

**THE**  
**DIGITALIST**  
**PAPERS**

**ESSAY**

ECONOMIC POSSIBILITIES FOR  
ARTIFICIAL INTELLIGENCE

# ECONOMIC POSSIBILITIES FOR ARTIFICIAL INTELLIGENCE

GABRIEL UNGER

Transformative AI has revolutionary potential. But Gabriel Ungers argues that whether we can fulfill it, and turn around decades of dissatisfaction, depends on answering three fundamental questions. Can we articulate a compelling shared vision of a future with TAI? Can we settle on a theory of economic growth and AI that allows us to think seriously and carefully about the effect of the latter? And can we redesign education and strengthen social connections in the age of TAI? These are new intellectual opportunities for the economics profession.



## I. Introduction

Over three million years ago, our ancestors first used stones to make tools out of other stones. Basic control over fire took another two million years. Another half million years to go from foraging to agriculture.

From 10,000 BC to about 1750, what little technological progress took place was haphazard and discontinuous. The life of a British serf in 1750, sowing his wheat with a plow in the fall, harvesting it with a scythe in the spring, would have been immediately recognizable to a Mesopotamian farmer from the earlier period, right down to the tools and the crop. Both would have had the same hard standard of living.

The First Industrial Revolution introduced not just the steam engine or the power loom, but economic growth itself, and a new regime of continuous progress. Since then, we as a species have made major advances in harnessing energy, processing material, and mastering physical transport. But throughout this period, the human mind has been a sovereign entity. These advances all came from the mind. They stood far apart from and outside of the mind. Even the IT Revolution of the past 50 years has largely respected this boundary. We created new ways of sharing and processing information. But it was still mostly the same content that one human might imagine and then share with another, only now via the internet instead of the phone or the mailman.

The significance of AI is that, for the first time in the history of the human species, this sovereignty is collapsing. We are creating a technology that seems to share many of our own cognitive abilities. It can pass school exams, talk to customers, design products, diagnose illness, provide personal advice, extend friendship, and write poetry. It can power robots that interact with the physical world. It can pursue scientific research, even on AI itself, in the improvement of its own capabilities. AI can do all of this right now, still in the first hours of its



infancy. The AI that can do all of this today is the weakest, least capable version of itself it will ever be.

Silicon Valley is now acutely aware of the revolutionary economic and social potential of this technology. Most of the rest of the world still does not grasp the scale of what is about to unfold.

It comes at a particularly significant moment. America has spent most of the past 50 years in a period of stagnant productivity growth.<sup>1</sup> Young Americans today no longer expect to enjoy a higher standard of living than their parents had.<sup>2</sup> They enter the workforce with more student debt, and a much longer and more uncertain road to marriage, home ownership, and economic security. Growth in real wages over the past few decades has been slow, by historical standards, particularly for those outside of the very top of the income distribution. Income inequality has risen. Industrial concentration has risen. Our trust in institutions has declined, as has our trust in each other.<sup>3</sup> There is both the economic fact of objectively diminished prospects, and the social fact of rising pessimism in opinion polls about our present and our future.

*The real opportunity here is not just about technology. It is about how we might be able to use a profoundly transformative technology to help rescue ourselves from our decades of increasing dissatisfaction and diminished expectations, in the service of a more promising economic and social life.*

All of this is to say: The real opportunity here is not just about technology. It is about how we might be able to use a profoundly transformative technology to help rescue ourselves from our decades of increasing dissatisfaction and diminished expectations, in the service of a more promising economic and social life.

But the translation of AI into a broader experience of strong and sustained economic growth, or into a better society, is not something that will happen quickly or on its own. The mechanics of economic growth are stubborn. And I argue below that successfully moving

from technology to economic prosperity is not just about technical benchmarks. It depends critically on answering three broad, related questions.

First, what is the vision of an AI future that ordinary people will find exciting and compelling, and not terrifying or hideous? Second, what is the theory of economic growth that helps us best realize the full economic potential of AI, in the service of this vision? Third, what must be done to ensure that as AI enables this economic progress, it also makes our experience of learning and connection stronger, not weaker?

In these early days of the AI revolution, there is a strange asymmetry between how large the stakes for humanity are, and how small and

*In these early days of the AI revolution, there is a strange asymmetry between how large the stakes for humanity are, and how small and confined the discussion of these issues still is, relative to their importance.*

confined the discussion of these issues still is, relative to their importance. Silicon Valley is aware of many of these questions (e.g., What do we ultimately want out of life, and how can AI help? What is our theory of economic growth and AI?) but not equipped to provide all the answers. Those who might be better equipped are mostly not asking the questions. We want to more directly engage the latter. There are plenty of AI optimists and AI pessimists. But optimism and pessimism are both forms of the same fatalism: They say the

future has already been decided. Hope is a categorically different kind of position. It says the future is open. Facing this openness, economists and other social scientists should turn more of their attention to the unresolved questions most important to a better future from AI.

