic: Supply/Production
Tags: peak oil, romania [list all tags]

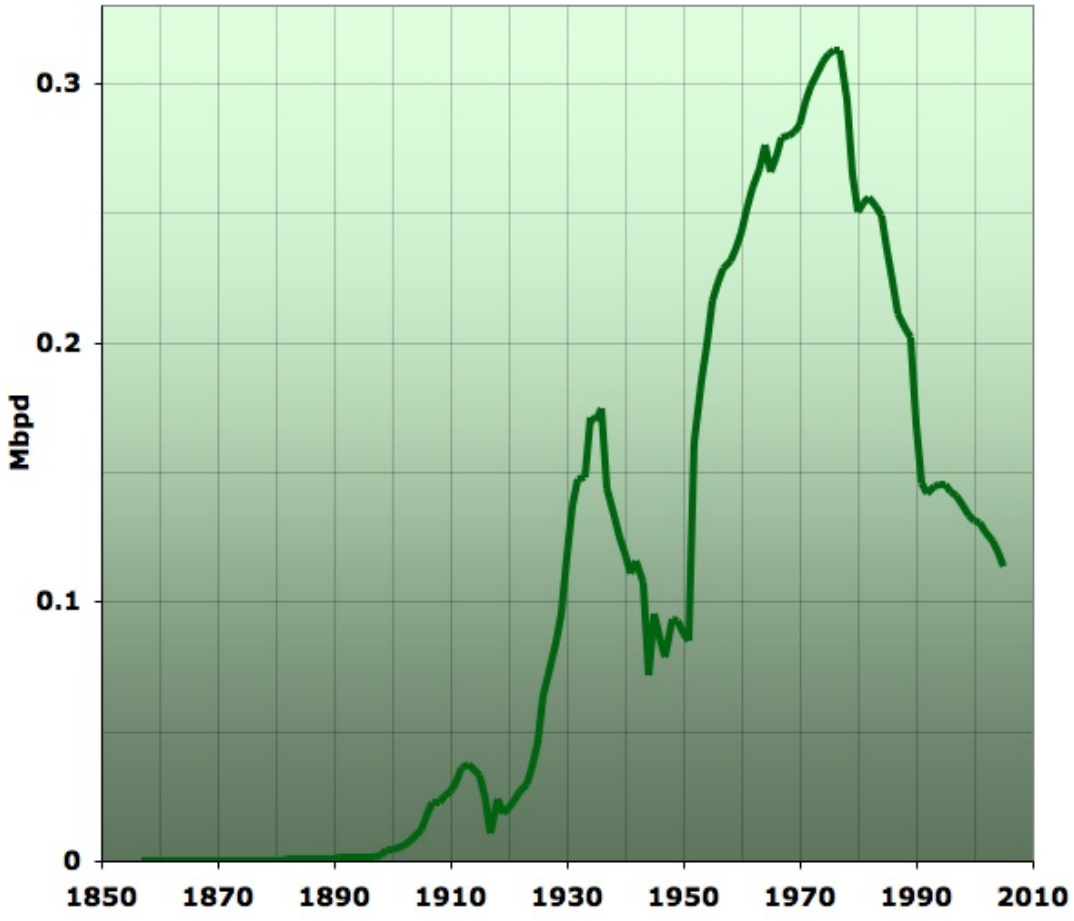

History of Romanian oil production versus time. Click for larger version. Source: American Petroleum Institute (courtesy J. Laharrere) for 1857-1958, and BP for 1965-2005 (includes NGLs). Data for 1959-1964 is linearly interpolated between the two data sets.

(For new folks, it might help to look at Super G's links to tutorial material on Hubbert Linearization. This post assumes readers have read prior linearization posts on TOD.)

Following our discussion of Hubbert Linearization of the Middle East, I was inspired to take a quick look at Romania. Romania is similar to the US in that production began very early (in fact the world's first commercial-scale oil production occurred in Romania in 1857), and peaked in the 1970s (1976 was the year of peak production). However it differs in that the amount of production has never been that large - peaking at just over 300,000 barrels a day, an amount

that a single large modern offshore platform might handle. It's also different in that the production profile (above) has been interrupted on a huge scale, especially by WWI and WWII, both of which caused an approximate 50% reduction in production.

