Food to 2050

Posted by Stuart Staniford on March 10, 2008 - 8:40am
Topic: Environment/Sustainability
Tags: 2050, agriculture, great transition, peak oil [list all tags]

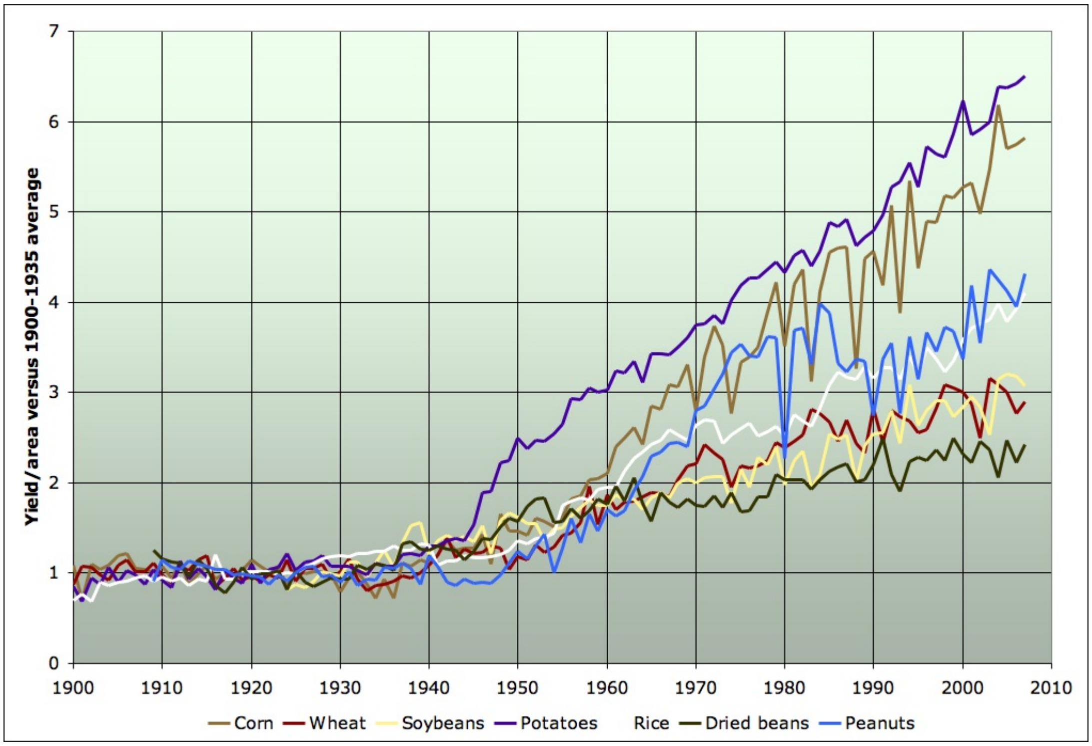

Average United States yields per unit area for various crops, 1900-2007. Yields are expressed as a multiplier of the 1900-1935 average. Source: National Agricultural Statistics Service.

This post continues an exercise I began a month or so ago of trying to figure out how civilization could be moved to a mostly sustainable footing by 2050, while still being recognizable as civilization, and in particular allowing some continued level of economic growth between now and then, especially in the developing countries. Let me remind you of the parameters of the exercise:

- **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.

I n Powering Civilization to 2050 I argued it was potentially feasible to transition to power civilization with a mix of solar, wind, and nuclear energy, with the transition well on the way to completion by 2050. (Luis de Sousa made a broadly similar argument in Olduvai Revisited 2008). This would require a period of belt tightening and conservation in the next couple of decades, but once the transition had overcome the critical threshold (as solar energy in particular became cheap), I suggested energy in general would get cheap again. I adopted the UN medium population projection which has population at about 9 billion by 2050, with growth slowing sharply. Making plausible assumptions for economic growth between now and 2050 if energy was available, we got to a world GDP of about $350 trillion in 2050 (in 2006 purchasing power parity dollars), versus about $70 trillion in 2007

If the average global citizen was significantly wealthier in 2050, they would undoubtedly want to drive more. The switch to primarily electrical energy sources for civilization would preclude doing this with all liquid fuels. In Four Billion Cars in 2050? I argued that, given that the average citizen will be living in a dense third world city by 2050, we can assume rates of ownership typical of the most car-free corners of western Europe at the moment (Holland), which gives rise to a few billion cars in 2050. I further argued that it seems feasible that this many plugin-hybrids could be built - there appears to be enough lithium for the batteries - and run on less than 10mbd of liquid fuels.

In this piece I want to look at another area that many people think is likely to be a critical bottleneck to civilization continuing - the area of food, agriculture, and soil. I am of course not an expert in these areas, but happily there is a lot of excellent scholarship and scenario building that I can lean on. My task is reduced to reporting of the existing science, with some modest adjustments to reflect where my assumptions differ from those of published scenarios (most especially the assumption of a near-term peak in oil supply, and a full-speed effort to convert society to carbon-free energy sources.)

Let's begin with two very helpful UN Food and Agriculture Organization reports: World agriculture: towards 2015/2030, and the sequel World Agriculture: Towards 2030/2050. What these reports do is basically look at projections for population and economic growth and then estimate how much food people would want in the future, and what quantity of agricultural commodities would be required to fulfill that demand. The first report focusses a lot more on the supply-side factors of how this could be done, while the second report extends the analysis out further in time but confines itself much more to demand side considerations.

The input assumptions about population and world GDP are slightly different than mine, but close enough that I am just going to adopt their food scenario wholesale, rather than trying to construct my own from first principles. The differences would be small - much smaller than the other uncertainties in the problem. Let me first summarize their scenario, and then we will start to explore the potential bottlenecks that might prevent achievement of this much food production. (However, I strongly encourage readers that care about where their food is going to be coming from in the future to take the time and read the FAO reports themselves.)

Let's start with a look at what the FAO scenario has for average nutrition. This next graph shows both history and projections to 2050 for daily dietary energy (in Kilocalories/day/person) in various regions of the world, as well as the global average.

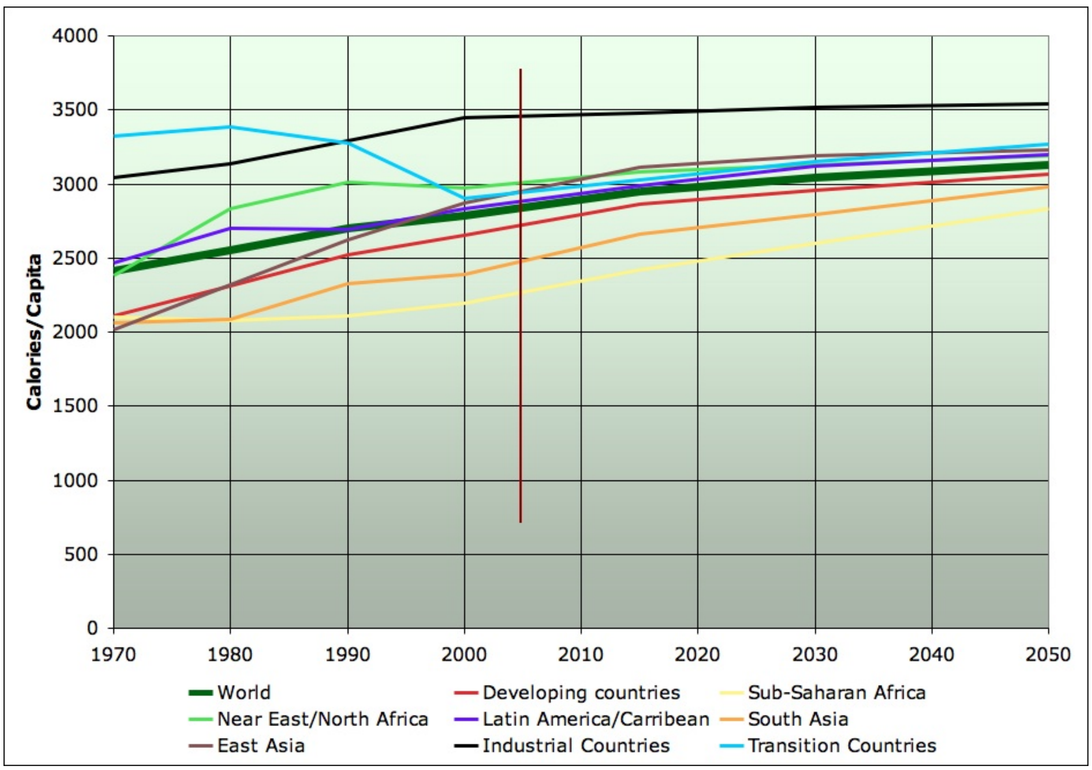

Per capita food availability 1970-2050 for various regions, together with world average. Values for 2000 and before are data (left of the vertical red line), 2010 onwards are projections (right of vertical red line). Source: Table 2.1 of UN Food and Agriculture Organization, World Agriculture: Towards 2030/2050.

