

February 26, 2024

Simon Johnson, Ronald A. Kurtz (1954) Professor of Entrepreneurship at MIT Sloan School of Management, and faculty co-director, MIT Shaping the Future of Work Initiative.¹

Testimony prepared for the Senate Armed Services Committee hearing on evolving workforce dynamics and the challenges for defense acquisition and defense industrial base personnel, Wednesday, February 28, 2024, 9:30am.

A. Summary and Recommendations

1) The U.S. has serious long-standing issues related to what people in the tech industry call “stranded expertise.” This term refers to people who have skills and experience that previously earned a high wage in productive activities, but which the market now does not highly value. With the rapid deployment of AI technology, we should expect this issue to become even more salient.

2) However, while the modern phenomenon of stranded expertise is associated with major economic and social problems, it also presents opportunities for the Defense Department, both for the defense industrial base in general and its human resources strategy in particular. Specifically, the DoD can and should think about strengthening its suppliers around the country by applying the following principles.

a. Identifying geographies where the tech sector is not otherwise booming and where there is sufficient housing (or willingness to build it) to accommodate the current and future workforce.

b. Seeking out cities and regions where potential hubs of interrelated expertise already exist and can be further strengthened (e.g., through apprenticeship programs which help specific suppliers while generating positive spillovers for related activities, and knowledge sharing programs that can be coordinated by universities and applied research institutes).

c. Focusing on developing expertise that is undervalued by the private sector, including skilled manual work and the tacit knowledge needed for modern manufacturing at scale.

d. Buying in on software, which is the main focus of private sector innovation, but also building next generation hardware with long development cycles and high capital costs.

e. Establishing incentives for the development of “pro-worker AI” – i.e., encouraging the development of artificial intelligence-based tools that are complementary to (boosting the productivity and effectiveness of) workers without a four-year college degree.

f. Scanning for up-and-coming technologies that have the potential to transform the organization of any economy, including (but not limited to) how it produces weapons and military equipment. Boosting productivity while sharing prosperity creates the soundest possible foundation for national defense. We should also aim to build and maintain a lead in all sectors and activities with fast-changing technology frontiers. This could happen, for example, in the

¹ This testimony draws on *Jump-Starting America* (co-authored with Jonathan Gruber) and *Power and Progress* (co-authored with Daron Acemoglu), as well as ongoing projects at the MIT Shaping the Future of Work Initiative (co-led with Daron Acemoglu and David Autor), including “Can We Have Pro-Worker AI? Choosing a Path of Machines in Service of Minds” (<https://shapingwork.mit.edu/wp-content/uploads/2023/09/Pro-Worker-AI-Policy-Memo.pdf>), and as part of the Industrial Strategy Seminar, co-organized with Brian Deese and Liz Reynolds. However, all views expressed here are personal and do not represent the views of any other person or organization.

context of a “Chips and Science Act 2.0” that included support for biomanufacturing, critical minerals, clean energy, additive manufacturing, quantum computing, and other cutting-edge technologies.

3) The U.S. private sector does some things very well. For example, we lead the world in software-centric projects that thrive based on venture capital finance, and our largest tech companies dominate the strategically important AI sector.²

4) However, in recent decades, the private sector has chosen to outsource a large part of its supply chains to other countries, without deeply considering the geopolitical risks that this creates. We have also lost a lot of the skills needed to design and build physical goods, both in normal times and when there is an emergency that requires rapid scale-up. If international trade is also disrupted, directly or indirectly, by a serious crisis, that makes the problem even worse.

5) Long-term, it makes sense for the DoD to think about how to support places and people that can remain world class at manufacturing while also offering the potential for rapid scale-up when needed.

6) In this context, apprenticeships are particularly appealing. The private sector underinvests in apprentices, because employers are concerned that workers will leave for better opportunities once trained. However, this should be much less of a concern to the DoD because everyone who completes an apprenticeship in the defense industrial base becomes a resource for that base.

7) In the U.S. today, we have deep and underutilized geographic pockets of technology-related expertise. The impact of AI should be tracked carefully, as it may create more such pockets. Matching available and relatively immobile skilled labor with current DoD requirements should not be difficult, and including workforce training into such a program makes even more sense when people are disinclined to move.

