

The Economics of Brain Emulations

Robin Hanson *

[a]Introduction

Technologists and economists both think about the future sometimes, but they each have blind spots.

Technologists think about specific future technologies, which they may foresee in some detail. Unfortunately, such technologists then mostly use amateur intuitions about the social world to predict the broader social implications of these technologies. This makes it hard for technologists to identify the technologies which will have the largest social impact.

Economists, in contrast, have a professional understanding of the social world, and are well-positioned to analyze the social implications of specific technologies. Using simple mathematical models based on powerful general concepts, economists could go well beyond simple trend projections.

Unfortunately, economists mostly rely on amateur intuitions about the feasibility of future technologies. Substantial technical innovations seem to

* Department of Economics, George Mason University, Washington, D.C.

them like "science fiction," being too silly to take seriously. Economists' future projections thus usually ignore specific future technologies.

As an economist (tenured professor) with a technology background (a physics masters and nine years of computer research), I try to avoid these blind spots. By applying economic theory to specific future technologies, I hope to go beyond trend projection to foresee the social consequences and relative importance of future technologies.

Of the many future technologies I have considered over the years, one stands out to me as likely to have the largest impact: brain emulations. This technology also happens to be relatively easy to analyze with standard economic tools. But before discussing this technology, let me outline an independent reason we have to expect a huge economic transition in the next century.

[a]Long Term Trends

A postcard summary of life, the universe and everything might go as follows. Our universe appeared and started expanding. Life appeared somewhere and then on Earth it began to make larger and smarter animals. Eventually humans appeared and became smarter and more numerous, by inventing language, then farming, then industry, and most recently computers.

The events in this summary are not evenly distributed in time. The first events are relatively evenly distributed: the universe started fourteen billion years ago, life appeared by four billion years ago, and on Earth animals started growing larger and smarter about half a billion years ago. But by comparison the other events are very closely spaced: our species appeared a few million years ago, farming started about ten thousand years ago, industry started about two hundred years ago, and computers started a few decades ago.

One might worry that this list of events is a biased sample of the globally important events, because we humans over-emphasize events that are about us. But I think not; I think these are in fact the globally important events, because they separate a chain of distinct key exponential growth modes. Exponential growth is where some quantity doubles after a certain time duration, and then continues to double again and again after similar durations. At each point in history some crucial quantity has been growing exponentially. And at a few rare transition points, the growth rate has suddenly increased.

The slowest growth mode started first. Our fourteen billion year old universe is expanding, and that expansion is now roughly exponential due to a mysterious "dark energy." The distance between the galaxies is predicted to double every ten billion years.

We don't know enough about the history of non-animal life in the universe to identify its growth rates, but we can see that for the last half billion years the

size of animals on Earth has been growing exponentially. While the size of the typical animal has changed little, the variation among animal sizes has greatly increased. Because of this, the mass of the largest animal has doubled about every seventy million years, and the mass of the largest brain has doubled about every third of a hundred million years. So the largest brains have doubled about three hundred times faster than the distance between galaxies.

Humans (really "our human-like ancestors") began with some of the largest brains around (relative to their bodies), and then tripled their brain size. Those brains, and the innovations they embodied, seem to have enabled a huge growth in the human niche - it supported about ten thousand humans about two million years ago, but about four million humans about ten thousand years ago.

While data is scarce, this growth seems roughly exponential, doubling about every two hundred thousand years. This is one hundred and fifty times faster than animal brains grew. (This growth rate for the human niche is consistent with faster growth for our ancestors, as some groups killed off others to take over the niche.)

About ten thousand years ago, those four million humans began to settle and farm, instead of migrating to hunt and gather. The human population on Earth then began to double about every nine hundred years. This farming growth rate is about two hundred and fifty times faster than hunting humans doubled.

Since the industrial revolution began a few hundred years ago, humans have grown even faster. Before the industrial revolution total human wealth grew so slowly that population quickly caught up, keeping wealth per person near a subsistence level. But in the last century or so wealth has grown faster than population, allowing for great increases in wealth per person.

