Innovation in Hard Times?

Posted by Stuart Staniford on March 15, 2007 - 12:00pm
Topic: Economics/Finance
Tags: collapse, innovation, peak oil [list all tags]

26
diggs
digg it

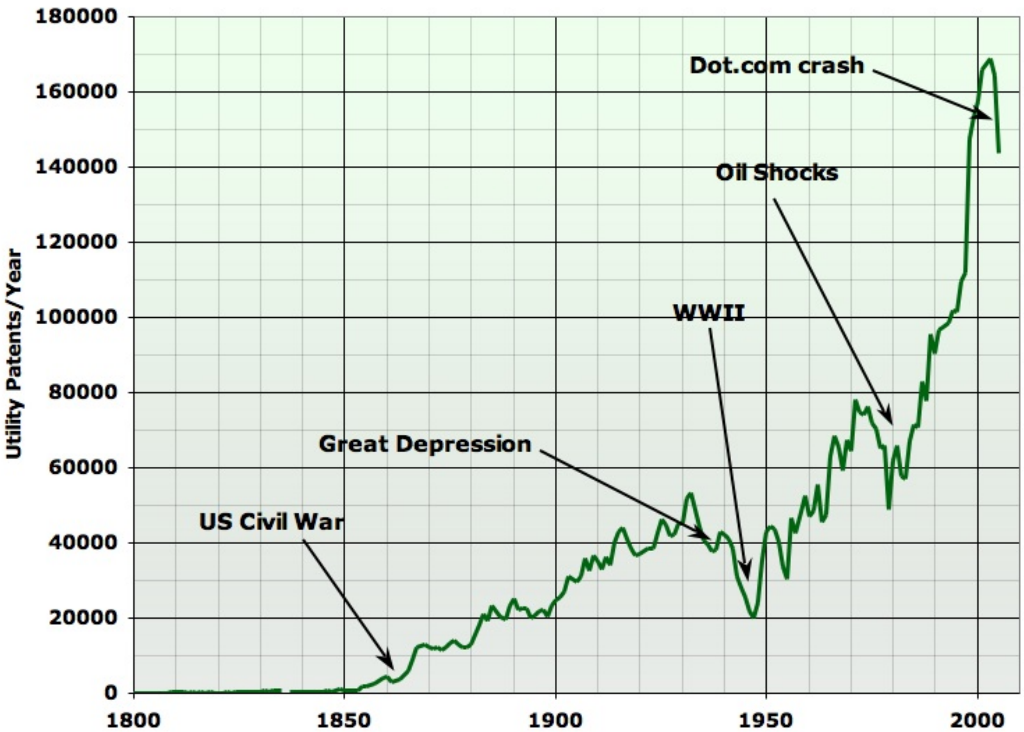

Utility Patents granted each year by the US Patent Office, with certain historical events added as annotations.

When thinking about what happens to society during a difficult time of some kind, one of the sharpest delineators between pessimists and optimists is their belief about the role of innovation. The optimists tend to assume that the can-do creative spirit of humanity will, always and everywhere, solve all problems, and thus the future will be ever brighter and brighter, with the possible exception of some brief and localized problems which will only serve to spur further innovation. The pessimists tend to assume that human innovation either a) doesn't occur at all, or b) generally makes things worse if it does.

I'm not quite sure where I fall on this spectrum yet, but it seems to me that, one way or another, as a society we are about to have a hard time here. Between the housing credit implosion, declines in Saudi oil production, the unknown but great damage that the Bush administration has done to the always fragile political arrangements in the economically critical Middle East, and monster hurricanes stomping on our cities, it's hard not to feel that the next decade is going to be one of the less fun ones in the historical record. I'm not going to commit myself on exactly how low the fun quotient is going to get - I really have no idea - but I'm sure not feeling good about the near to middle distance.

So it seems worth exploring further this question of the relationship between innovation and "hard times" of one kind and another. I am by no means an expert on either the history of innovation, or the economics of innovation, and one of the things I am hoping is to get links to the right things to read by those people who *are* such experts. However, I do have **some** perspective, as I have spent my working life in one way or another in "the innovation sector"; I have done stints in university research, venture capital funded startups, and also have some experience consulting on patent litigation. So I have a working knowledge of the way the US's formal institutions for innovation operate (and of course you could certainly argue that I have a vested interest in seeing them continue to operate).

