

Martin Wolf on robots

“If robots divide us, they will conquer”

Financial Times, February 4, 2014

“Enslave the robots and free the poor”

Financial Times, February 11, 2014

“Same as It Ever Was. Why the Techno-optimists Are Wrong”


Foreign Affairs, July/August 2015 Issue



If robots divide us, they will conquer

Martin Wolf, *Financial Times*, February 4, 2014

The rise of intelligent technologies may cost us dear – unless we understand the dangers



With one rub of his lamp, Aladdin could command an intelligent being able to fulfil all desires. His genie was a spirit. But the dream of powerful and intelligent artificial servants has also encompassed physical beings. Now, it is becoming a reality built of silicon, metal and plastic. But is it a dream or a nightmare? Will clever machines prove beneficial? Or will they be Frankenstein monsters?

This is [the question](#) raised by *The Second Machine Age*, a new book by Erik Brynjolfsson and Andrew McAfee of the Massachusetts Institute of Technology. This predicts that we will experience “two of the most amazing events in human history: the creation of true machine intelligence and the connection of all humans via a common digital network, transforming the planet’s economics. Innovators, entrepreneurs, scientists, tinkerers, and many other types of geeks will take advantage of this cornucopia to build technologies that astonish us, delight us, and work for us.”

What distinguishes the second machine age from the first is intelligence. The machines of the first age replaced and multiplied the physical labour of humans and animals. The machines of the second age will replace and multiply our intelligence. The driving force behind this revolution is, argue the authors, the exponential increase in the power (or exponential fall in the cost) of computing. The celebrated example is Moore’s Law, named after Gordon Moore, a founder of Intel. For half a century, the number of transistors on a semiconductor chip has doubled at least every two years. Similar progress has occurred elsewhere. (See chart.)

The authors argue that after half a century of progress we are seeing leaps in machine intelligence. As computing power grows exponentially, computers are managing tasks deemed beyond reach a few years ago. Soon, they predict, machine intelligence will be everywhere. They offer as a parallel the story of the inventor of chess who asked to be rewarded with one grain of rice on the first square of his board, two on the second, four on the third, and so forth. Manageable on the first half of the board, the reward reaches mountainous proportions on the second. Our reward will grow similarly.

Yet, to paraphrase a celebrated 1987 quip about computers by Robert Solow, a Nobel-laureate MIT economist, we see information technology everywhere except in the [productivity](#) statistics. Trends in output per hour in the US are quite mediocre. Indeed, after an encouraging surge in the 1990s and early 2000s, growth has subsided again. Recent performance in other big high-income economies is worse.

One possible explanation is that the impact of these technologies is overhyped. Not surprisingly, the authors disagree. Indeed, they argue that far from being exhausted, the possibilities are boundless: “digitisation makes available massive bodies of data relevant to almost any situation, and this information can be infinitely reproduced and reused”.

If so, why are measured increases in output so modest? The answers offered are: the plethora of cheap or free services (Skype or Wikipedia); the scale of do-it-yourself entertainment (Facebook); and the failure to account fully for all the new products or services. Before June 2007, an iPhone was out of the reach of even the richest man on earth. Its price was infinite. The fall from an infinite to a definite price is not reflected in the price indices. Again, the “consumer surplus” in digital products and services – the difference between the price and the value to consumers – is often huge. Finally, measures of gross domestic product also underestimate investment in intangible assets.

It seems quite plausible that the proliferation of new gadgets, and the rise of the digital economy with its uniquely low marginal costs, have had a far bigger effect on welfare and even GDP than current measures indicate.

Yet worries remain. The information age has coincided with – and must, to some extent, have caused – adverse economic trends: stagnation of median real incomes; rising inequality of labour income and of the distribution of income between labour and capital; and growing long-term unemployment.

Among the explanations are: fast-rising productivity in manufacturing; skills-biased technical change; the rise of global winners-take-all markets; and the role of rental income, particularly from intellectual property. Think of the difference between the cost of developing Google’s search algorithm and its value. Globalisation and financial liberalisation are also at work, both also boosted by new technologies.

