The Oil Drum
DISCUSSIONS ABOUT ENERGY AND OUR FUTURE

Powering Civilization to 2050

Posted by Stuart Staniford on January 28, 2008 - 10:00am
Topic: Alternative energy
Tags: 2050, agriculture, climate change, globalization, peak oil, photovoltaics, plateau, relocalization, renewable energy, solar power [list all tags]

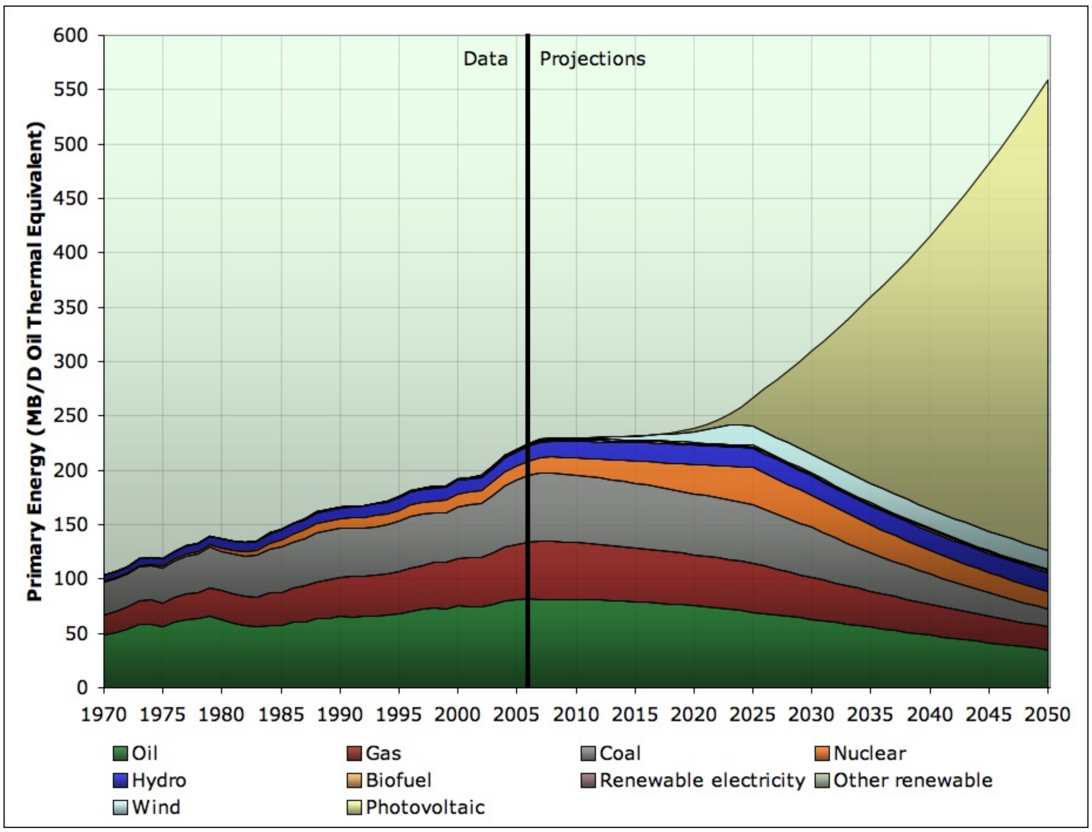

Global marketed primary energy production 1970-2050. Expressed in thermal equivalent of millions of barrels/oil day (ie electricity streams such as hydro or photovoltaic are treated as if they had been converted from fuel at 38% efficiency). Source: BP for fossil fuel, hydro, and nuclear data, EIA and IEA for renewable data, and author's calculations as described in the text for projections. This is a ***scenario not a forecast****.*

This post is the start of an attempt to sketch out what an integrated solution to the world's food, energy, land, climate, and economy problems might look like. My basic goal is to get to a somewhat defensible story of how civilization could get to 2050 in reasonable shape, despite the problems of climate change, peak oil, global population growth, etc.

Since it's not possible for me to entirely solve this problem in a week of part-time work, I put this out as a hasty straw-man. Feel free to point out the parts of this that don't work, or where my ignorance of some of the relevant issues shows particularly badly. Of course, I don't make the claim that I can predict what will happen forty years ahead. Nor do I expect the global population to pay much attention to what I think they should do. Instead, the value of a scenario is to try to think through the general issues that society faces, and the value of an **integrated** scenario is that we can think about how all the parts fit together holistically, whereas usually they get projected separately by specialists, and even the obvious interconnections get missed by decision-makers (if we try to solve our fuel problems by converting food to fuel, perhaps the price of food might go up).

With that said, for the remainder of the piece I'm arrogating to myself sole authorship of all relevant international treaties and implementing legislation at the national level. Here's how I'd go about it. In this first piece, I've analyzed the overall requirements for the problem, but only fleshed out any detail on the population, economy, and energy sectors; I did not have time to write up my analysis of transportation and agriculture/land issues. I will do so in a future piece.

Requirements

In engineering, there is a saying "requirements before design", which indicates that you should think through as comprehensively as possible what you want the system to do before you start trying to figure out how to build it to do that. Otherwise, there's a tendency to think of a subset of the requirements and then rush into a design, only later realizing that some were forgotten, which then must be added at much greater expense. So here are my list of requirements:

- **Population**: The global population is able to grow and go through its demographic transition with death rates continuing to go down. No die-offs.
- **Economy**: The world economy is able to grow on average over the period - modestly in developed countries, faster in developing countries.
- **Carbon emissions**: The global energy infrastructure will be mainly replaced with non-carbon-emitting energy sources by the end of the period, and residual emissions will be rapidly diminishing.
- **Fossil fuels**: I assume that peak oil is here about now but that declines will be governed by the Hubbert model (and thus will be gradual). I assume natural gas and coal are globally plentiful enough that climate policy is required to prevent their full use.
- **Technology**: I do not assume any massive breakthroughs - no technological miracles that solve problems in ways completely unknown or untested today. However, where technological sectors have long established rates of progress in key metrics, I extrapolate the metric to continue improving at the historic rate (eg the economics of solar power, or the yields/acre of agriculture are assumed to keep improving on the historical trajectory).
- **Impact on wild ecosystems**. Developed countries are assumed to maintain the protections they currently have in place (for national parks, wildernesses etc). Developing countries are assumed to exploit their unused land up to the point of best current practices for developed countries. Whatever impact on ecosystems arises from climate change due to past carbon emissions and the tail of emissions to 2050 is viewed as unavoidable.
- **Conservatism** Other than the above, I use the overarching principle of trying to assume as little change in the way the world works as possible - I assume it remains a more-or-less free market world, in which national governments regulate their own countries to temper the worst excesses of the free market and periodically enter into treaties on the more pressing global problems. I assume it remains full of highly imperfect humans mostly struggling to improve their own circumstances. I assume people are willing to come together and take collective action for the common good, but only when the need for that action has become so overwhelming and immediate as to be irrefutable.