As you can see, the history is that most regions of the world have been getting more and more food. The exceptions are some of the formerly communist countries which suffered a partial collapse of their societies as they attempted to transition to a different economic system. The FAO projects that as the developing countries continues to grow faster than the developed world, they will be able to afford more food, and thus they will continue to approach, but not completely achieve, developed world levels of (over)feeding.

I could quibble with a few things here - I might guess that wealthier developing countries will get closer to current developed country averages by 2050, and I wonder about the sharp trend break between the past and the projections in the developed world. Still, these are minor issues - I think this has to be in the right ballpark for any scenario that assumes continued improvement of economic conditions in the developing world, and no major societal collapses (which is what we are trying to figure out how to avoid).

If we take the FAO's scenario breakout of food groups (which they give by weight on a per-capita basis) and multiply by population, we get the following for total food demand:

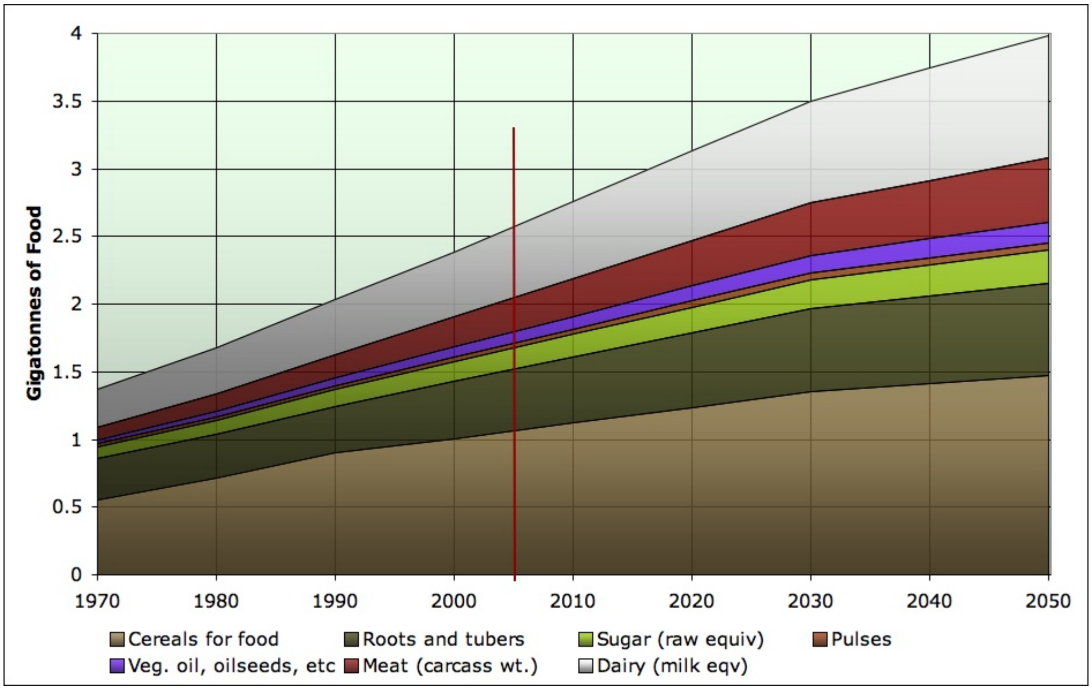

Total food requirement 1970-2050 by major food types. Values for 2000 and before are data (left of the vertical red line), 2010 onwards are projections (right of vertical red line). Source: Table 2.7 of UN Food and Agriculture Organization, World Agriculture: Towards 2030/2050 and UN Medium Population Scenario for population figures. Note that I did not include "Other food", which is only given in calorific terms in the table, and constitutes less than 10% of calories. Fruits and green vegetables would be included under that category.

As you can see, by 2050, the world would need to be producing about 50% more food than it is today (by weight - somewhat more in terms of energy in crops, since the meat component grows more than 50%). This contrasts with roughly doubling the planetary food production over the last 40 years. However, it's still an awful lot of extra food to produce - the required absolute increase in food production is similar in size to what has been achieved in the last forty years.

Let's now consider a variety of potential bottlenecks to achieving this kind of increase in food production. One major area of concern (water) I will reserve for its own future piece, but I address the other big potential constraints that I am aware of.

Land Use and Crop Yields

The doubling of global food production since the 1960s has not come about because of expanding cropland. The world has about 14.8 billion hectares of land area, and the uses of it over the last few decades are as follows:

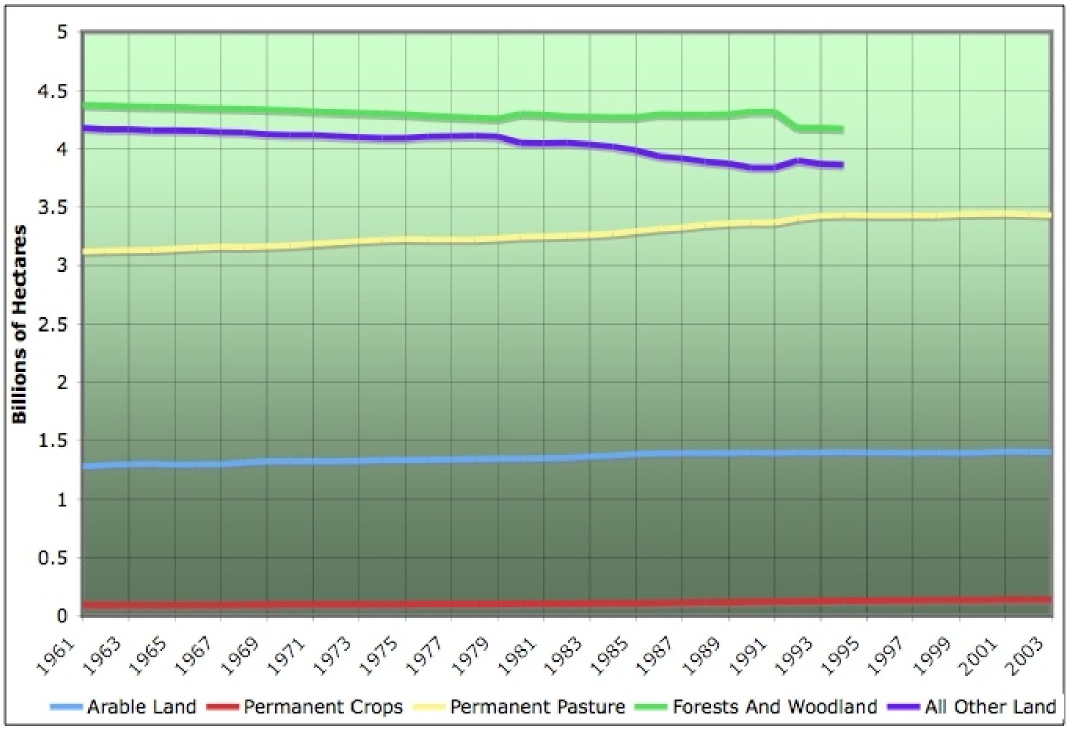

Major classes of global land use 1961-2003. Source: [FAO](#).

As you can see, the areas of cropland and pasture have increased slightly, at the expense of forests and other land, but the shifts are small percentage-wise. Instead, increased food production for the planet's extra billions of humans has largely come from big increases in agricultural yields.

I'm going to start with some yield data for the US, where we have long time series on yields for a number of crops. After that, we'll discuss the global situation. I have taken National Agricultural Statistics Service data on average US yields and reexpressed them on a common basis as a multiplier of the 1900-1935 average (or for those crops were the series doesn't start till after 1900, from whenever it does start until 1935).

