The Oil Drum
DISCUSSIONS ABOUT ENERGY AND OUR FUTURE

Four US Linearizations

Posted by Stuart Staniford on January 6, 2006 - 8:13am
Topic: Supply/Production
Tags: hubbert peak, oil prices, peak oil [list all tags]

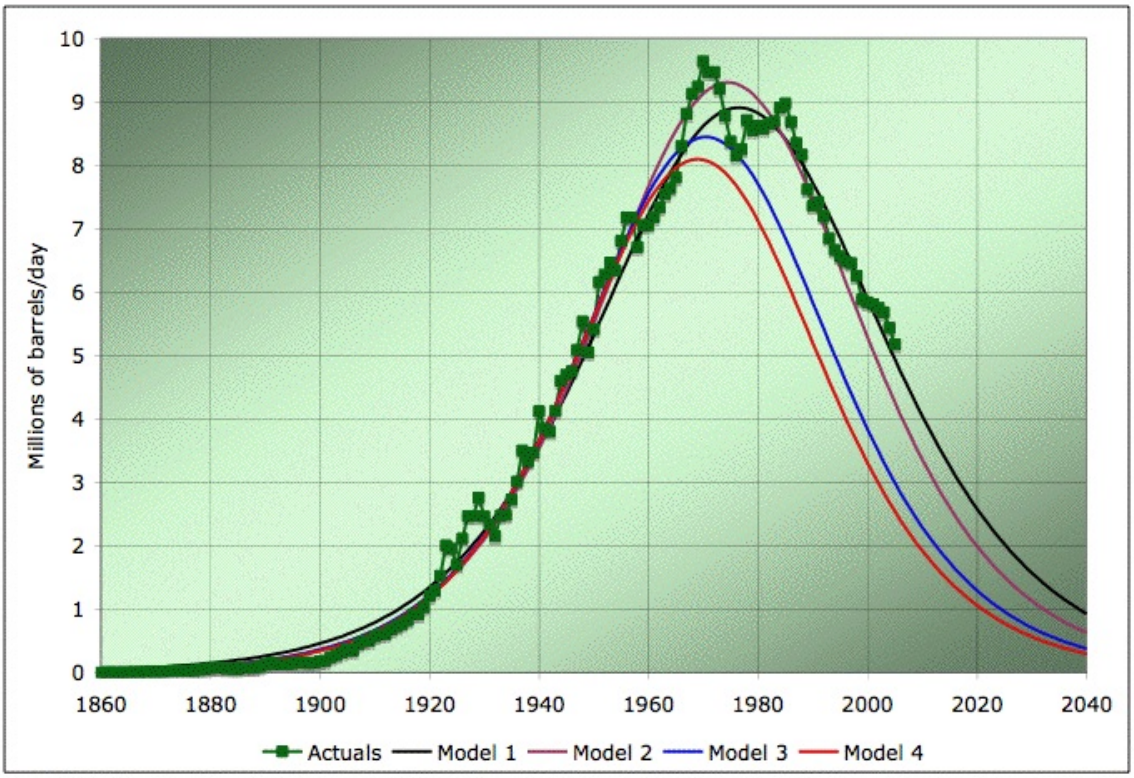

EIA Field production of crude, and four Hubbert models based on different linearizations. Source: EIA for the data, models as described in the text.

I'd like to begin an investigation into how much error we might expect to make by extrapolating Hubbert linearizations forward (for background on the method, try this post). It's going to take several posts, I think; there's a bit of work to do, and I'm not sure quite what schedule I'll be able to accomplish on it. Anyway, let's start gently with a look into the situation with the US, which is a popular target for linearization, because it's one of the most mature regions (though perhaps Romania is the best of all). Anyway, the EIA happily has US production data back to 1859, and to that I added one point for 2005, by summing the monthly data, and multiplying by 12/10. Note that this data includes Alaska - I prefer to work with whole countries.