8) Predicting exactly what technology skills the DoD will need in the future, even in the near future, is much more daunting. In the post-1945 period (in line with the “Endless Frontier” program) and in particular under the National Defense Education Act of 1958 (part of the policy reaction to Sputnik and the perceived engineering gap with the Soviet Union), the federal government successfully expanded access to higher education and specifically pulled more young people into what we now call STEM.

9) What we need now is to redouble our efforts in science, including more funding for basic research and more support for commercialization. This should be combined with complementary workforce development for people without a college degree. Undervaluing and underinvesting in manual skills has, for too long, undermined the strength of our defense industrial base.

10) The path of technology is never inevitable. It is and can be shaped by national policy, including how we think about and invest in the defense industrial base. A top priority now should be to encourage the development of AI-based tools that increase the effectiveness of workers without a college degree.

The remainder of this testimony is structured as follows.

- Section B: More details on what happened to jobs in the U.S. over recent decades, as well as lessons from relevant military-technology history.

² However, there are serious concerns that the digital ad-based model of social media damages young people, encourages extremism, and undermines our national capacity to sustain a full-range of innovation. AI-technologies may make all of these problems more severe. See *Power and Progress*, Chapters 10 and 11, for a thorough explanation of these points.

- Section C: Evidence on innovation and productivity growth in an economy where inter-regional mobility has become limited.
- Section D: Explanation of the potential workforce problems created by current and next generation AI, particularly if this focuses heavily on automation.
- Section E: Suggestions on how to respond to the business/competitive and national security threats posed by China.
- Section F: The national security benefits of funding science above currently appropriated levels.

B. Some Relevant History

When automation replaces people with machines, the obvious question is: what happens to the skilled workers who previously had those tasks? Historically, there have been episodes when the U.S. private sector created a large number of new tasks requiring expertise, and workers were able to benefit from these new opportunities (e.g., in the early 1900s and again after the 1940s) through employment at high relatively wages.

However, after the 1970s, the destruction of jobs through automation has outpaced the creation of new tasks for human experts, with this effect concentrated on blue-collar and white-collar tasks that had a significant routine element. Consequently, people previously in the middle of the wage distribution, particularly those who did not go to a four-year college, lost good jobs, and found themselves forced to compete in the low-skill (and low-wage) part of the labor market.

The most obvious manifestation of this issue is at the regional level. Over the past 40 years, innovation has become disproportionately concentrated in relatively few parts of the United States. Rigidities in local housing markets limited construction in cities that would otherwise have increased housing stock. Increased housing prices deter workers from moving to these new opportunities. The shock of cheap Chinese imports in the early 2000s accelerated the decline of manufacturing in some parts of the country.

The acceleration of Artificial Intelligence (AI) threatens to create even more stranded expertise as whole job categories may be wiped out. We have seen similar disruptions in U.S. labor markets previously, but the impact from this new challenge may be closer to that of the “China Shock” in the early 2000s than the arrival of automated telephone exchanges.³ We are not creating enough new tasks requiring expertise, and many people will have difficulty converting their expertise into different sectors and jobs if their current positions are eliminated.

³ Manual telephone exchanges employed a lot of young women. These jobs were eliminated with the arrival of automated exchanges. At the same time, however, women were able to find more and better opportunities across the economy. In contrast, the surge of imports from China after 2000s had major negative effects on employment in US manufacturing, and many communities and people have still not recovered (and perhaps never will recover fully, at least on our current technology path).

These developments pose major challenges for the U.S. macroeconomy, including our ability to sustain high wages, to boost productivity growth, and to maintain or continue the “compression” in wage inequality that occurred after the onset of the COVID-19 pandemic.⁴

At the same time, experience over the past century suggests that a robust defense industrial base needs three elements.

- a. Actual and potential capacity to meet predictable needs.
- b. The ability to pull in additional resources and scale up fast.
- c. Access to the greatest possible innovative depth.

The formative moment of the modern U.S. industrial base is instructive. In 1940, the U.S. was not prepared for war, either in terms of its stock of cutting-edge technology or its ability to scale up production of weapons, equipment, and other goods needed by the military. There was limited political support for anything that looked like intervention in (what was then regarded as just) another “European war”.