Note that the scale on the bottom totals 160 years. The peak looks sharp, but really the peak is 30 years ago.

Thus the profile is at least as choppy as the Middle East. But it turns out that this does **not** render the curve-fitting methods useless, as you might imagine. Firstly, if one supposed that production followed some underlying dynamic that led to a Gaussian shape, then the log of production would follow an upturned parabola (quadratic):

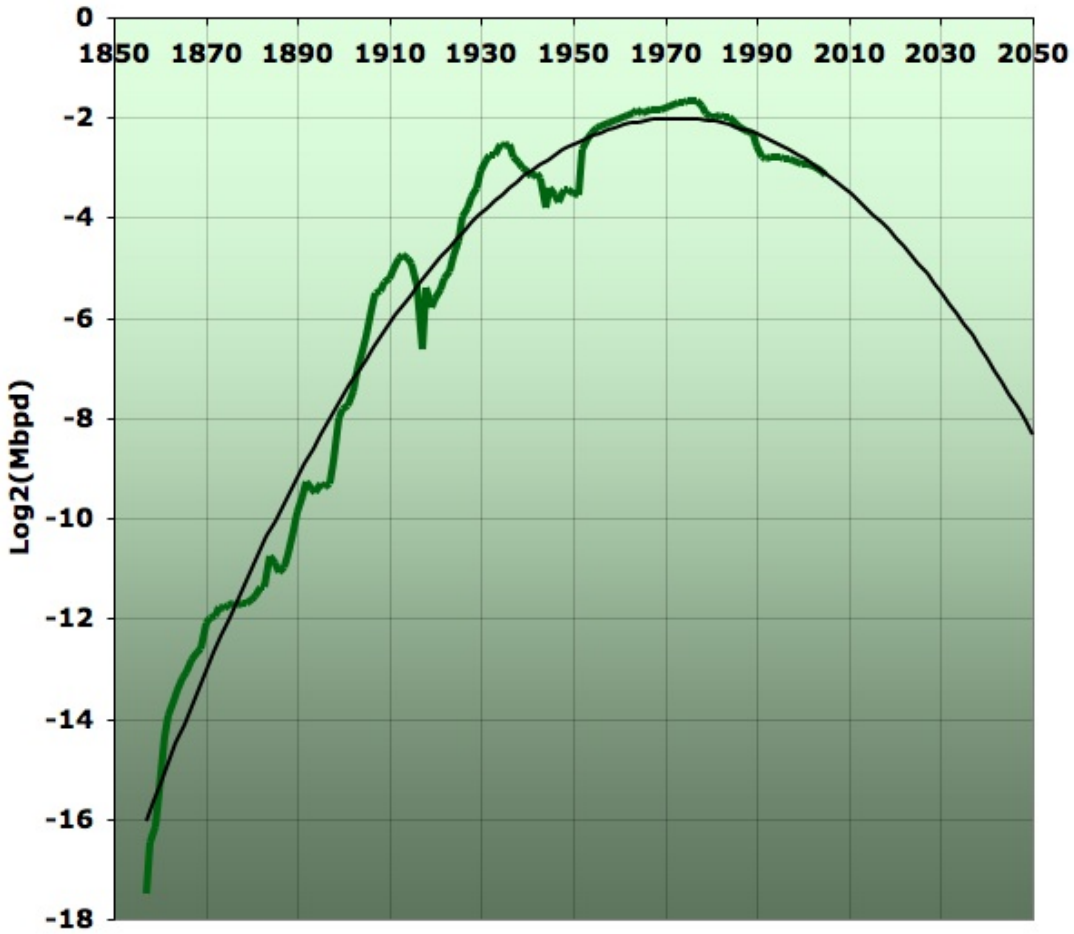

Log (base 2) of Romanian oil production versus time with a quadratic fit. Click for larger version. Source: American Petroleum Institute (courtesy J. Laharrere) for 1857-1958, and BP for 1965-2005 (includes NGLs). Production data for 1959-1964 is linearly interpolated between the two data sets.

