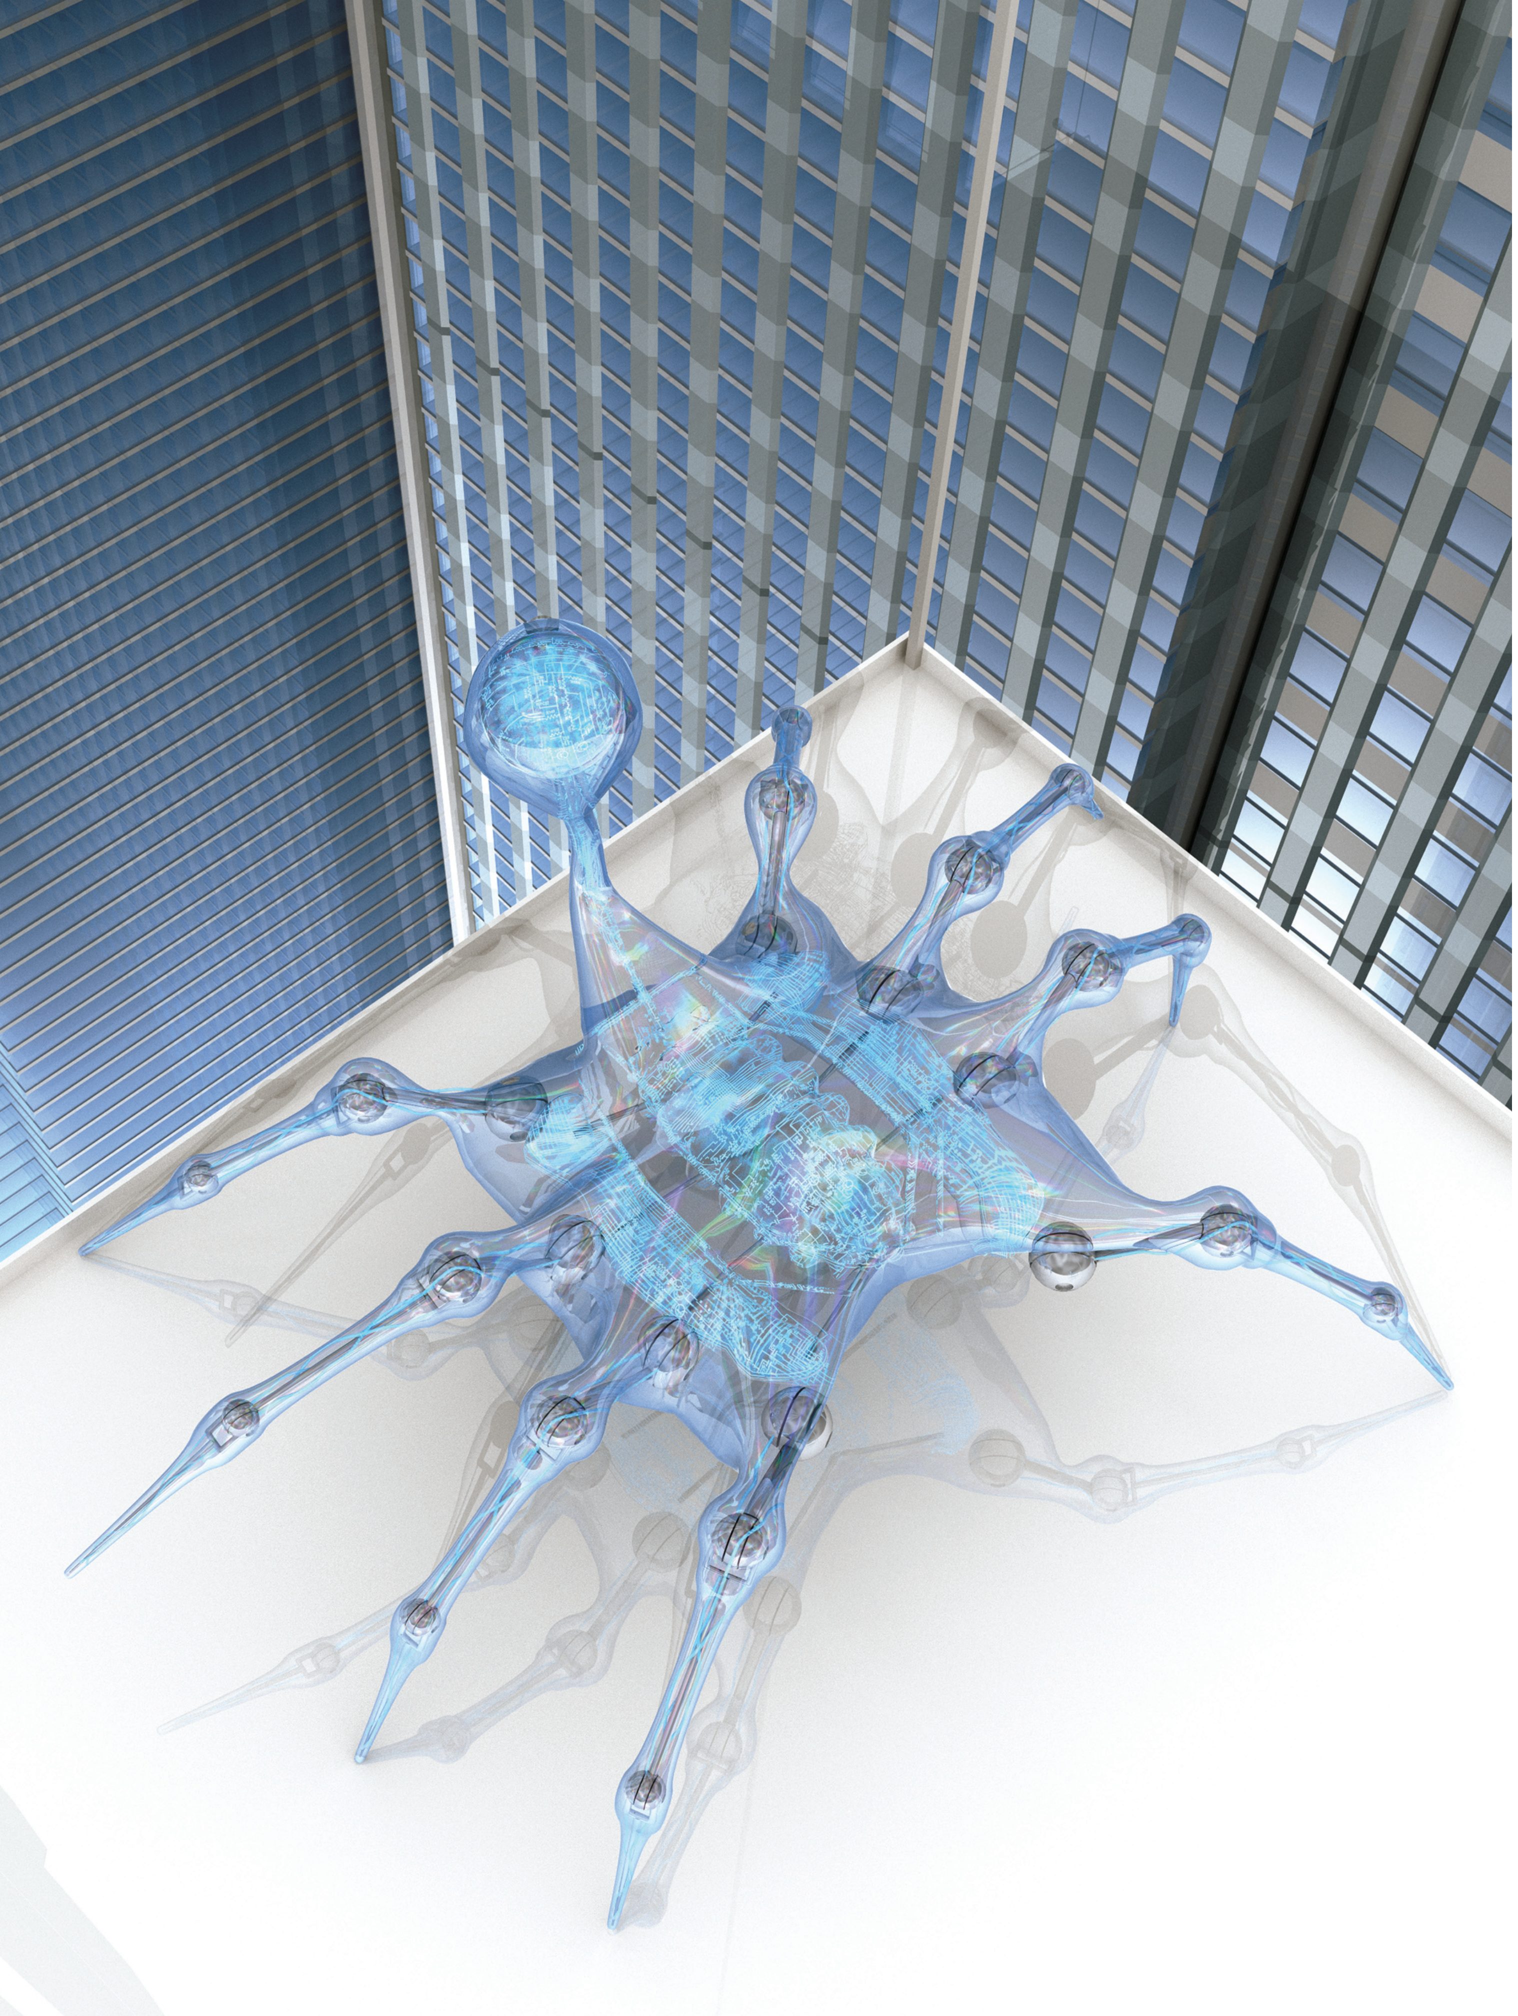

BRYAN CHRISTIE DESIGN





STUFFED INTO SKYSCRAPERS BY THE BILLION, BRAINY BUGBOTS WILL BE THE KNOWLEDGE WORKERS OF THE FUTURE BY ROBIN HANSON

# Economics Of The Singularity

OUR GLOBAL ECONOMY would stupefy a Roman merchant as much as the Roman economy would have confounded a caveman. But we would be similarly amazed to see the economy that awaits our grandchildren, for I expect it to follow a societal discontinuity more dramatic than those brought on by the agricultural and industrial revolutions. The key, of course, is technology. A revolutionary speedup in economic growth requires an unprecedented and remarkable enabling tool. Machine intelligence on a human level, if not higher, would do nicely. Its arrival could produce a singularity—an overwhelming departure from prior trends, with uneven and dizzyingly rapid change thereafter. A future shock to end future shocks.

Yes, this theory of mine is a social and economic one, and therefore not as unfailingly accurate or testable as one in the physical sciences. Nevertheless, social scientists routinely make short-term forecasts that hit the mark, and economists often offer insightful forecasts about unprecedented situations.

So indulge me as I outline how we economists view technological change. In so doing, I hope to explain why it's reasonable to view past history as a series of abrupt, seemingly unheralded transitions from one economic era to another, transitions marked by the sudden and drastic increase in the rate of economic growth. I will then show why another singularity is perhaps just around the corner. Finally, I will outline its possible consequences.