## II. The Vision of Society

Suppose the AI Revolution is maximally successful along its purely technological dimensions. What kind of broader society and economy do we hope for in this scenario?

At any moment, there will always be some balance in society of people who look to the future with more optimism or more pessimism. For every Luddite, a Whig. But what is particularly distinctive right now is not just that so many ordinary people feel pessimistic about the future of the world broadly, and about the future of technology specifically. It is that the *same* vision of the AI future that the tech industry thinks of as optimistic and utopian inadvertently comes off to so many other people as pessimistic and dystopian. This in turn has enabled a culture where most people now think of the future of technology with more dread than excitement. Every sci-fi movie is now a horror movie.

*What is particularly distinctive right now is not just that so many ordinary people feel pessimistic about the future of the world broadly, and about the future of technology specifically.*

This was not always our cultural attitude toward the technology of the future. From the First Industrial Revolution, up to as late as the early 1970s, the dominant attitude of the West to technological progress was passionate optimism. Prince Albert opened the planning for the 1851 World's Fair in London by announcing, "We are living at a period of most wonderful transition."<sup>4</sup>

This spirit was stronger than ever over a century later, as millions of Americans rapturously watched the Apollo missions on their new color televisions. In 1968, Pan Am announced a "Moon Flights" waiting list: For \$1, you could be at the top of the waiting list for the coming commercial space flights. Ninety thousand Americans signed up. Disney introduced a carnival ride at the 1965 New York World's Fair called the "Carousel of Progress" that showed the American family through successive eras of technological progress, into the future, all advertised with the motto, "There's a Great Big Beautiful Tomorrow." TV shows like *The Jetsons* warmly advertised a future of life in space and helpful robot servants. Companies like GE and Whirlpool introduced things like microwaves and washing machines in dramatic store stagings of the "Kitchen

of Tomorrow,” to show housewives how they could be freed from drudgery. The broad message was that we could count on the march of technological progress. It would bring abundance and empower ordinary people. It was a good servant, not a sadistic master. Curiously, the height of this optimism coincided with the nuclear arms race, almost certainly the greatest moments of existential risk humanity has ever faced. We were undeterred.

How many ordinary Americans today feel like they are still riding the carousel of progress? Who looks to the future of technology with hope? Not enough of us. Opinion polls now very consistently show only a tiny share of the American public thinks AI will make the world a better place. Will AI “do more good than harm?” Only 13% of people say yes.<sup>5</sup> Will AI have a positive impact over the next 20 years? Only 17% of people say yes.<sup>6</sup> A far larger share in both believe the impact will be negative.

*Who looks to the future of technology with hope? Not enough of us.*

But consider what is on offer.

The Silicon Valley vision for the AI future, from the perspective of an unsympathetic critic, might be described in the following terms: You won’t really have a job. You’ll stay at home, maybe living off some UBI welfare payments. At home, unemployed, you will be free to interact with your 12 AI friends and your AI-backed robots, or put on your VR goggles and glasses, where you can consume your AI slop (AI-created music and movies, all converging toward mediocre short-form video content). Your children—if you have any—will use AI to write their book reports—and then some teacher’s AI will grade them. More time for your children to also scroll through short-form video content! (You probably will not have children).

It is not strange to have critics. But what is strange is that with some changes of semantics and word choice, something like the above is difficult to distinguish from the same vision that many technologists are openly promoting for the future—even the people who most consider



themselves AI optimists. Policy papers about UBI, dreamy think pieces about the meaning of life after work, executives openly pushing the idea of AI friends. They do not seem to understand that their utopia is other people's dystopia. It fails on the two dimensions most fundamental to what we all want out of life: (1) a meaningful vocation, and (2) intense personal connection with other human beings.

Where is the vision that takes both of these imperatives seriously? It seems shockingly absent. We need to be able to articulate a future for the AI Revolution in which ordinary people have both exciting jobs and real relationships with each other, not unemployment and solitary lives at home, drugged out on dopamine. Meeting this challenge matters on some very tangible levels. First, articulating a better vision of the future will help companies and entrepreneurs build the right products, in the right direction. Second, particularly in a democracy, if Silicon Valley is unable to offer the rest of the country a more compelling vision of the future, it is placing itself in greater danger of regulatory control and punishment than it seems to realize.

*We need to be able to articulate a future for the AI Revolution in which ordinary people have both exciting jobs and real relationships with each other, not unemployment and solitary lives at home, drugged out on dopamine.*

To be clear, a more compelling vision is entirely possible. Deep within the American professional elite today, people have the experience of focusing their daily energies on difficult problems that require creative solutions. A neurosurgeon is presented with a complicated case. A lawyer confronts a new fact pattern in a high-stakes suit. An academic pursues an ambitious research program. In a sense, these people are all knowledge workers engaged in local acts of innovation. Outside of this elite, the dominant experience of work is routine, not innovative, and disempowering. Millions of Americans today are slotted into what they sometimes darkly call “fake email jobs,” doing tedious, repetitive things that can and should be automated.