Economists' best estimates of total world product (average income per person times the number of people) show it has been growing exponentially over the last century, doubling about every fifteen years, or about sixty times faster than under farming. And a model of the whole time series as a transition from a farming exponential mode to an industry exponential mode suggests that the transition is not over yet - we are slowly approaching an income doubling time of about six years, or one hundred and fifty times the farming growth rate.

A revised postcard summary of life, the universe, and everything, therefore, is that an exponentially growing universe gave life to a sequence of faster and faster exponential growth modes. First the largest animal brains grew slowly, and then the wealth of human hunters grew faster. Next farmer wealth grew much faster, and finally industry wealth grew faster still. Perhaps each new growth mode could not start until the previous mode had reached a certain enabling scale. That is, perhaps humans could not grow via culture until animal brains were large enough, farming was not feasible until hunters were dense enough, and industry was not possible until there are enough farmers near each other.

Notice how many important events are left out of this postcard summary. Fire, writing, cities, sailing, printing presses, steam engines, electricity, assembly lines, radio, and hundreds of other key innovations are not listed separately here. The reason is that most big changes are a *part* of some growth mode, but do not cause an *increase* in the mode's growth rate. While we do not know what exactly has made growth rates change, we do see that the number of such causes observed so far can be counted on the fingers of one hand.

While growth rates have varied widely, growth rate changes have been surprisingly consistent -- each mode grew from one hundred and fifty to three hundred times faster than its predecessor. Also, the recent modes have made a similar number of doublings before giving rise to a new mode. While the universe has barely completed one dark energy doubling time, and the largest animals grew through sixteen doublings, hunting grew through nine doublings, farming grew through seven and a half doublings, and industry has so far completed a bit over nine doublings.

This pattern explains event clustering - transitions between faster growth modes that double a similar number of times *must* cluster closer and closer in time. But looking at this pattern, we should wonder: are we in the last growth mode, or will there be more?

[a]A New Growth Mode?

If a new growth transition were to be similar to the last few, in terms of its number of prior doublings and its increase in the growth rate, then the remarkable consistency in the previous transitions allows a remarkably precise prediction. A new growth mode should arise sometime within about the next seven industry mode doublings (i.e., about the next seventy years) and give a new wealth doubling time of between roughly one and two weeks.

How sudden would such a transition be? We only have transition data on the last two transitions, and of those the industry transition was smoother. If the next transition happened around 2040, and was as smooth as the industry transition, then a simple model predicts the sequence of expected annual growth rates to be: 6.1%, 6.1%, 6.6%, 8.0%, 14%, 41%, 147%, 475%, 1025%, and so on. If growth rates fluctuate by about 0.5% per year, then growth rates would have doubled within two years of any noticeable change, and within two more years the world economy would be doubling more than yearly.

The suggestion that the world economy will soon double every week or two, after a transition lasting only a few years, seems so far from ordinary experience as to be, well, "crazy." Of course similar predictions made before the previous transitions would have seemed similarly crazy. Nevertheless, it seems hard to take this scenario seriously without at least some account of how it could be possible.

Now we should not expect to be able to get a very detailed account of a new growth mode. After all, most economics has been designed to explain the actual social worlds that we have seen so far, and not all the possible social worlds that might exist. And we are still pretty ignorant about the fundamental drivers of the previous modes. But we do want at least a sketchy account. Of the many future technologies that technologists have forecast, which could plausibly have anywhere near this impact on the economy?

One helpful hint is that innovations in larger economic sectors can produce larger social impacts. In the United States we spend about 1.5% of income on farming, 1.5% on mining, 2% on gas and electricity, 2.5% on communications, 3% on transportation, and 3.5% on construction. These small fractions make it hard to see how even dramatic innovations in these sectors could induce much faster growth. For such drama, we must look beyond the usual technology favorites, such space colonization, fusion energy, air cars, sea cities, or picture phones. We probably must even look beyond radical nanotechnology; while this might dramatically reduce the cost of capital for manufacturing, we only spend about 5% of income there.

