TACITUS LECTURE 2018

The Myth of The Innovation Menace

Why We Should Stop Worrying and Learn to Love AI

Nathan Myhrvold

22 February 2018

Are you about to be replaced by a machine? Will artificial intelligence (AI) or other new technology make your job obsolete?

It's hard not to wonder these things in a world that has become obsessed with technological progress. We love some progress, which is why lines form outside the Apple store a day ahead of a new iPhone release. But watch the news feed on your shiny new iPhone 10, and you'll see story after story predicting that robots, AI, or some other technology is coming for your job or for jobs in your community. They will disrupt the economy and all we hold dear.

This isn't a new story at all. In 1811, a group of Nottingham stockingers—people who wove stockings—revolted against new knitting machines, smashing them with sledgehammers. They even burned down several textile mills that used them.

Their movement came to be named for its fictional leader, Ned Ludd—we know them today as the Luddites. Although their name entered the language as a synonym for haters of technology, they were neither the first nor the last of their kind. Decades earlier in 1788, Parliament had passed the "Protection of Stocking Frames Act," which called for imprisonment or exile for damaging textile equipment. A later law in 1812 added the death penalty. Although some Luddites were hanged, the more common solution was penal transportation—as terrible as it may sound, many Luddites forcibly became Australians.

The idea that new technology was something to be feared is thus an old idea, but one originally confined to the working class. That is only natural: manual laborers were the ones who had the most to lose as new machinery made what few skills they had obsolete. The technology of 1811 could hardly replace the most skilled workers of that era, and that remains true even today. Instead, innovation was focused on increasing the speed and lowering the cost of incredibly rote and repetitive work like tying millions of tiny knots over and over again to knit a garment.

The modern heirs of Luddism are not uneducated workers at the bottom of the labor market. Instead the notion has become fashionable among highly educated intellectuals and even some economists. Open up The Guardian or the Financial Times, and you see headlines like "A.I. and robots threaten to unleash mass unemployment, scientists warn" and ominous alarms that "We're heading towards mass unemployment at the hands of technology."

A widely reported study done in 2013 at Oxford University argued that 47% of the US job market is "at high risk" to computerization within the next 15 years. MIT Technology Review, of all places, has run articles like "How Technology is Destroying Jobs." A report last year by a German IT association projected that robots and computers will destroy 10% of German jobs within five years.

The Innovation Menace

The common element to these claims is what I call the innovation menace: the fear that technological innovation will displace millions of workers, lead to widespread hardship, and throw society into crisis.

H.L. Mencken memorably defined puritanism as that haunting fear that somewhere, somebody is having fun. To borrow his formulation, the innovation menace is the fear that somewhere, somebody might have a good idea. Even worse, they might deploy that idea!

The striking thing about the whole notion of the innovation menace is that, despite intense study with the best available scholarship, it is an idea that throughout history has *never once proved correct*. Adoption of new technology has never been the primary cause of significant unemployment or widespread hardship.

Take nineteenth-century Nottinghamshire. It was hardly a cesspit of unemployment; to the contrary, it become one of the centers of the industrial revolution, the "Silicon Valley" of its day. The Nobel-prize-winning economist Robert Lucas pointed out that what really differentiated the industrial revolution is that "the living standards of the masses of ordinary people have begun to undergo sustained growth.... Nothing remotely like this economic behavior is mentioned by the classical economists, even as a theoretical possibility."

That doesn't sound like something to riot over, but the Luddites did. It's a safe bet that most of the stockingers that didn't either hang or help start Australia's wool industry wound up getting jobs in the new mills or elsewhere in the booming economy that helped define the industrial revolution.

Meanwhile there *were* terrible human crises in the 19th century that killed millions of people. The Napoleonic wars were contemporaneous with the Luddites. The Irish potato famine of 1845 happened soon after, and the American Civil War shortly after that. All of these hinged on intrinsically human issues—the ambition of narcissistic leaders, bad government policies, and dependence on the most radically un-technological form of labor: slavery.

Human failings like these really are menaces to society—they were then, and they are now. Similarly, it's mainly human mistakes, not new technologies, that cause recessions and depressions, with all the economic hardships those bring.

This Time Is Not Different

From our vantage point today, it's easy see that the Luddites got it all wrong; their fear of stocking frames now seems naïve and quaint. But a long line of prognosticators throughout the 19th, 20th and now 21st centuries has made essentially the same mistake.