Let me try to delineate the opposite ends of the spectrum on this question, starting with the pessimists. One of the people who has been most effective in getting people to at least think about the various crises facing society is Jim Kunstler whose pithy writing I envy every Monday morning. As far as I can tell, Kunstler doesn't think that innovation can play a significant role in solving society's problems, and is generally in favor of not doing further innovation on the technologies most important to society today, and instead wants to more-or-less roll back the various inventions of the 20th century, especially their impact on society, and roughly return to the 19th century:

> We have to produce food differently. The ADM / Monsanto / Cargill model of industrial agribusiness is heading toward its Waterloo. As oil and gas deplete, we will be left with sterile soils and farming organized at an unworkable scale. Many lives will depend on our ability to fix this. Farming will soon return much closer to the center of American economic life. It will necessarily have to be done more locally, at a smaller-and-finer scale, and will require more human labor. The value-added activities associated with farming -- e.g. making products like cheese, wine, oils -- will also have to be done much more locally. This situation presents excellent business and vocational opportunities for America's young people (if they can unplug their Ipods long enough to pay attention.) It also presents huge problems in land-use reform. Not to mention the fact that the knowledge and skill for doing these things has to be painstakingly retrieved from the dumpster of history. Get busy.
>
> We have to inhabit the terrain differently. Virtually every place in our nation organized for car dependency is going to fail to some degree. Quite a few places (Phoenix, Las Vegas, Miami....) will support only a fraction of their current populations. We'll have to return to traditional human ecologies at a smaller scale: villages, towns, and cities (along with a productive rural landscape). Our small towns are waiting to be reinhabited. Our cities will have to contract. The cities that are composed proportionately more of suburban fabric (e.g. Atlanta, Houston) will pose especially tough problems. Most of that stuff will not be fixed. The loss of monetary value in suburban property will have far-reaching ramifications. The stuff we build in the decades ahead will have to be made of regional materials found in nature -- as opposed to modular, snap-together, manufactured components -- at a more modest scale. This whole process will entail enormous demographic shifts and is liable to be turbulent. Like farming, it will require

> the retrieval of skill-sets and methodologies that have been forsaken. The graduate schools of architecture are still tragically preoccupied with teaching Narcissism. The faculties will have to be overthrown. Our attitudes about land-use will have to change dramatically. The building codes and zoning laws will eventually be abandoned and will have to be replaced with vernacular wisdom. Get busy.
>
> We have to move things and people differently. This is the sunset of Happy Motoring (including the entire US trucking system). Get used to it. Don't waste your society's remaining resources trying to prop up car-and-truck dependency. Moving things and people by water and rail is vastly more energy-efficient. Need something to do? Get involved in restoring public transit. Let's start with railroads, and let's make sure we electrify them so they will run on things other than fossil fuel or, if we have to run them partly on coal-fired power plants, at least scrub the emissions and sequester the CO_2 at as few source-points as possible. We also have to prepare our society for moving people and things much more by water. This implies the rebuilding of infrastructure for our harbors, and also for our inland river and canal systems -- including the towns associated with them. The great harbor towns, like Baltimore, Boston, and New York, can no longer devote their waterfronts to condo sites and bikeways. We actually have to put the piers and warehouses back in place (not to mention the sleazy accommodations for sailors). Right now, programs are underway to restore maritime shipping based on wind -- yes, sailing ships. It's for real. Lots to do here. Put down your Ipod and get busy.

Another contemporary pessimist is John Michael Greer, who operates under the title of "The Grand Archdruid of the Ancient Order of Druids in America". Many of you may wonder how on earth an Archdruid passed my initial credibility filters, but actually I read his blog avidly every week and consider him one of the most stimulating social critics I know of. Even though I often don't agree with him, I think he's worth the weight of ten more conventional religious leaders. And perhaps, like the Lorax, he speaks for an awful lot of trees, if not that many Druids.

