The Oil Drum
DISCUSSIONS ABOUT ENERGY AND OUR FUTURE

How fast should you boil a frog?

Posted by Stuart Staniford on February 2, 2006 - 11:57am
Topic: Environment/Sustainability
Tags: biofuel, carbon cycle, climate change, global warming, hubbert peak, kyoto, oil prices, peak oil [list all tags]

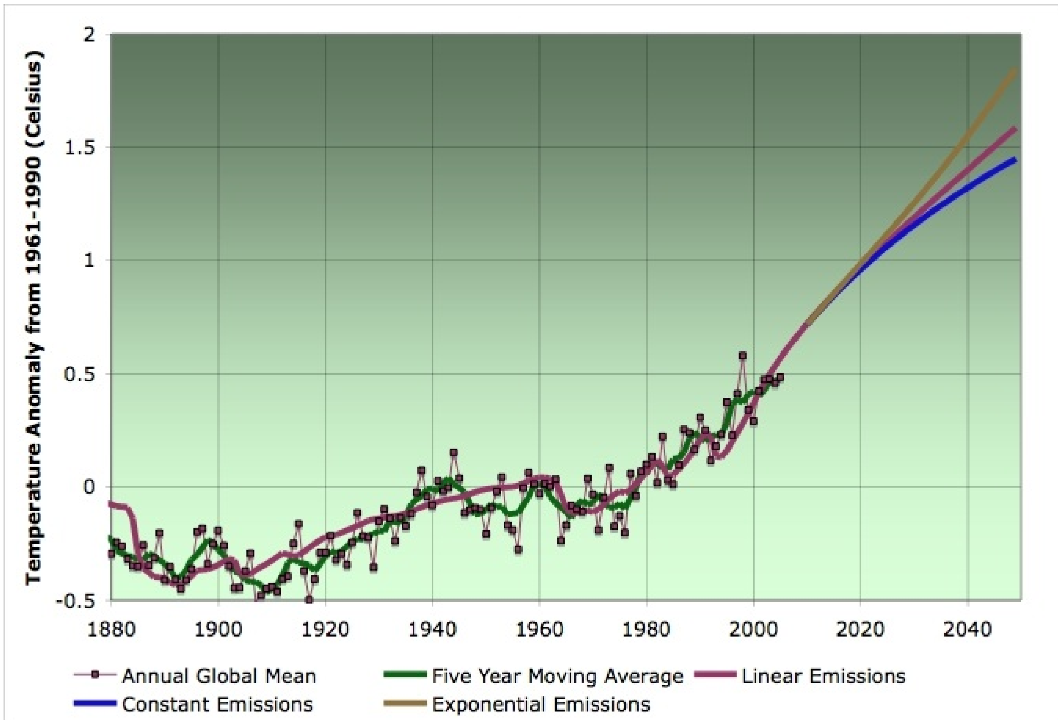

Global average temperature 1880-2005, together with one dimensional model fit (as described in the text) extrapolated to 2050 for the case of linear, constant, and exponential carbon emissions (all other forcings held constant after 2003). Data: UEA CRU.

This is the second post in a series on carbon in the economy. The first post was The Carbon Economy. I was going to do one post on implications of carbon emissions, but found myself obliged to split it into two, with this one being on translating CO_2 into temperature, and the second being on risks of that temperature. Again, I'm going to try and focus on developing rules of thumb that are decent enough approximations to be usable in reasoning, but simple enough to be comprehensible to non-climatologists.

The very basic physics of the greenhouse effect is probably familiar to most people, but let me take a second to review it for the rest of our readers. The sun, being very hot, emits enormous amounts of electromagnetic radiation. The spectrum of that radiation peaks in the range that our eyes can detect as visible radiation. The earth's atmosphere is mostly transparent to the sun's radiation, and so most of it reaches the ground and the parts that aren't reflected back upwards are absorbed and heat the earth.

Because the earth is much warmer than space, it also radiates electromagnetic radiation out in all directions. However, since the earth is much less warm than the sun, it emits much less energetic (longer wavelength) infrared radiation. If the earth's atmosphere were dry and free from greenhouse gases, it would be transparent and the earth would be quite a lot colder (permanently below freezing everywhere). Instead, certain gases absorb some of the infrared radiation, which warms the air more than it otherwise would be, and keeps the planet warmer.

The infrared absorption of the atmosphere is shown in the following figure which shows how much of the infrared gets through as a function of the wavelength (the blue areas under the curve are where radiation get through, in proportion to the height of curve). Also shown are labels for some of the notches where either CO_2 or water molecules are responsible for the absorption.

Transmittance of infrared radiation as a function of temperature. Source: Wikipedia.

So in general, if there is more CO_2, then it is harder for infrared radiation to get out of the atmosphere. That means the earth cannot shed heat as effectively, and thus will warm up. As it warms up, it will release more radiation. Eventually, it will get enough warmer that the outgoing radiation will balance the incoming sunlight (referred to as radiative equilibrium). Normally the earth is in radiative equilibrium to a very good approximation, but that is less true at the moment (as we shall see).

So the key question, in light of our explorations in CO_2 concentration the other day, is how much temperature rise do we get from any given increase in CO_2. This is known as the *climate sensitivity*, and is often expressed as what would be the effect of doubling CO_2 from the preindustrial value of 280ppm, but we will follow a newer and better convention.

This is a very complicated business because there are a lot of feedbacks in the system. Eg, if there is more CO_2 and things start to warm up, that means the warmer oceans will release more water vapor, which also blocks infrared, amplifying the effect of the CO_2. However, some of the water vapor might create more clouds, which reflect sunlight away from the earth and tend to reduce the temperature increase. On the third hand, the increased temperatures will melt some of the earth's ice and snow in some places, which would have reflected sunlight, but now won't, and so that enhances the temperature. On the fourth hand, the increased water vapor might cause more snowstorms and lead to *increased* accumulations of ice and snow in some places. Etc, etc. As you can imagine, this is a computer modeler's paradise and the climatologists have engaged in that big time.