Above all, insists the book, this is just the beginning. Much routine brain-work will be computerised, as happened to clerical skills. Middle-income jobs could hollow out far further. The outcome could be still more polarised incomes, with a tiny group of winners at the top and a vastly larger group struggling below. In 2012, for example, the top 1 per cent of Americans earned 22 per cent of all incomes, more than double their share in the 1980s.

There are good reasons why people should be disturbed by this. First, the lives of those at the bottom might get worse: the authors note that the life expectancy of an American white woman without a high school diploma fell five years between 1990 and 2008. Second, if income becomes too unequal, opportunities for young people dwindle. Third, the wealthy become indifferent to the fate of the rest. Finally, a vast inequality of power emerges, making a mockery of the ideal of democratic citizenship.

In the distant future, thinking machines might even overwhelm our sense of ourselves, just as the best human chess players now know they are not the best on earth. But well before that, the authors suggest that income inequality is likely to increase further, tarnishing the silver age of opportunity the book also promises.

Big challenges arise, then, both now and in the future, if we are to ensure the new machines do not become our Frankenstein monsters. These have big implications for public policy on property rights, education, taxation and other government measures that aim to promote human welfare. I will look at these controversial issues next week.

Enslave the robots and free the poor

Martin Wolf, *Financial Times*, February 11, 2014

The prospect of far better lives depends on how the gains are produced and distributed



In 1955, Walter Reuther, head of the US car workers' union, told of a visit to a new automatically operated Ford plant. Pointing to all the robots, his host asked: "How are you going to collect union dues from those guys?" Mr Reuther replied: "And how are you going to get them to buy Fords?" Automation is not new. Neither is the debate about its effects. How far, then, does what Erik Brynjolfsson and Andrew McAfee call *The Second Machine Age* alter the questions or the answers?

I laid out the core argument last week. I noted that the rise of information technology coincides with increasing income inequality. [Lawrence Mishel](#) of the Washington-based Economic Policy Institute challenges the notion that the former has been the principal cause of the latter. Mr Mishel notes: "Rising executive pay and the expansion of, and better pay in, the financial sector can account for two-thirds of increased incomes at the top." Changing social norms, the rise of stock-based remuneration and the extraordinary expansion of the financial sector also contributed. While it was a factor, technology has not determined economic outcomes.

Yet technology could become far more important. Professor Brynjolfsson and Mr McAfee also argue that it will make us more prosperous; and it will shift the distribution of opportunities among workers and between workers and owners of capital.

The economic impacts of new technologies are many and complex. They include: new services, such as Facebook; disintermediation of old systems of distribution via iTunes or Amazon; new products, such as smartphones; and new machines, such as robots. The latter awaken fears that intelligent machines will render a vast number of people redundant. A [recent paper](#) by [Carl Frey and Michael Osborne](#) of Oxford university concludes that 47 per cent of US jobs are at high risk from automation. In the 19th century, they argue, machines replaced artisans and benefited unskilled labour. In the 20th century, computers replaced middle-income jobs, creating a polarised labour market. Over the next decades, however, "most workers in transport and logistics occupations, together with the bulk of office and administrative support workers, and labour in production occupations, are likely to be substituted by computer capital". Moreover, "computerisation will mainly substitute for low-

skill and low-wage jobs in the near future. By contrast, high-skill and high-wage occupations are the least susceptible to computer capital.” This, then, would exacerbate inequality.

[Jeffrey Sachs of Columbia university and Laurence Kotlikoff](#) of Boston university even argue that the rise in productivity might make future generations worse off in aggregate. The replacement of workers by robots could shift income from the former to the robots’ owners, most of whom will be retired and are assumed to save less than the young. This would lower investment in human capital because the young could no longer afford to pay for it; and in machines because savings in this economy would fall.