A more promising fraction is the 70% of income we now pay to human labour, as opposed to other kinds of physical and social capital. Greatly lowering this cost could have a huge impact. And a robotics or artificial intelligence technology good enough to substitute wholesale for most human labour might just greatly lower such costs.

[a]Brain Emulations

For centuries now, people have been concerned about the possibility of machines replacing human labour. And many kinds of labour have in fact been replaced by machines. At first machines replaced humans at tasks needing physical strength, but more recently machines have replaced humans at mental tasks.

On the whole, however, machines have mainly helped humans be more productive at tasks that machines cannot do. By complementing humans, machines have so far greatly raised the value of most human labour. Because of this, most economists have not worried about machines replacing humans.

Previous trends need not continue, however. The key point to understand is that while *tasks* complement each other, *individuals* are substitutes for doing each task. There are many tasks that we want done, and machines are better suited to some tasks than to others. Thus slowly improving machines have two effects on human labour. First, machines get better at the tasks machines do best, which makes doing all the other tasks well more valuable. This complementary effect raises the demand for human labour. Second, some marginal tasks switch from humans to machines. This substitution effect lowers the demand for human labour.

So far humans still do most tasks worth doing, and so the net effect has been to raise human wages. But this picture would change dramatically if we had

machines that were good at almost all the tasks people now do. Human wages could then fall with the falling price of machines. And since the number of machines could grow as fast as the economy needed them, human population growth would no longer limit economic growth. Simple growth models can easily allow a new doubling time of a month, a week, or less.

Now admittedly, progress in robotics and artificial intelligence has been slow; it has been hard to write capable software. At current rates of progress it could be centuries before machines could get good at almost all tasks that people do. There is one approach to broadly capable artificial intelligence, however, that seems likely to succeed within the next century: brain emulations.

The idea here is to not "write" the relevant software, but to "port" it from a real human brain. Take a brain, and scan it in enough detail to see each neuron's type and its connections to other neurons. Study each type of neuron in enough detail to create a computer model of how its output signals depend on its input signals. Finally, create a computer model of the entire brain, connecting together models of each neuron, and connecting them to emulated eyes, ears, mouth, etc.

If the connection information and the neuron models are good enough, then the brain model should have roughly the same input-output behavior as the original brain. That is, you could talk to it and it would talk back. And if you could convince it to work for you, it could accomplish most of the same sort of

tasks as the original brain. It might even be conscious and enjoy its life (though this claim may long remain disputed). And once you had one such brain, you could as make billions more by just copying the software.

Three technologies are needed to make this work: enough neuron-type models, fast enough scanning, and large enough computers. These technologies have been steadily improving for many decades now, and they each seem likely to be ready by mid-century, if not by quarter-century. Of the three, progress in neuron modeling seems the hardest to forecast. But we already have good enough models of many neuron-types, we already have slow but accurate-enough scanners, and computers should be fast enough in a few decades. No grand breakthroughs seem required, just hard work and steady progress of the sort we have already seen.

Thus brain emulations seem likely to appear in time to cause the next big growth mode, and simple economic models suggest they are capable of producing such a mode. Within a few years human wages could begin falling dramatically, while economic growth rates skyrocket.

More precisely, such changes could happen if they were allowed. To keep our models simple and comparable, economists usually start by modeling peaceful low-regulation scenarios, such as where wages are set by supply and demand, and where people could make as many brain emulations as they wanted. One can imagine wealth transfers to ensure that humans do not starve due to falling

wages, and minimum wages or population controls to limit the number of brain simulations. Given the lack of a strong world government, it is not clear whether such regulations would be feasible, or if feasible whether they would be desirable. As important as these questions are, my space is up, and so they will have to wait for another essay.

For formal analysis and citations, see Hanson, *The Economics of Science*

Fiction: <http://hanson.gmu.edu/econofsf.html>