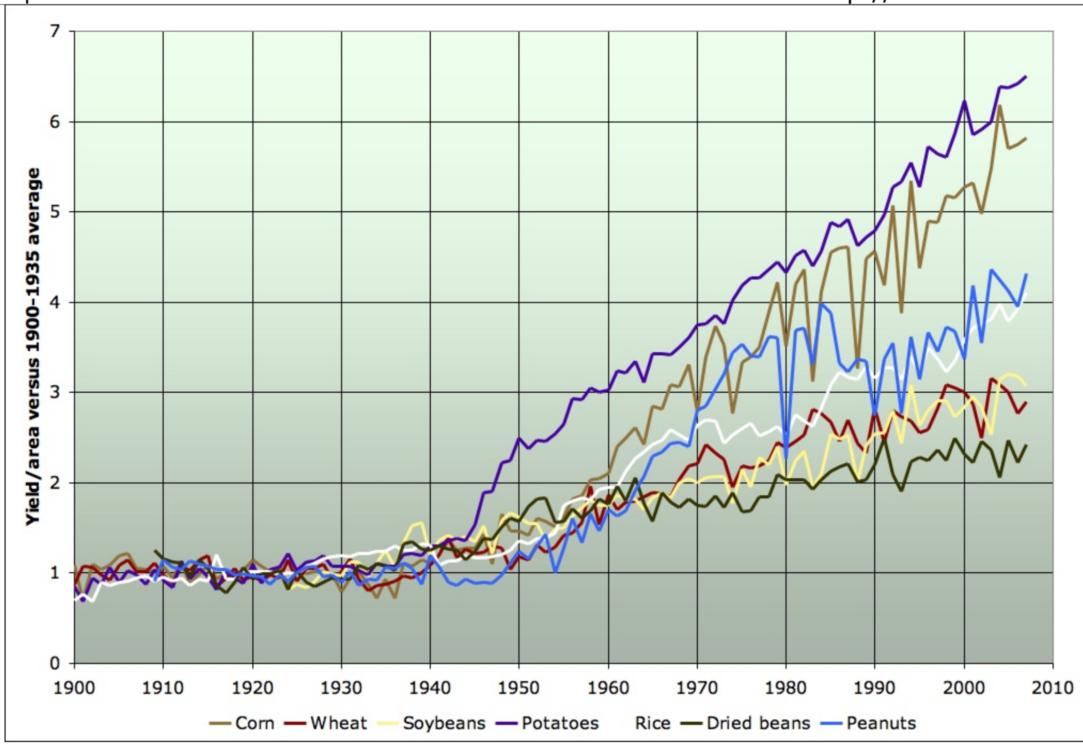

Average United States yields per unit area for selected crops, 1900-2007. Yields are expressed as a multiplier of the 1900-1935 average. Source: National Agricultural Statistics Service.

All the series show a roughly similar pattern. They were all fairly flat (with noise) until sometime in the late 1930s or 1940s. Then they all took off and began growing roughly linearly (again with noise). Modern yields are anywhere from 2.3 to 6.5 times greater than yields in the early twentieth century. Although some series have had periods of lagging for a decade or two (eg peanuts after 1983, dried beans - garbanzos and the like - after 1990), on the whole most of the series look like they are still increasing - there is no obvious pattern of yields flattening off yet. I encourage you to stare at this remarkable data for a long time. It's really worth thinking about the implications of it. Here are a few conclusions I draw.

Firstly, mechanization (and fossil-fuel powered machinery) are not the main cause of modern yields. Steam tractors were in widespread use in the late 1800s and early 1900s:

Steam Tractor in action in Ontario, 1916. Source: Ontario Govt Photo Archive.

The first gasoline powered tractor to be mass produced was introduced by Ford in 1917. Yet the yield take-off doesn't begin until 1940, and is almost certainly due to the agricultural innovations that comprise the green revolution. As The Future of Crop Yields and Cropped Area explains it:

> The Green Revolution strategy emerged from a surprising confluence of different lines of agricultural research (Evans, 1998) – the development of cheap nitrogenous fertilizers, of dwarf varieties of major cereals, and of effective weed control. Nitrogenous fertilizers increase crop production substantially, but make plants top-heavy, causing them to fall over. The development of dwarf varieties solves this problem, but at the cost of making plants highly susceptible to weeds, which grow higher than the dwarf plants, depriving them of light. The development of effective herbicides removed this problem. Further Green Revolution development focused on crop breeding to increase the harvest index – the ratio of the mass of grain to total above-ground biomass.

Secondly, anyone who wants to suggest that the world can be fed other than through industrial agriculture has some explaining to do about this data. Every crop shows yields prior to the green revolution that were flat and a small fraction of modern yields. If we returned to yields like that, either a lot of us would be starving, or we'd be terracing and irrigating most of the currently forested hillsides on the planet for food. While shopping for locally grown produce at your nearest organic farmer's market, stop and give a moment of thanks for the massive productivity of the industrial enterprise that brings you, or at least your fellow citizens, almost all of your calorie input.

Which raises a third important point. Food = Area Cropped x Average Yield. If average yields had not increased like this, humanity's impact on natural ecosystems would be much greater. It's true that industrial agriculture has a lot of impacts (nitrogen runoff and the like). However, the alternative would probably have been worse, since it would have required us to intensively exploit enormous areas of fragile, and currently less intensively exploited, land.

Fourthly, the period of greatest global warming, since 1950, coincides with the explosion of yields. I do not suggest that global warming *caused* increased yields. But at any rate, it would be hard to argue that industrial agriculture yields cannot grow rapidly in the face of the kind of warming we have seen to date: they just did

Well, is the global situation the same, or is this US data unrepresentative? I don't have access to as much data, but roughly, yes, it's the same:

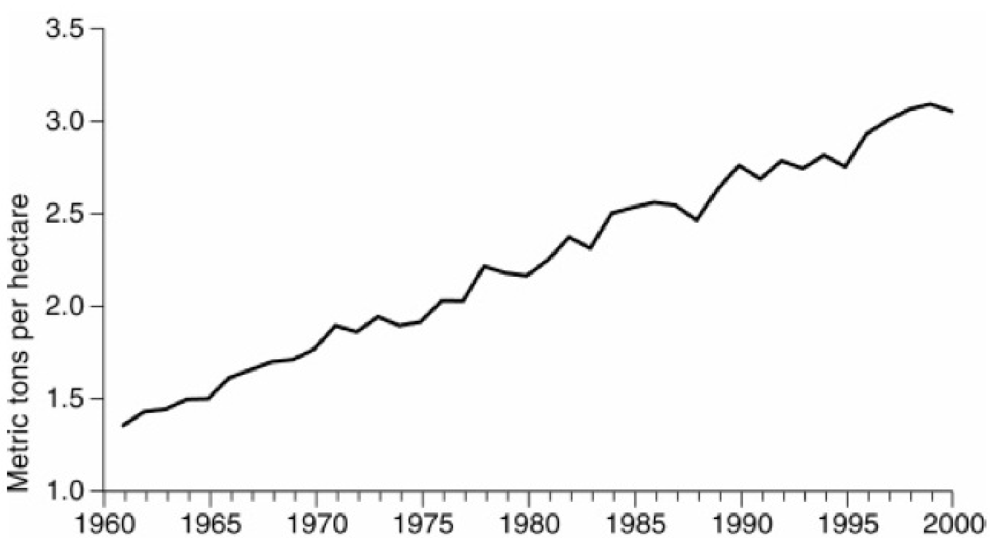

Average global cereal yields, 1961-2000. T. Dyson: World Food Trends: A Neo-Malthusian Prospect?, compiled from FAO data.

As you can see, global cereal yields are on the same roughly linear upward trajectory since 1961. Cereals are by far the most important food crop since not only do people eat a lot of them directly, but they also account for much of the input to the meat and dairy food groups that people eat, and thus are the base for the bulk of human calorie intake.

So obviously the critical question is whether or not yields can continue to increase in this manner? If we can just project out the linear increase than clearly a linearly increasing amount of food from a roughly constant amount of land is feasible, and humanity will be able to feed itself without having too much further impact on other ecosystems. On the other hand, if yields fail to increase, then we will be faced with unpleasant tradeoffs like trying to farm fairly unsuitable regions (think tropical rainforests, or the hilly parts of the western US), or not have enough food. So are we near some kind of theoretical yield limit?