The argument that a rise in potential productivity would make us permanently worse off is ingenious. More plausible, to me at least, are other possibilities: there could be a large adjustment shock as workers are laid off; the market wages of unskilled people might fall far below a socially acceptable minimum; and, combined with other new technologies, robots might make the distribution of income far more unequal than it is already.

So what should be done?

First, the new technologies will bring good and bad. We can shape the good and manage the bad.

Second, education is not a magic wand. One reason is that we do not know what skills will be demanded three decades hence. Also, if Mr Frey and Prof Osborne are right, so many low- to middle-skilled jobs are at risk that it may already be too late for anybody much over 18 and many children. Finally, even if the demand for creative, entrepreneurial and high-level knowledge services were to grow on the required scale, which is highly unlikely, turning us all into the happy few is surely a fantasy.

Third, we must reconsider leisure. For a long time the wealthiest lived a life of leisure at the expense of the toiling masses. The rise of intelligent machines makes it possible for many more people to live such lives without exploiting others. Today’s triumphant puritanism finds such idleness abhorrent. Well, then, let people enjoy themselves busily. What else is the true goal of the vast increases in prosperity we have created?

Fourth, we will need to redistribute income and wealth. Such redistribution could take the form of a basic income for every adult, together with funding of education and training at any stage in a person’s life. In this way, the potential for a more enjoyable life might become a reality. The revenue could come from taxes on bads (pollution, for example) or on rents (including land and, above all, intellectual property). Property rights are a social creation. The idea that a small minority should overwhelming benefit from new technologies should be reconsidered. It would be possible, for example, for the state to obtain an automatic share in the income from the intellectual property it protects.

Finally, if labour shedding does accelerate, it will be essential to ensure that demand expands in tandem with the rise in potential supply. If we succeed, many of the worries over a lack of jobs will fade away. Given the failure to achieve this in the past seven years, that may well not happen. But we could do better if we wanted to.

The rise of intelligent machines is a moment in history. It will change many things, including our economy. But their potential is clear: they will make it possible for human beings to live far better lives. Whether they end up doing so depends on how the gains are produced and distributed. It is possible that the ultimate result will be a tiny minority of huge winners and a vast number of losers. But such an outcome would be a choice not a destiny. A form of techno-feudalism is unnecessary. Above all, technology itself does not dictate the outcomes. Economic and political institutions do. If the ones we have do not give the results we want, we must change them.

Same as It Ever Was
Why the Techno-optimists Are Wrong
Martin Wolf, *Foreign Affairs*, July/August 2015



Belief in “the green light, the orgiastic future that year by year recedes before us,” as F. Scott Fitzgerald wrote in *The Great Gatsby*, is a characteristic American trait. But hope in a better future is not uniquely American, even if it has long been a more potent secular faith in the United States than elsewhere. The belief has older roots. It was the product of a shift in the temporal location of the golden age from a long-lost past to an ever-brighter future.

That shift was conceived and realized with the Enlightenment and then the Industrial Revolution. As human beings gained ever-greater control of the forces of nature and their economies became ever more productive, they started to hope for lives more like those of the gods their ancestors had imagined.

People might never be immortal, but their lives would be healthy and long. People might never move instantaneously, but they could transport themselves and their possessions swiftly and cheaply across great distances. People might never live on Mount Olympus, but they could enjoy a temperate climate, 24-hour lighting, and abundant food. People might never speak mind to mind, but they could communicate with as many others as they desired, anywhere on the planet. People might never enjoy infinite wisdom, but they could gain immediate access to the knowledge accumulated over millennia.

All of this has already happened in the world’s richest countries. It is what the people of the rest of the world hope still to enjoy.

Is a yet more orgiastic future beckoning? Today’s Gatsbys have no doubt that the answer is yes: humanity stands on the verge of breakthroughs in information technology, robotics, and artificial intelligence that will dwarf what has been achieved in the past two centuries. Human

beings will be able to live still more like gods because they are about to create machines like gods: not just strong and swift but also supremely intelligent and even self-creating.