A complex device, like a tractor or a building, can have thousands of parts, and each part can rely on dozens of technologies. Yet in most cases even a spectacular gain in the quality of one part bestows at best only a small improvement on the whole. Keep improving a part in successive increments of equal degree and you'll get ever smaller gains to the whole. This is the law of diminishing returns, and it applies not only to devices and organizations but to entire industries. Consider your personal computer: every couple of years its power-to-cost ratio has doubled, and yet as you go from one generation to the next, you probably notice only

## How Many Singularities Have There Been?

THE TWO SOLIDLY DEMONSTRATED singularities, the agricultural and industrial revolutions, came with little warning.

Were there any singularities before farming and industry? If we look back further in time, we can find even slower modes of growth that made sudden transitions to faster modes. For example, human hunter-gatherers vastly expanded their niche and spread throughout the world in a biologically short period of time. That transition apparently was made possible by special innovations in the unusually large protohuman brain. Before that transition, and after the emergence of animals some 500 million years earlier, the largest animal brains doubled in size roughly every 30 million years—less than 1 percent of the growth rate of human brains.

Looking further back, it is difficult to find long-term trends that may have paved the way for the emergence of animals. Still, it is interesting to note that the volume of our nearly 14 billion-year-old universe has been expanding exponentially due to a mysterious "dark energy," with a doubling time of 3 billion years—about 1 percent the growth rate of animal brain size.

Of course, because we have no accepted theory saying why various growth modes and transitions should be related, any similarities between them may be pure coincidence. But they do constitute precedents, for they show that vast changes can appear seemingly overnight. —R.H.

a small improvement as you plug away on your word processors and spreadsheets.

It turns out that most of these small, innovative gains come not from research labs but from hands-on builders and users. So the more a thing gets used, the more it tends to improve. It doesn't matter whether that thing is a physical device, such as a car, or a social organization, such as a corporation.

If any large system of interacting parts tends to improve by smooth gradations, then we should expect systems of systems, with their larger number of components and interactions, to improve even more smoothly. By this reasoning, the world economy should improve most smoothly of all. The world economy consists of the largest number of interacting parts of any man-made system, and everyone not stranded on an uncharted island contributes to the improvements in all those parts by using them. Finally, in each economic era the question of whether growth speeds up or slows down depends on two competing factors. Deceleration typically ensues as innovators exhaust the easy ideas—the low-hanging fruit. But acceleration also ensues as the economy, by getting larger, enables its members to explore an ever-increasing number of innovations.