Historically, the climate models didn't agree as closely as one might like. For example, here's a set of temperature projections for one particular CO_2 scenario (the A2 scenario used by the Intergovernmental Panel on Climate Change (IPCC), which lies between our linear and exponential emissions extrapolations from the other day).

Temperature increase over 2000 for eight global climate models in IPCC SRES A2 emissions scenario. Source: Wikipedia.

As you can see, the uncertainty is considerable. Indeed the IPCC, which is an international organization of government nominated scientists who write consensus reports about climate change, quotes the climate sensitivity to a doubling of CO_2 as 3.5 ± 0.9 ^oC, which is 6.3 ± 1.7 ^oF. Those are the one standard deviation error bars. This is a lot of uncertainty to add to the already considerable uncertainty that economics and resource constraints introduce into the emissions scenarios (though, as you can see from the figure, there's less uncertainty about the near future than the far future).

Another way to get at the situation is to look to the past. The climatologists have expended extraordinary levels of effort to drill sediments out of lake and ocean bottoms and analyze isotope ratios in them, examine tree rings, drill ice cores in ice sheets all over the world, etc, etc. Then they take all these time series, statistically munge them together, and try to come up with an estimate of global temperature. This graph is the result of a range of such exercises. The black line at right is special however - it depends on actual temperature records since regular meteorological observation began.

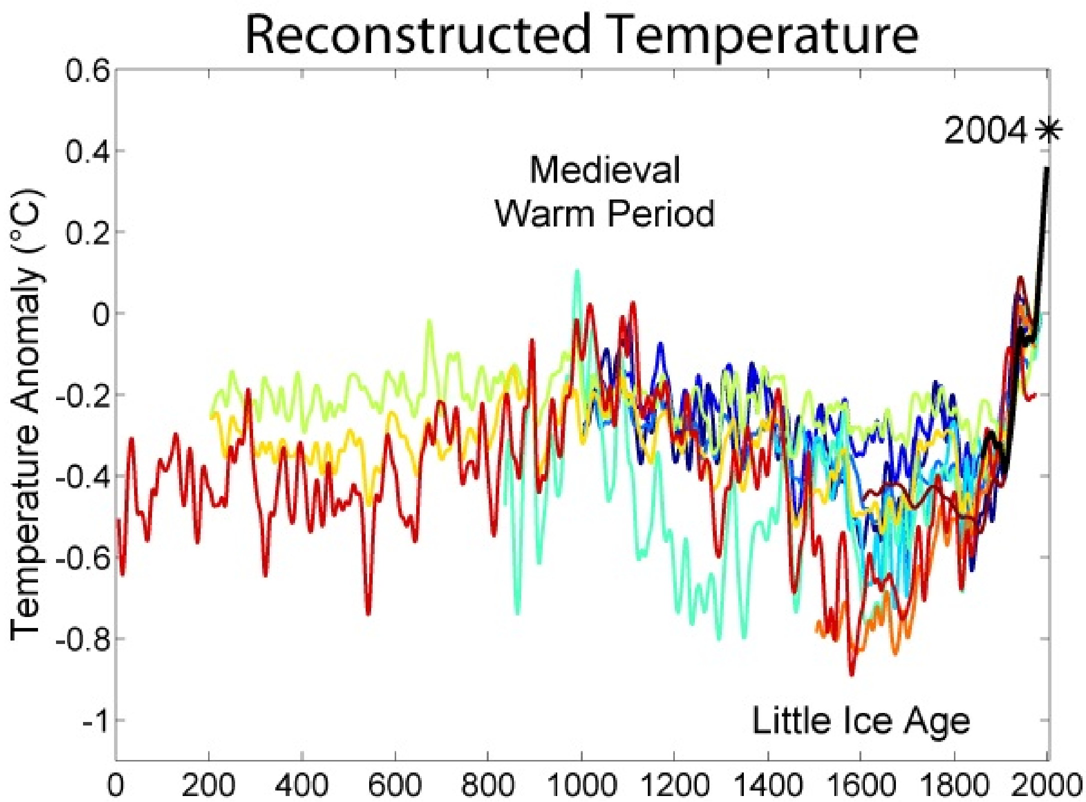

Ten different reconstructions of global temperature anomaly relative to the 1961-1990 temperature. Source and detailed references: Wikipedia.

Clearly, we run into significant uncertainty again. All of these efforts show recent temperatures being unprecedented in the last 2000 years. However, they differ considerably in how much natural variation there has been before (which of course must be considered in deciding how much of the recent temperature rise can be ascribed to greenhouse gases, and therefore used in estimating the climate sensitivity).

Let's take a look at that instrumental record (the global average data are courtesy of the University of East Anglia Climatic Research Unit who regularly update their time sequence).