Yet this is the optimistic version. Since Mary Shelley created the cautionary tale of Frankenstein, the idea of intelligent machines has also frightened us. Many duly point to great dangers, including those of soaring unemployment and inequality.

But are we likely to experience such profound changes over the next decade or two? The answer is no.

Small Change

In reality, the pace of economic and social transformation has slowed in recent decades, not accelerated. This is most clearly shown in the rate of growth of output per worker. The economist Robert Gordon, doyen of the skeptics, has noted that the average growth of U.S. output per worker was 2.3 percent a year between 1891 and 1972. Thereafter, it only matched that rate briefly, between 1996 and 2004. It was just 1.4 percent a year between 1972 and 1996 and 1.3 percent between 2004 and 2012.

On the basis of these data, the age of rapid productivity growth in the world's frontier economy is firmly in the past, with only a brief upward blip when the Internet, e-mail, and e-commerce made their initial impact.

Those whom Gordon calls "techno-optimists"—Erik Brynjolfsson and Andrew McAfee of the Massachusetts Institute of Technology, for example—respond that the GDP statistics omit the enormous unmeasured value provided by the free entertainment and information available on the Internet. They emphasize the plethora of cheap or free services (Skype, Wikipedia), the scale of do-it-yourself entertainment (Facebook), and the failure to account fully for all the new products and services. Techno-optimists point out that before June 2007, an iPhone was out of reach for even the richest man on earth. Its price was infinite. The fall from an infinite to a definite price is not reflected in the price indexes. Moreover, say the techno-optimists, the "consumer surplus" in digital products and services—the difference between the price and the value to consumers—is huge. Finally, they argue, measures of GDP underestimate investment in intangible assets.

These points are correct. But they are nothing new: all of this has repeatedly been true since the nineteenth century. Indeed, past innovations generated vastly greater unmeasured value than the relatively trivial innovations of today. Just consider the shift from a world without telephones to one with them, or from a world of oil lamps to one with electric light. Next to that, who cares about Facebook or the iPad? Indeed, who really cares about the Internet when one considers clean water and flushing toilets?

Over the past two centuries, historic breakthroughs have been responsible for generating huge unmeasured value. The motor vehicle eliminated vast quantities of manure from urban streets. The refrigerator prevented food from becoming contaminated. Clean running water and vaccines delivered drastic declines in child mortality rates. The introduction of running water, gas and electric cookers, vacuums, and washing machines helped liberate women from domestic labor. The telephone removed obstacles to speedy contact with the police, fire brigades, and ambulance services. The discovery of electric light eliminated forced idleness. Central heating and air conditioning ended discomfort. The introduction of the railroad, the steam ship, the motor car, and the airplane annihilated distance.

The radio, the gramophone, and the television alone did far more to revolutionize home entertainment than the technologies of the past two decades have. Yet these were but a tiny fraction of the cornucopia of innovation that owed its origin to the so-called general-purpose technologies—industrialized chemistry, electricity, and the internal combustion engine—introduced by what is considered the Second Industrial Revolution, which occurred between the 1870s and the early twentieth century. The reason we are impressed by the relatively paltry innovations of our own time is that we take for granted the innovations of the past.

Gordon also notes how concentrated the period of great breakthroughs was. As he writes:

Electric light and a workable internal combustion engine were invented in a three-month period in late 1879. The number of municipal waterworks providing fresh running water to urban homes multiplied tenfold between 1870 and 1900. The telephone, phonograph, and motion pictures were all invented in the 1880s.

And the benefits of these mainstays of the Second Industrial Revolution, Gordon points out, “included subsidiary and complementary inventions, from elevators, electric machinery and consumer appliances; to the motorcar, truck, and airplane; to highways, suburbs, and supermarkets; to sewers to carry the wastewater away.”