### The world economy, which now doubles in 15 years, would soon double in a week to a month

We have the tools to measure the world's economic product not only for today—it's about US $50 trillion per year—but also for times long past. A few years ago Angus Maddison, an economic historian at the University of Groningen, in the Netherlands, plotted a graph of world economic product—basically everything of value produced globally: bananas, submarines, magazine articles, you name it. It shows that from 1950 to 2003, growth was relatively steady. During that time, despite enormous technical change, no particular technology left much of a fingerprint on the data; no short-term accelerations in growth could be attributed to this or that technological development. Also, Maddison's data offer little support for the idea that innovation and growth have been accelerating recently.

Now look at the data for world product over the past 7,000 years, estimated by Bradford DeLong, an economic historian at the University of California, Berkeley. The data here tell a somewhat different story. For most of that time, growth proceeded at a relatively steady exponential rate, with a doubling of output about every 900 years. But within the past few centuries, something dramatic happened: output began doubling faster and faster, approaching a new steady doubling time of about 15 years. That's about 60 times as fast as it had been in the previous seven millennia.

WE CALL THIS transition the Industrial Revolution, but that does not mean we understand it well or even know precisely how and why it arose. But whatever the Industrial Revolution was, clearly it was an event worthy of the name "singularity."

If we look further back, we see what appears to be at least one previous singularity—the transition to an economy based on agriculture. And slow as economic growth during the agricultural era may seem in the aftermath of the Industrial Revolution, it was actually lightning fast compared with that of the economic era that came before, which was based on hunting and gathering.

In the roughly 2 million years our ancestors lived as hunters and gatherers, the population rose from about 10 000 protohumans to about 4 million modern humans. If, as we believe, the growth pattern during this era was fairly steady, then the population must have doubled about every quarter million years, on average. Then, beginning about 10 000 years ago, a few of those 4 million humans began to settle down and live as farmers. The resulting communities grew so fast that they quickly accounted for most of the world population. From that time on, the farming population doubled about every 900 years—some 250 times as fast as before.

Our understanding of the existence, nature, and relevance of these transitions clearly becomes more speculative the further back we look in time [see sidebar, "How Many Singularities Have There Been?"]. There may well have been two earlier singularities that started eras of this sort, although our ability to identify them and weigh their relevance is very speculative. I suggest an era defined by the growth of the brain from the emergence of animal life to the first protohumans and perhaps an earlier era defined by the growth of the universe from a time shortly after the big bang to the first animals.

So we have perhaps five eras during which the thing whose growth is at issue—the universe, brains, the hunting economy, the farming economy, and the industrial economy—doubled in size at fixed intervals. Each era of growth before now, however, has eventually switched suddenly to a new era having a growth rate that was between 60 and 250 times as fast. Each switch was completed in much less time than it had taken the previous regime to double—from a few millennia for the agricultural revolution to a few centuries for the industrial one. These switches constituted singularities.





Whatever may have been the key innovations behind these transitions, it is clear that they were far more potent than such familiar textbook examples of great innovations as fire, writing, computers, or plastics. Most innovations happen within a given growth era and do not change its basic nature, including its basic growth rate. A few exceedingly rare innovations, however, do suddenly change everything. One such innovation led to agriculture; another led to industry.

Therefore, we must admit that another singularity—at least the third one, and perhaps the fifth, depending on how you count—could lie ahead. Furthermore, data on these previous apparently similar singularities are some of the few concrete guides available to what such a transition might look like. We would be fools if we confidently expected all patterns to continue. But it strikes me as pretty foolish to ignore the patterns we see.

If a new transition were to show the same pattern as the past two, then growth would quickly speed up by between 60- and 250-fold. The world economy, which now doubles in 15 years or so, would soon double in somewhere from a week to a month. If the new transition were as gradual (in power-law terms) as the Industrial Revolution was, then within three years of a noticeable departure from typical fluctuations, it would begin to double annually, and within two more years, it might grow a million-fold. If the new transition were as rapid as the agricultural revolution seems to have been, change would be even more sudden.

Though such growth may seem preposterous, consider that in the era of hunting and gathering, the economy doubled nine times; in the era of farming, it doubled seven times; and in the current era of industry, it has so far doubled 10 times. If, for some as yet unknown reason, the number of doublings is similar across these three eras, then we seem already overdue for another transition. If we instead compare our era with the era of brain growth, which doubled 16 times before humans appeared, we would expect the next transition by around 2075.