Some people seem to think so. Lester Brown, who has been issuing alarming prognostications about food for several decades now, writes in Chapter 4 of his book *Outgrowing the Earth*

> Although the investment level in agricultural research, public and private, has not changed materially in recent years, the backlog of unused agricultural technology to raise land productivity is shrinking. In every farming community where yields have been rising rapidly, there comes a time when the rise slows and eventually levels off. For wheat growers in the United States and rice growers in Japan, for example, most of the available yield-raising technologies are already in use. Farmers in these countries are looking over the shoulders of agricultural researchers in their quest for new technologies to raise yields further. Unfortunately, they are not finding much.
>
> From 1950 to 1990 the world's grain farmers raised the productivity oftheir land by an unprecedented 2.1 percent a year, slightly faster than the 1.9 annual growth of world population during the same period. But from 1990 to 2000 this dropped to 1.2 percent per year, scarcely half as fast.

The argument in the second paragraph doesn't hold water to me. Population has been increasing pretty much linearly in recent decades, and agricultural yields have also been increasing pretty much linearly - I don't see any break from that pattern in the 1990-2000 decade. Of course, a linear rise will look like a dropping exponential growth rate, but Brown is careful to only point out the slowing in the yield growth rate. What he doesn't tell you is that world population growth had also dropped to only 1.4% during 1990-2000. In general, food prices until very recently were in a multi-decade secular decline, indicating that food production was not under serious supply-side constraint until the last few years:

Ratio of crude food/feed producer price index to all US consumer prices, Jan 1969-Dec 2007. Source: St Louis Fed.

And the argument in the first Brown paragraph I quoted doesn't seem to be how the agricultural scientists themselves are feeling. For example, Science reported last week:

> A decade ago, sequencing the maize genome was just too daunting. With 2.5 billion DNA bases, it rivaled the human genome in size and contained many repetitive regions that confounded the assembly of a final sequence. But last week, not one but three corn genomes, in various stages of completion, were introduced to the maize genetics community. In addition, researchers announced the availability of specially bred strains that will greatly speed tracking down genes involved in traits such as flowering time and disease resistance. These resources are ushering in a new era in maize genetics and should lead to tougher breeds, better yields, and biofuel alternatives. "We're sitting on very exciting times," says Geoff Graham, a plant breeder at Pioneer Hi-Bred International Inc.

The geneticists are well on the way to having complete genome sequences for thousands of corn varietals from all over the world. If I was a corn geneticist, I'd be pretty excited too.

A more grounded attempt to estimate the issue seems to be the FAO's discussion in World agriculture: towards 2015/2030:

> The slower growth in production projected for the next 30 years means that yields will not need to grow as rapidly as in the past. Growth in wheat yields is projected to slow to 1.1 percent a year in the next 30 years, while rice yields are expected to rise by only 0.9 percent per year.
>
> Nevertheless, increased yields will be required - so is the projected increase feasible? One way of judging is to look at the difference in performance between groups of countries. Some developing countries have attained very high crop yields. In 1997-99, for example, the top performing 10 percent had average wheat yields more than six times higher that those of the worst performing 10 percent and twice as high as the average in the largest producers, China, India and Turkey. For rice the gaps were roughly similar.
>
> National yield differences like these are due to two main sets of causes:
>
> Some of the differences are due to differing conditions of soil, climate and slope. In Mexico, for example, much of the country is arid or semi-arid and less than a fifth of the land cultivated to maize is suitable for improved hybrid varieties. As a result, the country's maize yield of 2.4 tonnes per ha is not much more than a quarter of the United States average. Yield gaps of this kind, caused by agro-ecological differences, cannot be narrowed.
>
> Other parts of the yield gap, however, are the result of differences in crop management practices, such as the amount of fertilizer used. These gaps can be narrowed, if it is economic for farmers to do so.
>
> To find out what progress in yields is feasible, it is necessary to distinguish between the

gaps that can be narrowed and those that cannot. A detailed FAO/IIASA study based on agro-ecological zones has taken stock of the amount of land in each country that is suitable, in varying degrees, for different crops. Using these data it is possible to work out a national maximum obtainable yield for each crop.

This maximum assumes that high levels of inputs and the best suited crop varieties are used for each area, and that each crop is grown on a range of land quality that reflects the national mix. It is a realistic figure because it is based on technologies already known and does not assume any major breakthroughs in plant breeding. If anything, it is likely to under-estimate maximum obtainable yields, because in practice crops will tend to be grown on the land best suited for them.

The maximum obtainable yield can then be compared with actual national average yield to give some idea of the yield gap that can be bridged. The study showed that even a technologically progressive country such as France is not yet close to reaching its maximum obtainable yield. France could obtain an average wheat yield of 8.7 tonnes per ha, rising to 11.6 tonnes per ha on her best wheat land, yet her actual average yield today is only 7.2 tonnes per ha.

For example:

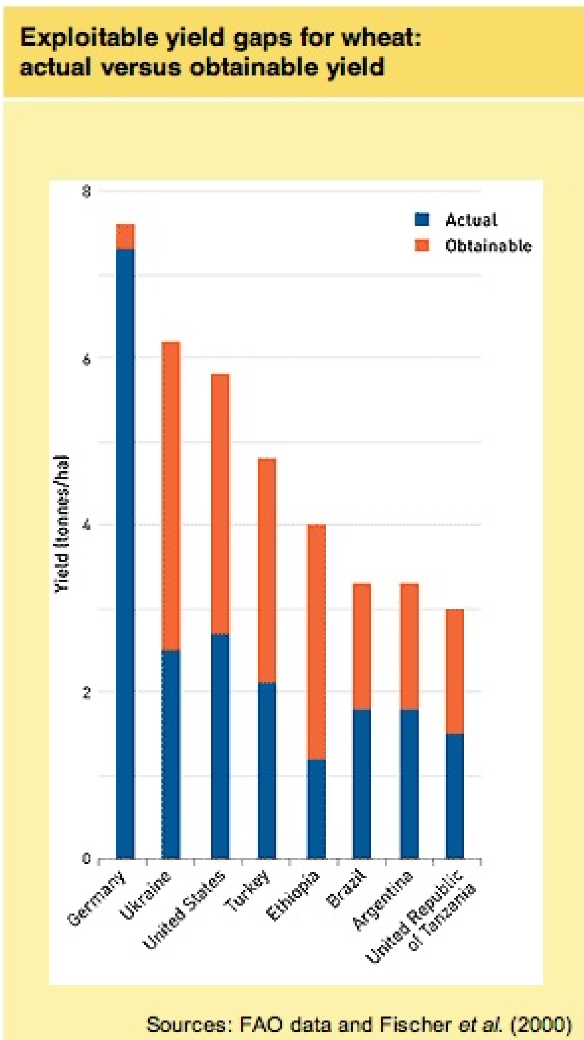

Gap between actual national yields and estimated yield with best currently known varietals and inputs. Source: FAO report, [World agriculture: towards 2015/2030](#)

And so,

> Similar yield gaps exist for most countries studied in this way. Only a few countries are actually achieving their maximum obtainable yield. When real prices rise, there is every reason to believe that farmers will work to bridge yield gaps. In the past, farmers with good access to technologies, inputs and markets have responded very quickly to higher prices. Argentina, for example, increased her wheat production by no less than 68 percent in just one year (1996), following price rises, although this was done mainly be extending the area under wheat. Where land is scarcer, farmers respond by switching to higher-yielding varieties and increasing their use of other inputs to achieve higher yields.
>
> It seems clear that, even if no more new technologies become available, there is still scope for increasing crop yields in line with requirements. Indeed, if just 11 of the countries that produce wheat, accounting for less than two-fifths of world production, were to bridge only half the gap between their maximum obtainable and their actual yields, then the world's wheat output would increase by almost a quarter.