Each time, some new technology plays the role of the villain. The predictions vary, but they are always dire. And even though they have never come true, new pundits keep taking up the cause of the innovation menace while dismissing the fact that it is an entirely failed line of argument.

Their reasoning is summed up neatly in the title of a book about the 2008 financial crisis, which has the property that once you have read the title, you may not need to read the book. It's called *This Time is Different* by Carmen Reinhart and Kenneth Rogoff.

Every time the intellectual menace comes up, the people involved flatter themselves to be, uniquely in human history, the first to recognize the danger that innovation poses in this new context. They dismiss past examples as irrelevant because *this time is different*. Tellingly, neo-Luddites are very eager to say that they are not Luddites. But after some throat clearing and chest thumping, they proceed to fall into the same intellectual traps.

It's worth looking at some examples. Forty years ago, in 1978, the average daily trading volume on the New York Stock Exchange was about 29 million shares a day. As with most exchanges for securities, the process was quite arcane. People would phone their brokers, who would then contact people at the exchanges, and those people would then run orders on little slips of paper down to the exchange floor.

These would go to men in loud-colored jackets, who would shout at each other and execute trades. I remember thinking, when I first saw the whole system in the 1980s, that it seemed ripe for computerization.

And that is what happened. New trading technologies did away with the little slips of paper. The technological upheaval followed on the heels of regulatory easing, like the 1986 Big Bang deregulation here in the UK.

The predictable result was that commission rates on buying and selling stocks plummeted. Retail investors who used to pay commissions of 1% or more now routinely trade stocks for free or for flat fees of a few dollars.

So should we mourn the poor displaced workers in the City who were impoverished by this technological onslaught? Must we protect Bloomberg terminals like stocking frames from the unemployed finance workers, and ship those who desecrate them off to Australia?

Of course not. By any metric, the financial industry has never been stronger. Average daily volume of stocks traded in New York is now over *a billion* shares a day. If you had told a trader in 1978 that this would happen, he would have said "that's impossible." There is no way that there could be enough little slips of paper or men in loud jackets to handle a billion shares a day. And even if there were, there is no way that investors could make that many decisions in a day.

What that 1978 trader couldn't imagine is that computers wouldn't just change the trading floor—they would change everything. On some days, two-thirds of market volume is automated trading between computers. The net result of this major technological shift was to strengthen the financial sector, not undermine it.

Here's another example: remember the so-called "digital divide?" When I was at Microsoft, our motto for many years was "a computer on every desk and in every home." In the mid 1980s, this was a radical concept, and many people told us we were foolish to think that everybody needed or wanted a computer. By the late 1990s, the combination of personal computers and the Internet became so compelling that our motto was much easier to sell.

But a group of innovation menacers decided that this was going to be bad for society because it would split people into the "digital haves and have-nots." The U.S. Congress even held hearings on this. This digital divide, we were told, would be a great inequity, widening the gap between rich and poor, and undermining economic mobility.

To anyone who understood the computer industry, this was absurd on its face. Computing power has dropped faster in price than any other technology in history. The Internet gave anyone—rich or poor—access to more information than the richest industrialist or most powerful government could have dreamed of just a few years prior. The rollout of the Internet to everyone was the most egalitarian and democratizing event in the history of communication.

So despite dire predictions of the digital divide by myriad experts, it never materialized. People both rich and poor today have more access to information and online services than ever before. In so many ways, this has saved lives, changed governance, and led to real improvements in day-to-day life. While there are still people on the planet who don't yet have Internet access, the global deployment has gone faster than most people imagined 20 years ago. And now an ambitious satellite company called OneWeb is making progress toward its goal of connecting every school on Earth by 2022 and, incredibly, every community by 2027.

Speed Bumps on the Road to the Future

You may have noticed that I am taking a much harder line on this subject than is usual, even for technology optimists. Normally somebody like me expresses some optimism, but then allows how the critics have some "valid points" and "raise important issues." That's code for "I don't believe a word of it, but I don't feel confident enough to challenge it."

The thing is, those peddling the innovation menace haven't actually made good points. All they offer is the same tired search for a downside just for the sake of being contrarian. I think it's important to stand up and say that.

