LISA MATHIAS

When the Economy Transcends Humanity

By Robin Hanson

What will our economy, workplaces, and society look like when we can copy our brains and build virtual workers to do our jobs? An economist looks at the next great era, a world dominated by robots.

What might a world full of robots as smart as humans look like? Experts in robotics and artificial intelligence have given a lot of thought to if and when such robots might appear. Most say it will happen eventually, and some say it will happen soon.

Knowing *when* advanced robots will appear doesn't tell us how they will change the world. For that, we need experts in social science, like economists.

Here, I outline a scenario of what a new robot-based society might look like. Some people say I shouldn't do this, because it's impossible, while others just say it is unscientific. Even so, I'm doing it anyway, because it seems useful—and it's fun.

Keep in mind, however, that I'm not arguing that this scenario is good; I'm just applying basic economics to make best guesses about what things would actually be like. While most of you have probably seen movies depicting worlds with smart robots, as an economist I intend to show you we can do a *lot* better by using careful economic analysis. We can actually say quite a lot about this new world.

Rise of the Em

There is a vast variety of possible robot minds, and a robot world will depend greatly on which minds inhabit it. Will these robots look like the Terminator or R2D2? Will they kill like HAL in *2001: A Space Odyssey?* These images of artificial intelligence are rooted in science fiction, but they don't represent what I think is the most likely form that future AI will take.

Picture a robotic mind of the future as though it were a digital version of a regular human, inhabiting a world that looks exactly like the one in which you actually live. The AI, in a very real sense, is you.

With computers today, we often "port" software running on old machines to get software that can run on new machines. To port software, we write an "emulator" to make the new machine look like the old machine to the software. Instead of writing new software for robot computers, we could instead try to port software from the old computer that is the human brain.

How could we take the memories, skills, hopes, dreams, and personality that are in a human brain and turn it into software for a computer? First, we need to take a particular human brain and scan it in great spatial and chemical detail, seeing which cell types are where, connected to what, and with what relevant internal states. Second, we need adequate models of each type of human brain cell—models saying how input signals are translated into output signals and internal state changes. Third, we need lots of parallel computer hardware.

If we fill up that computer hardware with emulations of each cell, emulations that are typed and connected just like the real cells in the scanned brain, then if the cell models are good enough the whole model must also be good enough. That is, it must have the same input-output signal behavior as the original brain. So, if you hooked it up to artificial eyes, ears, hands, etc., you could ask it to do a job and hope it might cooperate and do the job much like the original human.

Now, it isn't clear just how much cell detail we'll need in order to emulate brains. The more detail, the more computing hardware we'll need, and the longer it would take before computers are cheap enough to emulate brains cost-effectively. The best analysis so far—Anders Sandberg and Nick Bostrom's 2008 paper "Whole Brain Emulation: A Roadmap"—suggests that, if computer hardware continues to improve at past rates, we will have cheap emulations in about 30 to 75 years. And if computer hardware gains slow down by a factor of two, it will take 60 to 150 years.

Robots are likely to arrive first in the form of brain emulations, before smart software written by humans. Because they'd be digital copies of humans, we can say a lot about what their world will be like. I'll call them "ems" for short.

"Picture a robotic mind of the future as though it were a digital version of a regular human, inhabiting a world that looks exactly like the one in which you actually live. The AI, in a very real sense, is you."

Not All Robots Are Robotic. Ems Feel Human

When I say that ems are psychologically human, I mean that they remember a human life and retain the usual human inclinations. Ems love, laugh, fight, cry, joke, envy, work, play, and take pride. They have friends, lovers, bosses, and colleagues. Their world looks and feels to them much like our world looks and feels to us. Sometimes they see the same physical world that we see. At other times, they see an entirely virtual world running on computer hardware.

As in advanced economies today, most em jobs would be "desk" jobs. They would solve complex problems, develop strategies, publicize, and so on. Unlike today, however, most em offices would sit in virtual realities. Some ems, like plumbers, would need physical bodies to do their jobs. They may spend part of their day inhabiting a physical robot shell, but they would spend most leisure time in virtual realities. It should be cheap to give ems virtual realities that are spectacularly comfortable, engaging, beautiful, and inspiring. Ems also need never experience hunger, pain, or disease, unless they want to do so. In virtual reality, nothing needs cleaning or replacing, since there is no grime or wear. It would often be good to be an em.

On the downside, however, ems would spend most of their waking hours working. Because copying ems is easy, a single good em worker could take all the jobs in a profession. As we will see, this makes em labor markets very competitive. A few very productive em workers, each willing to work long hours at near-subsistence wages, would each be copied billions of times. Furthermore, most ems would be temporary copies created to do a task lasting several hours, and then end or retire.