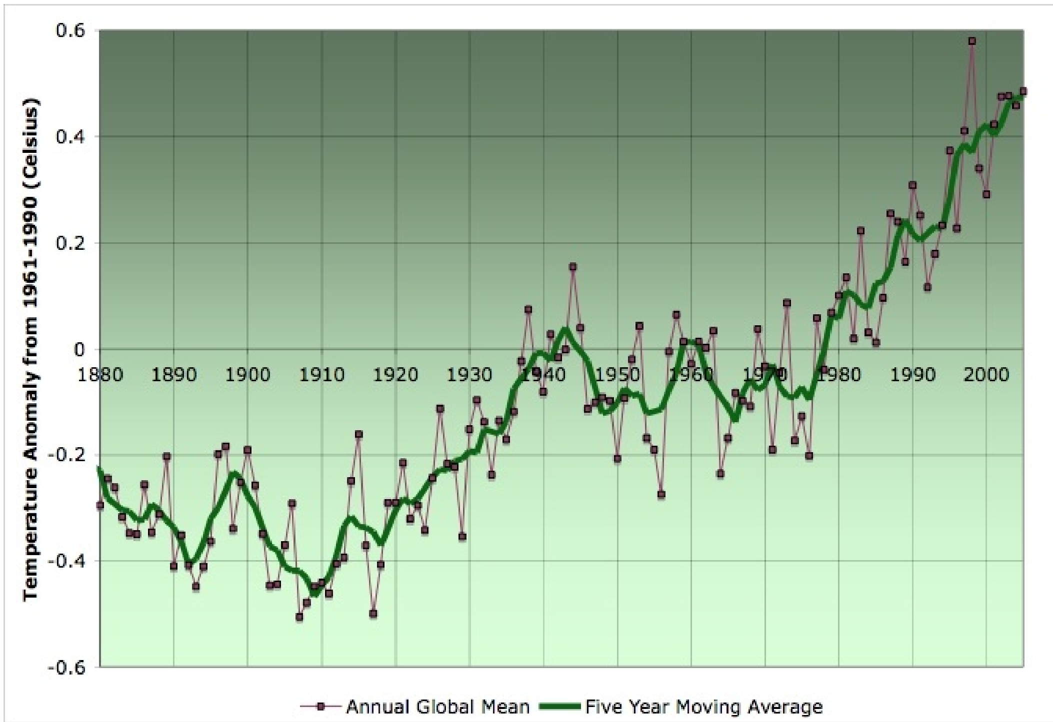

Global average temperature 1880-2005, together with five year moving average. Source: [UEA CRU](#).

Roughly speaking, you can see a cooling trend in the late nineteenth century, then a sharp rise in temperature from 1910-1940, approximate stability in the middle twentieth century, and then rapid rises after the late 1970s. Clearly, there is more going on in this than just CO_2. To get at what is happening, here's a figure from a paper of James Hansen *et al.* which is under submission at the moment and represents a huge array of modeling efforts with the GISS global climate model. The newest models are starting to get quite impressive - probably mainly because of better inputs.

In particular, this next picture shows estimates of all "climate forcings" that have changed from 1880. Climatologists classify all the external factors that affect the climate according to the equivalent change in sunlight at the top of the atmosphere that would have been needed to produce that change. The solar constant is around 1366 Watts/square meter (if the square meter in question faces the sun directly), so the variations of a few Watts/square meter are against that number. By putting all the different effects into a common framework, it's easier to compare their contributions to climate change.

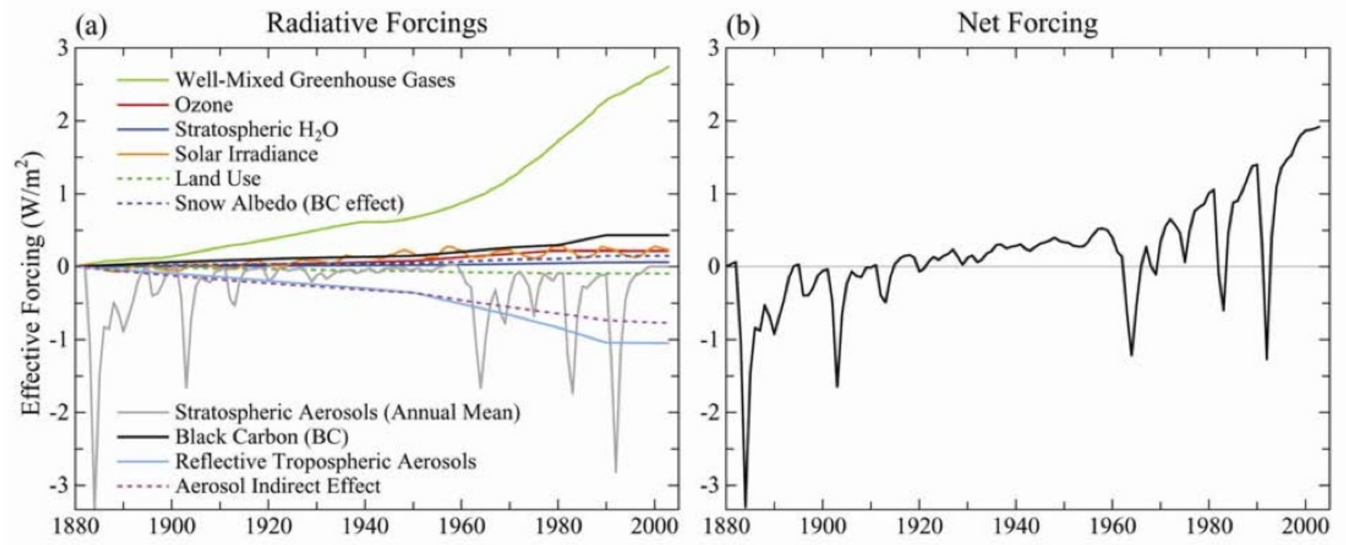

Estimates of main climate forcings over time (left), and net forcing (right). Source: [Hansen et al](#).

You can see that the most important effects heating the planet up are the global warming gases (reaching around 2.7 W/m^2), with black carbon (soot) in the atmosphere being a distant second. The most important compensating factors are aerosols (especially sulphates), which have both a direct effect reflecting sunlight, and also an indirect effect seeding clouds. Stratospheric aerosols (mainly from very large volcanic eruptions and early atmospheric nuclear tests) also play an important if intermittent role.