The promise of AI for the future of work here should be to take the elite's experience of creative problem-solving and transforming some

*Dario Amodei said that AI could give us “a nation of geniuses in a data center.” This inverts the goal: What we really want is to become closer to a nation of actual geniuses.*

small piece of the world around them through their work, and give it to everyone else. AI can do this by automating away all of the drudgery, and democratizing a broad set of technical capabilities. From this perspective, an increasing share of ordinary workers would effectively start to become knowledge workers. We would almost start to become a nation of scientists, in different ways, each of us taking on new challenges at work

with our AI assistants. These collaborative teams of AI assistants or agents and human knowledge workers could then continue to advance as the AI improved. Dario Amodei said that AI could give us “a nation of geniuses in a data center.”<sup>7</sup> This inverts the goal: What we really want is to become closer to a nation of actual geniuses.

Related but distinct features of the work today in the professional elite are autonomy and agency. An executive has a vision that a team of people is then commanded to enact. At the bottom of the labor force, the worker takes orders but never gives them. Creative work is a special case of this, but even when not particularly creative, most people have a broader interest in work that empowers them, in the very concrete sense of having power to coordinate resources and other people in the service of a project. A construction worker might prefer to run his own small independent business to working in someone else's much larger one, just to enjoy more of this. According to surveys, this is in fact by far the number one reason people start and run small businesses: not even for financial gain, but for the experience of autonomy in their work life. AI will be able to provide far more of these people with effective agents—AI models, AI-backed robots—at their command, to more broadly enact their will. It presents the opportunity for the small businesses of the

future, for the startups of the future, and for the most low-level workers at larger businesses, to fundamentally broaden their experience of agency.

Alongside the promise of more transformative work is the more simple promise of a return to higher economic growth and greater prosperity. At a basic level, economic growth is the promise of a better future: of less poverty, better health, higher standards of living, bigger and freer lives. Using AI to return to a world of higher economic growth would mean less zero-sum social thinking and pessimism in young Americans today. It would help our long-term fiscal problems, particularly our ability to pay for our increasingly concerning federal debt. This higher economic growth can exist without the transformative kind of work discussed above, but the two are not unconnected: An economy where everyone is effectively a knowledge worker is probably both the economy with the best chance at high aggregate productivity growth, and also the economy with a comparatively more equal distribution of income.

It is striking how often, in the history of economic thought, major thinkers have gotten the question of vocation wrong, so much so that we might forgive the technologists for now repeating their mistake. Both Keynes and Marx prophesized a world in which automation and growth would rescue us from all labor, and we could sit around painting and fishing. Today most discussion of these earlier works focuses on the failed empirical predictions. Omitted is any discussion of how these thinkers were also mistakenly *morally* in seeing endless leisure as the correct goal. Perhaps they confused the demeaning jobs most people were stuck in as they wrote, for the only kind of labor that would ever be possible for most people in any alternative economy. But most ordinary people today instinctively understand and share in one of the most important moral insights that Protestants like Luther and Calvin gave to the rest of the world: that almost all of us—from even the lowest ranks

of society—want to feel like we have been called in this world to do something of consequence with our lives. That a sacred calling is not just reserved for a tiny priesthood, but should extend to everyone.

(And in perhaps a more pagan vein: Even people who are otherwise apostates to our inherited ideas about vocation often still have some intense interest in “positional goods,” of status, elite jobs, elite colleges, and so on. Interest in these positional goods is probably only increasing. Note that 60 years ago, economists would discuss the “backward-bending labor supply curve”: At the highest income levels, labor supply would decrease, as executives traded away some income for more leisure. You no longer hear the term today. In the intensely competitive pseudo-meritocracy of American life, the top of the income distribution now takes less vacation time than everyone else, not more, as it would have been before. All of this simply extends the set of people and attitudes for which permanent leisure is the opposite of a shared goal.)

In our personal lives, AI could do for us what GE understood the first washing machines and dishwashers could do for housewives: relieve

*In our personal lives, AI could do for us what GE understood the first washing machines and dishwashers could do for housewives: relieve us of tedious burdens on our time.*

us of tedious burdens on our time. Somewhere tonight a parent with two children will spend several hours navigating a health insurance claim, or planning a grocery budget, or researching a medical condition, instead of spending time playing with the children. Another will spend hours trying to figure out why the air conditioner is broken, instead of getting some sleep or reading a book. According to surveys on time

use, men spend an average of about two hours a day on this kind of uncompensated domestic labor; women spend closer to three.<sup>8</sup> And the average American will spend almost an hour every workday driving to and from their job, focused at the wheel.<sup>9</sup>

These sorts of mundane cognitive tasks pollute our lives and rob us of the limited time we have on this earth and with each other. They are

all routine and absolutely amenable to AI (all the more so as AI-backed robots are increasingly integrated into the physical world, from self-driving cars to vacuum cleaners). One of the better uses of AI-based automation is to have it do all of this for us. And here we can be even more ambitious than GE was. A billionaire today will have an extensive home staff: a small army of people who cook, clean, and do any number of other domestic tasks, some planned, some spontaneous; some essential, some more frivolous. AI agents will increasingly be able to generalize this experience to everyone. At both home and at work, ordinary people will be able to effectively enjoy a large staff of extremely competent and sophisticated agents working on their behalf. It will be the first time in history that this will not be an experience reserved for a tiny segment of a society's aristocracy.