In fact, I think my fellow technologists have some culpability here. Everybody loves a bold vision for the future. They love it so much that they generally don't remember to check later whether the bold vision actually worked out. Silicon Valley has figured out that sincere predictions about future (i.e., those most likely to be true) always lose out in the short run to the crazy, bold (but almost certainly false) prediction. And in the long run, nobody cares. The tech world is full of people who have spouted utter nonsense for decades without anybody calling them on it.

It's important to keep in mind the self-interest involved here. Every tech visionary and CEO likes to brag about how their technology will "disrupt" a huge industry and destroy the competition. It's a tried-and-true way to boost your stock price. But while they are overhyping their own technologies—and how that will let their companies dominate their markets—they tend to paint the picture that technology marches rapidly though the rest of life.

If only it were so. I am a diehard technology optimist, but over decades in this business, I have come to realize that technological change is always slowed by limitations of various kinds.

One of the laws of technological change is that the bigger the ultimate impact on society, the longer it usually takes to happen. This is particularly true when the change touches on longstanding, deeply rooted aspects of society or in old industries where there are often lots of physical constraints. But the longer the change takes, the more time society has for adaptation.

The only fast transitions occur when technology is expanding into a vacuum—when it's building or competing in a new market unencumbered by big physical limitations. That's why a new tech company like Google or Facebook can grow so fast. That's also why a tech company that has lost its groove can fall quickly, like MySpace or Yahoo. They were recent entrants themselves. But the faster they grew, the harder they could fall. Yahoo users and advertisers could switch to Facebook, Google or others with no capital cost, no retraining, no new supply chain. You just click.

Contrast that to the automobile industry. Two different technologies claim to be disrupting this industry—electric vehicles and self-driving technology. They have turned the normally staid auto industry into a hotbed of competition.

They have also generated innumerable reports that innovation in automobiles is going to destroy jobs—some breathless headlines in recent years claim that self-driving cars will put up to a third of professional drivers out of work. That's an amazing claim. It would be far higher disruption, for example, than the Great Depression caused to the industry in the 1930s.

But a quick look at the numbers shows how truly implausible these projections are.

The average age of cars on the road in the United States is just under 12 years. Worldwide, the average passenger vehicle is even older. And because long-haul trucks cost so much, they tend to stay on the road even longer than cars do.

There are more than 1.2 billion cars and commercial vehicles worldwide. It would take at least $20 trillion of capital investment to replace all those vehicles.

Sorry, but you're not going to change the composition of that fleet overnight. The flip side of an average age of 12-plus years is that it the natural turnover rate for the fleet as a whole is at least 25 years. The production capacity of the automotive industry is sized to match that rate.

And actually, because the market for cars and trucks in China is growing so fast, a lot of the production capacity is needed just to satisfy new demand, rather than to replace old cars. The number of cars on the road is expected to top 2 billion by 2035.

At present, the self-driving car market has zero billion dollars in sales to end users. There's a pretty good market for self-driving car *companies*, if you have a startup to sell to one of the established players. But at the moment, there is no market for the cars themselves.

I have ridden in one of the test models out there. They do well in many traffic situations, but others completely baffle them. Earlier this week I was in Naples, and the man driving me gestured to the crazy rush-hour traffic careening around us and said, "Cars that drive themselves? Let's see them handle Italy!"

Those prototypes still cost way too much to be competitive on the market. Some recent estimates put the cost premium over a regular car at around $250,000 today. As a total technological optimist, I believe that clever people will eventually solve both the cost and performance issues. I look forward to the day I will own a self-driving car. But again: at the moment, sales for this technology is still stuck at zero.

Right off the bat, that should tell us something. It is characteristic of the innovation menace that it looms largest for things *that are not yet real*.

As long as self-driving cars remain a thing of the future, you can say whatever you want about them. You can imagine that they will be super-efficient, super-safe, highly cost-effective, and able to surmount any driving challenge. Once they finally come to market, we will find out their true strengths—but also their limitations, like higher cost, that undermine the assumptions implicit in sweeping statements.

But let's suspend disbelief for a moment and imagine that tomorrow all new cars, on every car lot everywhere, were all magically self-driving. They still wouldn't be the majority on the road until 2030 or so, and they wouldn't completely replace human-driven cars until something like 2043.

The same limitations apply to subsets, like long-haul trucks. Even though big trucks are only a fraction of the overall vehicle market, the existing factories are sized to fit an average vehicle age of more than a decade. Giant new factories could be built, but—as Tesla has discovered—that takes longer than you might expect, too.