We can see in Hansen et al's figure much of the explanation for the temperature variation. The cooling in the late nineteenth century is primarily due to volcanic eruptions, especially Krakatoa. The warming from 1910-1940 looks explicable as the earth warms up after that (it takes decades for the surface ocean to warm, and centuries for the deep ocean to equilibriate). Then the stable period in the mid twentieth century is due to the fact that global warming gases only slightly outweighed the cooling effects of tropospheric aerosol pollution. In the later part of the century, this breaks down as clean air measures are increasingly taken allowing global warming gases to dominate despite the best efforts of the nuclear weapons community and an assortment of volcanoes (led by Mt Pinatubo).

When Hansen et al. run these forcings through their model, they get this level of agreement:

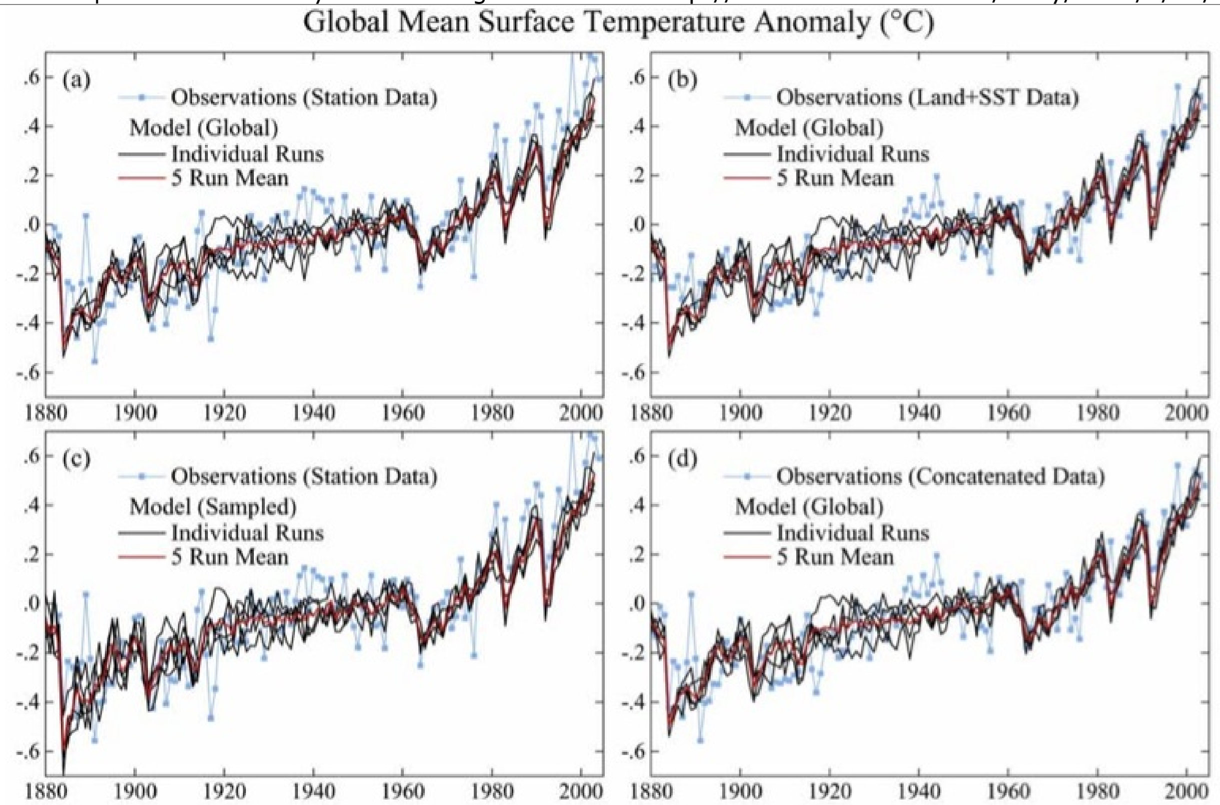

GISS model temperature versus observed temperature for four different schemes of global averaging. Source: Hansen et al.

Pretty good. However, they emphasize that the leading source of uncertainty in global warming is the forcing of the aerosols (both direct and undirect), which has a relative uncertainty of about 50% (ie huge), and the growth rate is so uncertain they can't even say if it's growing or shrinking. Hence considerable uncertainty propagates through into their estimate of the climate's sensitivity to greenhouses gases. They state they probably could produce equally good model fit with more climate sensitivity and less net forcing, or less climate sensitivity and more forcing. Thus the problem of uncertainty in the climate sensitivity persists, which makes extrapolating temperatures more difficult. Indeed, they estimate elsewhere based on paleoclimatalogical evidence that climate sensitivity is 0.75 ± 0.25 $^{o}C/W/m^{2}$ (1.35 ± 0.45 $^{o}F/W/m^{2}$). That is to say, if you increase the total climate forcing by 1 W/m^{2}, the world will eventually get 0.75 degrees Celsius warmer, except you could be wrong either way by 0.25 degrees Celsius (one standard deviation error bar).

However, you don't get this warming all at once. In their model, you get 50% of it in 25 years, 75% in 100 years, and it takes several hundred years to get to full equilibrium. That's the ocean; it's slow to warm up and it covers most of the planet.