Greer thinks that industrial society is pretty much fucked and there's nothing any of us can do. However, he believes we are likely to collapse back into a pre-industrial condition gradually over a period of a couple of centuries. His argument is mainly historical:

> Like modern industrial society, the Maya built their civilization on a nonrenewable resource base. In their case it was the fertility of fragile tropical soils, which couldn't support intensive corn farming forever. On that shaky foundation they built an extraordinary civilization with fine art, architecture, astronomy, mathematics, and a calendar more accurate than the one we use today. None of that counted when the crops began to fail. Mayan civilization disintegrated, cities were abandoned to the jungle, and the population of the Mayan heartland dropped by 90%.
>
> The parallels go deeper, for the Maya had other options. They could have switched from corn to more sustainable crops such as ramon nuts, or borrowed intensive wetland farming methods from their neighbors to the north. Neither of these happened, because corn farming was central to Maya political ideology. The power of the ahauob or "divine lords" who ruled Maya city-states depended on control of the corn crop, so switching crops or farming systems was unthinkable. Instead, Maya elites responded to crisis by launching wars to seize fields and corn from other city-states, making their decline and fall far more brutal than it had to be.

> Even so, the Maya decline wasn't a fast process. Maya cities weren't abandoned overnight, as archeologists of two generations ago mistakenly thought, but went under in a "rolling collapse" spread across a century and a half from 750 to 900. Outside the Maya heartland, the process took even longer. Chichen Itza far to the north still flourished long after cities such as Tikal and Bonampak were overgrown ruins, and Mayan city-states on a small scale survived in corners of the Yucatan right up to the Spanish conquest.
>
> Map the Maya collapse onto human lifespans and the real scale of the process comes through. A Maya woman born around 730 would have seen the crisis dawn, but the ahauob and their cities still flourished when she died of old age seventy years later. Her great-grandson, born around 800, grew up amid a disintegrating society, and the wars and crop failures of his time would have seemed ordinary to him. His great-granddaughter, born around 870, never knew anything but ruins sinking back into the jungle. When she and her family finally set out for a distant village, the last to leave their empty city, it would never have occurred to her that her quiet footsteps on a dirt path marked the end of a civilization.
>
> This same pattern repeats over and over again in history. Gradual disintegration, not sudden catastrophic collapse, is the way civilizations end. It usually takes somewhere between 150 and 350 years for a civilization to decline and fall. This casts a startling light on today's crisis. It took America two centuries of incremental change to transform itself from an agrarian society to its current status as an aging industrial behemoth. Now, with its resource base failing, it faces the common fate of civilizations. Yet if that fate follows its usual timeline, it could easily take two more centuries of incremental change to transform America to an agrarian society again.

And on technology, his view is that our main job now is to [figure out which ones we can salvage:]

> One of the most widely cited apocalyptic writers of my teen years, Roberto Vacca, argued in his book The Coming Dark Age that this extreme interdependence would prove to be the Achilles' heel of industrial society. His argument that too much interconnection among unstable systems would lead to cascading systems failures and the collapse of industrial civilization impressed the likes of Isaac Asimov, who contributed an introduction to the book. In retrospect, it proved to be embarrassingly wrong. Like so many others at that time, Vacca put the cart before the horse; the rising tide of interdependence and interconnection he saw moving through the industrial world was a reaction to improvements in information processing, not a force in its own right, and further developments along the same lines – especially the explosive growth in computer technology – proved more than adequate to keep the process moving.
>
> Still, Vacca was right to see the web of interconnections that unites today's industrial technology as a critical vulnerability. It's just that the vulnerability comes into play further along the arc of catabolic collapse. Many of today's technologies depend so completely on the support of an intact industrial system that they cannot operate without it. Many more could operate without it, at least in theory, but have been designed in a way that maximizes their dependence on other technologies and will have to be reengineered in a hurry as the fabric of the industrial system comes apart. A final set of technologies are largely or wholly independent of the system and can be expected to carry on without a hitch while industrial society comes apart around them.