Past, not prologue

The technologies introduced in the late nineteenth century did more than cause three generations of relatively high productivity growth. They did more, too, than generate huge unmeasured economic and social value. They also brought with them unparalleled social and economic changes. An ancient Roman would have understood the way of life of the United States of 1840 fairly well. He would have found that of 1940 beyond his imagination.

Among the most important of these broader changes were urbanization and the huge jumps in life expectancy and standards of education. The United States was 75 percent rural in the 1870s. By the mid-twentieth century, it was 64 percent urban. Life expectancy rose twice as fast in the first half of the twentieth century as in the second half. The collapse in child mortality is surely the single most beneficial social change of the past two centuries. It is not only a great good in itself; it also liberated women from the burden, trauma, and danger of frequent pregnancies. The jump in high school graduation rates—from less than ten percent of young people in 1900 to roughly 80 percent by 1970—was a central driver of twentieth-century economic growth.

All these changes were also, by their nature, one-offs. This is also true of the more recent shift of women entering the labor force. It has happened. It cannot be repeated.

Yet there is something else of compelling importance in the contrast between the breakthroughs of the nineteenth and early twentieth centuries and those of the second half of the twentieth and the early twenty-first century. The former were vastly broader, affecting energy; transportation; sanitation; food production, distribution, and processing; entertainment; and, not least, entire patterns of habitation. Yes, computers, mobile telecommunications, and the Internet are of great significance. Yet it is also essential to remember what has not changed to any fundamental degree. Transportation technologies and speeds are essentially the same as they were half a century ago. The dominant source of commercial energy remains the burning of fossil fuels—introduced with coal and steam in the First Industrial Revolution, of the late eighteenth and early nineteenth centuries—and even nuclear power is now an elderly technology. Although fracking is noteworthy, it does not compare with the opening of the petroleum age in the late nineteenth century.



The only recent connections between homes and the outside world are satellite dishes and broadband. Neither is close to being as important as clean water, sewerage, gas, electricity, and the telephone. The great breakthroughs in health—clean water, sewerage, refrigeration, packaging, vaccinations, and antibiotics—are also all long established.

The future's not what it used to be

The so-called Third Industrial Revolution—of the computer, the Internet, and e-commerce—is also itself quite old. It has already produced many changes. The armies of clerks who used to record all transactions have long since disappeared, replaced by computers; more recently, so have secretaries. E-mail has long since replaced letters. Even the Internet and the technologies that allow it to be searched with ease are now 15 years old, or even older, as is the e-commerce they enabled.

Yet the impact of all of this on measured productivity has been modest. The economic historian Paul David famously argued in 1989 that one should remember how long it took for industrial processes to adapt to electricity. But the computer itself is more than half a century old, and it is now a quarter of a century since David made that point. Yet except for the upward blip between 1996 and 2004, we are still—to adapt the Nobel laureate Robert Solow's celebrated words of 1987—seeing the information technology age “everywhere but in the productivity statistics.”

Meanwhile, other, more recent general-purpose technologies—biotechnology and nanotechnology, most notably—have so far made little impact, either economically or more widely.


The disappointing nature of recent growth is also the theme of an influential little book, *The Great Stagnation*, by the economist Tyler Cowen, which is subtitled *How America Ate All the Low-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better*. As Cowen writes:

The American economy has enjoyed . . . low-hanging fruit since at least the seventeenth century, whether it be free land, . . . immigrant labor, or powerful new technologies. Yet during the last 40 years, that low-hanging fruit started disappearing, and we started pretending it was still there. We have failed to recognise that we are at a technological plateau and the trees are more bare than we would like to think. That's it. That is what has gone wrong.

In considering the disappointing impact of recent innovations, it is important to note that the world's economies are vastly bigger than they used to be. Achieving a two percent economy-wide annual rise in labor productivity may simply be a much bigger challenge than it was in the past.