Let's now dig into what is inside that greenhouse gas forcing (just the green line from two figures up). Here are the main constituents:

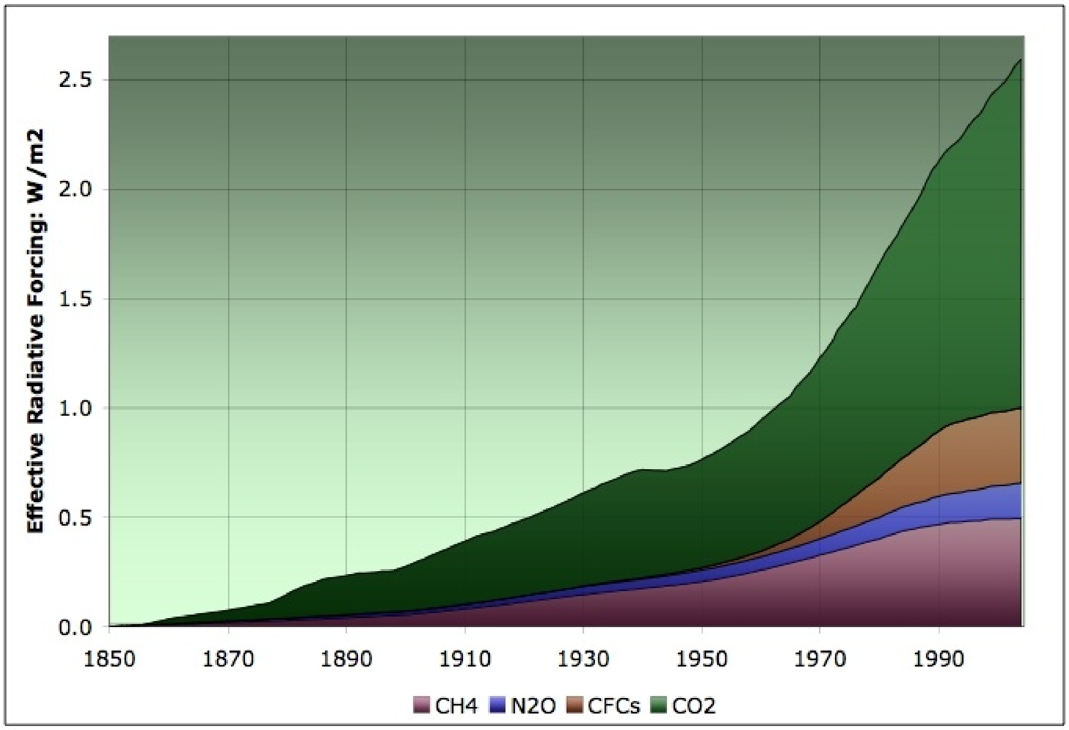

Excess radiative forcing over 1850 level due to various greenhouse gases. Source: [NASA, Goddard] for mixing ratios and Table 1 of [Hansen and Sato, 2000] for conversion formulae.

There are two important things I see here. The first is that the non-CO_2 components are a pretty important part of the total. However, they appear to be collectively stabilizing, whereas CO_2 is continuing to go through the roof. The reason for the methane emissions stabilizing doesn't seem to be very well understood and is in direct contradiction of the fears of permafrost methane release. The reason for the stabilization of CFC concentrations is the success of the Montreal protocol and its various follow-ons. However, these compounds are very long lived in the stratosphere, and there are still some emissions in developing countries, so there will not be a rapid decline in their concentration.

You might wonder why we don't have water in that graph, given that H_2O is an important and effective greenhouse gas. The reason is that water exchanges very fast with the ground/ocean, so there's no real memory in the water concentration (with a partial exception for the stratosphere). It varies rapidly in response to the weather. Thus it becomes one of many feedback loops that are all folded into the climate sensitivity, rather than being viewed as a long-lived greenhouse gas.

Ok. Now I'm going to do something that will make any real climatologists scream, but I'll argue it's defensible. Here's the thing. I need a model of the way carbon dioxide translates into temperature that's blog friendly. Later in this series, I want to be able to play with emissions scenarios, and see what they mean for temperature, and in the space of a blog post, where I might have 5-10 hours to work on it, I don't want to have to download, learn, and run a global climate model that probably needs a supercomputer to run on anyway (yep - you can get a real climate model and run it at home). I want something simple and lightweight enough to be usable in this format. But correct enough that it will keep us within the (significant) uncertainties of the problem. Remember, we have 33% uncertainty in the climate sensitivity overall, and the model spread in the 2001 IPCC report was even larger than that.

So here's my braindead blog-friendly one dimensional climate model, which turns out to work amazingly well. The idea is that we have an atmosphere coupled to a shallow ocean coupled to a deep ocean. We apply a forcing. The shallow ocean slowly starts to heat up. Newton's law of cooling suggests the heat flow should be proportional to the difference between the actual temperature of the ocean and the temperature that would be in equilibrium with the heat source (the sun as filtered through the atmosphere). That would imply that the shallow ocean would approach the new equilibrium temperature exponentially with a timescale of decades. But there's a complication. The shallow ocean is exchanging heat with a much larger deep ocean, which cools it and makes it take longer to reach equilibrium. So a simple model of that situation is that we take a sum of two exponentials - one with a short (decades) lifespan, and one with a long (centuries) lifespan, that together, when they asymptote, will give the full temperature rise specified by the climate sensitivity.

(In mathematical language, which you can safely ignore if it doesn't make sense to you, I'm going to convolute the forcing function with a sum of two $(1-e^{-t/l})$ type terms, which together add up to the climate sensitivity at their asymptote. One has a long life and one a short life.)

So, to calibrate this braindead model, I took a model run for GISS that only involved the greenhouse gas forcings since 1880 (ie they turned off volcanoes, aerosols, etc). Then I adjust the parameters to match it, given a climate sensitivity of 0.75 $^oC/W/m^2$. It turned out to work best to put 60% of the heat into a 15 year lifetime process, and the rest into a 500 year process (but the latter lifetime is not well constrained by the data as long as it's much greater than the first).

