

The great innovation debate

The Economist, Jan 12th 2013



Fears that innovation is slowing are exaggerated, but governments need to help it along

WITH the pace of technological change making heads spin, we tend to think of our age as the most innovative ever. We have smartphones and supercomputers, big data and nanotechnologies, gene therapy and stem-cell transplants. Governments, universities and firms together spend around \$1.4 trillion a year on R&D, more than ever before.

Yet nobody recently has come up with an invention half as useful as that depicted on our cover. With its clean lines and intuitive user interface, the humble loo transformed the lives of billions of people. And it wasn't just modern sanitation that sprang from late-19th and early-20th-century brains: they produced cars, planes, the telephone, radio and antibiotics.

Modern science has failed to make anything like the same impact, and this is why a growing band of thinkers claim that the pace of innovation has slowed. Interestingly, the gloomsters include not just academics such as Robert Gordon, the American economist who offered the toilet test of uninventiveness, but also entrepreneurs such as Peter Thiel, a venture capitalist behind Facebook.

If the pessimists are right, the implications are huge. Economies can generate growth by adding more stuff: more workers, investment and education. But sustained increases in output per person, which are necessary to raise incomes and welfare, entail using the stuff we already have in better ways—innovating, in other words. If the rate at which we innovate, and spread that innovation, slows down, so too, other things being equal, will our growth rate.

Doom, gloom and productivity figures

Ever since Malthus forecast that we would all starve, human ingenuity has proved the prophets of doom wrong. But these days the impact of innovation does indeed seem to be tailing off. Life expectancy in America, for instance, has risen more slowly since 1980 than in the early 20th century. The speed of travel, in the rich world at least, is often slower now than it was a generation earlier, after rocketing a century or so ago. According to Mr Gordon, productivity also supports the pessimists' case: it took off in the mid-19th century, accelerated in the early 20th century and held up pretty well until the early 1970s. It then dipped sharply, ticked up in late 1990s with computerisation and dipped again in the mid-2000s.

Yet that pattern is not as conclusively gloomy as the doomsayers claim. Life expectancy is still improving, even in the rich world. The productivity gains after electrification came not smoothly, but in spurts; and the drop-off since 2004 probably has more to do with the economic crisis than with underlying lack of invention. Moreover, it is too early to write off the innovative impact of the present age.

This generation's contribution to technological progress lies mostly in information technology (IT). Rather as electrification changed everything by allowing energy to be used far from where it was generated, computing and communications technologies transform lives and businesses by allowing people to make calculations and connections far beyond their unaided capacity. But as with electricity, companies will take time to learn how to use them, so it will probably be many decades before their full impact is felt.

Computing power is already contributing to dramatic advances far beyond the field of IT. Three-dimensional printing may cause a new industrial revolution. Autonomous vehicles, like the driverless cars produced by Google, could be common on streets within a decade. The performance of human prosthetics is rapidly catching up with that of natural limbs.

And although it is too soon to judge how big a deal these inventions will turn out to be, globalisation should make this a fruitful period for innovation. Many more brains are at work now than were 100 years ago: American and European inventors have been joined in the race to produce cool new stuff by Japanese, Brazilian, Indian and Chinese ones (see article).

Spend a penny—or two

So there are good reasons for thinking that the 21st century's innovative juices will flow fast. But there are also reasons to watch out for impediments. The biggest danger is government.

When government was smaller, innovation was easier. Industrialists could introduce new processes or change a product's design without a man from the ministry claiming some regulation had been broken. It is a good thing that these days pharmaceuticals are stringently tested and factory emissions controlled. But officialdom tends to write far more rules than are necessary for the public good; and thickets of red tape strangle innovation. Even many regulations designed to help innovation are not working well. The West's intellectual-property system, for instance, is a mess, because it grants too many patents of dubious merit.

The state has also notably failed to open itself up to innovation. Productivity is mostly stagnant in the public sector. Unions have often managed to prevent governments even publishing the performance indicators which, elsewhere, have encouraged managers to innovate. There is vast scope for IT to boost productivity in health care and education, if only those sectors were more open to change.