To shift a giant industry like automotive into a new technology, access to capital can throw up yet another limitation. Investors in Tesla are betting that something big will happen there: the company's market cap is about $550,000 per car sold in 2017. That's almost 100 times the same ratio for General Motors, where it's $5,700 per car. Tesla can get this multiple from investors because it is growing on a tiny base—it currently has a mere 0.3% or so of the U.S. car and light truck market. But as Tesla's sales grow, its access to capital must taper off. Certainly, the industry as a whole can't get access to capital at that rate.

At a cost of $250,000 extra per car, the self-driving varieties would be about 10 times more expensive than the average car on the road, and two to three times the cost of a long-haul truck. Where will the economic justification come from?

Presumably commercial owners could save the costs of driver wages. Professional drivers in the U.S. have a median income of about $40,000. A self-driving vehicle operated 24 hours a day could theoretically displace three human drivers, which means a self-driving car that cost $250,000 extra would pay

for itself in about two years. For a taxi, truck, or other commercial vehicle, that seems like a reasonable investment. But usage won't actually be 24 hours a day, seven days a week. And lower the vehicle usage, the harder it becomes to make a financial case for it. Certainly that $250,000 premium will eventually drop, but no one can really say how fast it will come down.

In a sense, that doesn't really matter. Even if self-driving catches on rapidly in some niches, it seems likely to remain a relatively rare phenomenon on the road for a long while, giving human drivers plenty of time to find new jobs. But if that's wrong, and the self-driving transition happens with surprising speed, that means that trillions of extra dollars will be flowing into the car and truck industry to accelerate the pace of transition and to pay for the extra cost. That flood of money will go into new factories and other industry realignment around the new technology.

Either way, human drivers still have years—probably a decade—to switch jobs. And when they do, they will be switching into an economy that is enjoying a rapid surplus caused by hundreds of billions of dollars a year in extra capital investment. In that kind of an economy, it generally isn't hard to find a job.

But, critics say, what *exactly* are those displaced drivers going to do?

The truth is, I don't know. And though it might sound strange, I think that is the crux of the issue.

The Folly of Futurism

Innovation menacers tend to fixate on one specific harm that they want to isolate. Their primary failure is one of imagination. They aren't able to see exactly how the economy will adapt—no one can—so they simply pretend that it won't. As a result, they end up extrapolating only the negative possibilities that suit their purpose, while ignoring the positive possibilities that, history suggests, are actually much more likely.

At last year's TED conference in Vancouver, I got into a discussion with an MIT professor who was espousing variations on the innovation menace. After he listened to my argument that technological adoption has never once caused widespread economic hardship, he finally blurted out, "But you can't prove it will always turn out well!"

Perhaps not, but so far in the long history of mass technology adaption *it always has*.

To claim that a rapid, willing adoption of technology must have a down side is to bet against everything in human nature. The requirement to prove up front what will happen is just silly. Nobody is smart enough to predict all of the things that will happen when technology causes a dramatic change to how we live and work. Those outcomes depend on the specifics of how society chooses to adapt the technology to our lives—and to adapt our lives to the technology. They also depend on what great new ideas people come up with once they are inspired by the pace of technological change.

The stockingers of Nottingham didn't know what jobs they would have in the industrial revolution that was about to explode in their region. But they got those jobs all the same. The stock brokers of the 1970s could not have guessed that commissions would drop when trading went digital, but that total volume would more than make up for it.

After so many consistent examples of technological progress, however, I have to grade MIT professors a bit more harshly. We ought to know better by now.

Thinking More Intelligently About AI

With self-driving cars, manufacturing is a particularly weak link. That's not the case for artificial intelligence. AI is mostly software, so it can be deployed rapidly on computer servers in the cloud. But even though production may easy for AI, the technology faces other important constraints.

Some people have been touting AI as the cure to all of our problems. Others are holding it up as the ultimate example of the innovation menace. And still others seem to be making both claims simultaneously. AI, we are told, will steal our jobs, rob us of what it means to be human, and even—in the most extreme scenarios—eliminate the human race.

Both the hype and the fear are based on the same underlying reasoning: if computers can be as smart as we humans are, then they will displace us from tasks that today are ours alone.

As a technological optimist, I am convinced that the human brain is a biological machine and that eventually scientists will figure out how that machine works. That knowledge should allow us to eventually build machines that we will recognize as genuinely intelligent.