So it seems like the quadratic does about as good a job of modeling this as any smooth curve is likely to do with such a choppy time series. But could it have predicted anything, or does the noise just destabilize things too much? Well, I didn't have time to do a thorough investigation, but I checked at the peak year of 1976, and the Gaussian knew very well it was at the peak:

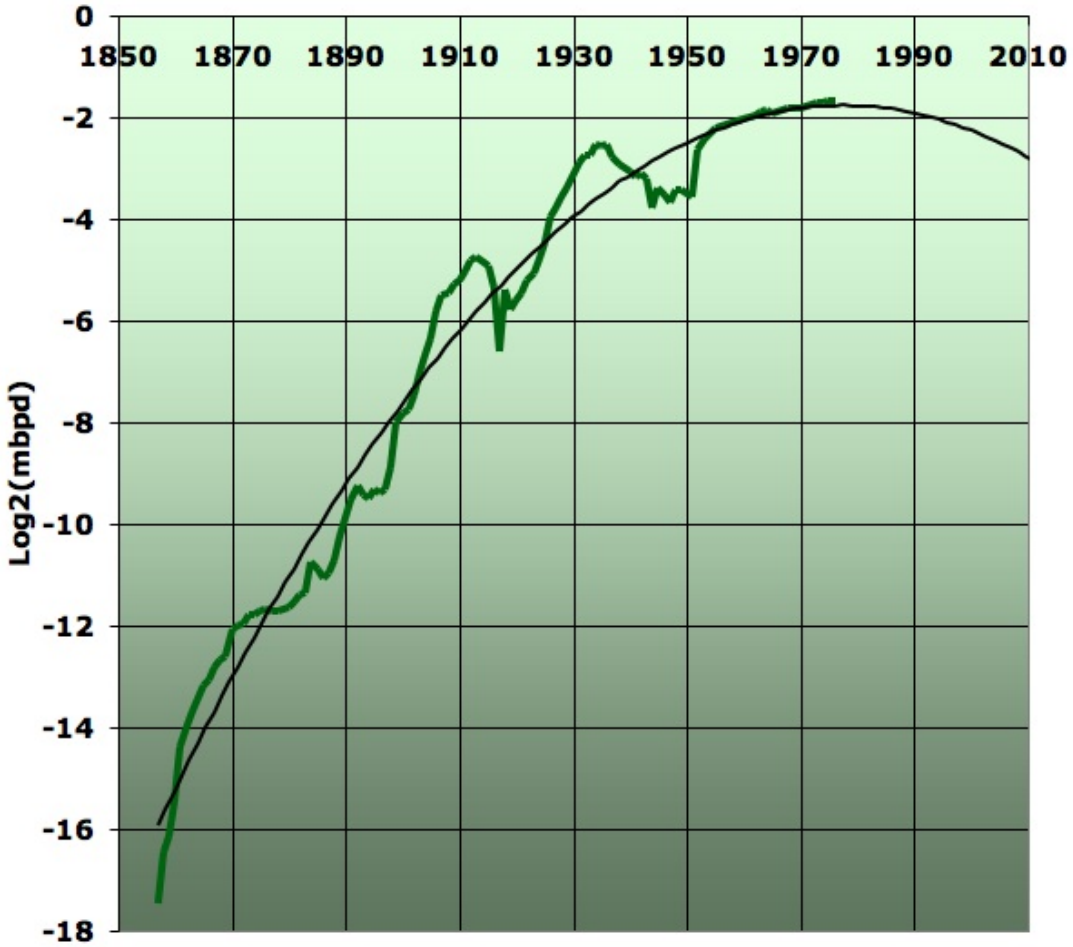

Log (base 2) of Romanian oil production versus time (through 1976 only) with a quadratic fit. Click for larger version. Source: American Petroleum Institute (courtesy J. Laharrere) for 1857-1958, and BP for 1965-2005 (includes NGLs). Production data for 1959-1964 is linearly interpolated between the two data sets.