To kick us off, I started with a regression of the data between 1958 (which is the point Deffeyes chooses) and today. That gives me this picture (Model 1):

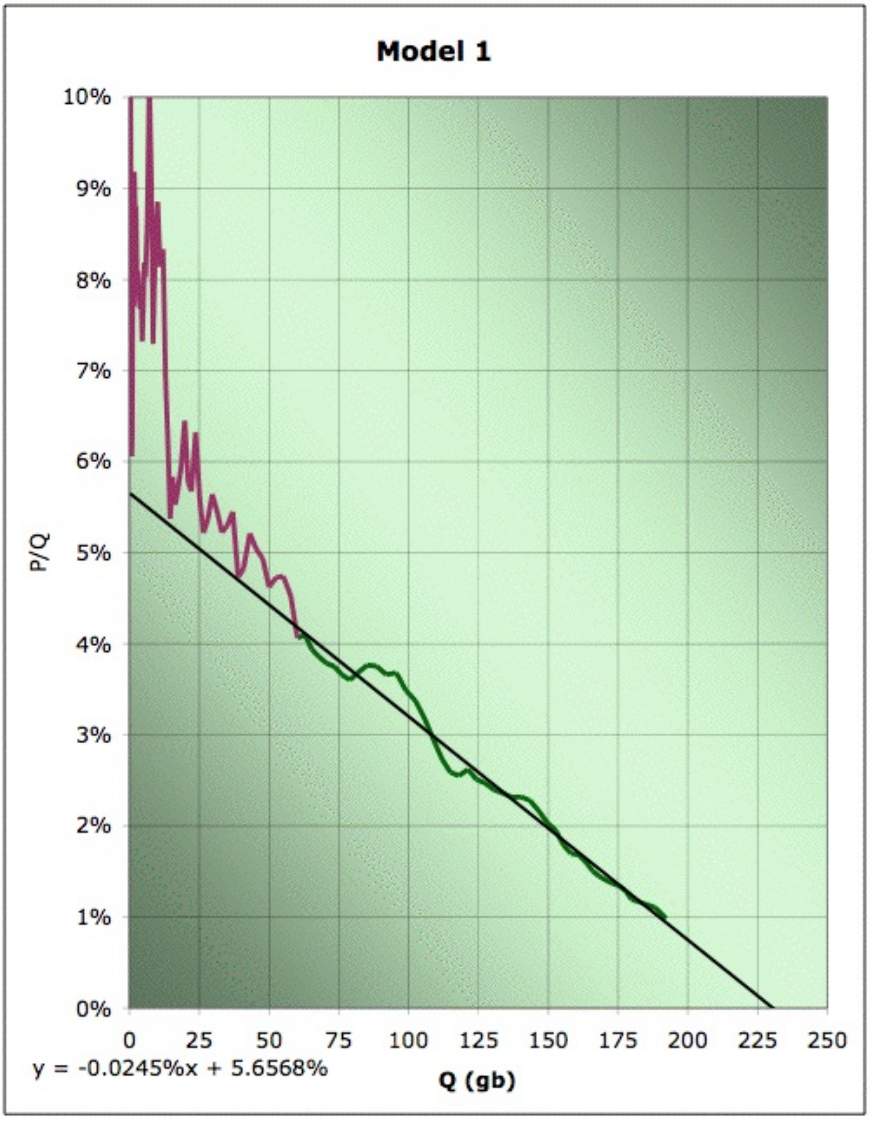

Linearization plot of EIA Field production of crude, together with linear fit for the data from 1958 to 2005. Green data is used for line fitting, plum data is not. Source: EIA for the data. 2005 data is extrapolated from data through October.

Clearly, we got a pretty good fit on recent data, and we ended up with a URR (ultimately recoverable resource) estimate of about 230gb, and a K of 5.7%. Remember that K is best thought of as the initial growth rate in the model, or equivalently, the final decline rate (with change rates gradually interpolating from one to the other over the history).