But to the extent that AI otherwise intersects with our social lives, we want an AI future that somehow expands and deepens our experience of connection to each other, not one that shows up to replace the people in our lives with AI. I present the further working out of such a vision as a challenge of the highest importance, to both Silicon Valley and the rest of the country. We will know someone has succeeded when the popular attitude toward the future once again more closely resembles the joyful optimism of American schoolchildren in the 1960s.

*To the extent that AI otherwise intersects with our social lives, we want an AI future that somehow expands and deepens our experience of connection to each other, not one that shows up to replace the people in our lives with AI.*

### III. The Theory of Economic Growth and AI

#### The structural context

**T**he economic promise of the AI Revolution is that it can fundamentally change our experience of cognitive labor, in the way that the Industrial Revolution changed our experience of



physical labor. Over a sufficiently long horizon, this might prove to be an even more radical change to our individual experiences of work, and to the economy.

Technologists, acutely aware of how revolutionary this long-term potential is, have started to make dramatic short-term economic predictions in a way that often exasperates economists. A common tendency within Silicon Valley is to go from very narrow discussions of purely technical progress (benchmarks, scaling laws, etc.) to very dramatic, high-level economic claims (about GDP growth or unemployment), and skip the 20 steps that connect them. But economic

*To tell a convincing story about a technology leading to economic growth, you need to have some account of which sectors and firms are going to use the technology, and in which ways.*

growth is not just about replacing gadgets with better gadgets. The reorganization of production requires all sorts of structural changes. To tell a convincing story about a technology leading to economic growth, you need to have some account of which sectors and firms are going to use the technology, and in which ways. Which current human tasks might be a substitute for AI, and which might be a complement? Will the sectors

most effected by AI simply decline as a share of GDP, if demand is inelastic, and economic activity simply be reallocated elsewhere?

At the same time, economists, more equipped to think through these kinds of issues, should take the long-term transformative potential of AI more seriously. Part of the price of saying something compelling about AI and economic growth is developing a more serious structural view of the economy, and how AI then enters into it. It would not be enough to have some lab of scientists off to the side, powered by AI, doing amazing things, and self-recursively improving. It has to somehow translate into a broad mass of sectors and firms using AI to achieve higher productivity growth.

Over the past several decades, data reveals a rising dispersion of productivity and income in the US economy. The gap between the wealthiest states and cities in the US versus the poorest states and cities is now growing. Income per capita in the wealthiest US states like Massachusetts or Connecticut is now about \$100,000; in the poorest US states like Mississippi or West Virginia, it is about \$50,000. In earlier periods of US history, the poorer regions were catching up to the richer ones. Since the 1980s, the tendency has been reversed: The richest states and cities are now pulling further away in income and GDP per capita.<sup>10</sup> But a similar pattern exists now throughout the US economy along many other basic economic dimensions, such as the rising productivity gap between the leading firms and other firms, or in income between the richest workers and the rest of the labor market.

The US economy has been characterized by a fundamental and increasing kind of dualism. There is a small number of fancy firms with fancy workers in fancy cities (New York City, San Francisco, Boston, etc.), and a long unproductive tail of everyone else. In one part of this dual economy, a professional elite has access to advanced technology, high salaries, and very capital-intensive forms of production. The other part does not. Much of the IT Revolution of the past 50 years disproportionately helped this elite become even more productive. It did much less to help everyone else. But it is difficult to sustain a high level of aggregate productivity growth with only a small share of exciting firms and workers. The whole economy needs to be involved. In the post-war period, productivity growth was more uniformly distributed across firms and workers throughout the economy. Today it is not.

*It is difficult to sustain a high level of aggregate productivity growth with only a small share of exciting firms and workers. The whole economy needs to be involved.*

From this perspective, the decisive structural question about AI and productivity growth may be: Will AI disproportionately help only elite firms and elite workers, or help the long tail? But there is reason for hope: AI can democratize capabilities and empower ordinary workers and firms in very dramatic ways. Maybe you're a student in Kansas and now have access to a truly elite education. Maybe you're a worker at a small firm in Kansas and suddenly have access to a virtual army of very powerful coworkers. Maybe you're an entrepreneur in Kansas and want to start some amazing new firm, and AI agents can act as your first 10 employees.