A few comments on these requirements. The first two on population and economy I view as minimally necessary to ensure the world's political and social stability. If the population or economy were to contract substantially, I believe a number of countries would become unstable and there would be revolutions, mass population migrations, etc. All predictability in the world would likely be lost at that point. The first, on population, also seems to me a non-negotiable moral imperative. Of course, these conditions alone will not be enough to prevent a normal amount of human mistake and misery, but I believe they are required to prevent mistakes and misery from becoming globally endemic.

The carbon emission requirement seems essential to have some hope of stabilizing the climate. There is a considerable possibility that the climate is already destabilized to an important degree. Nonetheless, it's likely possible to make it much worse by continuing to emit exponentially increasing amounts of carbon. The best we can hope for is to start now, as energetically as possible, substituting our energy sources and sinks as best we can, until we don't emit carbon.

On fossil fuels, I have explicated my best guesses on peak oil timing (here and here) and on decline rates, and will let those stand for now.

On technology - of course, it's not likely that 40+ years of human ingenuity will pass without some big breakthroughs (think of what has happened since 1968 - personal computers, the Internet, statin-class drugs, mobile phones, etc, etc). But there is no hope at all of predicting specifically what those breakthroughs will be, or whether we will make one important to our climate/energy problems. So I assume nothing here - any breakthroughs will hopefully serve to offset whatever negative surprises the world may pose to us. However, I think it would be unduly pessimistic to assume that well established trends of technical progress will not continue, unless there is a clear game-changing theoretical barrier in the way of further progress.

As everyone knows, the rest of nature has taken a real beating from humans, and it's going to get worse. Climate change means that species must move, and habitat fragmentation means they can't in many cases. I recognize the tragedy in this, but I do not see what can be done about it given the other requirements above. I don't think it's realistic to assume that developing countries will do any better than the examples that developed countries have set in terms of set-asides - I certainly support efforts to conserve as much as can be conserved.

And then, finally, I am still a liberal on the contemporary American political scale, but a rather centrist and very independent one. In particular, at this point in my life I have a very strong sympathy for the strain of thinking in the Anglo-American conservative tradition which says that humans are complex and highly imperfect, that it has taken a long time to get society to work as well as it does, unsatisfactory as that may be, and that radical experiments to improve the way it works are dangerous. One should attempt change by making improvements step-by-step on the most pressing issues of the day, and one should be prepared to accept that society will never be utopia. In this case, we have no choice but to make some radical changes in where we get energy from, but let's at least not try to change any more things than we have to. I'm more than happy to settle for a world that will be a minimally acceptable place for my children and grandchildren to grow up and live out their destinies. I take the requirements above to be a description of that minimal level. The world will not become a utopia (whether you prefer your utopia agrarian, socialist, libertarian, spiritual, or transhumanist).

Outline of Solution

My basic approach is as follows. Over the next fifty years, we're going to phase out most burning of fossil fuels, but they will still be used for petrochemicals and fertilizer (manufacture of which will be mainly in the Middle East). We will cope with short term energy problems by efficiency improvements, but in the long term we will power society predominantly by massive amounts of solar PV, with smaller amounts of wind, and legacy hydro. We will use a global transmission grid to balance supply and demand between the nighttime and cloudy areas and the areas in the sun that generate power. Nuclear is avoided in the long term out of proliferation and waste concerns but is used in the short and medium term. Owners of fossil fuel infrastructure will be compensated at fair market value.

Ground transportation will be by a mix of electric cars and electrified public transport (in areas of high enough density). The car fleet will be moved through hybrids to plug-ins to full electrics as storage technology slowly improves. Developing countries will be encouraged to urbanize and develop as rapidly as feasible to reduce pressure on remaining wild ecosystems and to build public transport systems in their very dense cities.

Building heating and cooling will be transitioned predominantly to ground source heat pumps powered by electricity instead of burning fossil fuels.

Agriculture will remain predominantly industrialized, and ongoing yield improvements, particularly in the lower-yielding poor countries, are assumed to be able to feed the world. The residual oil production and modest and regulated amounts of biofuels will be used for certain applications where the advantages of liquid fuels are indispensible (predominantly heavy construction and agricultural machinery, shipping, and aviation). There is considerable scientific uncertainty on the extent of soil depletion, but the assumption here is that at-risk areas will be placed in conservation reserves, and that, later in the century when energy becomes cheap again, restoration and remediation will be attempted.

The overall economic approach for implementation will be a hybrid "markets-within-a-plan" approach. A pure free market approach is likely to be disastrous (eg starving the poor to make biofuels for the rich, which will result in riots and revolutions). However, markets are very powerful drivers of innovation and efficiency when well designed. We will set general goals with binding targets by treaty, and then use a combination of subsidy auctions, rights auctions, and reverse auction retirements of fossil fuel infrastructure to meet the binding targets. Market competition will improve the technology and drive down the required subsidies over time.

In general, this will require a massive global infrastructure project. It will be expensive, but it's not impossible. It seems very cheap compared to further uncontrolled experiments with the climate, or to allowing the world to descend into starvation and chaos by adopting dysfunctional approaches to our energy challenges. It will place civilization on a tolerably sustainable footing for the longer term.

With that, let's turn to looking at the major sectors of the global economy in a little more detail. Again, I'm going to go through population, economy, and energy sectors this time. The rest will have to await future pieces.