When choosing which humans to scan, or which ems to train and copy for new tasks, firms would select the smartest, hardest working, most flexible, and cooperative folks with the best-matched attitudes. They'd also select the most productive "tweaks" to the emulation process itself, to make ems especially attentive, inspired, energetic, etc.

So, while ems would be recognizably human in having mental features that fall near the usual range of human variation, ems would be distributed quite differently within those ranges. The typical em would be more productive than the best humans today. Ems might sympathize with, and be grateful to, their human ancestors, but ems would feel and be superior in most ways they care about.

The human-to-em selection process would be highly unequal. While the em economy would quickly grow to have billions and then trillions of ems (since virtual space is limited only by computer memory size), most of these would be copies of a dozen to a thousand of the most-suitable humans. While those best humans could each have a "clan" of billions of em copies, most humans would give rise to just a few ems. These clans would represent towns or nations of the same ems (think Agent Smith in *The Matrix,* but nicer).

The copy clan would be a natural unit of em organization. Because copies are similar and share many interests, ems would tend to trust their clan's advice on love, work, finances, and so on. Each em might feel like a visitor from a planet full of folks just like her. Just as politics in ancient cities often consisted of shifting coalitions of great families, em politics might consist of shifting coalitions of copy clans.

With only a few hundred important clans, the em social world becomes more like the social world of our forager ancestors, where people only ever met one or two hundred people in a lifetime and knew each one very well. An em meeting with a George might know George's basic personality and style well, even if she had never met this particular George before.

Slow Human Computers and Fast Em People

Computers can be fast or slow; faster computers are more expensive, but do more calculations per second. This ability to pay more for speed works especially well when calculations can be done in parallel, with many processors running at the same time. For parallel calculations, it costs about twice as much for a machine that calculates twice as fast.

The fact that the human brain calculates in parallel implies a very wide range of speeds where emulation calculations can be done twice as fast if one uses a computer that costs twice as much. This range includes kilo-ems and milli-ems, who think and execute tasks a *thousand* times faster or slower than humans, and may include mega-ems and micro-ems, who run a *million* times faster or slower.

Human brains have a reaction time of about a tenth of a second, because human body parts can't change position faster than that. Since body parts half as long oscillate twice as fast, you'd need a double-speed em to react fast enough for a body that is half of a human's height. For a body 2 millimeters tall, you'd need a kilo-em.

Most ems, however, would live in virtual reality without a physical body. Such ems could meet in the same virtual room even when their brains (i.e., emulation hardware) are far apart. This would greatly cut travel, and hence travel congestion, which is the biggest limit to city size today. Thus, em cities may be much bigger and denser, perhaps holding populations in the trillions. Em cities may instead be limited in size by difficulties in cooling lots of dense, fast computer servers. Such cities might look somewhat like the dense tall cities of *The Fifth Element* or *Star Wars: Episode II—Attack of the Clones*.

Em minds would run at speeds matched to the jobs they perform. Ems controlling fast machines, like nanotech factories, must run fast, while those controlling slow machines, like boats, can run more slowly. Ems in product-development races might run as fast as possible, while poor em retirees might run as slowly as possible. Ems who take their leisure time much faster than work can offer near 24/7 service to clients, and bosses who run faster than subordinates can better coordinate big organizations.

Ems faster than kilo-ems, and who meet together virtually but actually have brains on opposite sides of a city, would notice signal delays due to the speed of light. On the other hand, for ems slower than kilo-ems, their subjective century-or-so careers would have hardly started before job practices change, wasting their job-training investments. Together, these delay and career effects suggest that kilo-ems are near the typical em speeds.

You might wonder, would these virtual people ever die? And if not, how long would they work? Why not keep them working forever?

I estimate that an em career would last only a subjective century or so. After that, the em would "retire" and spend the rest of his or her virtual days playing virtual shuffleboard (or doing anything else he or she liked). Why such a low job half-life? Human minds are like most complex adapting systems: As they adapt to an environment, they get more complex and fragile, and less able to adapt later to very different environments. We see this effect in software "rot," in manufactured products, in biological cells and species, and in human minds today. So we expect this pattern to continue with em minds. Even though they could last forever, old em minds would still become obsolete.

Faster ems would have many features we associate today with higher status. They would tend to be richer, be bosses, win arguments, host meetings, and sit at central locations. Em speeds would also clump, because ems running at different speeds couldn't interact naturally, unless one of them temporarily ran at the others' speeds. Thus, ems would be ranked into a discrete hierarchy of social classes running at different speeds.