An overriding goal should be that AI helps us overcome the extreme dualism of the economy in its current form. The only form of aggregate productivity growth that is both high and sustainable is one that is broadly based and involves the entire economy. Economists, AI companies, and policymakers at both national and local levels should consider what can be done to make sure that the commercial benefits of AI are not disproportionately confined to a small part of the economy.

*We need stronger and more confident views about the mechanics of growth as it plays out across different kinds of sectors and different kinds of firms. Without them, it is very difficult to have deep insight into economic growth, or to develop any kind of usable pro-growth policy agenda.*

Seventy years after Solow discovered that empirically, most of US economic growth comes from TFP (total factor productivity) growth,<sup>11</sup> not labor or capital accumulation, discussions of economic growth too often still have technological progress speeding up or slowing down as their bottom line. The theory of economic growth still requires more causal and structural content. We need stronger and

more confident views about the mechanics of growth as it plays out across different kinds of sectors and different kinds of firms. Without them, it is very difficult to have deep insight into economic growth, or to develop any kind of usable pro-growth policy agenda.

## Labor market policy

One of the first tasks for a pro-growth policy agenda is clearly the labor market. The near-term unemployment threatened by AI seems to us to be substantially overstated. AI will augment some tasks and automate others; the task composition of specific jobs will change sooner than specific jobs are eliminated. The workers most “at risk” of the latter will be, in an sense, particularly well-suited to weather AI disruption. The young McKinsey consultant in New York makes an easier transition to a new industry than the assembly-line worker in Gary, Indiana. The history of the US labor market is generally one of resiliently absorbing massive supply increases, as it did when women entered the workforce after WWII, without causing unemployment for anyone else. And to the extent that any jobs are in fact at risk right now, there is always a kind of ascertainment bias around these technological transitions: We always know the names of the old jobs being destroyed far more clearly than we do the names of the new jobs being created.

*To the extent that any jobs are in fact at risk right now, there is always a kind of ascertainment bias around these technological transitions: We always know the names of the old jobs being destroyed far more clearly than we do the names of the new jobs being created.*

Preliminary research now suggests AI may already be impacting the US labor market, along precisely these lines. A number of papers now suggest that young workers in the most AI-exposed sectors (like software engineers in tech) may be experiencing higher unemployment rates now than their peers in less exposed sectors.<sup>12</sup> Even if this is true, there is still no evidence so far of AI causing substantial employment effects in the US economy at the aggregate level. If anything, the more dominant tendency of the US economy in the past several decades has been not enough job churn, which is a basic feature of a healthy and dynamic economy.

But to the extent that higher transitional unemployment from AI later becomes a more real thing, consider how little the US spends right now

on unemployment insurance (UI) as it is. UI averages a cost of something like 0.3% of GDP over the past several decades (and more like 0.1% to 0.2% of GDP in the past couple of years). It is, from a broad perspective, quite cheap; we could quadruple it, and still easily afford it.<sup>13</sup> The basic implication here is that if part of the price of technological progress were periods of transitional unemployment, there are a set of very feasible policies to accommodate this.

In a sense, this would be a move toward the Nordic model of labor market policy—“flexicurity” (flexibility for firms, security for workers). In the Nordic model, it is easy for firms to hire or fire workers, but workers then benefit from extremely generous UI and a set of other “active labor market policies” around search, training, apprenticeships, etc. The Nordic countries spend something like 1% to 2.5% of GDP on UI and associated policies, an order of magnitude more than the US does today.<sup>14</sup> (And all of this in sharp contrast not only to the current US approach, but even more to the corporatist model of continental Europe, of higher job protection but lower UI security).

The prospect of necessary labor market disruption comes at a particularly inauspicious time: The US labor market has suffered a serious decline in “dynamism” over the past 50 years. Workers moving

*You would be forgiven for thinking that dynamism went up over this period of time (the old model of someone spending 40 years working at the same factory giving way to a more dynamic economy). In fact, according to the data, the dominant tendency has been the opposite one.*

from old jobs into new jobs, new industries, new cities and states, at a sufficiently high rate, is up to a point a normal part of a healthy economy. It is very directly connected, both theoretically and empirically, to higher economic growth. It is also strongly associated with higher wages for workers (labor economists call this climbing “the job ladder”—substantial wage increases come more from switching jobs than within-job raises, especially at the start of a career).



You would be forgiven for thinking that dynamism went up over this period of time (the old model of someone spending 40 years working at the same factory giving way to a more dynamic economy). In fact, according to the data, the dominant tendency has been the opposite one. More than ever, people are persistently stuck, in the same jobs and industries or the same cities they happen to be born in. Any forces of labor market disruption will run up against whatever forces are causing all of this.<sup>15</sup> Part of successfully returning to a higher growth economy will be returning to higher dynamism.