> These three classes have an uncomfortable similarity to the three categories used by battlefield medics in the process known as triage. Triage — the word comes from French and means "trying" or "testing" – is a care-rationing process used when the number of wounded overwhelms the people and resources available to treat them. Incoming wounded are sorted out into three classes. The first consists of those who will die even if they get care. The second consists of those who will survive even if they receive no care. The third consists of those who will live if they get help but will die without it. In a triage situation, all available resources go to the third category. When the need for care outruns the available time and resources, this harsh but necessary logic maximizes the number of survivors.
>
> The coming of deindustrial society will require us to approach technology in much the same way. Technological triage requires more complex judgments than the battlefield variety, however. Not all technologies are of equal value for human survival; it won't do us any good to preserve video game technology, let's say, if by doing so we lose the ability to grow food. Some technologies necessarily depend on other technologies —firearms, for example, presuppose a certain level of metalworking ability. Finally, technological triage involves four categories, not three. Alongside technologies that can't be saved no matter what we do, technologies that are certain to be saved even if we do nothing, and technologies that will be saved if we act and lost if we do not, there are technologies that have gone out of existence but could be brought back and put into use if action is taken now.

Moving now to the optimists, the quintessential cornucopian was the late Professor Julian Simon, an economist whose book The Ultimate Resource 2 is must reading for anyone wanting to understand both sides of the debate. Essentially, Simon had the view that the condition of humanity had always improved everywhere, taken over any length of time, and that it always would in the future too. (Strangely enough, he didn't look at the Mayan's :-) To get a feeling for his thought, take this Cato institute essay:

> People have since antiquity worried about running out of natural resources--flint, game animals, what-have-you. Yet, amazingly, all the historical evidence shows that raw materials--all of them--have become less scarce rather than more. It is beyond any doubt that natural resource scarcity--as measured by the economically meaningful indicator of cost or price--has been decreasing rather than increasing in the long run for all raw materials, with only temporary and local exceptions. And there is no reason why this trend should not continue forever. The trend toward greater availability includes the most counterintuitive case of all--oil.
>
> Food is an especially important resource. The evidence is particularly strong that the trend in nutrition is benign despite rising population. The long-run price of food is down sharply, even relative to consumer products, as a result of increased productivity. And per person food consumption is up over the last 30 years. The increase of height in the West is another mark of improved nutrition.
>
> (Africa's food production per person is down, but in the 1990s, few people any longer claim that Africa's suffering has anything to do with a shortage of land or water or sun. Hunger in Africa clearly stems from civil wars and government interference with agriculture, which periodic droughts have made more murderous.)

> Only one important resource has shown a trend of increasing scarcity rather than increasing abundance. It is the most important and valuable resource of all--human beings. Certainly, there are more people on earth now than ever before. But if we measure the scarcity of people the same way that we measure the scarcity of other economic goods--by how much we must pay to obtain their services--we see that wages and salaries have been going up all over the world, in poor countries as well as in rich countries. The amount that one must pay to obtain the services of a barber or a professor has risen in India, just as the price of a barber or professor has risen in the United States over the decades. That increase in the price of people's services is a clear indication that people are becoming more scarce even though there are more of us.

And he viewed innovation as the central fount from which this bounty sprouted:

> How can it be that economic welfare grows over time along with population, instead of humanity's being reduced to misery and poverty as population grows and we use more and more resources? We need some theory to explain this controversion of common sense.
>
> The process operates as follows: More people and increased income cause problems in the short run--shortages and pollutions. Short-run scarcity raises prices and pollution causes outcries. Those problems present opportunity and prompt the search for solutions. In a free society solutions are eventually found, though many people seek and fail to find solutions at cost to themselves. In the long run the new developments leave us better off than if the problems had not arisen. This theory fits the facts of history.
>
> Technology exists now to produce in virtually inexhaustible quantities just about all the products made by nature--foodstuffs, oil, even pearls and diamonds--and make them cheaper in most cases than the cost of gathering them in their natural state. And the standard of living of commoners is higher today than that of royalty only two centuries ago--especially their health and life expectancy, and their mobility to all parts of the world.
>
> The extent to which the political-social-economic system provides personal freedom from government coercion is a crucial element in the economics of resources and population. Skilled persons require an appropriate social and economic framework that provides incentives for working hard and taking risks, enabling their talents to flower and come to fruition. The key elements of such a framework are economic liberty, respect for property, and fair and sensible rules of the market that are enforced equally for all.