More important, the share of total output of the sectors with the fastest growth in productivity tends to decline over time, while the share of the sectors where productivity growth has proved hardest to increase tends to rise. Indeed, it is possible that productivity growth will essentially cease because the economic contribution of the sectors where it is fastest will become vanishingly small. Raising productivity in manufacturing matters far less now that it generates only about an eighth of total U.S. GDP. Raising productivity in caring for the young, the infirm, the helpless, and the elderly is hard, if not impossible.

Yet perhaps paradoxically, recent technological progress might still have had some important effects on the economy, and particularly the distribution of income, even if its impact on the size of the economy and overall standards of living has been relatively modest. The information age coincided with—and must, to some extent, have caused—adverse economic trends: the stagnation of median real incomes, rising inequality of labor income and of the distribution of income between labor and capital, and growing long-term unemployment.



Information technology has turbocharged globalization by making it vastly easier to organize global supply chains, run 24-hour global financial markets, and spread technological know-how. This has helped accelerate the catch-up process of emerging-market economies, notably China. It has also allowed India to emerge as a significant exporter of technological services.

Technology has also brought about the rise of winner-take-all markets, as superstars have come to bestride the globe. Substantial evidence exists, too, of “skills-biased” technological change. As the demand for and rewards offered to highly skilled workers (software programmers, for example) rise, the demand for and rewards offered to those with skills in the middle of the distribution (such as clerks) decline. The value of intellectual property has also risen. In brief, a modest impact on aggregate output and productivity should not be confused with a modest impact across the board.

No crystal ball required

The future is, at least to some extent, unknowable. Yet as Gordon suggests, it is not all that unknowable. Back in the nineteenth and early twentieth centuries, many had already realized the changes that the recent inventions might bring. The nineteenth-century French novelist Jules Verne is a famous example of such foresight.

The optimistic view is that we are now at an inflection point. In their book *The Second Machine Age*, Brynjolfsson and McAfee offer as a parallel the story of the inventor of chess, who asked to be rewarded with one grain of rice on the first square of his board, two on the second, four on the third, and so forth. Manageable in size on the first half of the board, the reward reaches mountainous proportions toward the end of the second. Humanity’s reward from Moore’s law—the relentless doubling of the number of transistors on a computer chip every two years or so—will, they argue, grow similarly.

These authors predict that we will experience

two of the most amazing events in history: the creation of true machine intelligence and the connection of all humans via a common digital network, transforming the planet’s economics. Innovators, entrepreneurs, scientists, tinkerers, and many other types of geeks will take advantage of this cornucopia to build technologies that astonish us, delight us, and work for us.

In the near term, however, the widely mentioned possibilities—medicine, even bigger data, robots, 3-D printing, self-driving cars—look quite insignificant.

The impact of the biomedical advances so far has been remarkably small, with pharmaceutical companies finding it increasingly difficult to register significant breakthroughs. So-called big data is clearly helping decision-making. But many of its products—ultra-high-speed trading, for example—are either socially and economically irrelevant or, quite possibly, harmful. Three-D printing is a niche activity—fun, but unlikely to revolutionize manufacturing.

Making robots replicate all the complex abilities of human beings has proved extremely difficult. Yes, robots can do well-defined human jobs in well-defined environments. Indeed, it is quite possible that standard factory work will be entirely automated. But the automation of such work is already very far advanced. It is not a revolution in the making. Yes, it is possible to imagine driverless cars. But this would be a far smaller advance than were cars themselves.

Inevitably, uncertainty is pervasive. Many believe that the impact of what is still to come could be huge. The economist Carl Benedikt Frey and the machine-learning expert Michael Osborne, both of Oxford University, have concluded that 47 percent of U.S. jobs are at high risk from automation. In the nineteenth century, they argue, machines replaced artisans and benefited unskilled labor. In the twentieth century, computers replaced middle-income jobs, creating a polarized labor market.