Population

Here, from an old piece of mine on population, is a graph of the various UN population scenarios:

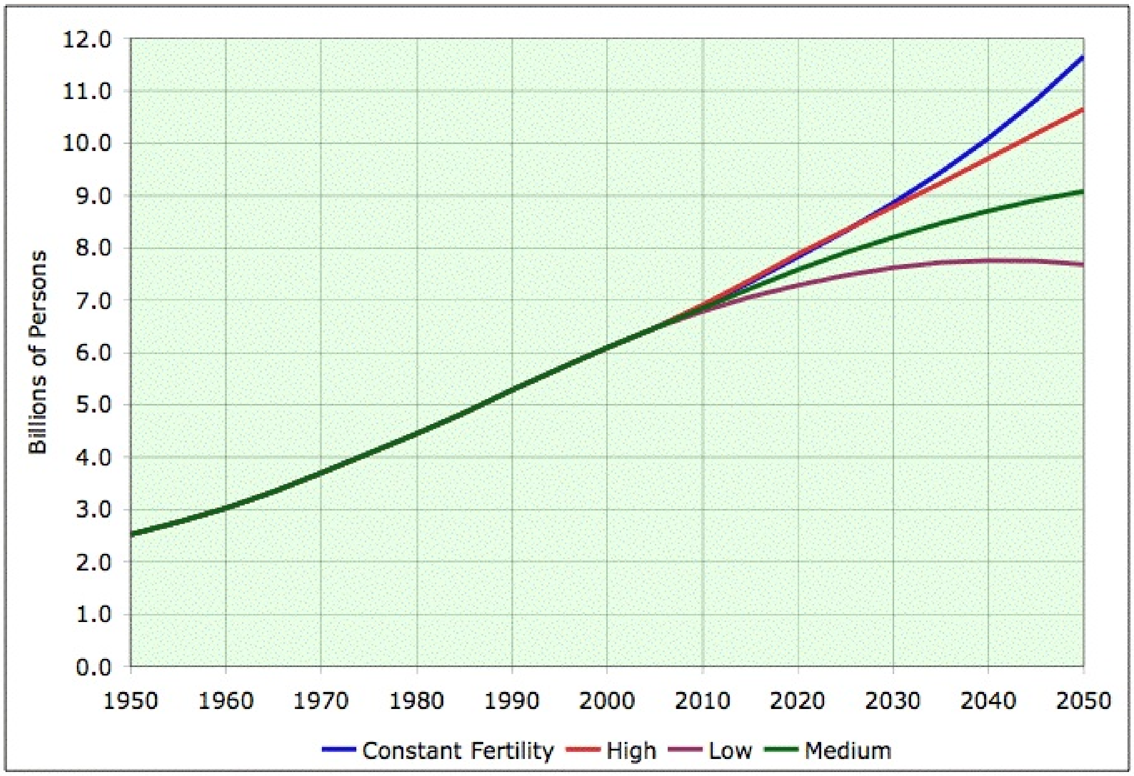

UN population projections through 2050. Click to enlarge. The medium scenario (dark green) is the UN's best guess as to what will happen. High and Low represent their best estimates of the range of reasonably likely outcomes. The Constant Fertility line is their estimate of what would happen if world average fertility did not decline any further. Source: World Population Prospects: The 2004 Revision.

In this piece, we will take the medium population scenario as our goal to be supported by the planetary economy. It's worth looking for a moment at the assumptions the demographers are making in the less developed countries (where most of the population growth will be occurring):

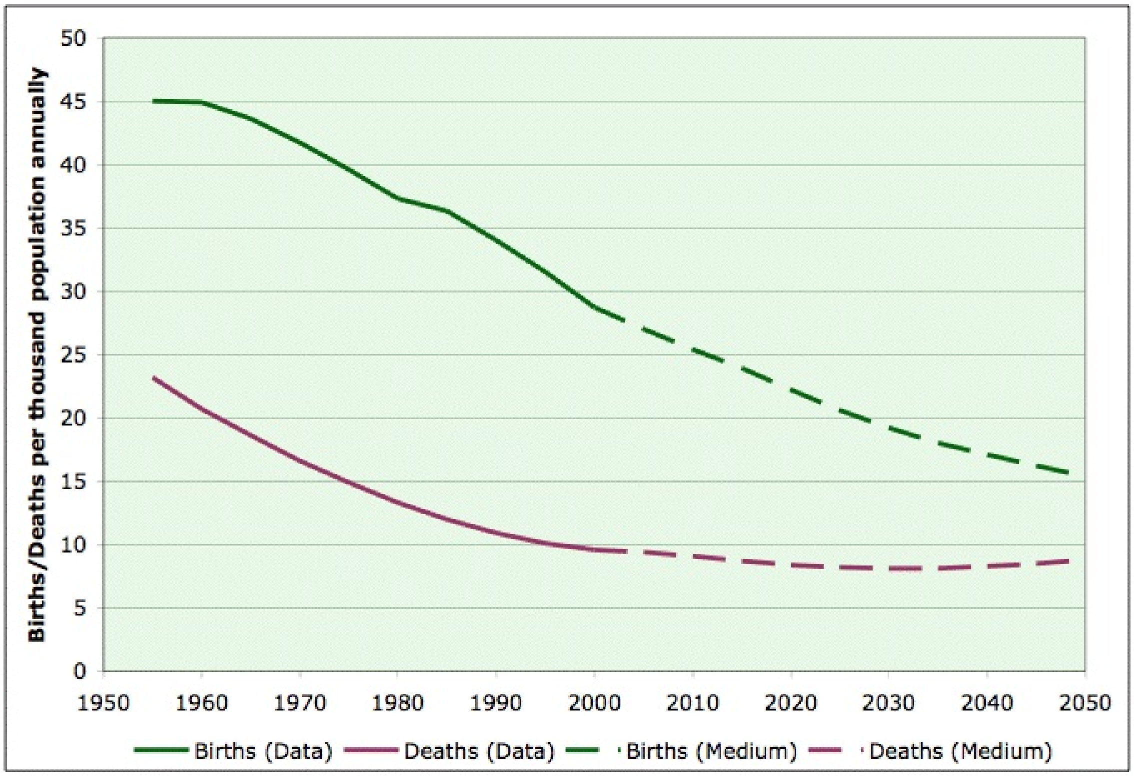

*Birth and Death rates for UN **less** developed countries (excluding least developed). The lines through 2000 are data, and after that the lines are the UN's medium projection. Source: World Population Prospects: The 2004 Revision.*

As you can see, death rates have been falling, and they are assumed to fall further as countries continue to develop and get wealthier. Birth rates follow death rates down, but with a lag (this is the demographic transition). Least developed countries are following a similar trajectory but are not as far along. The assumption I make here is that if the economy and food supply are allowed to grow at reasonable rates, then population will follow the UN's assumptions.