Now, recall that there's this slightly black art in the linearization process. They hardly ever fit well early on in the history, so one has to choose some point from which to start fitting. Obviously, choosing different points gives slightly different results. How sensitive are we to that? Well, if we picked 1930, which is about the earliest vaguely plausible place to try and start a straight line, we'd end up looking like this picture, which I'm calling Model 2:

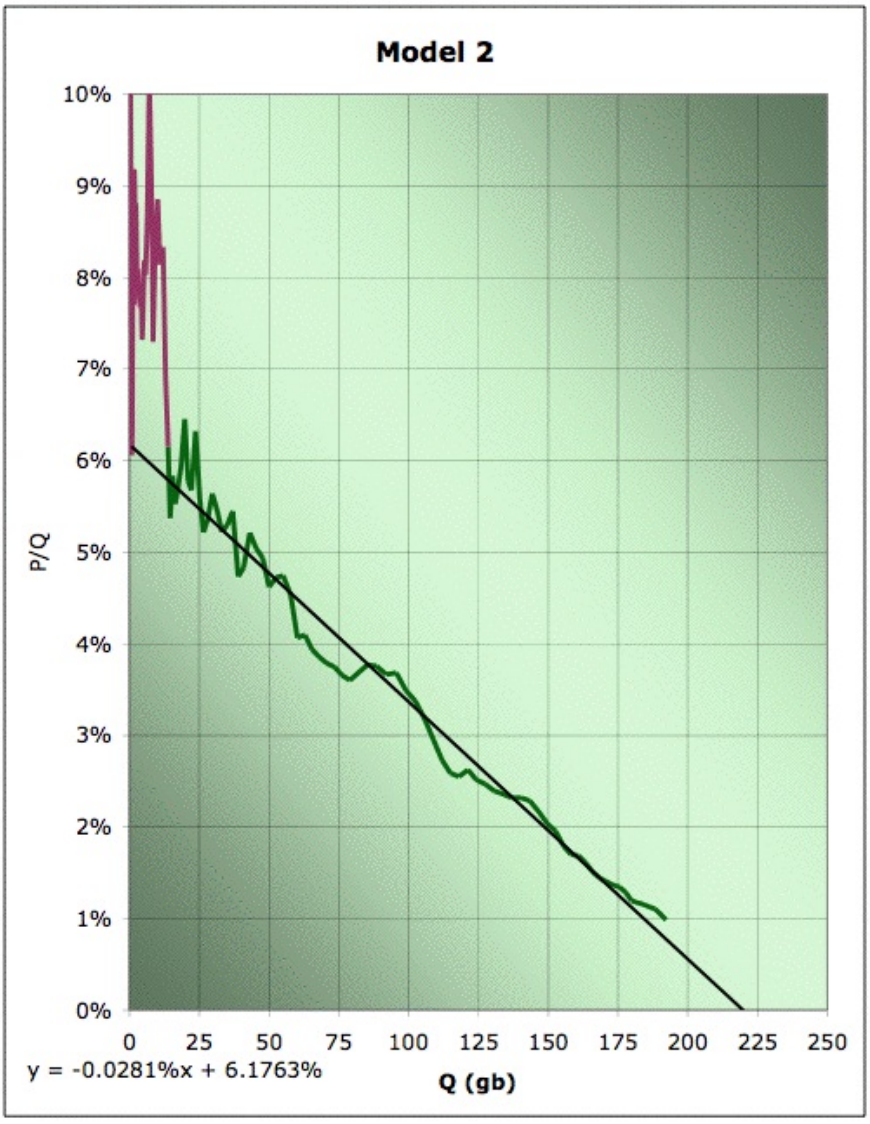

Linearization plot of EIA Field production of crude, together with linear fit for the data from 1930 to 2005. Green data is used for line fitting, plum data is not. Source: EIA for the data. 2005 data is extrapolated from data through October.

Clearly, this doesn't do *quite* as well on the recent data. Our URR has changed to about 220gb, and our K is now 6.2%. So the URR changed by around 5%, and K changed by around 10% of it's value. Those are probably decent rough estimates of our uncertainty in these quantities with the full history known to us now.