Indeed, computer science has already found approaches that, at least for some subset of interesting problems, allow computers to be as smart or smarter than we are by using techniques that work very differently than human brain does. Put those two together, and I believe that, yes, one day we will create machines that are smarter than us.

But taken broadly across human abilities, that day certainly isn't today. Nor next Thursday. Nor even in the foreseeable future.

The history of AI research is that a small breakthrough is made and causes lots of excitement and hype, which in turn leads to disillusionment when the hype is not realized. This has happened in several cycles since the late 1950s, when AI was first conceived as an application for digital computers. Each cycle has made some progress, but it turns out that human intelligence is not just one trick or technique—it is many.

It's a useful analogy to look at a specific area where computers are already vastly better than humans: mathematics.

For a long time, complex mathematical computations were done by computers—not machines, but people who had the job title "computer." That profession actually dates back four centuries—the first known mention is from 1613. Human computers made, for example, navigational tables, initially using little more than pencil and paper. These jobs were still around in the 1960s, when the workers had slide rules and mechanical calculators along with plenty of pencils and paper. The 2016 feature film "Hidden Figures" told the story of Katherine Johnson and other African-American women who worked in a computer team at NASA doing calculations for the early space missions.

Of course by now even a smart watch has the capability to add and subtract numbers billions of times faster than a human being could. You might assume that NASA doesn't need any more human computers because the digital kind is so superior.

But you'd be wrong. NASA employs far more programmers, mathematicians, and computational physicists than the human computers it employed in the 1960s. Rather than being replaced, the jobs for humans *multiplied*.

You see, mathematical problems come in many degrees of difficulty. Simply adding up a column of numbers is hard for a human, but trivial for a computer. Anybody—even a middle school child—can enter numbers into Excel and hit return.

But not all math is simple arithmetic. Some math problems are difficult for even the most advanced computers; they require enormous amounts of computing time. In order to do problems like that, you need a lot of humans who can think up ways to program computers to do those calculations as efficiently as possible.

It's a classic example of something else that the innovation menacers routinely forget: in almost all areas where we have deployed computers, the more capable the computers have become, the more demand there has been for them. Meeting that demand requires lots of human effort and jobs.

Normally in economics, a big increase in the supply of a commodity lowers prices because demand is fixed. Computer power—especially for mathematics—has greatly increased in supply, but that counter-intuitively stimulates demand for more. The reason is that each new increase in computer power—whether it's greater raw computational ability or clever new software algorithms that do useful things with it—makes another class of problems accessible to computer solution. All with human help, of course.

We know from theoretical work that some mathematical problems are so complicated that they will *always* be challenging or even impossible for computers to solve. This means that there is (currently at least) no end date in sight for employing humans to push forward the boundary of computationally hard problems.

So even though computers are already enormously "smarter" than we are for arithmetic, and nobody doing scientific, engineering, or financial calculations uses pencil and paper anymore, there are many other problems where computers struggle and may continue to need human help for some time.

This tells us something important about AI. Like mathematics, intelligence is not just one simple kind of problem, such as pattern recognition. It's a huge constellation of tasks of widely differing complexity. As AI gets more capable, more of these tasks will be done better by computers, while others are likely to remain elusive for some time.

In particular, if you ask: 'What are the biggest successes of AI today in a practical sense?' the answer is quite underwhelming.

Neural networks and machine learning are by now fairly old technologies, and they are incorporated in many software systems but typically only in very minor roles. Here's one example: a few months ago, in preparation for the next cookbook I am writing, I ordered a bunch of books on pizza from the Italian version of Amazon.com.

American Express promptly shut off my credit card.

Credit-card fraud detection was one of the early commercial successes for neural networks. In my case, the algorithm came to the wrong conclusion, but overall this application has added a valuable feature to banking. One would be hard pressed to find any humans that were thrown out of work by fraud detection—except criminals. It's not like AMEX had banks of people scanning credit-card transactions in real time. In fact, neural networks for fraud detection almost certainly caused a lot of humans to be *hired* because it was a nice human who took my call to get my card to work again.

Two of the most widely reported AI announcements have come from Google. A few years ago, the Google Brain project created a massive neural network that analyzed a week's worth YouTube videos and taught itself to identify cat videos. More recently, DeepMind (a different group within Google) created a program that learned to play the strategy game Go well enough to defeat the world's best player.