*Part of successfully returning to a higher growth economy will be returning to higher dynamism.*

The optimal labor market policy both for individual workers and for the acceleration of the broader process of technological progress and economic growth may look quite different from the one we have today. But the promise of higher UI, and perhaps other more active labor market policies, is that we can advance both interests at once, instead of setting them against each other. All of this will be even more compelling if it is done in the service of the vision sketched above, of a movement toward an economy of less drudgery and more exciting vocations for ordinary workers. From this perspective, active labor market policy is about accelerating the search of the worker for this opportunity and accelerating the reorganization of firms toward this new economy.

## IV. Social and Cognitive Risks

The promise of an AI future of greater economic growth and more transformative work is in a sense all downstream of the future of our cognitive development. At the center of any future society in which AI has truly helped economic and social flourishing must be young people for whom AI made the experiences of education and social connection stronger, not weaker.

## Education

First, consider the formation of the human mind, as it plays out over the basic stages of formal education. Students read books. They write essays. They learn to perform simple mathematical operations. They learn to speak a second language. They cultivate a set of broad analytical capabilities that will equip them later, both as workers and as people.

There are growing anecdotal reports from teachers that the ability of young American students to do the most basic reading and writing, in a sustained and unassisted way, is already declining. LLMs present the threat of letting students outsource the basic tasks that cognitive development demands. Ominously, in the background, the Flynn effect—the documented trend of average IQ to rise each decade for the general human population—now shows signs of slowing dramatically, or even outright reversing.<sup>16</sup> More narrowly, literacy and math scores for American students are now at their lowest in several decades.<sup>17</sup>

Education can and should change with technology. I do not cling here to some Humboldtian fantasy of humanist, screen-free tutoring for everyone as the pedagogical ideal. It is probably true that

*AI seems particularly promising in potentially democratizing access to all kinds of education to students who may not otherwise have access to the best teachers and schools. But it is also probably true that AI, if misused, can have destructive consequences here, particularly along the path of least resistance.*

each generation of pedagogically useful technologies has first endured some combination of ignorance and dismissal. And AI seems particularly promising in potentially democratizing access to all kinds of education to students who may not otherwise have access to the best teachers and schools. But it is also probably true that AI, if misused, can have destructive consequences here, particularly along the path of least resistance. What should the

model of education look like now? How can it become better with AI, and not worse?

Part of the promise of AI here is obviously that it will be able to help provide broader access to the frontier of knowledge for more students. There are clearly some number of classes where some kind of AI instructor will be as good, if not better, than a human professor. AI instructors will also be able to precisely customize content and very real interaction with each student. This is a crucial difference between AI and the MOOCs (“massive open online courses”) that briefly inspired excitement 15 years ago and are now widely regarded as a disappointment. Here, as in other industries, AI allows both greater scale and greater customization. As AI instructors become increasingly sophisticated, there will probably still be room for star academics at the leading universities, but what becomes of most human instructors teaching introductory courses at most universities is a very open question.

*As AI instructors become increasingly sophisticated, there will probably still be room for star academics at the leading universities, but what becomes of most human instructors teaching introductory courses at most universities is a very open question.*

Now, more than ever, the very basic model of higher education will be open for revision. This is fundamentally a good thing. Throughout the 20<sup>th</sup> century, would-be reformers like Dewey and Piaget pleaded for more alternatives to the ancient model of the professor lecturing to passive students, an anachronism from a time before the printing press when books were literally scarce, and the lecture was simply a way to share one book with many people. The human lecture might still make sense as the live development of a series of powerful and original views to confront a student with, if delivered by a compelling interlocutor. It will make increasingly less sense as the uncharismatic regurgitation of the content of assigned readings by a human lecturer, if the alternative

is an AI instructor that is capable of greater command of the material, more interaction, and more customization to the individual student.

Educators should take this as a real opportunity to rethink the basic model. An overriding question should be: How can education better develop the creative and analytical powers of the student? AI needs to somehow be in the service of this, instead of becoming an increasing substitute for it in the minds of the student, particularly in the most important stages of cognitive development.

## Mental health

Second, consider mental health, and the experience of social connection more broadly. Rates of depression, anxiety, and other mental health problems are substantially higher for Gen Z and Millennials, compared to earlier generations. There is also the clear empirical evidence, across many different dimensions, of declining

*Smart phones, and the internet more generally, are an obvious candidate explanation for the dimensions here that bear most directly on social life: people at home alone on screens, instead of being with each other in person.*

social connection and social capital. Rates of marriage are down. Rates of childbirth are down. Rates of social events, parties, even casual intimacy, are all down. Smart phones, and the internet more generally, are an obvious candidate explanation for the dimensions here that bear most directly on social life: people at home alone on screens, instead of being with each other in person.

The economic research on the connection between the former (the rise of mental health problems) and the latter (smart phones and the internet) is still unresolved; it would be dishonest to claim we can confidently say today exactly how much one led to the other. But it would be even more dishonest to confidently say there is no connection at all, no cause for concern, no need for further study. The growing use of AI characters as personal friends for their human users (and in some cases, it must be

said frankly, as more than friends) makes these concerns even more pressing.