What innovation could possibly induce so fabulous a speedup in economic growth? It is easier to say what could not. Because of diminishing returns, no change that improved just one small sector of the economy could do the trick. In advanced countries today, farming, mining, energy, communications, transportation, and construction each account for only a small percentage of economic activity. Even so extraordinary an innovation as radical nanotechnology would do no more than dramatically lower the cost of capital for manufacturing, which now makes up less than 10 percent of U.S. GDP.

No, the next radical jump in economic growth seems more likely to come from something that has a profound effect on everything, because it addresses the one permanent shortage in our entire economy: human time and attention. They are by far the most productive components of today's economy. About two-thirds of all income in the rich countries is paid directly for wages, and much of the remaining third represents indirect costs of labor. (For example, corporate income largely reflects earlier efforts by entrepreneurs.) So any innovation that could replace or dramatically improve human labor would be a very big deal.

ONE OF THE pillars of the modern singularity hypothesis in its many forms is that intelligence is a general elixir, able to cure many if not all economic ailments. Typically, this belief is expressed in the form of an argument that the arrival of very intelligent machines will produce the next singularity. Some people hope this arrival will follow a new Einstein, who will discover a powerful general theory of intelligence applicable to those machines. Others envision an "intelligence explosion" via a series of powerful design innovations, beginning with one that would make machines smart enough to help us quickly find a second innovation, allowing even smarter machines, and so on. A few even imagine innovations so unprecedentedly potent that a single machine embodying the first innovation could go through the entire innovation series by itself, unnoticed, within a week, and then take over the world.

There are many views on how intelligence might arise in a machine. One argument holds that hardware is the critical limiting factor and predicts that human-level machine intelligence will come soon after we have computer hardware whose performance is comparable with that of the human brain.

Another argument focuses on knowledge as the true limiting factor. This view is behind several huge artificial-intelligence database projects, including Cyc, under construction for 23 years

JOHN CASTI

EXPERT VIEW:
## John Casti

WHO HE IS
Senior Research Scholar, the International Institute for Applied Systems Analysis, in Laxenburg, Austria, and cofounder of the Kenos Circle, a Vienna-based society for exploration of the future. Builds computer simulations of complex human systems, like the stock market, highway traffic, and the insurance industry. Author of popular books about science, both fiction and nonfiction, including *The Cambridge Quintet*, a fictional account of a dinner-party conversation about the creation of a thinking machine.

SINGULARITY WILL OCCUR
Within 70 years

MACHINE CONSCIOUSNESS WILL OCCUR
Questionable

MOORE'S LAW WILL CONTINUE FOR
20 more years with current technology

THOUGHTS
"I think it's scientifically and philosophically on sound footing. The only real issue for me is the time frame over which the singularity will unfold. [The singularity represents] the end of the supremacy of *Homo sapiens* as the dominant species on planet Earth. At that point a new species appears, and humans and machines will go their separate ways, not merge one with the other. I do not believe this necessarily implies a malevolent machine takeover; rather, machines will become increasingly uninterested in human affairs just as we are uninterested in the affairs of ants or bees. But in my view it's more likely than not that the two species will comfortably and more or less peacefully coexist—unless human interests start to interfere with those of the machines."