But how would we have done in the past? Based on model 1, the 50% point would have passed in about 1976. So let's try and do a linearization analysis if we only have data up to 1976 (but sticking with the 1930 starting point). In the world case, this is roughly analagous to trying to extrapolate linearizations today. That gets us Model 3:

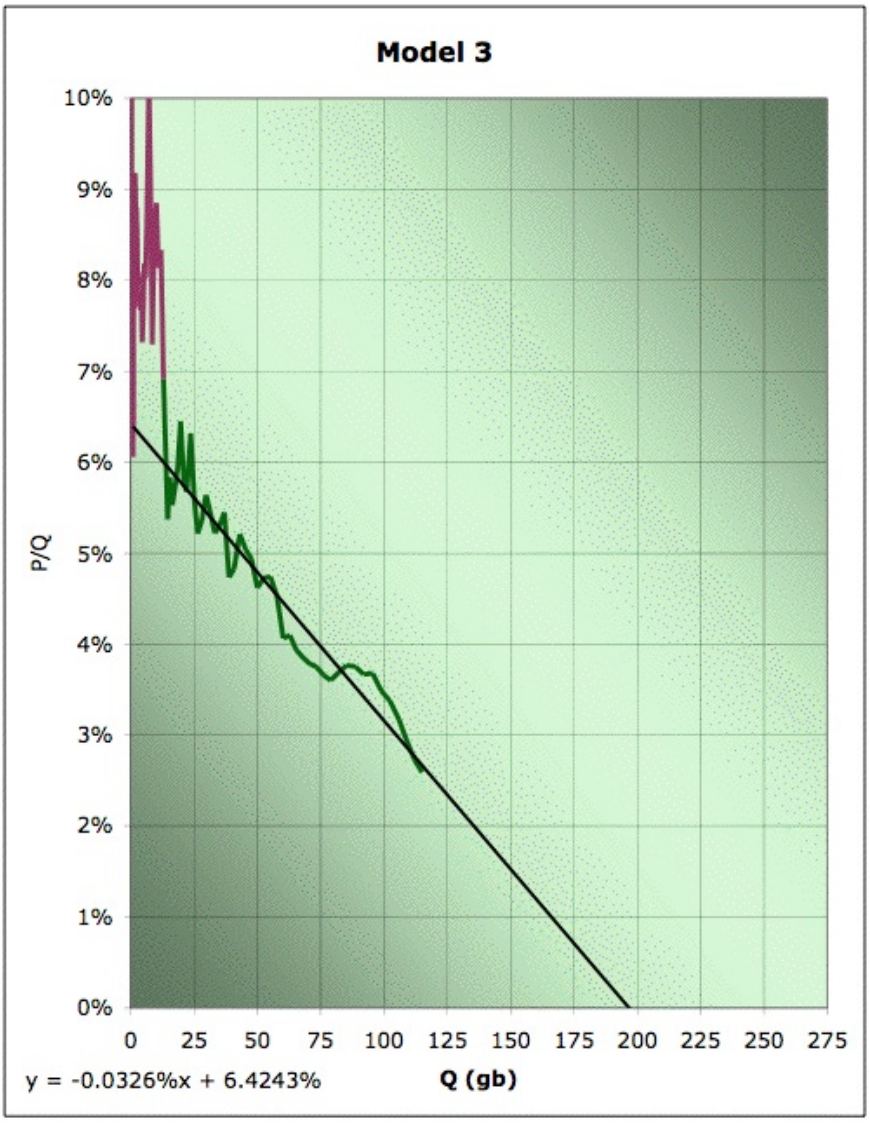

Linearization plot of EIA Field production of crude, together with linear fit for the data from 1930 to 1976. Green data is used for line fitting, plum data is not. Source: EIA for the data. 2005 data is extrapolated from data through October.