By the 1960s, executives in the tobacco industry began to understand that their product might be both addictive and causally linked to cancers. Depositions later revealed that many of them then privately discouraged their own children from touching cigarettes. Fifty years later, Steve Jobs was prohibiting his children from using various Apple products at home (when asked by a reporter how his kids liked the iPad—four years after its public release!—he replied: “They haven’t used it. We limit how much technology our kids use at home.”).<sup>18</sup> This is very widely and publicly reported to be the norm in Silicon Valley today, not the exception. As Mark Zuckerberg told a journalist in 2019, “I don’t generally just want my kids to be sitting in front of a TV or a computer for a long period of time.”<sup>19</sup> That many Silicon Valley CEOs today appear to strictly limit the screen time of their own children, while trying to maximize the screen time of your children, is a rather curious thing.

Either way, the analogy to the tobacco industry suggests the need for serious and independent research into the potential cognitive, social, and mental health consequences of AI. A number of the leading AI companies have, to their real credit, already started research on some of these themes. But logic and experience with a range of products with potentially dangerous side effects, from tobacco to cars, suggests that the sole source of safety research should not be the industries themselves: It should be pursued more broadly. Concerns about the health effects of tobacco inspired independent research efforts at the leading public health universities from the 1960s forward.<sup>20</sup> In tandem to the

*A number of the leading AI companies have, to their real credit, already started research on some of these themes. But logic and experience with a range of products with potentially dangerous side effects, from tobacco to cars, suggests that the sole source of safety research should not be the industries themselves: It should be pursued more broadly.*



corporate research initiatives from the AI industry on this theme, we may want to consider an analogously serious and independent research program about the consequences of AI on mental health, cognitive formation, and social connection.

## V. Conclusion

In Silicon Valley today, there are frequent discussions about the “existential risk” AI presents to humanity. What if bad actors use AI to advance destructive ends? What if rogue AI systems take on a life of their own?

But as concerning as some of these scenarios may be, these discussions typically overlook a broader set of more challenges that, in our view, are just as significant. That the US, and with it most of the advanced economies around the world, have been economically stagnant for the past several decades, in a regime of diminished productivity growth, is in a very real sense “existential.” We want a future of high economic growth, and a society where ordinary people can enjoy exciting vocations and rich experiences of social and family life. We want a path in which AI genuinely sparks sustained economic growth, empowers normal people to have more transformative careers, and makes our social and cognitive experiences deeper and better. This is as existential as anything else. We need to focus now on making the most of this opportunity fate has given us. This can be aided by a better sense of the current structural problems of the US economy, more thought about education and mental health, and above all, a vision for what we want out of the future.

Keynes concluded his essay “Economic Possibilities for Our Grandchildren” with a call for economists “to get themselves thought of as humble, competent people, on a level with dentists,” as they presided over another hundred years of normal economic growth—a

**Economics as a discipline now has vast new intellectual opportunities. Beyond asking what the unemployment rate from AI might be in five years, or whether some part of the stock market will burst in three years, there are bigger, longer-term questions about how we want to reorganize our societies over the next fifty years, and the next hundred.**

rather deflationary view of the work that remained.<sup>21</sup> But another implication of the views developed here is that economics as a discipline now has vast new intellectual opportunities. Beyond asking what the unemployment rate from AI might be in five years, or whether some part of the stock market will burst in three years, there are bigger, longer-term questions about how we want to reorganize our societies over the next fifty years, and the next hundred. The intellectual opportunity here extends far beyond economics, to the rest of the social sciences. There may be a fundamental openness to the rest of the 21<sup>st</sup> century, an expanded set of economic and social possibilities, that did not characterize the rigidity and institutional convergence of the late 20<sup>th</sup> century. In the service of exploring this set of possibilities, we may need to leave the dentist's office and assume a higher task.