To prove that not all cornucopians are economists, we can look at the case of Ray Kurzweil, a well-known computer visionary who believes that the rate of technological progress, which he views essentially as an extension of biological evolution by other means, has been accelerating at ever greater rates since the Cambrian. His book The Singularity is Near is also a must read for a well-rounded view of the subject. But he lays out his basic thesis here.

> Consider that the price-performance of computation has grown at a superexponential rate for over a century. The doubling time (of computes per dollar) was three years in

1900 and two years in the middle of the 20th century; and price-performance is now doubling each year. This progression has been remarkably smooth and predictable through five paradigms of computing substrate: electromechanical calculators, relay-based computers, vacuum tubes, transistors, and now several decades of Moore's Law (which is based on shrinking the size of key features on a flat integrated circuit). The sixth paradigm—three-dimensional molecular computing—is already beginning to work and is waiting in the wings. We see similar smooth exponential progressions in every other aspect of information technology, a phenomenon I call the law of accelerating returns.

Where is all this headed? It is leading inexorably to the intelligent universe that Jim Gardner envisions. Consider the following: As with all of the other manifestations of information technology, we are also making exponential gains in reverse-engineering the human brain. The spatial resolution in 3D volume of in-vivo brain scanning is doubling each year, and the latest generation of scanners is capable of imaging individual interneuronal connections and seeing them interact in real time. For the first time, we can see the brain create our thoughts, and also see our thoughts create our brain (that is, we create new spines and synapses as we learn). The amount of data we are gathering about the brain is doubling each year, and we are showing that we can turn this data into working models and simulations.

Already, about 20 regions of the human brain have been modeled and simulated. We can then apply tests to the simulations and compare these results to the performance of the actual human brain regions. These tests have had impressive results, including one of a simulation of the cerebellum, the region responsible for physical skill, and which comprises about half of the neurons in the brain. I make the case in my book (The Singularity is Near) that we will have models and simulations of all several hundred regions, including the cerebral cortex, within 20 years. Already, IBM is building a detailed simulation of a substantial portion of the cerebral cortex. The result of this activity will be greater insight into ourselves, as well as a dramatic expansion of the AI tool kit to incorporate all of the methods of human intelligence.

By 2029, sufficient computation to simulate the entire human brain, which I estimate at about 10^{16} (10 million billion) calculations per second (cps), will cost about a dollar. By that time, intelligent machines will combine the subtle and supple skills that humans now excel in (essentially our powers of pattern recognition) with ways in which machines are already superior, such as remembering trillions of facts accurately, searching quickly through vast databases, and downloading skills and knowledge.

But this will not be an alien invasion of intelligent machines. It will be an expression of our own civilization, as we have always used our technology to extend our physical and mental reach. We will merge with this technology by sending intelligent nanobots (blood-cell-sized computerized robots) into our brains through the capillaries to intimately interact with our biological neurons. If this scenario sounds very futuristic, I would point out that we already have blood-cell-sized devices that are performing sophisticated therapeutic functions in animals, such as curing Type I diabetes and identifying and destroying cancer cells. We already have a pea-sized device approved for human use that can be placed in patients' brains to replace the biological neurons destroyed by Parkinson's disease, the latest generation of which allows you to download new software to your neural implant from outside the patient.