But maybe it would think that at previous points too? How about the pre-WWII peak in 1936? Well, it wasn't fooled. It knew perfectly well that wasn't the real peak, though it had a very poor idea at that stage of when the peak would be:

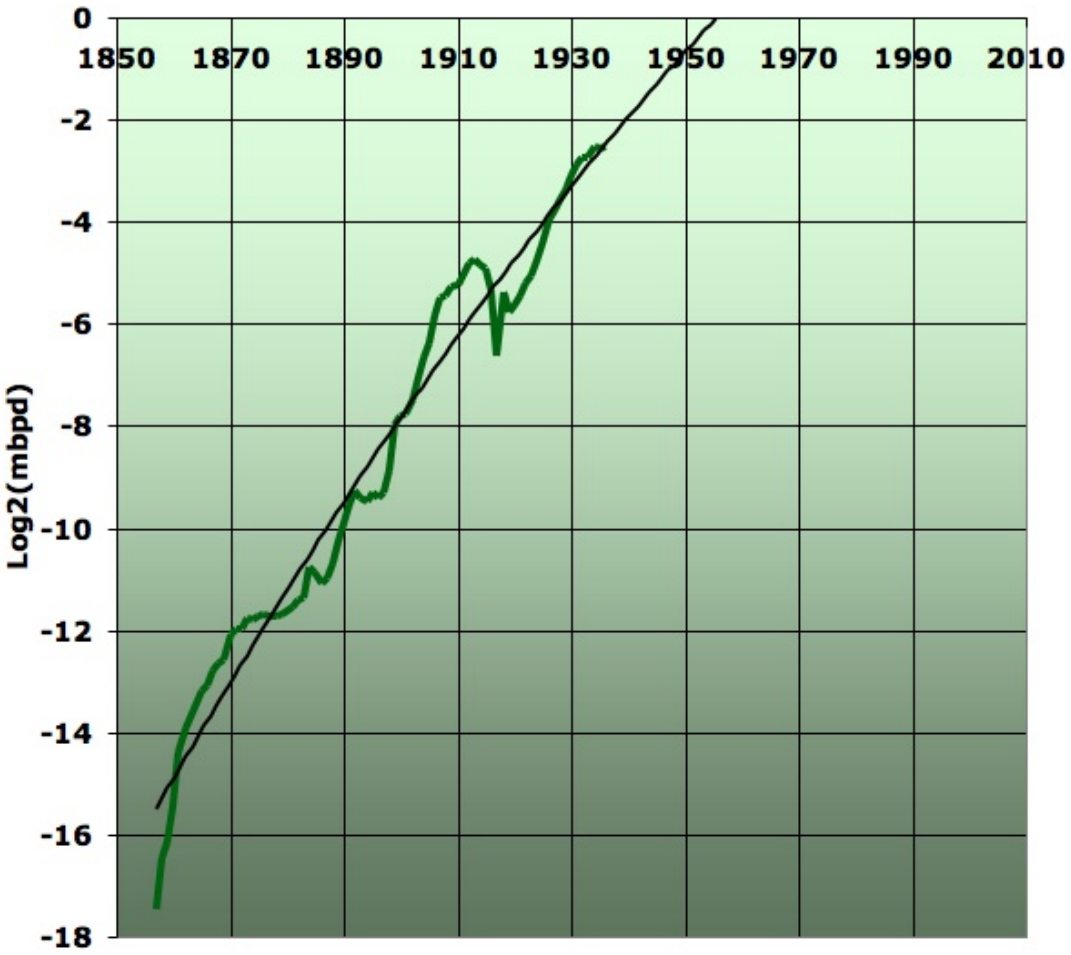

Log (base 2) of Romanian oil production versus time (through 1936 only) with a quadratic fit. Click for larger version. Source: American Petroleum Institute (courtesy J. Laharrere) for 1857-1958, and BP for 1965-2005 (includes NGLs). Production data for 1959-1964 is linearly interpolated between the two data sets.