Another way to try to get at the issue is to look at how current yields compare to the theoretical potential of photosynthesis. This is generally expressed as *net primary productivity* (NPP) - the amount of carbon that plants can fix, exclusive of that used to power their own respiration. The net primary productivity is the photosynthetic product that is available to be eaten by people and other animals, rot into the soil, etc. Here is a map of the fraction of net primary productivity appropriated by humans published by Haberl et al last year in the Proceedings of the National Academy of Sciences, which I take to be a decent representative of the state-of-the-art in this kind of calculation:

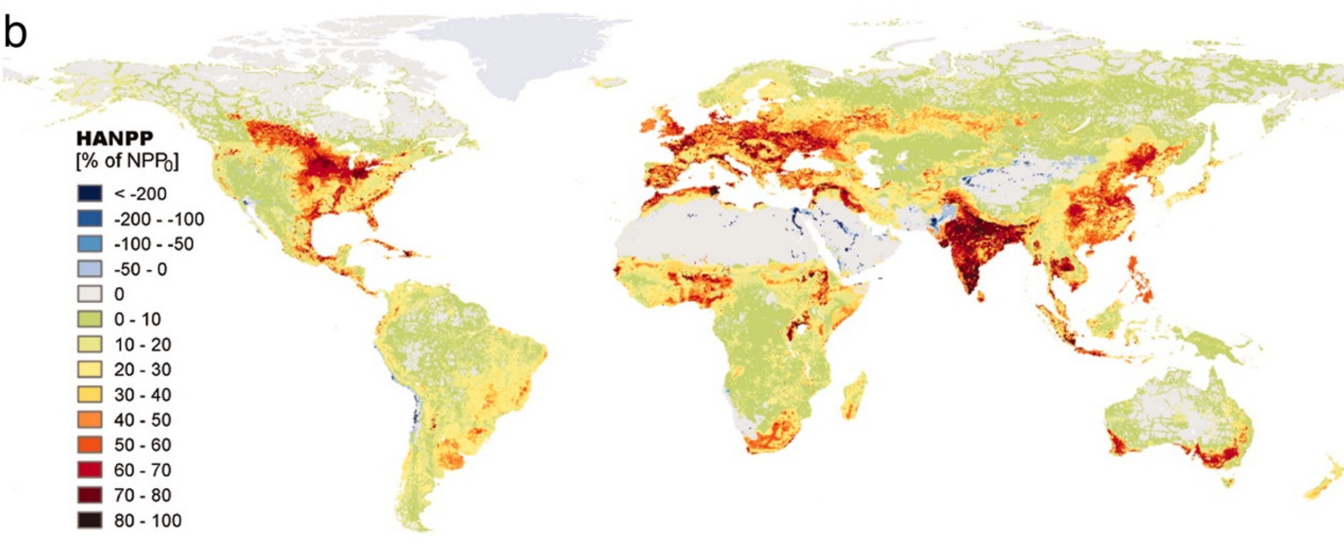

Global distribution of fraction of potential net primary productivity appropriated by humans. Source: Haberl et al: Quantifying and mapping the human appropriation of net primary production in earth's terrestrial ecosystems

You might look at the red - 60%-80% appropriation of NPP in many of the world's key crop growing areas, and think there wasn't enough head room for another 50%+ increase in yield in those areas. However, it's important to understand exactly how the accounting in these calculations is done. Let's consider a piece of the US midwest that used to be tall-grass prairie and is now under corn. What Haberl et al would do is first use a vegetation model (specifically, this one) to establish that it would be a prairie there absent human intervention, and figure out how much carbon the prairie would have fixed as NPP. That quantity they call NPP_0 (for that particular area - they compute NPP_0 for every cell in a global grid). So this is an estimate of the theoretical carbon fixation in the absence of any human influence. In particular, this is with the rainfall that falls naturally - carbon fixation in actual use could potentially exceed this if the crop was irrigated.

Then they would run the model again, but constrained to have cornfields rather than prairies. The carbon fixed by the model in that scenario would be NPP_{act}. Thus a model estimate of the actual carbon fixation in the actual human use of the area.

Next, they would figure out NPP_h which would be basically the carbon in the harvested corn based on national agricultural statistics (and in agricultural residues if those were harvested and statistically tracked also, but not likely in the case of corn). So NPP_h is the part that we humans really use (either by eating or feeding to our animals).

Given the actual NPP_{act}, and the NPP_h they would then compute the difference, NPP_t - basically the carbon in the corn stover which gets returned to the ground, eaten by mice, or whatever happens to it.

So then the human appropriation of net primary productivity (HANPP) is defined as $1 - NPP_t/NPP_0$. That is to say, if you look at the carbon that the prairie would have fixed, and then the carbon in the corn-stover, the difference is what is considered to be human appropriated. And that's the thing in the map that's 60-100% in the midwest (and other heavily utilized major cropland areas). However, this is **not** the same as the theoretical yield. In particular, a lot of the appropriated carbon comes about due to the difference between NPP_0 and NPP_{act} - the corn field doesn't fix as much carbon as the prairie, probably mainly because it starts the season out as bare soil and has to grow an annual crop from seed, instead of being a set of perennial grasses that can sprout from last year's roots and cover the available area in chlorophyll much faster.

Let's look at their Table 2 to make this clearer. This table shows the global breakdown of HANPP by food class. If we look at the "Cropping" category, we can see the different figures.

Table 2. Breakdown of global HANPP (excluding human-induced fires) in the year 2000 to land-use classes

Land use category	NPP_0, gC/m²/yr	NPP_{act}, gC/m²/yr	NPP_h, gC/m²/yr	NPP_t, gC/m²/yr	HANPP on this area, %	ΔNPP_{LC}, %	Contribution to total HANPP, %
Cropping	611	397	296	101	83.5	35.0	49.8
Grazing land	486	433	41	392	19.4	11.0	28.5
Forestry	720	720	48	673	6.6	0.0	10.6
Infrastructure areas	586	221	63	158	73.0	62.3	3.7
Wilderness	229	229	None	229	None	None	0.0
Global average or total	502	454	63	391	22.1	9.6	92.7*

*The remaining 7.3% are caused by human-induced fires (see Table 1).

Summary of human appropriation of net primary productivity. NPP_0 is modeled carbon fixation in wild condition. NPP_{act} is carbon fixation in actual human usage. NPP_h is carbon harvested or unfixed by harvest. NPP_t is residual carbon flowing into ecosystem. Source: Haberl et al: Quantifying and mapping the human appropriation of net primary production in earth's terrestrial ecosystems

As you can see, the average m^2 of cropfield (worldwide) would fix 0.6kg of carbon if it wasn't actually a field, but instead was covered in whatever the climactic climax vegatation is in that location. As a square meter of a field instead, it fixed 0.4kg of carbon, and of that humans got, on average 0.3kg as food and straw etc, leaving 0.1kg to go to the ground. So the HANPP is considered to be 5/6 (1 - 0.1/0.6). (The authors insist on three significant figures (83.5%), but I'm skeptical that the calculations are really that accurate). However, hopefully it should be clear by now that that doesn't mean there's a theoretical limit of only increasing yield by a further 1/5. Instead, there are multiple targets for the agronomists and geneticists to go after. The gap between the 0.4kg of NPP_{act} and the 0.6kg NPP_0 could be addressed with plants that had a longer growing season, covered the ground earlier, etc. To the extent some cropland is water-limited, irrigation could potentially increase the total NPP feasible. To the extent the 0.3kg of NPP_h is showing up as straw rather than food, then potentially that could be increased further.

A few decades down the road, one imagines heat-loving genetic mutant corn plants that pop up in the spring from perennial roots, promptly cover the ground with leaves that flatten themselves to the soil, and then start spitting out corn kernels, which can be harvested several times a year. It might not look much like a corn plant, but made into Doritos, people would probably still eat it (well, Americans would, anyway).

In short, another factor two of global cropland yields seems not to be ruled out on theoretical grounds. However, much more than that would appear to require the geneticists to come up with better photosynthesis (black plants basically - on which there has been no progress, as far as I understand).

Finally, it's worth mentioning that the FAO thinks there is considerable potential to use more land for agriculture:

> There is still potential agricultural land that is as yet unused. At present some 1.5 billion ha of land is used for arable and permanent crops, around 11 percent of the world's surface area. A new assessment by FAO and the International Institute for Applied Systems Analysis (IIASA) of soils, terrains and climates compared with the needs of and for major crops suggests that a further 2.8 billion ha are to some degree suitable for rainfed production. This is almost twice as much as is currently farmed.