In the next picture, I plot this model, based on the greenhouse gas forcings above, and superimposed on the GISS model plot of the same thing. You can see the braindead model cannot quite match the GISS line - there was no parametrization that would bend quite that sharply. However, it's pretty close, and certainly it's well within the bounds of the uncertainty in the climate sensitivity. The upper and lower lines are what the same model with the same parameters does if you just change the climate sensitivity up or down by 0.25 $^oC/W/m^2$

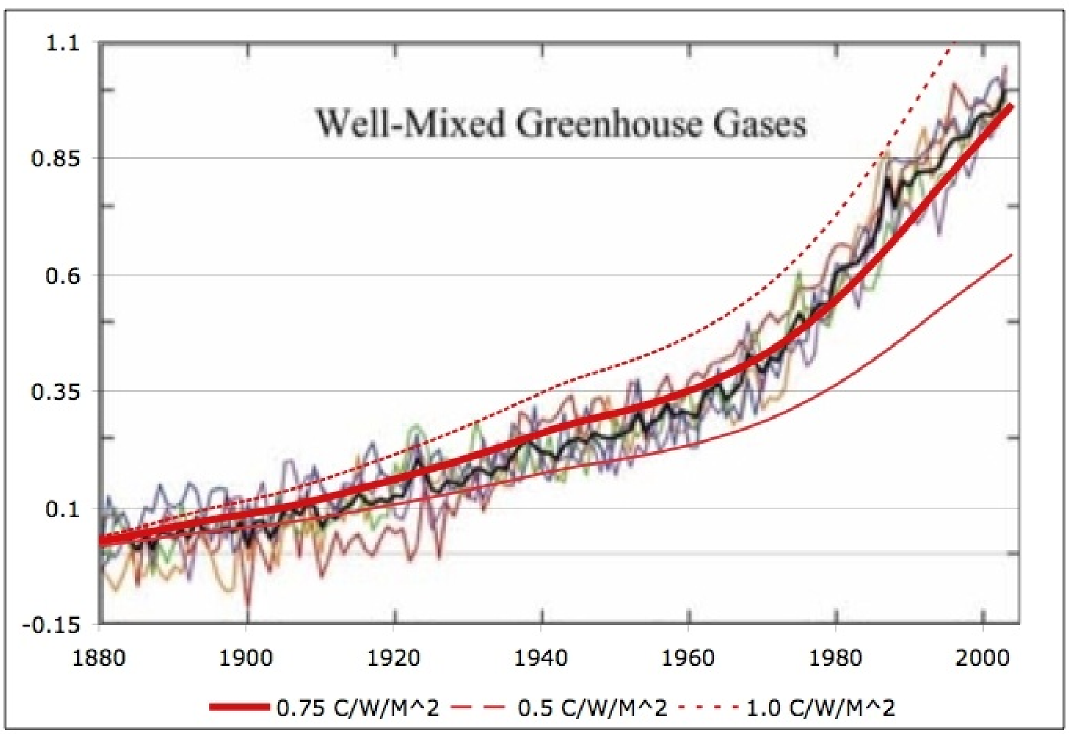

Fit of simple one-dimensional climate model applied to greenhouse gas forcings only, as compared to GISS model run in Hansen et al (top right panel of Figure 10). Center red line is model for climate sensitivity of 0.75

$^0C/W/m^2$, while thinner lines above and below represent what the model would produce for climate sensitivities one standard deviation above and below the estimate. Simple model agrees with GISS to within the uncertainties in the problem.

Ok. So I'm not claiming this is a **great** climatological model. However, I claim it's an adequate one - it will get us by given how large the other uncertainties are in the forcings and how the climate responds to them, as long as we don't try to extrapolate too far into the future, outside the domain it's calibrated on. And it's blog-post friendly.

If we now turn to the full net forcing trace (which I acquired the hard way), my 1D model does this, when compared to the UAE global temperature data:

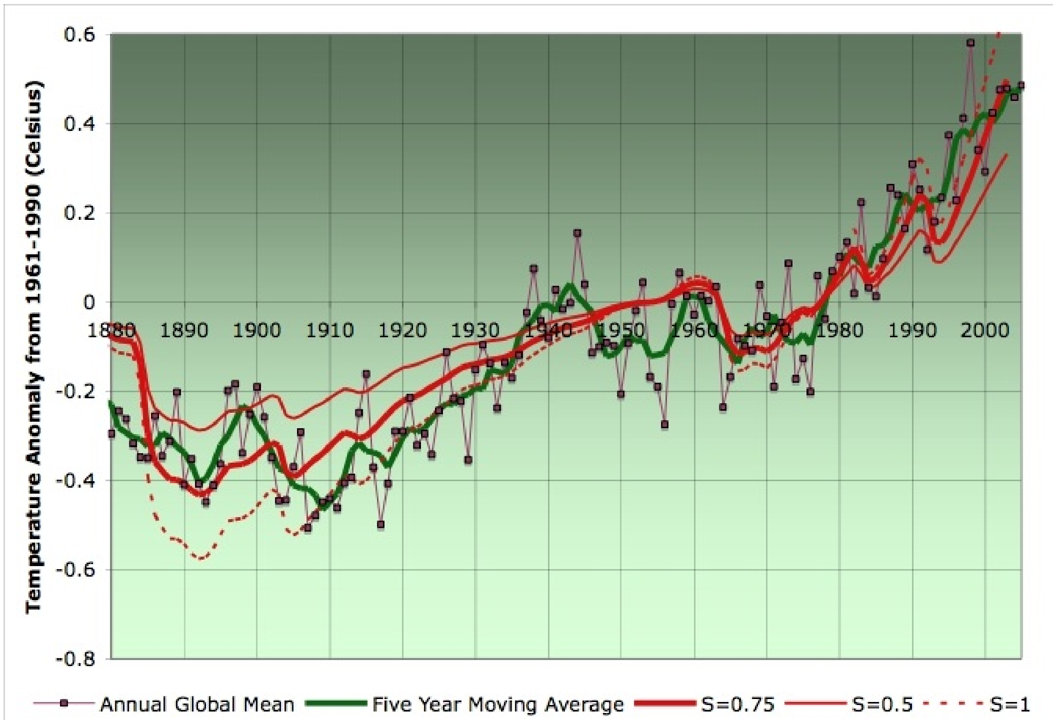

Global average temperature 1880-2005, together with five year moving average, and 1D model fits for climate sensitivity of 0.75 ± 0.25 $^0C/W/m^2$. Source: UEA CRU.

Not bad for a Model with Very Little Brain! I did tweak my parameters somewhat in the face of this more demanding trace. The model is now putting 70% of the heat into a 11 year lifetime exponential (instead of 60% into a 15 year), and the rest into a 200 year exponential. Also, I had no forcing data before 1880 to initialize the state of the model with, so I had to fake some up to get it started off in a reasonable manner. Still and all, the degree of fit suggests the model is capturing the first order physics of what is going on adequately. It's not quite as good as GISS, but it sure was a hell of a lot cheaper to develop :-)

Alright! **Finally** (and believe me when I say that this post has been harder on me than on you), we get to figure out what our exponential, linear, and constant carbon emissions mean in temperature terms. Recall that in The Carbon Economy post I made this plot:

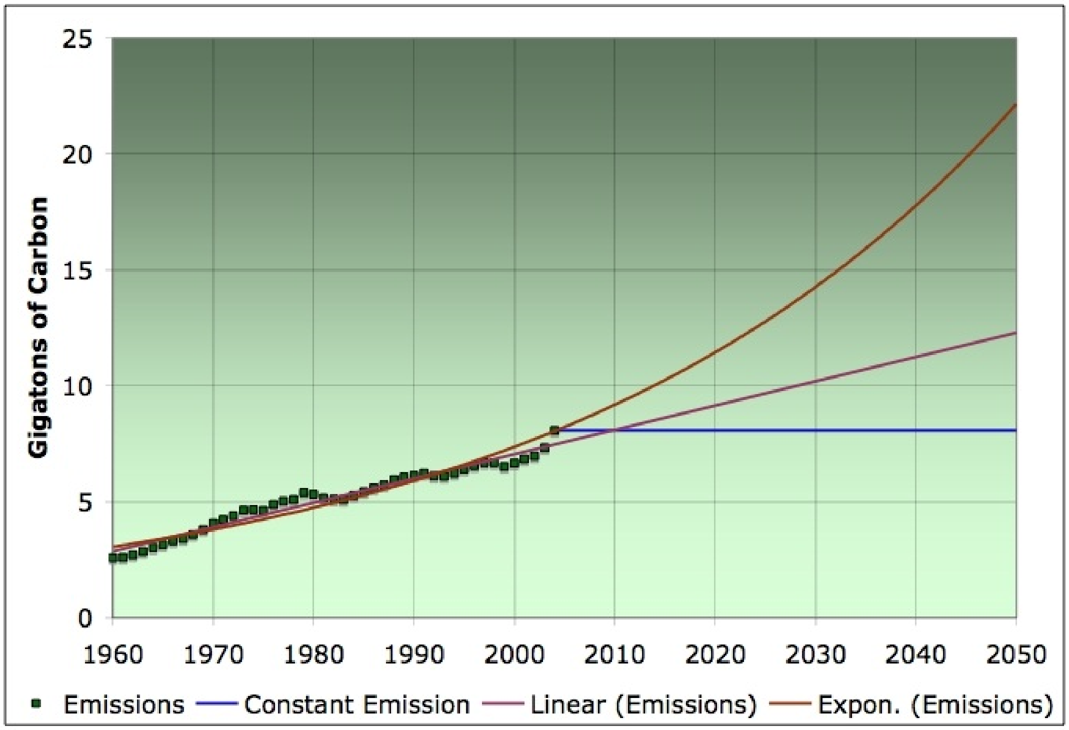

Carbon emissions in Gt/year 1960-2004, together with linear, exponential, and constant extrapolations through 2050. Click to enlarge. Source: [ORNL](#) through 2002. 2003-2004 were estimated by scaling the 2002 numbers by the appropriate percentage increases in coal, oil, and natural gas from the [BP annual production numbers](#).

Then, I [estimated the resulting CO_2 concentrations.](#) What I'm now going to do is hold all other forcings constant:

- aerosols get neither better or worse,
- no major volcanic eruptions or nuclear wars
- non-CO_2 greenhouse gases stay flat.

and just look at where temperature goes as a result of the different carbon emission scenarios. First, here's the additional climate forcing over the 2004 level in each scenario. This can be contrasted to the approximately 2 W/m^2 of net forcing to date.