The rapid growth in the rich world before the 1970s was encouraged by public spending on infrastructure (including in sewage systems) and basic research: the computer, the internet and the green revolution in food technology all sprang out of science, where there was no immediate commercial aim. Wars provide the sharpest example of the innovative power of government spending: astounding new developments in drone and prosthetic technology—let alone the jet engine—are a bittersweet testament to that. Even in these straitened times, money should still be found for basic research into areas such as carbon capture and storage.

For governments that do these things well—get out of the way of entrepreneurs, reform their public sectors and invest wisely—the rewards could be huge. The risk that innovation may slow is a real one, but can be avoided. Whether it happens or not is, like most aspects of mankind's fate, up to him.

Has the ideas machine broken down?

The idea that innovation and new technology have stopped driving growth is getting increasing attention. But it is not well founded



BOOM times are back in Silicon Valley. Office parks along Highway 101 are once again adorned with the insignia of hopeful start-ups. Rents are soaring, as is the demand for fancy vacation homes in resort towns like Lake Tahoe, a sign of fortunes being amassed. The Bay Area was the birthplace of the semiconductor industry and the computer and internet companies that have grown up in its wake. Its wizards provided many of the marvels that make the world feel futuristic, from touch-screen phones to the instantaneous searching of great libraries to the power to pilot a drone thousands of miles away. The revival in its business activity since 2010 suggests progress is motoring on.

So it may come as a surprise that some in Silicon Valley think the place is stagnant, and that the rate of innovation has been slackening for decades. Peter Thiel, a founder of PayPal, an internet payment company, and the first outside investor in Facebook, a social network, says that innovation in America is “somewhere between dire straits and dead”. Engineers in all sorts of areas share similar feelings of disappointment. And a small but growing group of economists reckon the economic impact of the innovations of today may pale in comparison with those of the past.

Some suspect that the rich world’s economic doldrums may be rooted in a long-term technological stasis. In a 2011 e-book Tyler Cowen, an economist at George Mason University, argued that the financial crisis was masking a deeper and more disturbing “Great Stagnation”. It was this which explained why growth in rich-world real incomes and employment had long been slowing and, since 2000, had hardly risen at all (see chart 1). The various motors of 20th-century growth—some technological, some not—had played themselves out, and new technologies were not going to have the same invigorating effect on the economies of the future. For all its flat-screen dazzle and high-bandwidth pizzazz, it seemed the world had run out of ideas.

Glide path



The argument that the world is on a technological plateau runs along three lines. The first comes from growth statistics. Economists divide growth into two different types, “extensive” and “intensive”. Extensive growth is a matter of adding more and/or better labour, capital and resources. These are the sort of gains that countries saw from adding women to the labour force in greater numbers and increasing workers’ education. And, as Mr Cowen notes, this sort of growth is subject to diminishing returns: the first addition will be used where it can do most good, the tenth where it can do the tenth-most good, and so on. If this were the only sort of growth there was, it would end up leaving incomes just above the subsistence level.


Intensive growth is powered by the discovery of ever better ways to use workers and resources. This is the sort of growth that allows continuous improvement in incomes and welfare, and enables an economy to grow even as its population decreases. Economists label the all-purpose improvement factor responsible for such growth “technology”—though it includes things like better laws and regulations as well as technical advance—and measure it using a technique called “growth accounting”. In this accounting, “technology” is the bit left over after calculating the effect on GDP of things like labour, capital and education. And at the moment, in the rich world, it looks like there is less of it about. Emerging markets still manage fast growth, and should be able to do so for some time, because they are catching up with technologies already used elsewhere. The rich world has no such engine to pull it along, and it shows.