These are huge technical accomplishments over the prior state of the art. It used to be impossible to use a computer to distinguish cats from other animals, and although computers could beat humans at chess, Go proved much harder to approach computationally.

Many commentators have touted these Google announcements as reasons for concern about an innovation menace. Yet it's hard to see who would ever be thrown out of work by either one.

I have no doubt that the underlying technologies will find important applications. My company, for example, has developed a machine learning system that can recognize early warning signs of cervical

cancer better than any human doctor can. It offers the chance to save the lives of millions of women, particularly in the developing world. We were able to do this using machine learning that is far less sophisticated than Google Brain. It turns out to be easier to find cancer than cats.

We have also created AI-enabled microscopes that can count malaria-infected cells in a blood drop and are working on an AI-powered test for tuberculosis. The possibilities for applying AI to improve medicine are pretty remarkable.

DeepMind's champion-level Go-playing AI is also very impressive — but it's hardly evidence that AI can now do almost anything, as many media stories about the achievement absurdly suggested. Think about it: Is the human that the computer beat the smartest, richest, most powerful, or most feared person on earth? No, he is 34-year-old Lee Sedol of Korea. A phenomenally good human Go player is not a master of all intellectual pursuits, and current experience suggests that is true for computers, too.

So far, computers have had the easiest time showing superhuman performance at tasks that humans struggle for years to learn and master, like mathematics or Go. The more challenging tasks for computers have always been the ones that are so trivial for us that we do not think about them—those that require common sense or dealing with the messy ambiguities of the real world—like visually identifying cats or other familiar things.

At the moment AI is in one of its cyclic high-hype phases and many of its envisioned applications are not yet real. That makes it perfect fodder for invoking the innovation menace. But all of our experience with AI and with computers in other areas like mathematics suggests that the transition will not happen overnight and will not generate massive unemployment.

While the threat of AI to employment is overdrawn, even that is too boring for some prognosticators. At a recent conference, one technology executive suggested that AI is becoming so powerful that it threatens humanity. Another has taken to Twitter to warn repeatedly that AI poses "vastly more risk" than nuclear weapons.

This last idea is so absurd, it's hard to know where to start. In one corner we have the nukes. Although only two nuclear weapons have ever been detonated in war, they killed somewhere between 130,000 and 250,000 people. The accounts of the survivors are horrific. And afterward, during the Cold War, the power of these terrible weapons was increased by more than a thousand-fold from those that destroyed Hiroshima and Nagasaki.

In the other corner we have AIs that detect cat videos, play Go, and deny my credit card.

To suggest that these are comparable is insulting to both our intelligence and the very real dangers that nuclear weapons still pose.

To invoke the innovation menace in this way, you have to make four completely unfounded, counterfactual assumptions. First, you must assume that AI will start working massively better than it does today before we learn how to control it.

Next, you have to assume that the AI can somehow develop motivations totally different from those it was programmed to have.

Third, you must assume the AI will become intent on killing or enslaving us all, even though that would seem hugely counterproductive. Why, pray tell, should supersmart beings that we created bear us a grudge? It's not like they compete with us for food or a place to live.

And lastly, you have to assume the AI could acquire the means to pull off its dastardly plan, despite the best efforts of its creators to stop it.

Dreaming up god-like powers paired with malevolent intent is fun if you are creating a superhero movie or a Games of Thrones plot line. You can always postulate a demon like the Norse god Loki or the Night King's army of white walkers that poses a bigger threat to humanity than nuclear weapons.

The scenario is of course insane, but it has an internal logic that follows from the starting assumption of a supernatural god bent on destroying mankind. Likewise, if you assume that computers will be vastly smarter and incredibly more powerful that we are, and at the same time that they will have a malevolent outlook on humanity, then yes, logically that would be a bad combination. But please, let's not confuse it with reality.

At the moment, computers are smarter than us in a few very limited ways, like at mathematics and Go playing. They barely equal us in spotting cat videos. They will undoubtedly become more capable in the coming decades. But we will have ample time in those decades to figure out how to put them to good use while avoiding existential catastrophes.

I was discussing this with a friend of mine, one of the great geniuses of our time, who said, "but really, Nathan, wouldn't you be at least a little afraid of a super-intelligence?" I said, "look, that's how the rest of us think about you!"

And my strategy for survival with both is the same. I am really hoping that they like me.