Ok, so how about the linearization? Well, as Laharrere has shown before, this is about the best poster child there is:

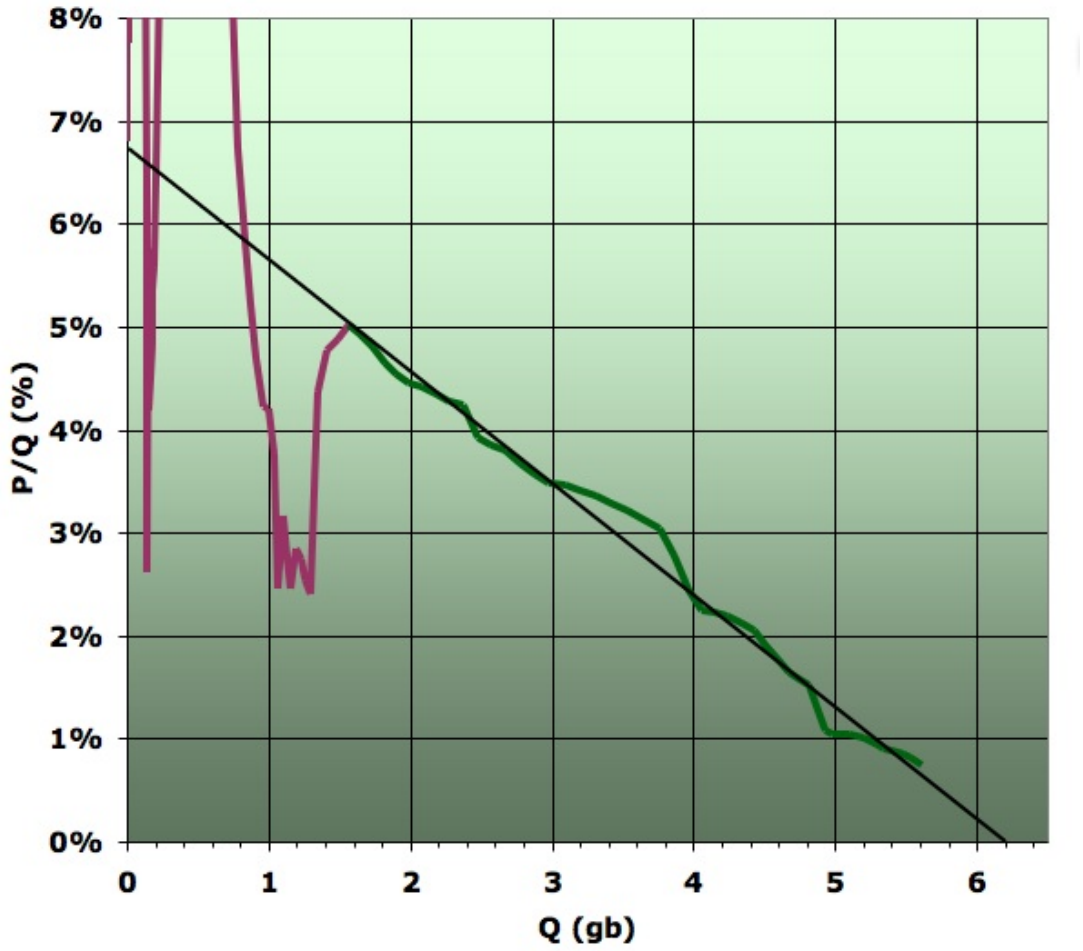

Linearization of Romanian production (the y-axis the ratio of annual production to cumulative production, x-axis is cumulative production. Linear fit is to the region for 1955 on (green data - plum data is prior to 1955). Click for larger version. Source: American Petroleum Institute (courtesy J. Laharrere) for 1857-1958, and BP for 1965-2005 (includes NGLs). Production data for 1959-1964 is linearly interpolated between the two data sets.

Not only would the linearization have given you about the right answer (or at least about the same answer it would give you today) at the peak in 1976, it would have been telling you pretty much the same answer from some point in the early 1960s on (I hedge a little since that falls in my missing data gap that I linearly interpolated). Also, this is the most mature country profile in the world, both because it started so early and also because the peak is narrower than the US peak (since K is around 6.7% versus the 5.6% for the US).