This is hardly unusual. For most of human history, growth in output and overall economic welfare has been slow and halting. Over the past two centuries, first in Britain, Europe and America, then elsewhere, it took off. In the 19th century growth in output per person—a useful general measure of an economy’s productivity, and a good guide to growth in incomes—accelerated steadily in Britain. By 1906 it was more than 1% a year. By the middle of the 20th century, real output per person in America was growing at a scorching 2.5% a year, a pace at which productivity and incomes double once a generation (see chart 2). More than a century of increasingly powerful and sophisticated machines were obviously a part of that story, as was the rising amount of fossil-fuel energy available to drive them.



But in the 1970s America’s growth in real output per person dropped from its post-second-world-war peak of over 3% a year to just over 2% a year. In the 2000s it tumbled below 1%. Output per worker per hour shows a similar pattern, according to Robert Gordon, an economist at Northwestern University: it is pretty good for most of the 20th

century, then slumps in the 1970s. It bounced back between 1996 and 2004, but since 2004 the annual rate has fallen to 1.33%, which is as low as it was from 1972 to 1996. Mr Gordon muses that the past two centuries of economic growth might actually amount to just “one big wave” of dramatic change rather than a new era of uninterrupted progress, and that the world is returning to a regime in which growth is mostly of the extensive sort (see chart 3).



Mr Gordon sees it as possible that there were only a few truly fundamental innovations—the ability to use power on a large scale, to keep houses comfortable regardless of outside temperature, to get from any A to any B, to talk to anyone you need to—and that they have mostly been made. There will be more innovation—but it will not change the way the world works in the way electricity, internal-combustion engines, plumbing, petrochemicals and the telephone have. Mr Cowen is more willing to imagine big technological gains ahead, but he thinks there are no more low-hanging fruit. Turning terabytes of genomic knowledge into medical benefit is a lot harder than discovering and mass producing antibiotics.

The pessimists' second line of argument is based on how much invention is going on. Amid unconvincing appeals to the number of patents filed and databases of “innovations” put together quite subjectively, Mr Cowen cites interesting work by Charles Jones, an economist at Stanford University. In a 2002 paper Mr Jones studied the contribution of different factors to growth in American per-capita incomes in the period 1950-93. His work indicated that some 80% of income growth was due to rising educational attainment and greater “research intensity” (the share of the workforce labouring in idea-generating industries). Because neither factor can continue growing ceaselessly, in the absence of some new factor coming into play growth is likely to slow.

The growth in the number of people working in research and development might seem to contradict this picture of a less inventive economy: the share of the American economy given over to R&D has expanded by a third since 1975, to almost 3%. But Pierre Azoulay of MIT and Benjamin Jones of Northwestern University find that, though there are more people in research, they are doing less good. They reckon that in 1950 an average R&D worker in America contributed almost seven times more to “total factor productivity”—essentially, the contribution of technology and innovation to growth—that an R&D worker in 2000 did. One factor in this may be the “burden of knowledge”: as ideas accumulate it takes ever longer for new thinkers to catch up with the frontier of their scientific or technical speciality. Mr Jones says that, from 1985 to 1997 alone, the typical “age at first innovation” rose by about one year.

A fall of moon dust

The third argument is the simplest: the evidence of your senses. The recent rate of progress seems slow compared with that of the early and mid-20th century. Take kitchens. In 1900 kitchens in even the poshest of households were primitive things. Perishables were kept cool in ice boxes, fed by blocks of ice delivered on horse-drawn wagons. Most households lacked electric lighting and running water. Fast forward to 1970 and middle-class kitchens in America and Europe feature gas and electric hobs and ovens, fridges, food processors, microwaves and dishwashers. Move forward another 40 years, though, and things scarcely change. The gizmos are more numerous and digital displays ubiquitous, but cooking is done much as it was by grandma.

Or take speed. In the 19th century horses and sailboats were replaced by railways and steamships. Internal-combustion engines and jet turbines made it possible to move more and more things faster and faster. But since the 1970s humanity has been coasting. Highway travel is little faster than it was 50 years ago; indeed, endemic congestion has many cities now investing in trams and bicycle lanes. Supersonic passenger travel has been abandoned. So, for the past 40 years, has the moon.