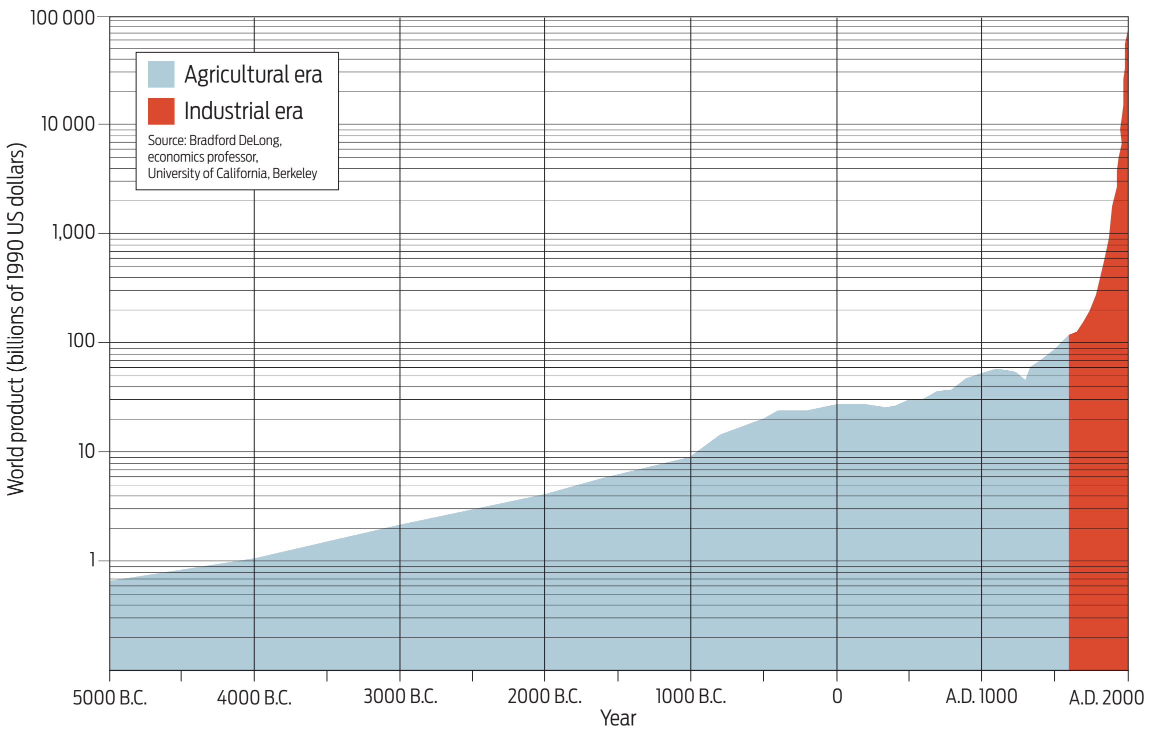


**SPOT THE TRANSITION:** After millennia of slow growth tied to rising population, the world economy took off. For the first time ever, output per capita rose above mere subsistence levels. It's been rising ever since.

and now at Cycorp in Austin, Texas. Cyc now possesses millions of pieces of commonsense knowledge, added mostly by hand. Eventually, Cyc may know enough to begin to read and assimilate all written knowledge, and the more it knows, the faster it should be able to learn. So it is possible, though hardly inevitable, that Cyc will eventually undergo a rapid knowledge explosion.

I find those scenarios interesting but unlikely to come to pass anytime soon. Regarding advanced machine intelligence, my guess is that our best chance of achieving it within a century is to put aside the attempt to understand the mind, at least for now, and instead simply focus on copying the brain.

This approach, known as whole brain emulation, starts with a real human brain, scanned in enough detail to see the exact location and type of each part of each neuron, such as dendrites, axons, and synapses. Then, using models of how each of these neuronal components turns input signals into output signals, you would construct a computer model of this specific brain. With accurate enough models and scans, the final simulation should have the same input-output behavior as the original brain. It would, in a sense, be the "uploaded mind" of whoever served as the template. Whether the emulation indeed constitutes a person and whether that person has rights is another story, to which I will return later.

If current trends continue, we should have computer hardware and brain scans fast and cheap enough to support this scenario in a few decades. What may well take longer are input-output models in sufficient detail for every relevant type of human neuron part. But I think those details will accrue in time. We already have sufficient models for some types of neuronal components, gathered after only a modest effort. And we have no reason to expect the other types to be harder.

> **Brain emulation would simulate the "uploaded mind" of whoever served as the template**

Project Blue Brain, a joint effort by IBM and the Ecole Polytechnique Fédérale de Lausanne, in Switzerland, has made some impressive progress: in December 2006, the project finished mapping and modeling the 10 000-odd neurons and 30 million synapses in a rat's neocortical column. Similarly impressive, in 2004 a Stockholm University team observed realistic behavior in a simulation of 8 million neurons and 4 billion synapses. But we still have far to go.

An emulation of a brain could merely do what that brain can already do, although if done in sufficiently powerful hardware, the cognition might occur faster. Still, even if all we were able to achieve was a computer with the mental powers of a particular human, that would be more than just interesting—it would also be incredibly useful.

Though it might cost many billions of dollars to build one such machine, the first copy might cost only millions and the millionth copy perhaps thousands or less. Mass production could then supply what has so far been the one factor of production that has remained critically scarce throughout human history: intelligent, highly trained labor.

Okay, so might these machines be conscious, with wills of their own, and if so, could they be selfish, even malevolent? Yes, yes, yes, and yes. More on that later; for now, let's get back to the economic argument.

Creating human-level intellect in a machine would be an astounding achievement, but it is not immediately obvious that it would launch a new era of much faster growth, with doubling times measured in months or less. After all, more and more capable machines have been replacing and aiding humans for centuries without sparking such an explosion. To answer that objection, we've got to start with the fundamentals: what economic theory says about growth rates.