Economy

The way the economy develops in my scenario is as follows (I'm going to summarize it here, but some of the justification comes later). This graph shows global GDP on a $2007 purchasing power parity basis. The data through 2007 are from the IMF. My scenario basically has global growth dropping sharply for the next few years due to an assumed US recession caused by the credit collapse. Then growth starts increasing again, but slowly due to high energy prices. Finally, as our energy problems get thoroughly solved around 2025, growth returns to about the current long term trend in developed countries. Developing countries are assumed to gradually slow down until by 2050 they are not growing much faster than developed countries (since they will be fairly developed by then).

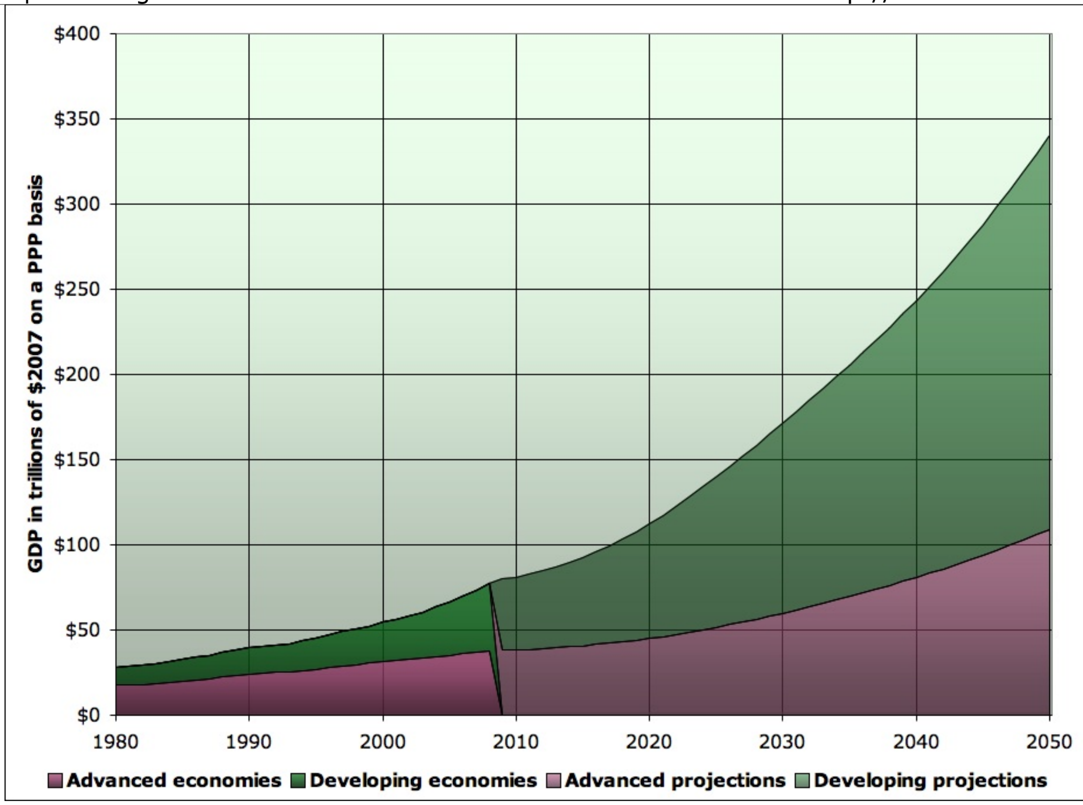

Actual and projected global GDP 1980-2050. Expressed in $2007 dollars on a PPP basis. Source: IMF for historical data, and author's calculations as described in the text for projections.

The global economy, about $72 trillion in 2007, will be several hundred trillion dollars by 2050 under my assumptions.

In addition to GDP data, the IMF has also kept track of the level of investment in the global economy. I show this data, and also how much investment there would be in the future, assuming that the fraction of GDP invested averages the same in the future as in the past. I place this graph here to give some idea of the world's investment budget for when we start spending some of it on a new global energy infrastructure.

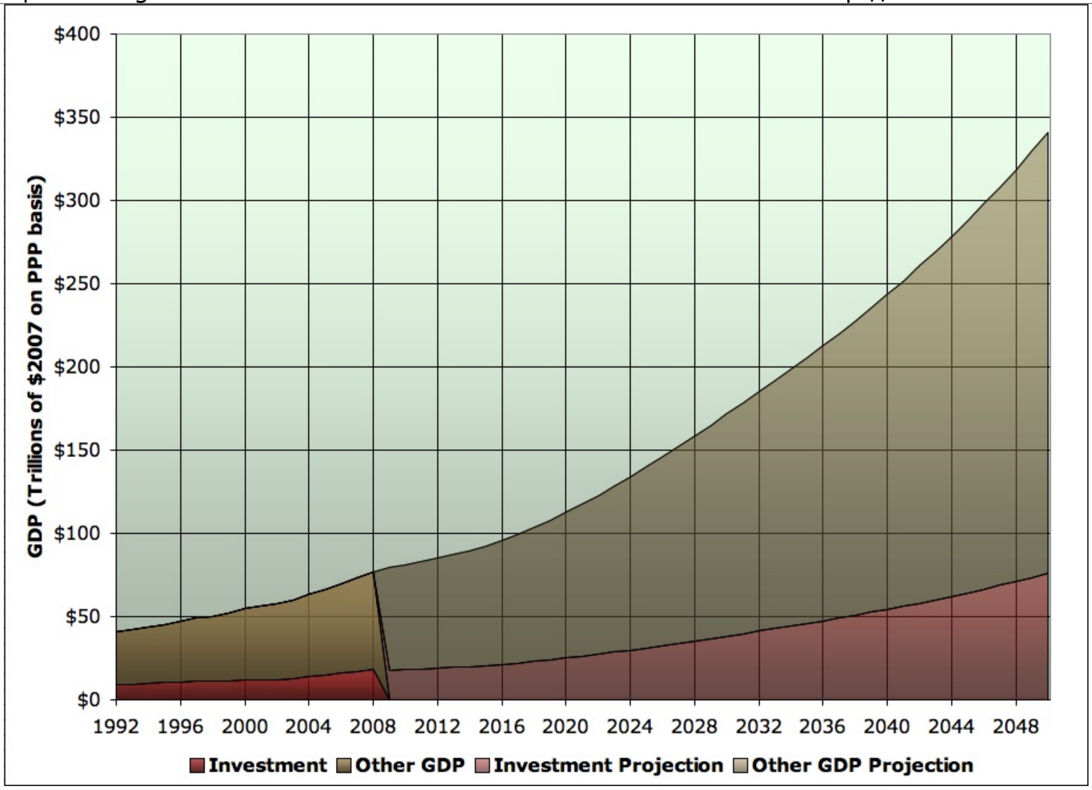

Actual and projected global GDP 1980-2050, broken out into investment component versus immediate consumption. Expressed in $2007 dollars on a PPP basis. Source: IMF for historical data, and author's calculations as described in the text for projections.