1. E.g., John G. Fernald, “Productivity and Potential Output Before, During, and After the Great Recession,” *NBER Macroeconomics Annual* 29, no. 1 (2015): 1–51, <https://doi.org/10.1086/680580>.
2. Richard Wike, Moira Fagan, Christine Huang, Laura Clancy, and Jordan Lippert, “Views of Children’s Financial Future,” in *Economic Inequality Seen as Major Challenge Around the World* (Pew Research Center, January 2025), <https://www.pewresearch.org/global/2025/01/09/views-of-childrens-financial-future/>; Megan Brenan, “Americans Less Optimistic About Next Generation’s Future,” Gallup News, October 25, 2022, <https://news.gallup.com/poll/403760/americans-less-optimistic-next-generation-future.aspx>.
3. Laura Silver, Scott Keeter, Stephanie Kramer, Jordan Lippert, Sofia Hernandez Ramones, Alan Cooperman, Chris Baronavski, Bill Webster, Reem Nadeem, and Janakee Chavda, “Americans’ Trust in One Another,” May 8, 2025, <https://www.pewresearch.org/politics/2024/06/24/public-trust-in-government-1958-2024/>; Pew Research Center, “Public Trust in Government: 1958–2024,” June 24, 2024, <https://www.pewresearch.org/politics/2024/06/24/public-trust-in-government-1958-2024/>.
4. Prince Albert, Speech at the Mansion House, March 21, 1850, in *The Principal Speeches of His Royal Highness the Prince Consort* (John Murray, 1862).
5. Julie Ray, “Americans Express Real Concerns About Artificial Intelligence,” Gallup News, August 26, 2024, <https://news.gallup.com/poll/648953/americans-express-real-concerns-artificial-intelligence.aspx>.
6. Colleen McClain, Brian Kennedy, Jeffrey Gottfried, Monica Anderson, and Giancarlo Pasquini, “Public and Expert Predictions for AI’s Next 20 Years,” in *How the U.S. Public and AI Experts View Artificial Intelligence* (Pew Research Center, April 2025), <https://www.pewresearch.org/internet/2025/04/03/public-and-expert-predictions-for-ais-next-20-years/>.
7. Dario Amodei, “Machines of Loving Grace: How AI Could Transform the World for Better,” *Dario Amodei* (blog), October 2024, <https://www.darioamodei.com/essay/machines-of-loving-grace>.
8. Bureau of Labor Statistics, *American Time Use Survey: 2023 Results* (U.S. Department of Labor, 2023), [https://www.bls.gov/news.release/archives/atus\\_06272024.pdf](https://www.bls.gov/news.release/archives/atus_06272024.pdf).
9. Charlynn Burd, Michael Burrows, and Brian McKenzie, *American Community Survey Report: Travel Time to Work in the United States, 2019* (United States Census Bureau, March 2021), <https://www.census.gov/content/dam/Census/library/publications/2021/acs/acs-47.pdf>.
10. Peter Ganong and Daniel Shoag, “Why Has Regional Income Convergence in the U.S. Declined?,” *Journal of Urban Economics* 102 (November 2017): 76–90, <https://doi.org/10.1016/j.jue.2017.07.002>.
11. Robert M. Solow, “Technical Change and the Aggregate Production Function,” *Review of Economics and Statistics* 39, no. 3 (August 1957): 312–320, <https://doi.org/10.2307/1926047>.

12. Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen, “Canaries in the Coal Mine? Six Facts About the Recent Employment Effects of Artificial Intelligence,” Working Paper (Stanford Digital Economy Lab, 2025), <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/>; Guy Lichtinger and Seyed Mahdi Hosseini Maasoum, “Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data,” SSRN Working Paper No. 5425555 (November 5, 2025), <http://dx.doi.org/10.2139/ssrn.5425555>.
13. See Congressional Budget Office, “Unemployment Insurance: Budgetary History and Projections,” accessed November 10, 2025, <https://www.cbo.gov/publication/61179>.
14. Organisation for Economic Co-operation and Development (OECD), “Public Spending on Labour Markets,” accessed November 10, 2025, <https://www.oecd.org/en/data/indicators/public-spending-on-labour-markets.html>.
15. Ryan Decker, John Haltiwanger, Ron S. Jarmin, and Javier Miranda, “Declining Business Dynamism: What We Know and the Way Forward,” *American Economic Review Papers & Proceedings* 106, no. 5 (May 2016): 203–207, <https://doi.org/10.1257/aer.p.20161050>; Raven Molloy, Christopher Smith, and Abigail Wozniak, “Internal Migration in the United States,” *Journal of Economic Perspectives* 25, no. 3 (Summer 2011): 173–196, <https://doi.org/10.1257/jep.25.3.173>.
16. Brent Bratsberg and Ole Rogeberg, “Flynn Effect and Its Reversal Are Both Environmentally Caused,” *Proceedings of the National Academy of Sciences* 115, no. 26 (2018): 6674–6678, <https://doi.org/10.1073/pnas.1718793115>.
17. “NAEP Long-Term Trend Assessment Results: Reading and Mathematics, 2022,” Nation’s Report Card, accessed November 10, 2025, <https://www.nationsreportcard.gov/highlights/ltr/2022/>.
18. Nick Bilton, “Steve Jobs Was a Low-Tech Parent,” *New York Times*, September 10, 2014, <https://www.nytimes.com/2014/09/11/fashion/steve-jobs-apple-was-a-low-tech-parent.html>.
19. Cory Steig, “How Mark Zuckerberg Lets His Toddlers Use Their Screen Time,” CNBC, October 23, 2019, <https://www.cnbc.com/2019/10/23/how-mark-zuckerberg-manages-kids-screen-time.html>.
20. Allan M. Brandt, *The Cigarette Century: The Rise, Fall, and Deadly Persistence of the Product that Defined America* (Basic Books, 2009).
21. John Maynard Keynes, *Essays in Persuasion* (W.W.Norton & Co., 1963).