To keep a modern economy thriving, we must accomplish many mental tasks. Some people (we call them engineers) have to design new products, systems, and services. Other people have to build, market, transport, distribute, and maintain them, and so on. These myriad tasks are mostly complements, so that doing one task better increases the value of doing other tasks well. But for each task, humans and machines may also be substitutes; it can be a wasted effort to have them both do the same task.

The relative advantages of humans and machines vary from one task to the next. Imagine a chart resembling a topographic cross section, with the tasks that are "most human" forming a human advantage curve on the higher ground. Here you find chores best done by humans, like gourmet cooking or elite hairdressing. Then there is a "shore" consisting of tasks that humans and machines are equally able to perform and, beyond them an "ocean" of tasks best done by machines. When machines get cheaper or smarter or both, the water level rises, as it were, and the shore moves inland.

This sea change has two effects. First, machines will substitute for humans by taking over newly "flooded" tasks. Second, doing machine tasks better complements human tasks, raising the value of doing them well. Human wages may rise or fall, depending on which effect is stronger.

For example, in the 1920s, when the mass-produced automobile came along, it was produced largely by machines, with human help. So machines dominated that function—the assembly of cars. The resulting proliferation of machine-assembled cars raised the value of related human tasks, such as designing those cars, because the financial stakes were now much higher. Sure enough, automobiles raised the wages of machinists and designers—in these cases, the complementary effect dominated. At the same time, the automobile industry lowered the pay of saddle makers and stable hands, an example of the substitution effect.

So far, machines have displaced relatively few human workers, and when they have done so, they have in most cases greatly raised the incomes of other workers. That is, the complementary effect has outweighed the substitution effect—but this trend need not continue.

In our graph of machines and humans, imagine that the ocean of machine tasks reached a wide plateau. This would happen if, for instance, machines were almost capable enough to take on a vast array of human jobs. For example, it might occur if machines were on the very cusp of human-level cognition. In this situation, a small additional rise in sea level would flood that plateau and push the shoreline so far inland that a huge number of important tasks formerly in the human realm were now achievable with machines. We'd expect such a wide plateau if the cheapest smart machines were whole-brain emulations whose relative abilities on most tasks should be close to those of human beings.

In such a scenario, the economy would start growing much faster, for three reasons. First, we could create capable machines in much less time than it takes to breed, rear, and educate new human workers. Being able to make and retire machine workers as fast as needed could easily double or quadruple growth rates.

Second, the cost of computing has long been falling much faster than the economy has been growing. When the workforce is largely composed of computers, the cost of making workers will therefore fall at that faster rate, with all that this entails for economic growth.

Third, as the economy begins growing faster, computer usage and the resources devoted to developing computers will also grow faster. And because innovation is faster when more people use and study something, we should expect computer performance to improve even faster than in the past.

Together these effects seem quite capable of producing economic doubling times much shorter than anything the world has ever seen. And note that this forecast does not depend on the rate at which we achieve machine intelligence capabilities or the rate at which the intelligence of machines increases. Merely having computer-like machines able to do most important mental tasks as well as humans do seems sufficient to produce very rapid growth.

CYPRESS SEMICONDUCTOR

## EXPERT VIEW: T.J. Rodgers

**WHO HE IS**
Founder and CEO of Cypress Semiconductor Corp., in San Jose, Calif., known for his brash opinions about the business world and politics. Owner of the Clos de la Tech winery and vineyards, in California, where he's trying to make the best American pinot noir.

**SINGULARITY WILL OCCUR**
Never

**THOUGHTS**
"I don't believe in technological singularities. It's like extraterrestrial life—if it were there, we would have seen it by now. However, I do believe in something that is more powerful because it is real—namely, exponential learning. An exponential function has the property that its slope is proportional to its value. The more we know, the faster we can learn.

"Technological transitions are required to maintain an exponential rate of learning. The first airplanes were certainly not as good as well-appointed trains in moving masses comfortably, but the transition later proved essential to maintaining our progress in human mobility. Gene splicing is a breakthrough technology but has not yet done (or been allowed to do) a lot for mankind. That will change.

"I don't believe in the good old days. We will be freer, more well-educated and even smarter in the future—but exponentially so, not as a result of some singularity."

LIFE IN A robot economy would not be merely a sped-up version of our lives today. When I apply basic economic theory and some common sense to this scenario, I conclude that humans would probably be neither the immortal, all-powerful gods that some hope for nor the hated and hunted prey that some fear.