Next, I looked at the idea we explored the other day at the suggestion of WebHubbleTelescope of plotting P (annual production) against Q (cumulative production). Since the logistic differential equation is $P = dQ/dt = KQ(1-Q)$, a plot of P versus Q should be a quadratic if the production was exactly Hubbertian. Again, this is very roughly true in Romania, though the production profile is so noisy that the fit is not as pretty as in the US case:

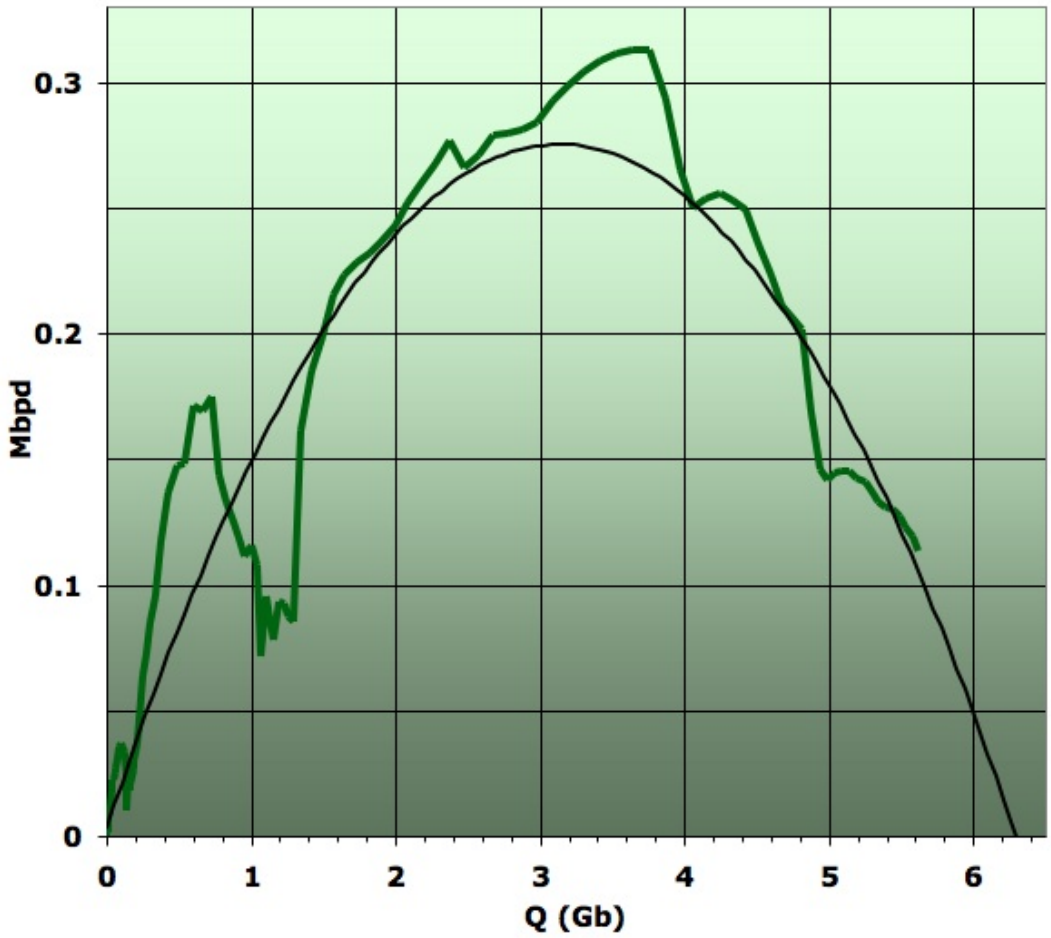

Romanian annual oil production (in million barrels/day) versus cumulative production (in billions of barrels) with a quadratic fit. Click for larger version. Source: American Petroleum Institute (courtesy J. Laharrere) for 1857-1958, and BP for 1965-2005 (includes NGLs). Production data for 1959-1964 is linearly interpolated between the two data sets.