Clearly, we've dropped our URR down to about 190gb - off by a little more than 15% from a current estimate. Our K estimate is now 6.4% - about 11% higher than we would estimate today. Finally, let's suppose we were doing this back in 1955, and trying to scoop Hubbert. I would have come up with something like Model 4 here:

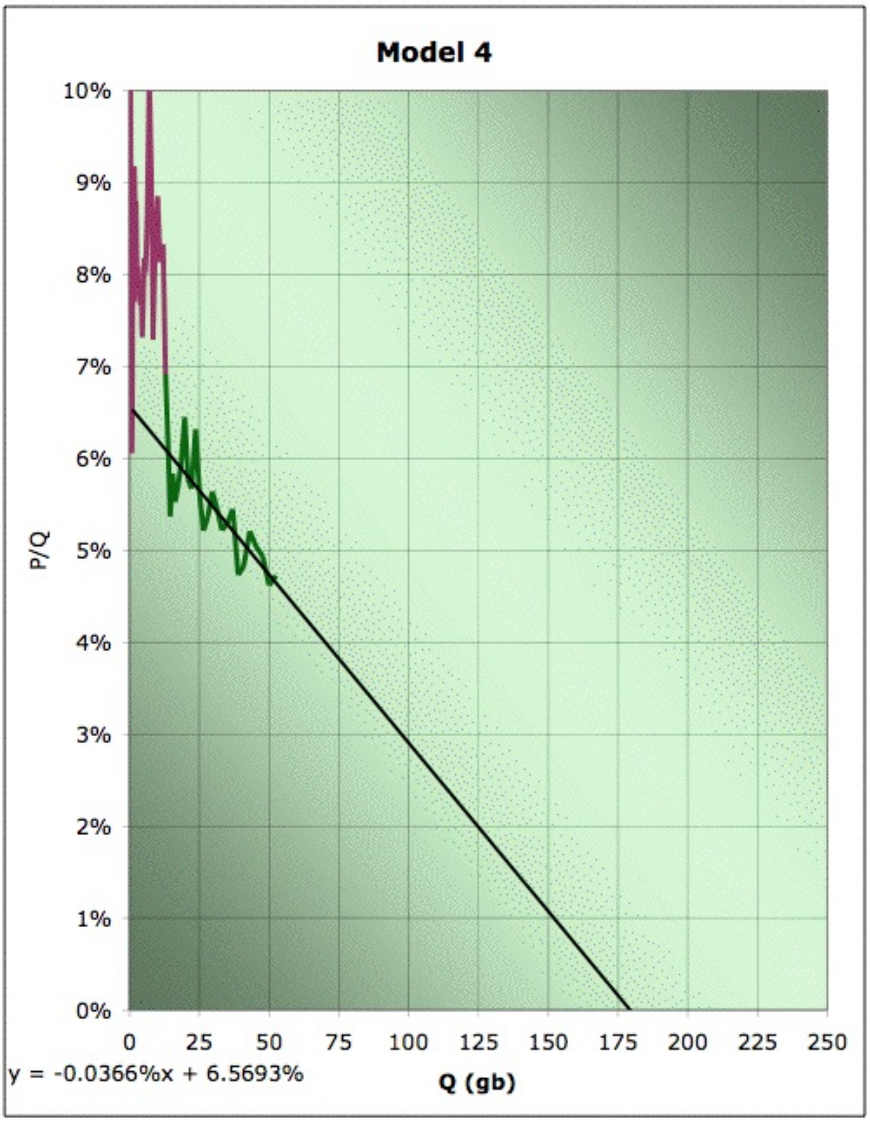

Linearization plot of EIA Field production of crude, together with linear fit for the data from 1930 to 1955. Green data is used for line fitting, plum data is not. Source: EIA for the data. 2005 data is extrapolated from data through October.

Our estimate of URR is now down to about 180gb - over 20% low. Our estimate of K has gone up to 6.6% (and truth be told, we're lucky to do that well - there's a lot of countries where we'd go hopelessly wrong trying to extrapolate from this early on in the history).