> If you consider what machines are already capable of, and apply a billion-fold increase in price-performance and capacity of computational technology over the next quarter century (while at the same time we shrink the key features of both electronic and mechanical technology by a factor of 100,000), you will get some idea of what will be feasible in 25 years.
>
> By the mid-2040s, the nonbiological portion of the intelligence of our humanmachine civilization will be about a billion times greater than the biological portion (we have about 10^{26} cps among all human brains today; nonbiological intelligence in 2045 will provide about 10^{35} cps). Keep in mind that, as this happens, our civilization will be become capable of performing more ambitious engineering projects. One of these projects will be to keep this exponential growth of computation going. Another will be to continually redesign the source code of our own intelligence. We cannot easily redesign human intelligence today, given that our biological intelligence is largely hard-wired. But our future—largely nonbiological—intelligence will be able to apply its own intelligence to redesign its own algorithms.
>
> So what are the limits of computation? I show in my book that the ultimate one-kilogram computer (less than the weight of a typical notebook computer today) could perform about 10^{42} cps if we want to keep the device cool, and about 10^{50} cps if we allow it to get hot. By hot, I mean the temperature of a hydrogen bomb going off, so we are likely to asymptote to a figure just short of 10^{50} cps. Consider, however, that by the time we get to 10^{42} cps per kilogram of matter, our civilization will possess a vast amount of intelligent engineering capability to figure out how to get to 10^{43} cps, and then 10^{44} cps, and so on.
>
> So what happens then? Once we saturate the ability of matter and energy to support computation, continuing the ongoing expansion of human intelligence and knowledge (which I see as the overall mission of our human-machine civilization), will require converting more and more matter into this ultimate computing substrate, sometimes referred to as “computronium.”
>
> What is that limit? The overall solar system, which is dominated by the sun, has a mass of about 2×10^{30} kilograms. If we apply our 10^{50} cps per kilogram limit to this figure, we get a crude estimate of 10^{80} cps for the computational capacity of our solar system. There are some practical considerations here, in that we won’t want to convert the entire solar system into computronium, and some of it is not suitable for this purpose anyway. If we devoted 1/20th of 1 percent (.0005) of the matter of the solar system to computronium, we get capacities of 10^{69} cps for “cold” computing and 10^{77} cps for “hot” computing. I show in my book how we will get to these levels using the resources in our solar system within about a century.

About now, you may be thinking that the Archdruid sounds like a sane and moderate individual. And if I tell you that Kurzweil takes 250 different supplements daily because he's trying to keep his 56 year old body around long enough to become immortal when the technology become available in a couple of decades, you might really start to wonder. But there's no question that he has an enviable track record as an inventor, and is a frequent headline speaker where computer scientists congregate.

So what to think? The truth lies in the middle somewhere? For myself, I try to go where the data lead, and follow them whether they go left, right, or up the center of the garden path.

As a first cut at looking at the question of how much innovation happens in crises, I started by looking at the number of patents on inventions issued by the US patent office. Admittedly, this is a crude metric. Patents vary in quality and significance, and those variations may have trends not captured by the sheer quantity. Also, the things that the patents are *about* very likely vary over time, and those variations could be highly significant. Still, it's a place to start and reveals some interesting things, I think:

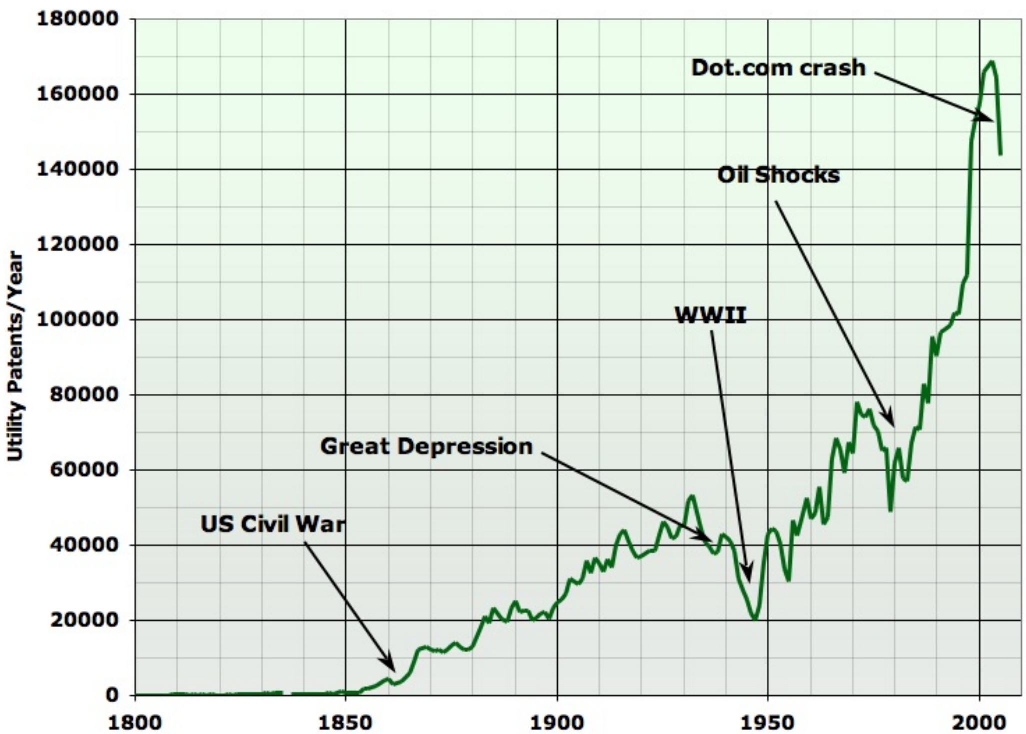

Utility Patents granted each year by the US Patent Office, with certain historical events added as annotations.

I think one could draw some support for both sides of the argument here. On the one hand, it's quite clear that the level of innovation goes **down** when society comes under stress. I have pointed to a few of the more prominent drops in the curve. So, presumably there is some level of societal stress so great that society's innovative institutions would cease to function. The worst case in the record was the combination of the depression in the 30s and WWII, which between them caused a roughly 60% drop in the rate of patent applications. The seventies oil shocks were around a 35% reduction, and the US civil war about a 25% reduction. So far, the post 2000 tech crash is about a 15% event, but one wonders what a housing crash and peak oil are going to add onto that in coming years (possibly after some intervening rebound - the application rate started shooting up again in 2005, but it often takes 2-4 years for an application to turn into a patent, if it is going to).

On the other hand, nothing in the last 200 years has been nearly enough to cause the patent application rate to drop to zero. And it seems likely that particularly relevant technologies may be invented in those eras of low overall patent productivity - one thinks of radar, jet engines, and nuclear power during WWII. I think there is also some support in the data for Professor Simon's idea that following a crisis, the innovation rate increases.

Indeed, I have to say that the overall rate of patent applications looks somewhat super-exponential, as Kurzweil would argue it should be. If we look just at the periods between crises, the interval from 1865 to 1932 had a combined annual growth rate (CAGR) of 3.29%/yr. Then from 1947 to 1974, the CAGR was 5.06%/yr, but it increased further to 5.6% between 1983 and 2003. That seems like evidence for a non-linear positive feedback loop.

But on the third hand, I am a lot less sanguine than Kurzweil that runaway positive feedback processes usually end well. Anyone who's been following the US housing market for the last few years, and especially the last few weeks, knows exactly what I mean.

On the fourth hand, I think those who want to argue that innovation will cease need to explain in far more detail **why** the institutions for it will stop operating. The Mayans didn't have venture capitalists, or patent offices, or universities, as far as we know. As we speak, those institutions are rapidly reorienting themselves towards our energy and climate problems:

> Out of the ashes of the Internet bust, many technology veterans have regrouped and found a new mission in alternative energy: developing wind power, solar panels, ethanol plants and hydrogen-powered cars.
>
> It is no secret that venture capitalists have begun pouring billions into energy-related start-ups with names like SunPower, Nanosolar and Lilliputian Systems.
>
> But that interest is now spilling over to many others in Silicon Valley — lawyers, accountants, recruiters and publicists, all developing energy-oriented practices to cater to the cause.
>
> The best and the brightest from leading business schools are pelting energy start-ups with résumés. And, of course, there are entrepreneurs from all backgrounds — but especially former dot-commers — who express a sense of wonder and purpose at the thought of transforming the $1 trillion domestic energy market while saving the planet.
>
> "It's like 1996," said Andrew Beebe, one of the remade Internet entrepreneurs. In the boom, he ran Bigstep.com, which helped small businesses sell online. Today, he is president of Energy Innovations, which makes low-cost solar panels. "The Valley has found a new hot spot."

I think a realistic argument that innovation will not be an enormous factor in our response to the crises we face needs to come to grips with the institutional differences between us and earlier societies.

SOME RIGHTS RESERVED This work is licensed under a Creative Commons Attribution-Share Alike 3.0 United States License.