However, it predicts the same URR as the linearization. But it turns out that back at the peak it would have given you a much less clear idea of what was going on than either the linearization or the Gaussian fit to logP:

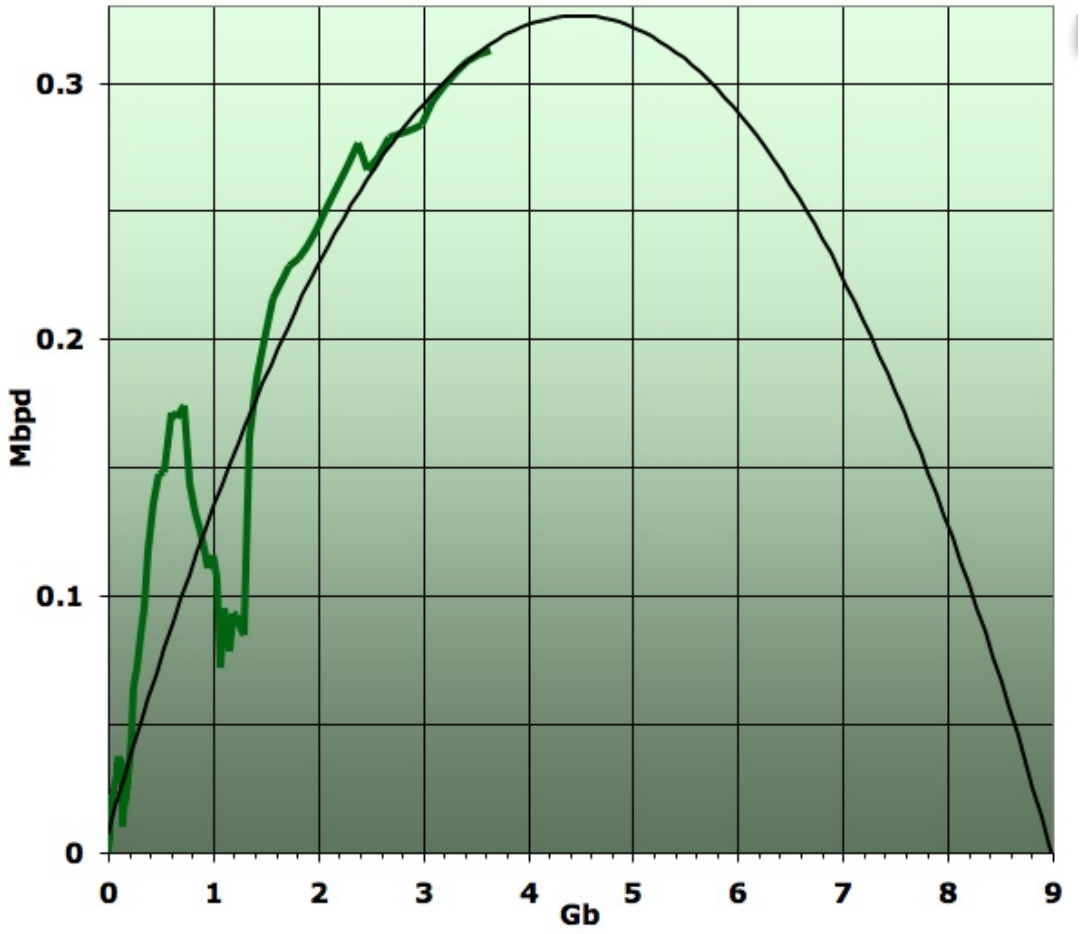

Romanian annual oil production (in million barrels/day) versus cumulative production (in billions of barrels), only through 1976, with a quadratic fit. Click for larger version. Source: American Petroleum Institute (courtesy J. Laharrere) for 1857-1958, and BP for 1965-2005 (includes NGLs). Production data for 1959-1964 is linearly interpolated between the two data sets.

It's 50% higher on the URR estimate than a modern estimate would be, and also thinks the peak is still a decade away when in fact it has arrived.

So in Romania, linearization works despite a very choppy profile, and again it seems to be because the profile is roughly Gaussian (decorated with major ugly noise) and the linearization technique serves as a useful robust predictor in that situation. However, directly fitting the Gaussian in logP versus t space works quite well here also. Fitting the differential equation in P versus Q space works eventually but is less stable to noise early on.

< warning type="self promotion">
I make my living consulting for a high enough rate that I can afford to support a family and nonetheless contribute these posts to The Oil Drum pro-bono. If you have a project that you think could benefit from my skills and you would like to support this arrangement, talk to me -- Stuart Staniford.
</warning>