Medicine offers another example. Life expectancy at birth in America soared from 49 years at the turn of the 20th century to 74 years in 1980. Enormous technical advances have occurred since that time. Yet as of 2011 life expectancy rested at just 78.7 years. Despite hundreds of billions of dollars spent on research, people continue to fall to cancer, heart disease, stroke and organ failure. Molecular medicine has come nowhere close to matching the effects of improved sanitation.

To those fortunate enough to benefit from the best that the world has to offer, the fact that it offers no more can disappoint. As Mr Thiel and his colleagues at the Founders Fund, a venture-capital company, put it: "We wanted flying cars, instead we got 140 characters." A world where all can use Twitter but hardly any can commute by air is less impressive than the futures dreamed of in the past.

The first thing to point out about this appeal to experience and expectation is that the science fiction of the mid-20th century, important as it may have been to people who became entrepreneurs or economists with a taste for the big picture, constituted neither serious technological forecasting nor a binding commitment. It was a celebration through extrapolation of then current progress in speed, power and distance. For cars read flying cars; for battlecruisers read space cruisers.

Technological progress does not require all technologies to move forward in lock step, merely that some important technologies are always moving forward. Passenger aeroplanes have not improved much over the past 40 years in terms of their speed. Computers have sped up immeasurably. Unless you can show that planes matter more, to stress the stasis over the progress is simply a matter of taste.

Mr Gordon and Mr Cowen do think that now-mature technologies such as air transport have mattered more, and play down the economic importance of recent innovations. If computers and the internet mattered to the economy—rather than merely as rich resources for intellectual and cultural exchange, as experienced on Mr Cowen's popular blog, *Marginal Revolution*—their effect would be seen in the figures. And it hasn't been.

As early as 1987 Robert Solow, a growth theorist, had been asking why "you can see the computer age everywhere but in the productivity statistics". A surge in productivity growth that began in the mid-1990s was seen as an encouraging sign that the computers were at last becoming visible; but it faltered, and some, such as Mr Gordon, reckon that the benefits of information technology have largely run their course. He notes that, for all its inhabitants' Googling and Skypeing, America's productivity performance since 2004 has been worse than that of the doldrums from the early 1970s to the early 1990s.

The fountains of paradise

Closer analysis of recent figures, though, suggests reason for optimism. Across the economy as a whole productivity did slow in 2005 and 2006—but productivity growth in manufacturing fared better. The global financial crisis and its aftermath make more recent data hard to interpret. As for the strong productivity growth in the late 1990s, it may have been premature to see it as the effect of information technology making all sorts of sectors more productive. It now looks as though it was driven just by the industries actually making the computers, mobile phones and the like. The effects on the productivity of people and companies buying the new technology seem to have begun appearing in the 2000s, but may not yet have come into their own. Research by Susanto Basu of Boston College and John Fernald of the San Francisco Federal Reserve suggests that the lag between investments in information-and-communication technologies and improvements in productivity is between five and 15 years. The drop in productivity in 2004, on that reckoning, reflected a state of technology definitely pre-Google, and quite possibly pre-web.

Full exploitation of a technology can take far longer than that. Innovation and technology, though talked of almost interchangeably, are not the same thing. Innovation is what people newly know how to do. Technology is what they are actually doing; and that is what matters to the economy. Steel boxes and diesel engines have been around since the 1900s, and their use together in containerised shipping goes back to the 1950s. But their great impact as the backbone of global trade did not come for decades after that.

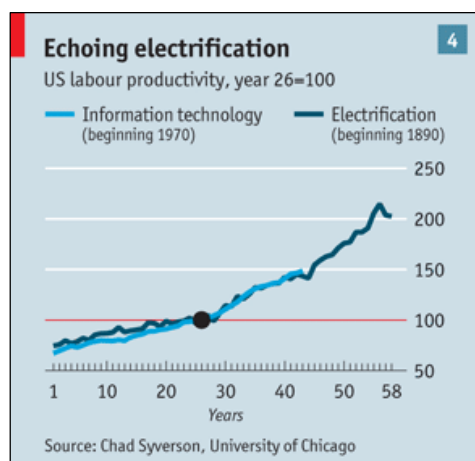
Roughly a century lapsed between the first commercial deployments of James Watt's steam engine and steam's peak contribution to British growth. Some four decades separated the critical innovations in electrical engineering of the 1880s and the broad influence of electrification on economic growth. Mr Gordon himself notes that the innovations of the late 19th century drove productivity growth until the early 1970s; it is rather uncharitable of him to assume that the post-2004 slump represents the full exhaustion of potential gains from information technology.