Over the next decades, they write, “most workers in transportation and logistics occupations, together with the bulk of office and administrative support workers, and labour in production occupations, are likely to be substituted by computer capital.” Moreover, they add, “computerisation will mainly substitute for low-skill and low-wage jobs in the near future. By contrast, high-skill and high-wage occupations are the least susceptible to computer capital.” That would exacerbate already existing trends toward greater inequality. But remember that previous advances also destroyed millions of jobs. The most striking example is, of course, in agriculture, which was the dominant employer of humanity between the dawn of the agricultural revolution and the nineteenth century.

The economists Jeffrey Sachs and Laurence Kotlikoff even argue that the rise in productivity generated by the coming revolution could make future generations worse off in the aggregate. The replacement of workers by robots could shift income from the former to the robots’ owners, most of whom will be retired, and the retired are assumed to save less than the young. This would lower investment in human capital because the young could no longer afford to pay for it, and it would lower investment in machines because savings in this economy would fall.

Beyond this, people imagine something far more profound than robots able to do gardening and the like: the “technological singularity,” when intelligent machines take off in a rapid cycle of self-improvement, leaving mere human beings behind. In this view, we will someday create machines with the abilities once ascribed to gods. Is that imminent? I have no idea.

Been there, done that

So how might we respond now to these imagined futures?

First, new technologies bring good and bad. We must believe we can shape the good and manage the bad.

Second, we must understand that education is not a magic wand. One reason is that we do not know what skills will be demanded three decades hence. Also, if Frey and Osborne are right, so many low- to middle-skilled jobs are at risk that it may already be too late for anybody much over 18 and for many children. Finally, even if the demand for creative, entrepreneurial, and high-level knowledge services were to grow on the required scale, which is highly unlikely, turning us all into the happy few is surely a fantasy.

Third, we will have to reconsider leisure. For a long time, the wealthiest lived a life of leisure at the expense of the toiling masses. The rise of intelligent machines would make it possible for many more people to live such lives without exploiting others. Today's triumphant puritanism finds such idleness abhorrent. Well then, let people enjoy themselves busily. What else is the true goal of the vast increases in prosperity we have created?

Fourth, we may need to redistribute income and wealth on a large scale. Such redistribution could take the form of a basic income for every adult, together with funding for education and training at any stage in a person's life. In this way, the potential for a more enjoyable life might become a reality. The revenue could come from taxes on bads (pollution, for example) or on rents (including land and, above all, intellectual property). Property rights are a social creation. The idea that a small minority should overwhelmingly benefit from new technologies should be reconsidered. It would be possible, for example, for the state to obtain an automatic share of the income from the intellectual property it protects.

Fifth, if labor shedding does accelerate, it will be essential to ensure that demand for labor expands in tandem with the rise in potential supply. If we succeed, many of the worries over a lack of jobs will fade away. Given the failure to achieve this in the past seven years, that may well not happen. But we could do better if we wanted to.

The rise of truly intelligent machines, if it comes, would indeed be a big moment in history. It would change many things, including the global economy. Their potential is clear: they would, in principle, make it possible for human beings to live far better lives. Whether they end up doing so depends on how the gains are produced and distributed.

It is also possible that the ultimate result might be a tiny minority of huge winners and a vast number of losers. But such an outcome would be a choice, not a destiny. Techno-feudalism is unnecessary. Above all, technology itself does not dictate the outcomes. Economic and political institutions do. If the ones we have do not give the results we want, we will need to change them.

As for the singularity, it is hard to conceive of such a state of the world. Would a surpassed humanity live happily ever after, tended, like children, by solicitous machines? Would people find meaning in a world in which their intellectual progeny were so vastly superior to themselves?

What we know for the moment is that there is nothing extraordinary in the changes we are now experiencing. We have been here before and on a much larger scale. But the current and prospective rounds of changes still create problems—above all, the combination of weak growth and significant increases in inequality. The challenge, as always, is to manage such changes. The only good reason to be pessimistic is that we are doing such a poor job of this.

The future does not have to be a disappointment. But as Gatsby learned, it can all too easily be just that.