Yes, robot-human wars would be possible, but it is important to remember that few differences between humans ever lead to war. We do not fear that the short will conspire to murder the tall in their sleep, nor that the right-handed will exterminate the left-handed. Short, tall, left-handed, and right-handed people all trade with, befriend, and marry one another with abandon, making such wars almost unthinkable. Instead, wars today happen between largely separate nations and ethnic groups. Similarly, robots well-integrated into our economy would be unlikely to exterminate us.

Wages could fall so far that most humans could not live on them

Would robots be slaves? Laws could conceivably ban robots or only allow robots "born" with enough wealth to afford a life of leisure. But without global and draconian enforcement of such laws, the vast wealth that cheap robots offer would quickly induce a sprawling, unruly black market. Realistically, since modest enforcement could maintain only modest restrictions, huge numbers of cheap (and thus poor) robots would probably exist; only their legal status would be in question. Depending on local politics, cheap robots could be "undocumented" illegals, legal slaves of their creators or owners, "free" minds renting their bodies and services and subject to "eviction" for nonpayment, or free minds saddled with debts and subject to "repossession" for nonpayment.

The following conclusions do not much depend on which of these cases is more common. For example, in any of these cases human wages would rise or fall rapidly, depending on the shape of the human advantage landscape. After the flood of the plateau, there might still be some mountain peaks of human tasks left. Some rich people might still want to be served and entertained by real human beings. So for those jobs, human wages could rise. But if in the end the machine ocean completely inundated all of Task-Land, then wages would fall so far that most humans would not, through their labor alone, be able to live on them, though they might work for other reasons.

In either case, human labor would no longer earn most income. Owners of real estate or of businesses that build, maintain, or supply machines would see their wealth grow at a fabulous rate—about as fast as the economy grows. Interest rates would be similarly great. Any small part of this wealth should allow humans to live comfortably somewhere, even if not as all-powerful gods.

Because copying a machine mind would be cheap, training and education would cost no more than a software update. Instead of long years to train each worker, a few machines would be trained intensely, and then many copies would be made of the very best trainees. Presumably, strong security would prevent bootleg copies.

Organizational decision cycles would shorten, favoring streamlined, decentralized processes run by fast machine minds in key positions of authority. Fast minds could be whole-brain emulations sped up relative to human brains. This scenario would marginalize slow bureaucratic human committees, regulators, and the like. Fast growth rates would likely discourage slow long-distance transport and encourage local production.

Some robots responsible for administration, research, law, and other cognitive work might live and work entirely in virtual environments. For others, crude calculations suggest that tiny bodies a few millimeters tall, with sped-up minds to match their faster body motions, might allow insectlike urban densities, with many billions living in the volume of a current skyscraper, paying astronomical rents that would exclude most humans.

As emulations of humans, these creatures would do the same sorts of things in their virtual realities and skyscrapers that humans have done for hundreds of thousands of years: form communities and coalitions, fall in love, gossip, argue, make art, commit crimes, get work done, innovate, and have fun. Just as farming was more alien to our human nature than hunting and gathering, and industry was more alien still, their world would be even more distant from human origins. But human nature seems flexible enough to accommodate such changes.

The population of smart machines would explode even faster than the economy. So even though total wealth would increase very rapidly, wealth per machine would fall rapidly. If these smart machines are considered "people," then most people would be machines, and per-person wealth and wages would quickly fall to machine-subsistence levels, which would be far below human-subsistence levels. Salaries would probably be just high enough to cover the rent on a tiny body, a few cubic centimeters of space, the odd spare part, a few watts of energy and heat dumping, and a Net connection.

While copying would make robot immortality feasible in principle, few robots would be able to afford it. And when reproduction via copying dominates, few robots would be able to afford robot versions of human children.

While whole-brain-emulation robots would be copies of particular humans, we should expect vast inequality in copy rates. Investors who paid the high costs for scanning a human brain would carefully select the few humans most likely to be flexible, cooperative, and productive workers, even while living a short, hardscrabble, childless, and alien life in robotic bodies or virtual offices. Investors who paid for copying existing machine minds would select robots with a track record of achieving this ideal. As a result, there would be large first-mover advantages and winner-take-all effects. For example, if docile minds turned out to be the most productive, then the robot world might consist mainly of trillions of copies each of a few very docile human minds.

In this case, the meek would indeed inherit the Earth. ❑

*TO PROBE FURTHER For additional resources on reconstructing the deep economic past, speculations on a rapid intelligence explosion, and the likely effects of machine intelligence on economic growth, see http://spectrum.ieee.org/jun08/singularityprobe.*