Here's the breakdown for where the alleged potential cropland is:

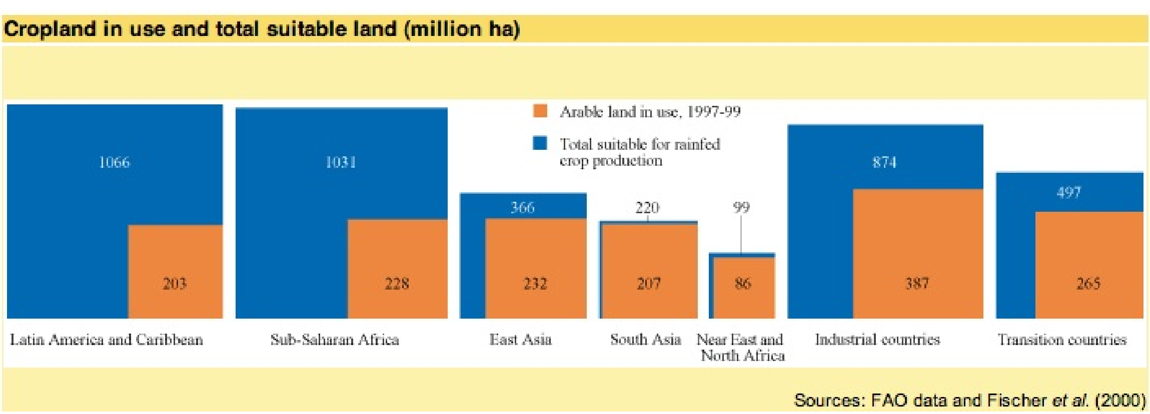

Regional breakdown of land considered available for cropping, compared to land in present use for that purpose. Source: FAO report World agriculture: towards 2015/2030

However, "much of the land reserve may have characteristics that make agriculture difficult, such as low soil fertility, high soil toxicity, high incidence of human and animal diseases, poor infrastructure, and hilly or otherwise difficult terrain." Caveat emptor!

If you look carefully at this figure - with the available land mainly in South America and Sub-saharan Africa, and the HANPP map above, you'll realize that much of what the FAO is talking about is cutting down the remaining tropical rainforests and using them for agriculture. I don't think that's a very good idea for a host of different reasons - better that we eat mutant corn, I think. The great bulk of the best land is almost certainly in production already.

Soil Loss

It appears to me that until recently, there has been a good deal of scientific confusion on the seriousness of soil erosion, estimates of the rate of erosion vary by more than an order of magnitude, and the overall data situation make global oil reserves look like a model of precision. As such, I don't think it's possible to make a clear evaluation of how near term the threat is globally. My best impression is that it's regionally quite severe, especially on fragile and marginal lands (dry, steep, or thin-soiled), but is probably not a near-term (next few decades) threat on the core agricultural regions from which most food comes (which tend to be flatter places with deep soils that don't erode quickly). It is certainly a major concern on the century timescale. However, there are many cultural practices that can help while still allowing good yields and, if I'm reading the literature correctly, erosion appears to be controllable, even within the context of fairly industrial styles of agriculture. Let me quickly sketch some of the debate.

The last global evaluation appears to have been GLASOD done by Oldeman et al and published in 1990. They produced a map which looks like this:

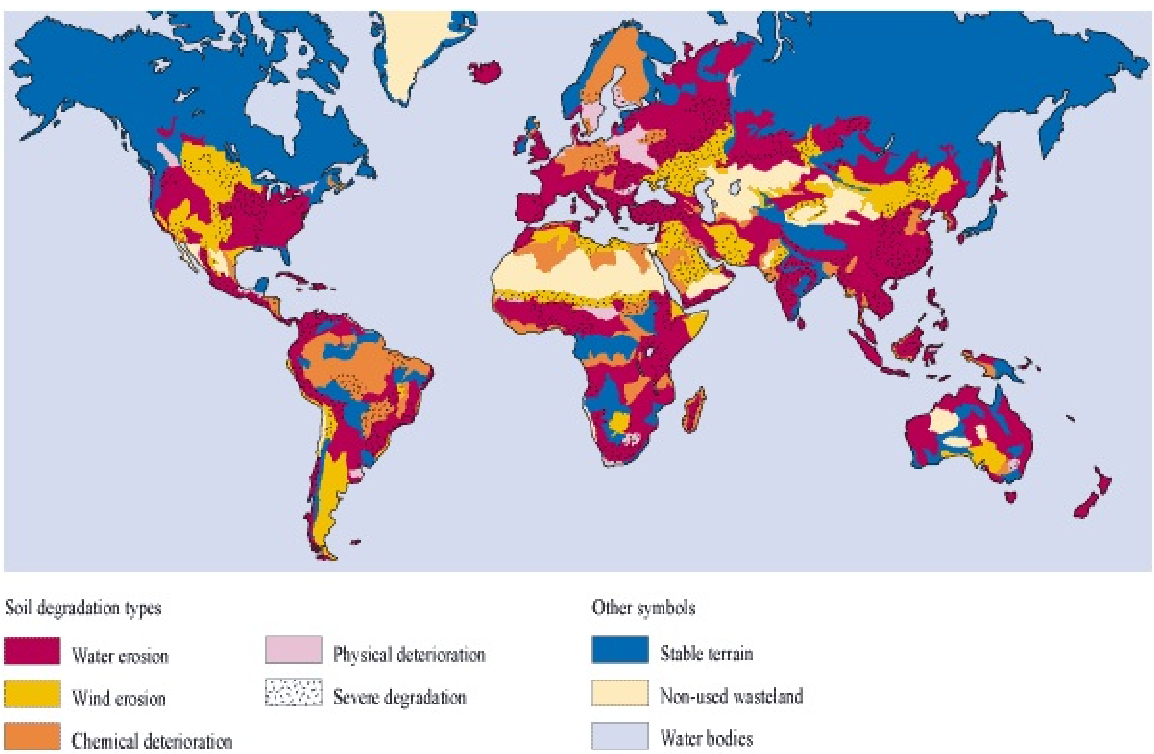

Global map of soil degradation. Source: GLASOD map, as shown in FAO report World agriculture: towards 2015/2030

This looks really bad - everywhere humans are, the soil is degraded, and much of the world's core crop land is in the "severely degraded" category. However, that did not yet have much noticeable effect on global yields, which have continued to increase by leaps and bounds since then. Moreover, this map was produced by what amounts to a survey of soil scientists, who used their subjective judgement. The instructions for filling out the questionaire describe how to set up the map cells, and then say:

> The next step involves evaluation of the degree, relative extent, recent past rate and causative factors for each type of human-induced soil degradation, as it may occur in the delineated physiographic unit. This evaluation process should be carried out in close cooperation with national and/or international experts with local knowledge of the region. The evaluation process results in a list of of human-induced soil degradation types per physiographic unit, ranking them in order of importance.

So this doesn't sound like a precise, quantitative sort of estimate. And more quantitative estimates are dogged with problems. A central issue is that most soil eroded from place A (let's say a steep field on the side of a valley) isn't necessarily lost to cultivation. Instead, it may end up in place B (let's say the flood plain of the river in the bottom of the valley) where it may still be of use in cultivation.

The US is the best measured place, in that we at least have a national agency charged with regular quantitative assessments of soil erosion (a legacy of the dustbowl years). The last assessment was the 2003 National Resources Inventory.

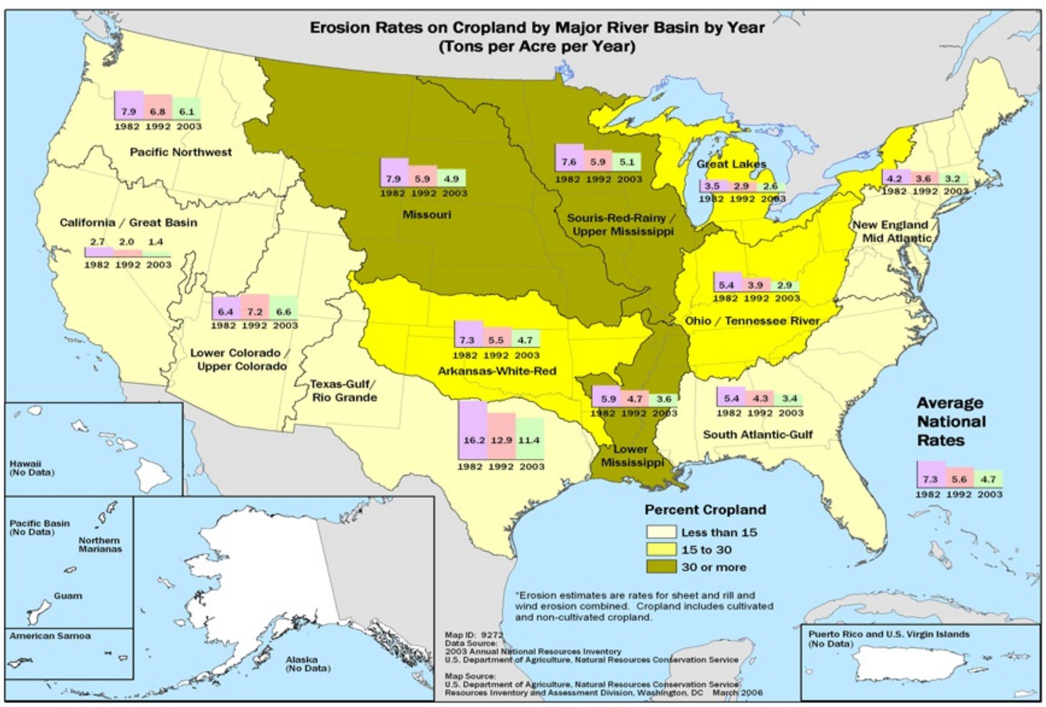

NRCS maps of US soil erosion in 2003. Source: US National Resources Conservation Service 2003 National Resources Inventory