Total investment is about $17 trillion in 2007, and rises to somewhere around $75 trillion/year by 2050. This includes spending on houses, factories, offices and infrastructure of all forms, but excludes current consumption.

Again, left alone with a sufficiency of resources, a capitalist economy will tend to grow. Inventors will invent more efficient ways of doing things and more desirable end-consumer goods; entrepreneurs will bring them to market; people will get more productive; bigger and better houses, factories, cars, etc will be built, and the whole thing will be bigger next year than it was last year. Thus the basic analysis to justify my scenario consists in showing that none of the apparent resource bottlenecks are necessarily fatal to growth. By far the most important of these is energy.

Energy Sector

There are basically three choices for non-carbon-emitting, non-biofuel, energy production technologies that are already in practical commercial use and have potential for major expansion: nuclear, wind power, and solar. Let me briefly discuss the trade-offs.

Nuclear is the furthest along in that it already provides a material fraction of global primary energy (about 5.6% of marketed energy on a thermal basis). See Is Nuclear Power a Viable Option for Our Energy Needs? for a good summary of the case for nuclear. Nuclear has a high energy return on energy invested and there is sufficient uranium for a long time (though there are short term price issues associated with the ending of the burning of the stockpile of uranium from old nuclear warheads). New nuclear energy will probably be reasonably cheap as long as interest rates aren't too high (since most of the cost is the upfront capital to build the plant, the cost of finance is critical to nuclear economics).

Nonetheless, I do not favor nuclear power as the long-term major solution to powering the global economy. Pervasive presence of a nuclear power industry throughout all or almost all countries of the world has three major issues: developing countries are frequently corrupt with a tendency for their building subcontractors to do things like leaving the rebar out of concrete to save money. Nuclear power is an industry were mistakes and short-cuts cannot be tolerated. Furthermore, the nuclear fuel chain always has the potential to be diverted into weapons use by its owners (the basis for western concerns about Iran). Nuclear weapons remain the only way humanity has come up with to not just end our civilization but end our species and most life on the planet and they ought not to be proliferated further. And finally, after fifty years of operation of the industry there is still no commercially demonstrated permanently satisfactory solution for what to do with the waste. I share these concerns with many members of the public, and, even if I didn't, the presence of broad political objections to nuclear power would make its future problematic.

That said, the issues with nuclear power seem to me somewhat less pressing than climate change, and in scenario building it turns out to be very valuable to allow the nuclear industry to grow for a while, and then buy them out at fair market value in due course. In my scenario, I allow global nuclear to grow until 2025 (at the historical 5.4% annual growth rate that obtained globally from 1980-2006 according to BP data) and then start buying out older and less-profitable plants to reduce the nuclear contribution by 3% per year.

See The Most Frequently Asked Questions About Wind Energy for the case for wind. Wind power is already cost-competitive with fossil fuel power, emits no carbon in operation, and has a very satisfactory EROEI of around 18 on average. Wind power delivered has grown on average 23.7% annually from 1990 to 2005 according to the IEA, and the new capacity delivered has been growing 30% - 40% annually in the last couple of years, suggesting that the growth rate is accelerating. I feel wind is an excellent option.

However, many people differ as soon as someone tries to put a wind turbine near them - objecting to the noise, visual impact, and harm to bird-life. Local political opposition to wind plants has frequently been very strong, and because wind is extremely diffuse, one needs to put an awful lot of turbines in place to garner it's full potential. For this reason, I assume wind continues to grow, but eventually saturates the politically acceptable sites. The two countries that have the most wind are Denmark with about 1.1 TWhr/year per million people of wind production in 2005, and Germany with about 0.33 TWhr/year per million people. New wind installation in Denmark has largely stopped, and it appears to be slowing in Germany. I treat the global politically acceptable maximum as around 0.5 TWhr/year/million people, and grow wind up to that point at the historical 23.7% growth rate. (This threshold is highly uncertain and my calculation should only be treated as qualitatively indicative. If I'm wrong, we'll end up with more wind and less solar, but the issues are basically the same either way). Wind power still produced only 0.2% of global marketed primary energy in 2005 (treated on a thermal oil equivalent basis with an assumed power plant fuel efficiency of 38% - what BP does for hydro and nuclear electricity statistics). Thus it will take a long time still for it to grow to providing an appreciable fraction of our power.

My feeling is that photovoltaics are the right answer for the long term future. They don't harm wildlife, don't pollute, people around them don't seem to object to them much, they don't critically depend on anything in ultimately short supply, and they have outstanding energy payback. For example, Nanosolar claims payback of manufacturing energy in less than one month for their state-of-the-art product, but more conventional options still have energy payback in the single digit years, implying an EROEI in the tens to hundreds as we go from current to future PV products.

There are two major issues to overcome: economics, and intermittency. PVs are not yet economically competitive with fossil fuel energy, but there is a long cost history that lends itself to a fairly stable extrapolation that is quite encouraging:

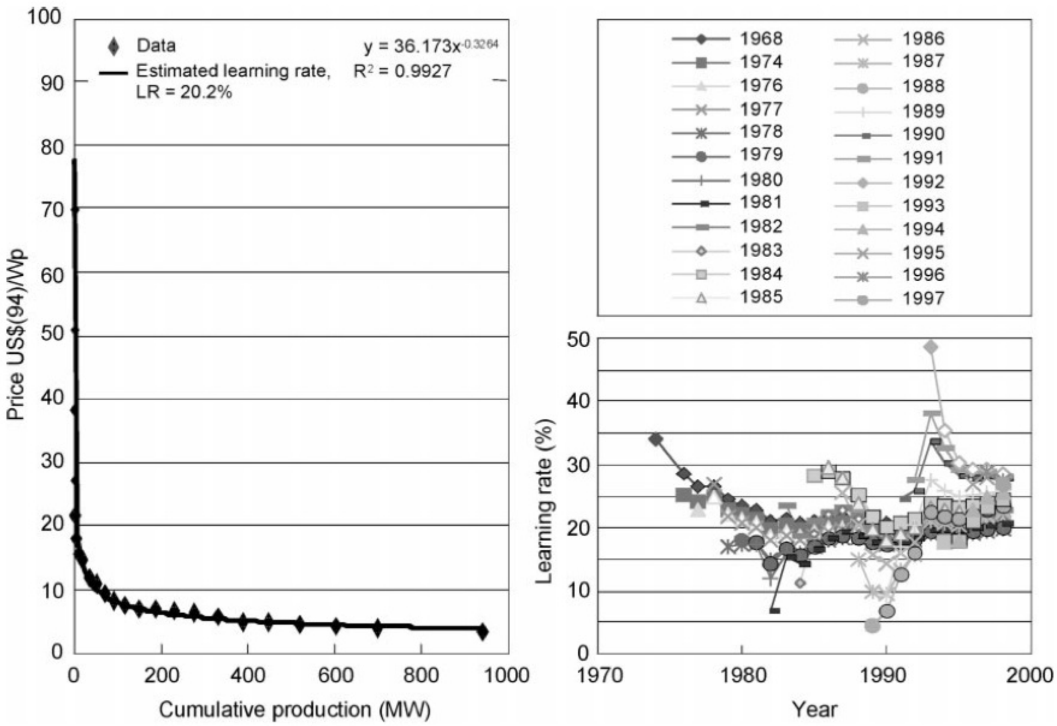

Left panel shows cost of PV panels versus versus cumulative installed capacity. Right panel shows a sensitivity analysis for the learning rate (the percentage drop in the cost due to a doubling of the installed capacity). Source: Fig 3 of McDonald and Schrattenholzer, Learning Rates for Energy Technologies.

The *learning rate* is the percentage by which some technology drops in cost per doubling of installed capacity. Solar has been dropping at about 22% for each doubling, and this is fairly stable, give or take a few percent. If anything, the learning rate is improving slightly over time. According to the IEA, the installed base of PVs grew at a combined average growth rate (CAGR) of 34.9% from 1990 to 2005. Thus we would expect costs to drop by about 9% per year, which would correspond to halving every eight years. At current costs of about $4/peak watt, unsubsidized PV power costs about 10c/kWhr in sunny places like Los Angeles and about 16c/kWhr in a cloudy place like Seattle (from Solar Revolution, p 110). Thus it is probably already competitive with retail electricity in many sunny places, and will become competitive with wholesale prices of about 5c/kWhr in less than a decade in sunny places and in about 15 years in cloudy places.

With serious policy help, PV installed capacity can grow much faster than the 35% global average. Eg in Germany, PV has grown at a CAGR of 61% over the same 1990-2005 period. However, the global installed base of PV is miniscule - in 2005 it only comprised 0.0033% of marketed primary energy (on a thermal equivalent basis).

The effect of all these trends - tiny current installed base, rapid growth, very fast learning curve, high EROEI tends to mean that PV can be of almost no meaningful benefit to the global situation in the short term, but in a couple of decades from now reaches critical mass, and then will potentially be in a position to provide almost all of society's power within a couple more decades from that. Since PV can be readily fit into all kinds of otherwise unused surfaces on buildings, and also spread out over otherwise low-value desert, and can be applied in installations from a single panel up to thousands of acres or more, it can be ramped up very quickly - there are few barriers to deployment. This is the basis for my selecting it as the backbone of long-term sustainable power for society in my scenario.

The remaining problem that needs to be solved is the intermittency (PV provides no power when it's dark and not much when it's very cloudy). There are basically two possible approaches to this. The first is that we would install enough storage everywhere that the energy stored during the day would be enough to power usage at night. I have not been able to construct a believable story about how current electric storage technology can scale to the required magnitude in a timely way, and thus this approach, as far as I can see at present, faces a critical bottleneck. It's one thing to have a battery that will power a plug-in hybrid for an hour commute. It's another to have enough batteries to get a region through a week of clouds and rain.

The second approach is to construct a global electricity grid. As far as I'm aware, this approach was first proposed by Sanyo under the rubric Project Genesis. Their idea was to install PVs throughout the world's deserts, and connect them up via superconducting cables to the world's cities - they estimated 4% of the world's desert area would be required. To get some feel for the issues, you might want to stare for a while at this screenshot of the planet:

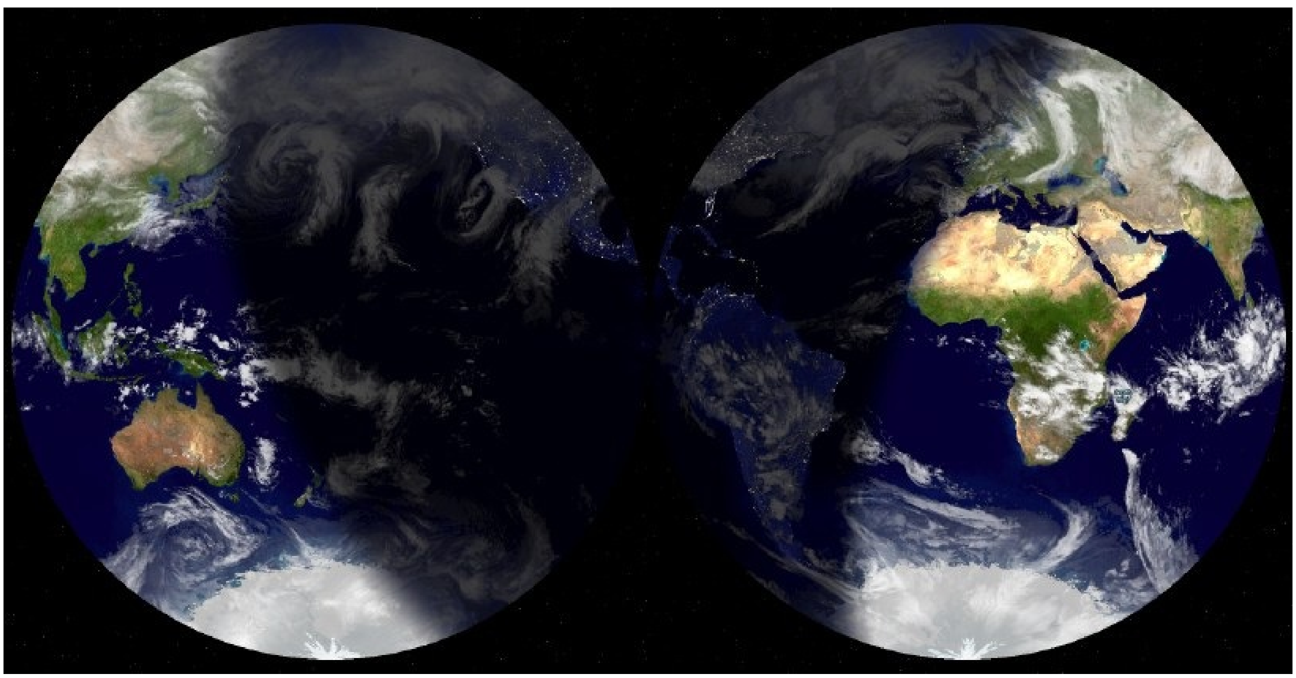

View of the earth as of about 11pm Pacific time on Friday 1/25/08. Source: Screenshot of OS X Planet.

Obviously, half the planet is dark at any given time, and about half the rest is under cloud. We would have to be generating enough power in Africa, the Middle East, and the non-cloudy portions of Asia to power the globe, and then shipping it from those sources to wherever it was required all over the world. 12 hours later, those areas would be returning the favor. Nighttime electricity use is only about 30-40% of the daily peak, but it's still a lot of energy to move. The worst case seems to be early evening here in California in the northern hemisphere winter:

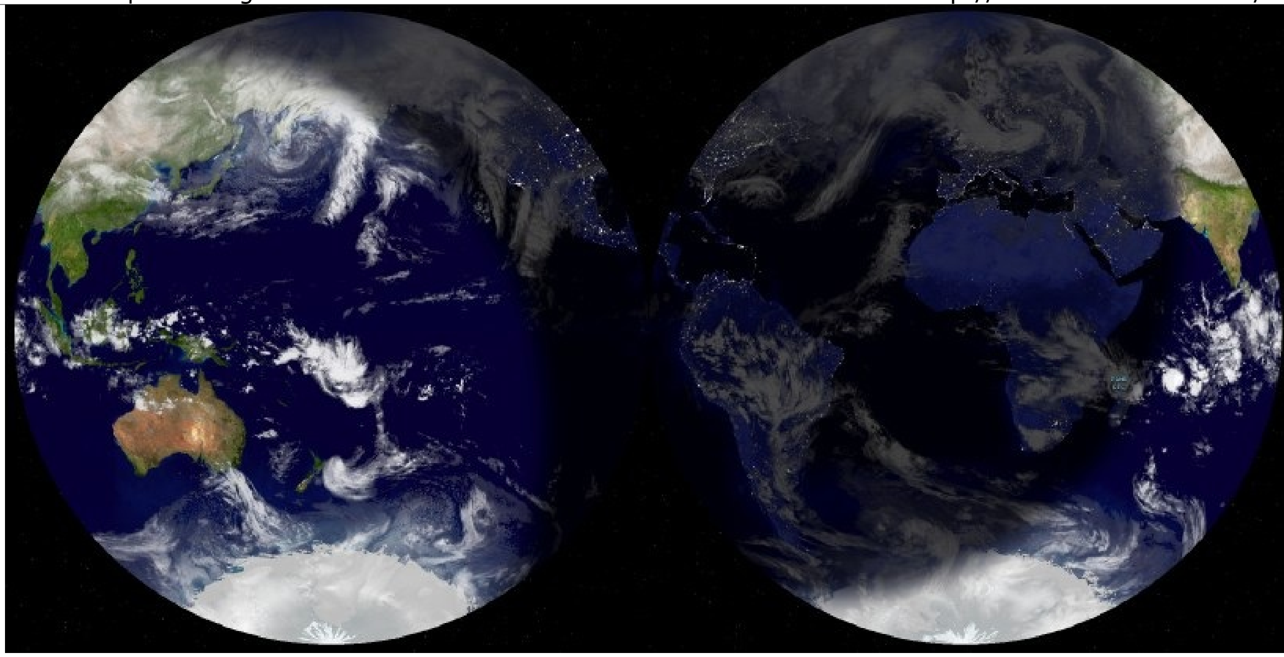

View of the earth as of about 6:30pm Pacific time on Saturday 1/26/08. Source: Screenshot of [OS X Planet](#).

In that case, Asia and Australia have to power their own daytime use, as well as evening use in the Americas, and residual night time use in Europe and Africa. I imagine that the Australians would clean up financially with big PV arrays in the interior.

I have two changes I think are required to Project Genesis. One is that most PV already is being installed on buildings in the urban environment, and I expect this to continue (albeit supplemented by utility scale plants in waste areas). The second is that I don't want to rely on superconducting cables, since they are still on the drawing board, which violates my "no breakthrough" requirement.

However, existing technology for high energy transmission lines appear to be able to do the job, albeit with significant losses. High voltage direct current lines lose about 3% per 1000km. The earth has 6378km radius, so it's 20,000km to go to the exact opposite point on the other side. However, if we figure the average electron only needs to go about 2/3 of the way around the planet to get to its customer, then average losses will be 1- 0.97^13 = 33%. In short, by shipping PV around the world, we lose about a third of the part that's shipped. However, if we figure on average two thirds of PV goes locally to the awake/bright side of the earth with little loss, and one third goes to the asleep/dark side at 1/3 lost, overall losses are a fairly manageable 1/9 of total power generated. This is in the same ballpark as the losses of natural gas in LNG shipping (about 15% of the natural gas). In short, existing technology appears able to get the job done - it's an enormous global infrastructure project, but doesn't require breakthroughs. Any breakthroughs in electric transmission can only make it better.

Before I analyze the costs required to build a global renewable electricity grid, let me summarize my total energy generation scenario as it stands at present (I reserve the right to improve it over time). I assume PV capacity grows at the historic rate of 35% until 2012, when the replacement for the Kyoto treaty comes into force. Since I will be writing the treaty, in my scenario, it will call for a much accelerated rate of growth in renewables, and so PV then grows at the German rate of 61% until 2025. At that point, it is reaching critical mass and will then slow down and just grow overall top line societal energy usage by 3%/year. But after about 2025, energy will get cheap again and it's not clear how fast usage will grow - it will be constrained by demand, not supply.