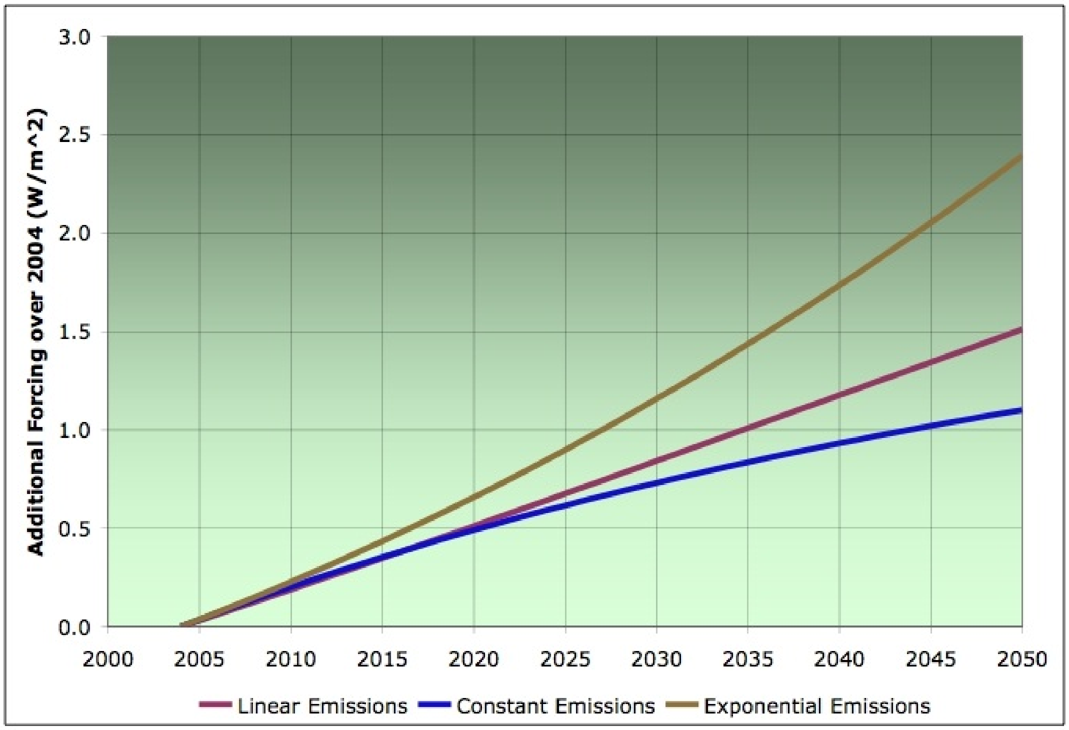

Excess radiative forcing over 2004 level due to various carbon emissions scenarios. All other forcings held constant. Uses Table 1 of [Hansen and Sato, 2000](#) for conversion formulae from CO_2 mixing ratios to radiative forcings.

And if we run those into my model with the same parameters as last time (70% into the 11 year exponential, and 30% into the 200 year one), we get:

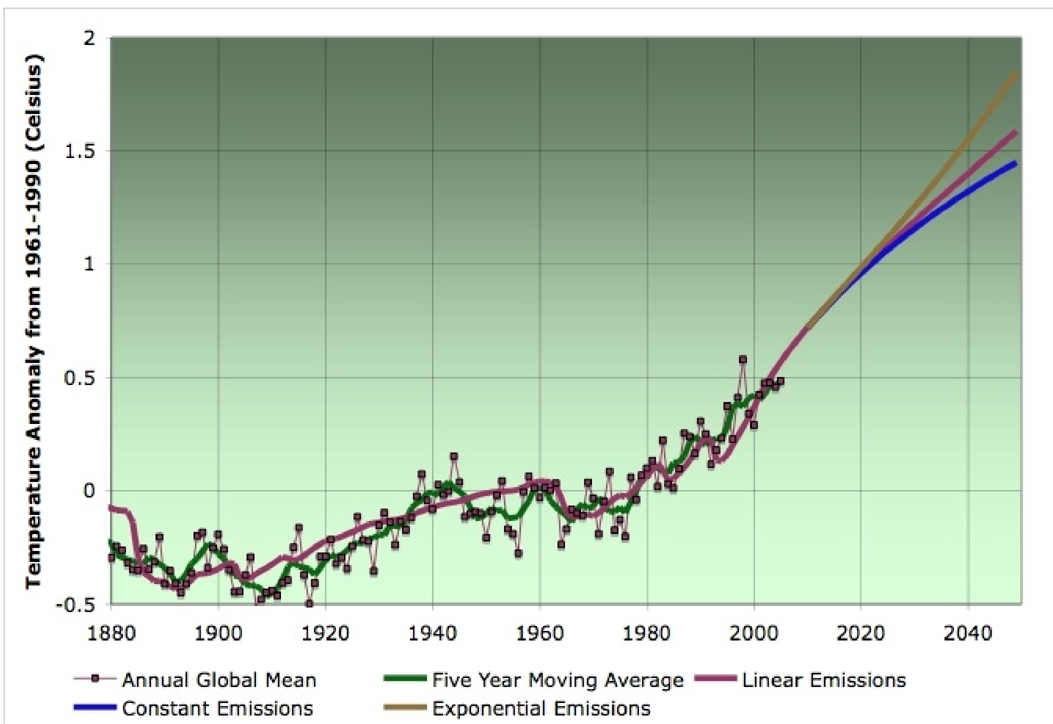

Global average temperature 1880-2005, together with one dimensional model fit (as described in the text) extrapolated to 2050 for the case of linear, constant, and exponential emissions (all other forcings held constant after 2003). Data: [UEA CRU](#).

A comparison with this figure from the Wikipedia (repeated from above for your viewing convenience) suggests the simple model is well within the range of global climate model projections from the 2001 IPCC report, giving us further confidence that it is a reasonable tool for use in scenario experiments.

Temperature increase over 2000 for eight global climate models in IPCC SRES A2 emissions scenario. Source: Wikipedia.

Obviously, this model can't tell us detail on any other aspect of climate than global temperature, but that will probably serve our purposes here.

In the next post, we'll talk about the emerging implications of that kind of warming, before moving on to talk about biomass flows in the economy (which has certainly been given a new surge of interest by the President's remarks yesterday evening).