The Copy Economy and Runaway Economic Growth

The most important feature that all these artificial brains share is being easy to copy. This has many big implications. First, it makes training much cheaper. You can train one robot and then make millions of copies having that skill. Second, easy copying would greatly reduce the wages of robots and of the humans who compete with them. This is simply a matter of supply and demand. When anyone can easily make many robots, the cost to rent robot workers must fall to near the cost of renting robot hardware.

Human workers competing with very cheap and productive robots could not live on the resulting wages. Humans will need income from charity, or from stocks, bonds, and real estate, whose values would skyrocket. Otherwise, they'd starve. Today, we don't kill unproductive retirees and take their stuff, partly from fear of destroying the financial, legal, and political institutions we rely on to keep the peace among the rest of us. Humans might similarly hope to survive as marginalized "retirees" in a robot economy.

"The population of robotic minds might quickly reach trillions, greatly outnumbering humans."

How much faster could an economy of easily copied robots grow? Standard economic theory says that even a one-week doubling time isn't crazy. It takes both workers and tools together to make things; because we are slow at growing and improving workers, our economy now grows mainly by inventing better tools. While our factories can make more tools quickly, those don't help much without more workers to use them. A robot economy, however, can quickly add both more tools and more workers to use them.

Compared to growth today, robot growth depends less on increasing quality and more on increasing quantity. The population of robotic minds might quickly reach trillions, greatly outnumbering humans. Note, however, that this need not imply philosophical, political, or military conflict with humans. Remember the type of robots we are talking about here. Psychologically speaking, ems feel human.

An economy that grows faster would be different in some predictable ways. For example, faster economic growth implies higher interest rates, which is bad for projects and activities that take a long time to produce an economic return. These may include huge construction projects, space travel, or shipping goods long distances. A faster economy would thus have faster construction, less space travel, and more local production. Long regulatory delays would also be expensive; fast, agile governance would be necessary.

A fast-growing economy would quickly grow big, and bigger economies are also different in predictable ways. For example, they have more room for specialization, and they could pay for better and more varied movies, music, novels, and art. A larger total population can hold bigger versions of organizations like firms and cities.

A bigger future economy can also contain more inequality: bigger differences in nations, firms, and cities; between the richest and poorest; and between the largest and smallest.

How the Previous Eras of Man Suggest a New Robot Era

Is the dawning of a robot economy at all believable in the foreseeable future? It may be if we consider the robot economy to be the next step in a pattern of previous eras of human activity, following the eras of foragers, farmers, and industry. First, foragers wandered the land in search of food, then farmers grew their food while staying in one place, and finally industrial folks had machines to help them make food and much more.

These eras are distinguished by how fast we grew in our capacity and power. Each era is characterized by the speed at which innovations could be shared. Starting about 2 million years ago, foragers learned to share their innovations via culture; their economy (which tracked their population) doubled every quarter million years. Starting about 10,000 years ago, farmers learned to share innovations much faster, perhaps via long-distance trading networks. The farming economy doubled every thousand years.

Finally, a few hundred years ago, we learned to share innovations even faster, via communication among networks of technical experts. Our industrial economy has since doubled about every 15 years. This growth has been faster even than our population grew, allowing a great increase in wealth per person for the first time in history. Each new era grew faster but lasted a shorter time, to encompass a similar amount of total growth. Each era has seen seven to nine doublings in capacity, and similar jumps in growth rates relative to preceding eras.

If we try to extend this pattern to eras before humans, we find something that roughly fits: Animal brains doubled roughly 16 times every 30 million years, starting a half billion years ago. Might this pattern also extend into the future? If another era comes after ours, and fits this historical pattern, it would appear sometime in the next century or so; within a few years, its economy would double every few weeks. That new economy might be a robot economy.

I warned at the start of this article that I'm just applying basic economics to make best guesses about what a future robot world would actually be like. I'm not saying this scenario is good or bad.

If this future actually happens, it will, of course, be very important to think carefully about which aspects of this scenario are good and to be encouraged, and which are lamentable and to be minimized. Surely, the first task is to figure out what is likely to happen if we do nothing. This is the task I have begun. ❑

About the Author

Robin Hanson is an associate professor of economics at George Mason University, hanson.gmu.edu. He is known as a founder of the field of prediction markets, and also has interests in signaling, health, politics, law, disagreement, hypocrisy, and the future. He is writing a book on the subject of this article.