If we translate these four models back into production versus time graphs, we get the following picture:

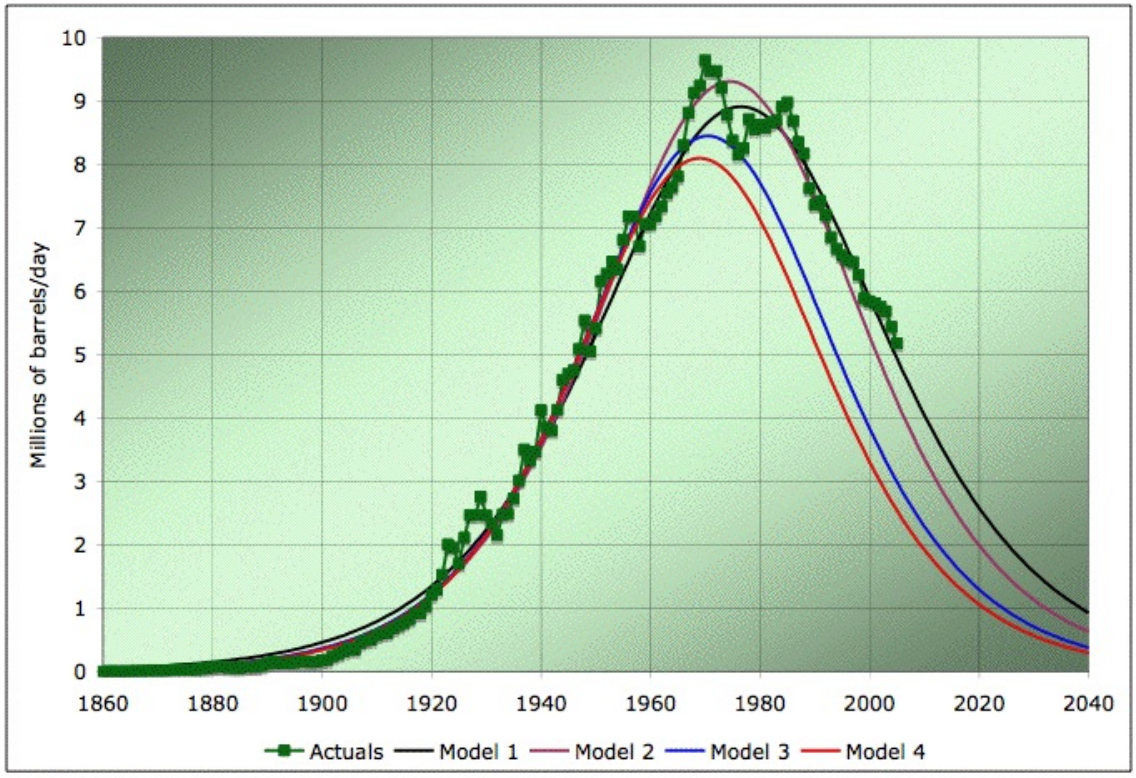

EIA Field production of crude, and four Hubbert models based on different linearizations. Source: EIA for the data, models as described in the text.

Again, Model 1 is our best current fit. It is indeed a remarkably good fit for such a very simple model applied to what would appear to be a very complex process (I confess that I still don't fully understand why this model works as well as it seems to in practice). Model 2 is the one where we started fitting in 1930 instead of 1958. It doesn't work quite as well, but it's not bad either.

Model 3 is if we only have data to 1976. You can see pretty clearly where it goes wrong: it misses the Alaska North Slope production. Of course, if we had really been doing this in 1976, we would have **known** that there was a super-giant field about to come on stream much later than the others and change the picture. And of course model 4 is a little worse again - that's the one where we only have the history to 1955.

The conclusion in the US seems to be this: doing the linearization in the early days causes you to underestimate URR and overestimate K. However, it's by no means useless. The errors even in 1956 are much less than a factor of 2 - it gets us in the right ballpark. And I think it's particularly striking that even when you get the URR somewhat wrong, the shape of the downslope is still about right. Very early on in the history, you could tell with about a 10% relative error what the post peak decline rates were going to be. That's remarkable.

What I'd like to do next is write a little program to more systematically estimate the uncertainties in the linearization. I think we can also make some improvement by underweighting the early data and overweighting the later part of the history. Then I want to get into people's concerns about whether the world is fundamentally different enough from individual countries that we cannot safely use the same process. But those will have to wait for future posts.