And information innovation is still in its infancy. Ray Kurzweil, a pioneer of computer science and a devotee of exponential technological extrapolation, likes to talk of "the second half of the chess board". There is an old fable in which a gullible king is tricked into paying an obligation in grains of rice, one on the first square of a chessboard, two on the second, four on the third, the payment doubling with every square. Along the first row, the obligation is minuscule. With half the chessboard covered, the king is out only about 100 tonnes of rice. But a square before reaching the end of the seventh row he has laid out 500m tonnes in total—the whole world's annual rice production. He will have to put more or less the same amount again on the next square. And there will still be a row to go.

Erik Brynjolfsson and Andrew McAfee of MIT make use of this image in their e-book "Race Against the Machine". By the measure known as Moore's law, the ability to get calculations out of a piece of silicon doubles every 18 months. That growth rate will not last for ever; but other aspects of computation, such as the capacity of algorithms to handle data, are also growing exponentially. When such a capacity is low, that doubling does not matter. As soon as it matters at all, though, it can quickly start to matter a lot. On the second half of the chessboard not only has the cumulative effect of innovations become large, but each new iteration of innovation delivers a technological jolt as powerful as all previous rounds combined.

The other side of the sky

As an example of this acceleration-of-effect they offer autonomous vehicles. In 2004 the Defence Advanced Research Projects Agency (DARPA), a branch of America's Department of Defence, set up a race for driverless cars that promised \$1 million to the team whose vehicle finished the 240km (150-mile) route fastest. Not one of the robotic entrants completed the course. In August 2012 Google announced that its fleet of autonomous vehicles had completed some half a million kilometres of accident-free test runs. Several American states have passed or are weighing regulations for driverless cars; a robotic-transport revolution that seemed impossible ten years ago may be here in ten more.



That only scratches the surface. Across the board, innovations fuelled by cheap processing power are taking off. Computers are beginning to understand natural language. People are controlling video games through body movement alone—a technology that may soon find application in much of the business world. Three-dimensional printing is capable of churning out an increasingly complex array of objects, and may soon move on to human tissues and other organic material.

An innovation pessimist could dismiss this as “jam tomorrow”. But the idea that technology-led growth must either continue unabated or steadily decline, rather than ebbing and flowing, is at odds with

history. Chad Syverson of the University of Chicago points out that productivity growth during the age of electrification was lumpy. Growth was slow during a period of important electrical innovations in the late 19th and early 20th centuries; then it surged. The information-age trajectory looks pretty similar (see chart 4).


It may be that the 1970s-and-after slowdown in which the technological pessimists set such store can be understood in this way—as a pause, rather than a permanent inflection. The period from the early 1970s to the mid-1990s may simply represent one in which the contributions of earlier major innovations were exhausted while computing, biotechnology, personal communication and the rest of the technologies of today and tomorrow remained too small a part of the economy to influence overall growth.

Other potential culprits loom, however—some of which, worryingly, might be permanent in their effects. Much of the economy is more heavily regulated than it was a century ago. Environmental protection has provided cleaner air and water, which improve people's lives. Indeed, to the extent that such gains are not captured in measurements of GDP, the slowdown in progress from the 1970s is overstated. But if that is so, it will probably continue to be so for future technological change. And poorly crafted regulations may unduly raise the cost of new research, discouraging further innovation.