These estimates are made by applying models (the Universal Soil Loss Equation and the Wind Erosion Equation) to topographical and climate data. The model inputs are things like the rainfall data, the slope of the field, the erodibility of the particular soil, etc. The overall amounts of erosion are decreasing, and the amount is not imminently scary. The current national average of 4.7 tons/acre/year corresponds to a little more than 1 kg/m^2/yr, which in turn is about 1mm/year, or an inch in twenty five years. That's not good, but doesn't sound like a likely disaster before 2050, particularly given that the rate of erosion is dropping quite rapidly.

However, these estimates in one way overstate the problem because the USLE and WEE are designed to assess how much soil is removed from its original location, but not where that soil goes. Most of it is unlikely to make it all the way out to the ocean, but instead end up somewhere else where it may be put to use. An extraordinary paper by [Trimble in 1999](#) assessed the details of where soil went in a single valley in Wisconsin by doing detailed samples and cross sections of the alluvial plains. His estimates of the trends and disposition of soil is as follows:

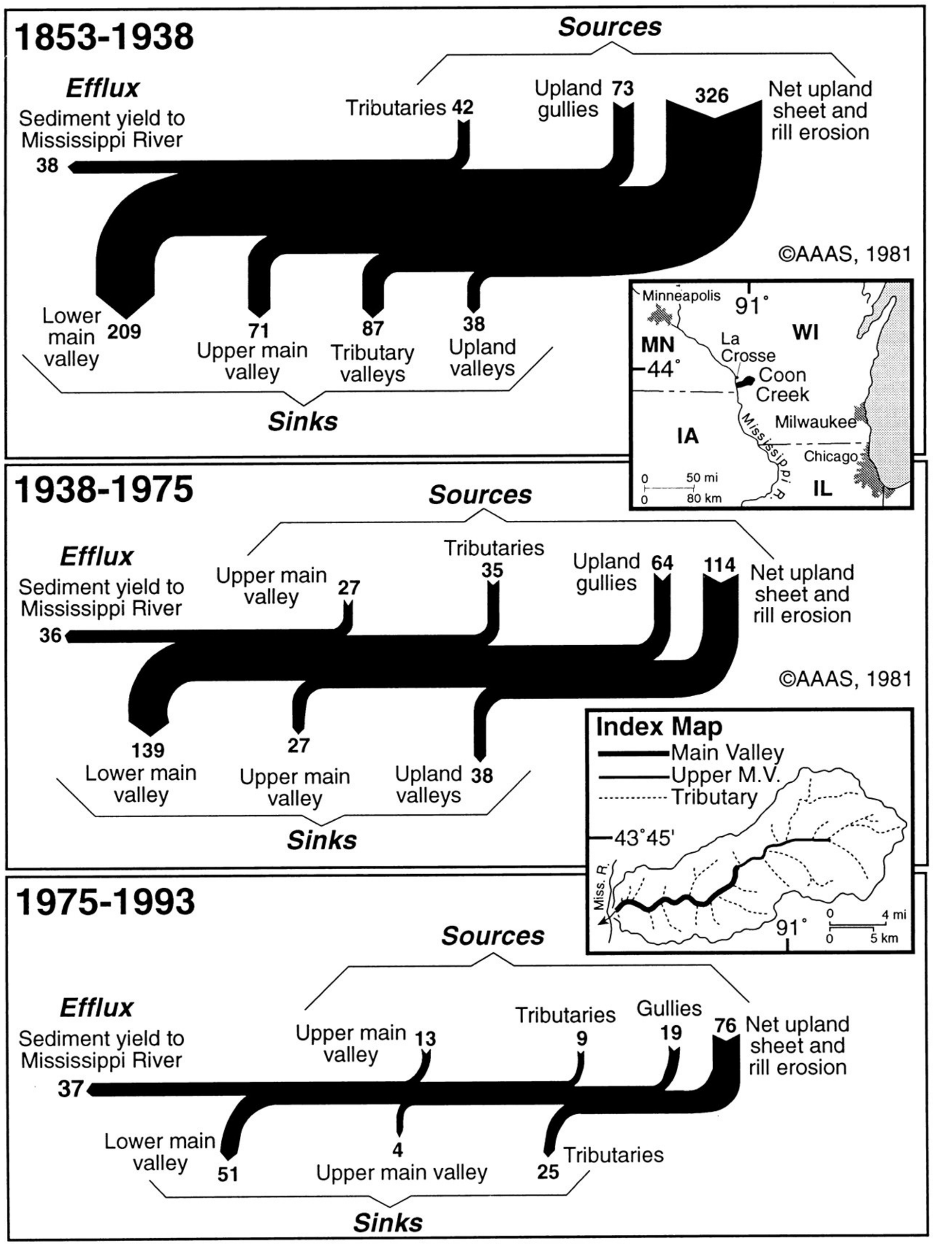

Disposition of soil erosion in Coon Creek watershed, Wisconsin. Source: S. Trimble [Decreased Rates of Alluvial Sediment Storage in the Coon Creek Basin, Wisconsin, 1975-93](#)

Clearly, the soil erosion is decreasing, but also, most of it hasn't gone that far, and, therefore, could potentially be put back at some point in the future if that becomes economically desirable.

Still, in the long term, it seems that eroding an inch every few decades from upland areas is certainly not sustainable, though it's not an imminent crisis either. In an important meta-analysis last year, D. Montgomery compiled erosion rates for a wide variety of situations and plotted the following cumulative density function for the probability of different erosion rates:

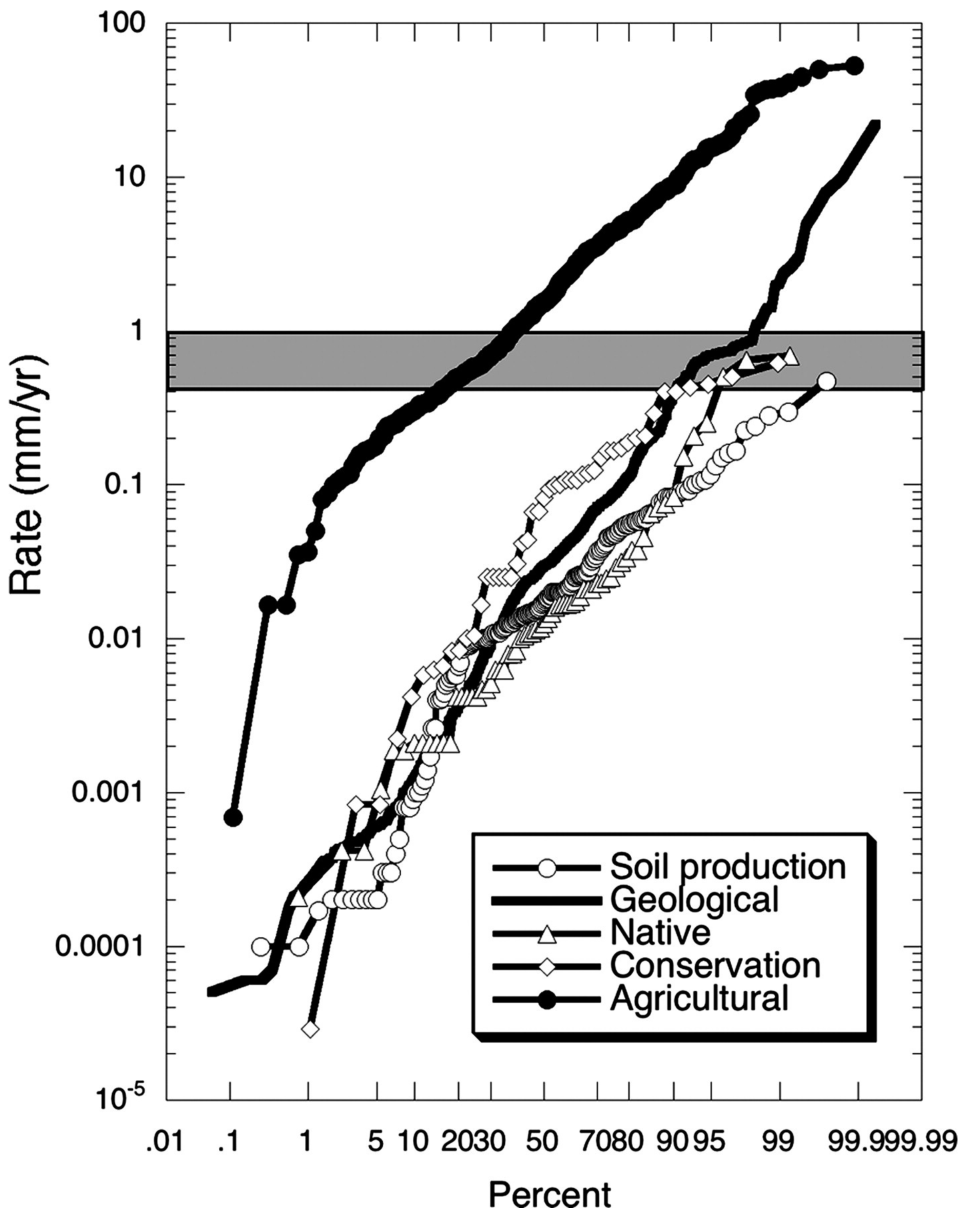

Cumulative distribution function of soil erosion and formation rates from numerous studies around the world. Hollow circles represent rates of soil formation, solid line is geological erosion rates, triangles are soil erosion rates under native vegetation, while diamonds are soil erosion rates under various conservation tillage methods (terracing or no-till agriculture). Solid circles represent plough-based agriculture. Source: D. Montgomery, Soil erosion and agricultural sustainability

The key things to note are these:

- Rates of soil production and erosion under native vegetation are roughly similar, suggesting soil depths are naturally in equilibrium.
- Rates of "agricultural" erosion are a couple of orders of magnitude higher, suggesting that ploughing is not a long-term proposition.
- Rates of "Conservation" erosion are roughly comparable to to natural erosion rates under native vegetation. This covers more sustainable management regimes such as terracing and no-till agriculture.

This suggests that the long-term sustainability of industrial agriculture requires the use of no-till farming systems in which ploughing is not done, crop residues are left on the field, and weeds are managed another way (primarily via herbicides today).

Fertilizer

The three major fertilizer nutrients applied in industrial agriculture are Nitrogen (N), Phosphorus (P), and Potassium (K). None appear to be a critical constraint on agriculture to the 2050 timeframe, though there are significant issues with nitrogen in the short term.

Nitrogen fertilizer is manufactured via the Haber-Bosch process in which nitrogen gas (which forms almost 80% of the atmosphere) is heated with hydrogen over an iron catalyst at high temperatures and pressures to form ammonia (NH_3) which is subsequently reacted with other compounds to form urea, ammonium sulphate, and other compounds used as fertilizer. Presently, almost all the hydrogen input to this process is produced by steam reformation of natural gas, and this is the cause of the short term problem since natural gas supplies are problematic, and likely to worsen with both Europe and North America probably at or past peak natural gas. Fertilizer manufacture is exiting these regions and moving to the Middle East, Trinidad, and other places with more natural gas.

However, in the long term, there's no reason nitrogen fertilizer has to be made from natural gas. I n my scenario in which energy production is dominated by renewable/nuclear electricity by 2050, the natural source of hydrogen for Haber-Bosch is by electrolyzing water. Producing nitrogen fertilizer is unproblematic as long as society has ample energy.

The reserves and reserve-base for phosphorus are enormous. According to the USGS, 2006 global production of phosphate rock was 145 million tons, while reserves were 18 billion tons, and the reserve base was 50 billion tons. For the 2050 timeframe, I consider reserve base to be the more appropriate number for the same reasons discussed under lithium. The reserve base for phosphate rock is 350 times larger than 2006 production, so there is no evidence of a problem at present.

Some bloggers are concerned that the Hubbert linearization suggests peak phosphorus has already past. However, Hubbert linearization is not very reliable if there is no independent evidence to suggest peak is at hand, due to the problem of dual peak structures giving rise to misleading linear regions (eg see the UK oil linearization). In this case, with enormous reserves, and stable phosphorus prices (they haven't varied outside the range of $27-$28/ton from 2002-2006), it seems very unlikely that phosphorus is in trouble. JD has made a similar point (snark warning).

Potassium comes from the mining of potash. The USGS estimates the global reserve base to be 550 times larger than current usage. So potassium is unlikely to limit civilization any time soon.

Fuel use in Farming and Food Transport

I don't have global statistics, but at least in the US, agriculture is a minor user of oil. In total, it only used 2.2% of oil in 2000. This contrasts with cars and light trucks, which used 40%, heavy trucks which used 12.7%, air travel at 6.7% etc. Since agriculture is such a critical industry, we can ensure it is preferred for oil usage.

Furthermore, all shipping trade only uses 2.5% of US oil use. Most of that is shipping things other than food, but the bulk of food transportation is in there too. Amongst critics of globalism, the image of strawberries being flown from Chile is a popular thing to pick on. However, things like strawberries form a miniscule fraction of our diet. A more representative image of global food trade would be a grain ship like this one:

Grain ship docked in Australia.

Shipping is extremely energy efficient - two orders of magnitude better per ton-mile than air freight. Thus, long-haul shipping of food will be cost effective long after oil has peaked. Ships can also be run on nuclear power, as the US navy has been demonstrating for decades.

In Conclusion

There seems to be reason for cautious optimism that if other global problems can be solved, food production will not be a critical constraint on civilization to 2050. If industrial agricultural yields maintain their historical trajectory, there will be enough food without needing much more land. In case yields fail to continue increasing, more land is potentially available globally, though likely of poor quality. Soil erosion is an important problem, but not a critical emergency, and can seemingly be solved permanently with no-till farming methods. Fertilizer does not appear to be seriously constrained in the long-term, though nitrogen fertilizer needs to be transitioned away from reliance on natural gas. Agriculture only needs a tiny fraction of global liquid fuel use to operate, and this can be maintained for a long time, since food production is a critical infrastructure.

However, if we were to keep growing the conversion of food into biofuels, all bets would be off.

Other sources

In addition to the sources linked directly above, I consulted the following references

- Smil, V. The Earth's Biosphere: Evolution, Dynamics, and Change (book).

- Wild, A. Soils, Land and Food: Managing the Land during the Twenty-First Century (book)
- Brown, L. Plan B 2.0: Rescuing a Planet Under Stress and a Civilization in Trouble (book)
- McKibben, B. Deep Economy: The Wealth of Communities and the Durable Future (book)
- T. Dyson: World food trends and prospects to 2025
- Millenium Ecosystem Assessment, Ecosystems And Human Well-Being: Desertification Synthesis
- P. Waggoner and J. Ausubel: How Much Will Feeding More and Wealthier People Encroach on Forests?
- Great Transition: The Promise and Lure of The Times Ahead.
- IPCC Special Report on Emissions Scenarios, Chapter 4 (An Overview of Scenarios).
- Climate Change 2007: Impacts, Adaptation and Vulnerability, Chapter 5 (Food, fibre and forest products)
- J. Pretty et al. Resource-Conserving Agriculture Increases Yields in Developing Countries
- J. Pretty and R. Hine Reducing Food Poverty with Sustainable Agriculture: A Summary of New Evidence
- D. Pimentel et al Environmental and Economic Costs of Soil Erosion and Conservation Benefits