Fossil fuel usage is assumed to be impacted in the near term by a significant US recession and global slowdown, and in the medium term by strenuous conservation efforts inspired both by high energy prices and the ongoing legislation/regulation process responding to climate concerns. From 1979 to 1983, overall primary energy consumption dropped as a result of the second set of 70s oil shocks. I assume that, this time around, developing country demand growth is too robust to cause an actual drop in overall energy consumption, but I assume it essentially stays flat until the renewable energy explosion starts to really hit with wind making increasing contributions after 2015, and solar PV starting to completely swamp the situation from the early 2020s on. After 2025, we can easily grow PVs and we can begin to aggressively retire the remaining fossil fuel and nuclear infrastructure.

This graph summarizes the scenario at present:

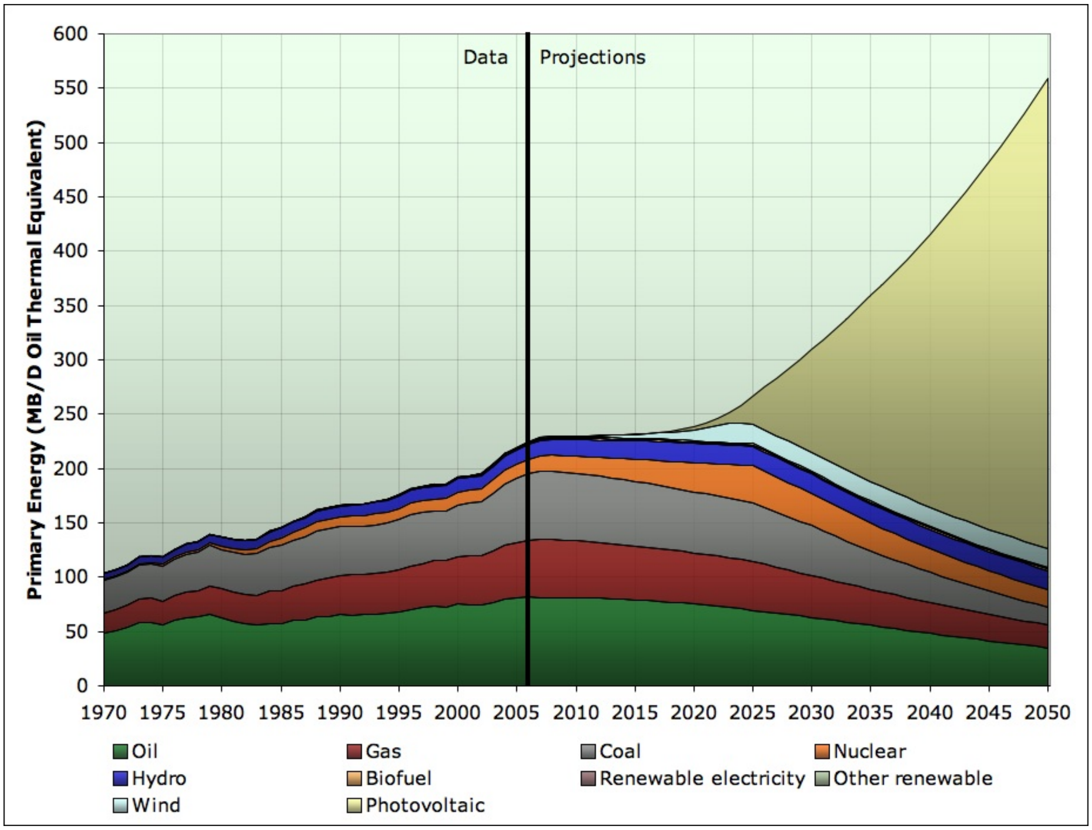

Global marketed primary energy production 1970-2050. Expressed in thermal equivalent of millions of barrels/oil day (ie electricity streams such as hydro or photovoltaic are treated as if they had been converted from fuel at 38% efficiency). Source: BP for fossil fuel, hydro, and nuclear data, EIA and IEA for renewable data, and author's calculations as described in the text for projections. This is a ***scenario not forecast****.*

Of course, I wouldn't claim for a moment that my back of the envelope calculations will actually be the way things play out quantitatively. However, it does seem to me that the qualitative features of this graph are likely **if** we do the required international public policy groundwork to install a global renewables grid. The qualitative features I mean are a period over the next 15-20 years of high energy prices and slower global growth, followed by a period after that when renewables take off and become the main power source for society and energy becomes cheap again.

The remaining questions are around costs - can we afford to to do this? Well, if you look at my scenario, we need about 550mbd of primary energy capacity, which under my assumptions corresponds to about 125 PWhr/year, or about 14 TW of electricity. Of that, we probably need worst case capacity to move about half of it around the globe. So we need 7TW capacity cables circling the globe in a roughly east-west arrangement, but let's add 50% for path deviations due to the complexity of land topography. Therefore, we need 7TW x 60,000km of cable capacity by 2050.

Now, the cost of HVDC lines is ballpark $1m/km/GW. So given that we need 7000 GW x 60,000km, we will have to spend about 400 trillion dollars ($2007) between now and 2050 to achieve that. That's a lot (learning curve may reduce it somewhat - I'm not assuming any). However, if by this means we keep economic growth going, then the system will certainly be affordable. Recall that above I gave the IMF investment data and a projection of it. In my scenario, GDP from 2008 to 2050 totals about $7700 trillion, and investment at the historical ratio is $1700 trillion. So the cost of the renewable grid is about 25% of investment, or 5% of GDP. **However**, we save on all the fuel. For example, in 2006, the global fuel bill for oil, coal, and natural gas (at commodity prices) was about $3.6 trillion, which was 5.4% of 2006 global GDP according to the IMF. Presumably fuel prices going forward are not likely to be much better than 2006. Thus, although a global renewables grid would require a major investment over the course of the next forty - fifty years, it's only comparable to what we would be spending on fuel if we stick with our current course of action. And our current course leaves us with no idea what kind of climate we'll be living in, and whether it permits civilization or not.

Well, with that, I'm about out of time for this piece. Next week, I'll try to flesh out what some of the other sectors of the economy might look like. In the meantime, feel free to poke holes in this in comments.