Another thing which may have changed permanently is the role of government. Technology pessimists rarely miss an opportunity to point to the Apollo programme, crowning glory of a time in which government did not simply facilitate new innovation but provided an ongoing demand for talent and invention. This it did most reliably through the military-industrial complex of which Apollo was a spectacular and peculiarly inspirational outgrowth. Mr Thiel is often critical of the venture-capital industry for its lack of interest in big, world-changing ideas. Yet this is mostly a response to market realities. Private investors rationally prefer modest business models with a reasonably short time to profit and cash out.

A third factor which might have been at play in both the 1970s and the 2000s is energy. William Nordhaus of Yale University has found that the productivity slowdown which started in the 1970s radiated outwards from the most energy-intensive sectors, a product of the decade's oil shocks. Dear energy may help explain the productivity slowdown of the 2000s as well. But this is a trend that one can hope to see reversed. In America, at least, new technologies are eating into those high prices. Mr Thiel is right to reserve some of his harshest criticism for the energy sector's lacklustre record on innovation; but given the right market conditions it is not entirely hopeless.

Perhaps the most radical answer to the problem of the 1970s slowdown is that it was due to globalisation. In a somewhat whimsical 1987 paper, Paul Romer, then at the University of Rochester, sketched the possibility that, with more workers available in developing countries, cutting labour costs in rich ones became less important. Investment in productivity was thus sidelined. The idea was heretical among macroeconomists, as it dispensed with much of the careful theoretical machinery then being used to analyse growth. But as Mr Romer noted, economic historians comparing 19th-century Britain with America commonly credit relative labour scarcity in America with driving forward the capital-intensive and highly productive "American system" of manufacturing.



The view from Serendip

Some economists are considering how Mr Romer's heresy might apply today. Daron Acemoglu, Gino Gancia, and Fabrizio Zilibotti of MIT, CREi (an economics-research centre in Barcelona) and the University of Zurich, have built a model to study this. It shows firms in rich countries shipping low-skill tasks abroad when offshoring costs little, thus driving apart the wages of skilled and unskilled workers at home. Over time, though, offshoring raises wages in less-skilled countries; that makes innovation at home more enticing. Workers are in greater demand, the income distribution narrows, and the economy comes to look more like the post-second-world-war period than the 1970s and their aftermath.

Even if that model is mistaken, the rise of the emerging world is among the biggest reasons for optimism. The larger the size of the global market, the more the world benefits from a given new idea, since it can then be applied across more activities and more people. Raising Asia's poor billions into the middle class will mean that millions of

great minds that might otherwise have toiled at subsistence farming can instead join the modern economy and share the burden of knowledge with rich-world researchers—a sharing that information technology makes ever easier.

It may still be the case that some parts of the economy are immune, or at least resistant, to some of the productivity improvement that information technology can offer. Sectors like health care, education and government, in which productivity has proved hard to increase, loom larger within the economy than in the past. The frequent absence of market pressure in such areas reduces the pressure for cost savings—and for innovation.

For some, though, the opposite outcome is the one to worry about. Messrs Brynjolfsson and McAfee fear that the technological advances of the second half of the chessboard could be disturbingly rapid, leaving a scourge of technological unemployment in their wake. They argue that new technologies and the globalisation that they allow have already contributed to stagnant incomes and a decline in jobs that require moderate levels of skill. Further progress could threaten jobs higher up and lower down the skill spectrum that had, until now, seemed safe.

Pattern-recognition software is increasingly good at performing the tasks of entry-level lawyers, scanning thousands of legal documents for relevant passages. Algorithms are used to write basic newspaper articles on sporting outcomes and financial reports. In time, they may move to analysis. Manual tasks are also vulnerable. In Japan, where labour to care for an ageing population is scarce, innovation in robotics is proceeding by leaps and bounds. The rising cost of looking after people across the rich world will only encourage further development.

Such productivity advances should generate enormous welfare gains. Yet the adjustment period could be difficult. In the end, the main risk to advanced economies may not be that the pace of innovation is too slow, but that institutions have become too rigid to accommodate truly revolutionary changes—which could be a lot more likely than flying cars.