In late August 1940, Sir Henry Tizard arrived in Washington, DC, as head of an expert team bearing information about some of Britain’s most important technological discoveries. The Battle of Britain raged through the summer and fall of 1940, with the German Luftwaffe first crippling critical airfields and then switching to the Blitz bombing of civilian areas, including London, Coventry, Birmingham, and other major cities. In this moment of great national desperation, Tizard and a few others persuaded Churchill’s government to put aside all conventional notions of secrecy, with the goal of receiving greater material assistance from the United States. Of all Tizard’s offerings, without question the most immediately consequential was a small mechanical device, about the size of a hockey puck: the resonant cavity magnetron. This simple and even elegant piece of equipment created the possibility—with a lot of additional work—of smaller, more powerful, and more accurate radar sets.

To the great credit of the U.S., the wartime scientific organization under Vannevar Bush immediately recognized the potential of this technology, created the Radiation Lab at MIT (bringing together physicists from around the country), and brought this technology to fruition and into scale production by early 1941 – in time to play a decisive role in the Battle of the Atlantic (where plane-based radar was used to spot German submarines).⁵ The success of this

⁴ As documented by MIT colleague David Autor and coauthors, “The labor market tightness following the height of the COVID-19 pandemic led to an unexpected compression in the US wage distribution that reflects, in part, an increase in labor market competition. Rapid relative wage growth at the bottom of the distribution reduced the college wage premium and counteracted nearly 40% of the four-decade increase in aggregate 90-10 log wage inequality.” (<https://shapingwork.mit.edu/research/the-unexpected-compression-competition-at-work-in-the-low-wage-labor-market/>)

⁵ Experience at Pearl Harbor on December 7, 1941, should remind us that there is no necessary or automatic link from technology development to effective application in the defense sphere. Just after 7am on that morning, two radar operators reported a large number of unidentified aircraft heading towards Oahu. The officer to whom they reported told them to “Forget it.” The official congressional report determined, “The real reason, however, that the information developed by the radar was of no avail was the failure of the commanding general to order an alert commensurate with the warning he had been given by the “War Department that hostilities were possible at any moment” (pp.141-142, Investigation of the

effort, as well as the Manhattan Project and other technology development and scale-up efforts, was the basis for Bush’s 1945 *Endless Frontier* report. This provided a framework for post-war federal government efforts to support the development of breakthrough science, with both civilian and military applications. This effort received a further major impetus after the Soviet Union launched Sputnik, the first artificial satellite, in October 1957.

The story of military aviation development in the 1920s and 1930s is also instructive, although in part because of its cautionary elements. In 1923, France had the largest and most modern airforce in the world. By the late 1930s, however, all of this equipment was completely antiquated (or already out of service).

The French were not short of planes in May 1940. Instead, they lacked (a) modern fighter planes, capable of intercepting German bombers and defeating German fighters; (b) any kind of modern air defense system integrating human observers (and preferably radar); and (c) the strategy and tactics to use air defense effectively against a Luftwaffe that had been built and trained to practice close air support.⁶

Of course, 1940 was a long time ago, and the weapons of war continue to change fast. The more general lesson from history is that stockpiles are never enough. For robust national defense, we need strong national innovation across a wide range of technologies, most of which must presumably be rooted in the non-military economy. Of course, finding or imagining military and defense applications is entirely appropriate and a good idea.

The most important goal is to build a coherent vision for how technology is developed and applied to national defense. A key component of this vision must be the skills required at all levels of the future workforce.

C. Innovation, Productivity Growth, and Regional Issues⁷

In the decades that followed World War II, the U.S. led the world in innovation, creating entirely new sectors such as jet aircraft, life-saving drugs and vaccines, microelectronics, satellites, and

Pearl Harbor Attack, Joint Committee on the Investigation of the Pearl Harbor Attack, Congress of the United States, July 20, 1946.)

⁶ The key breakthrough by German armored forces across the Meuse at Sedan was made possible by intense support from dive bombers, acting in place of towed artillery. The Germans then shifted focus to disrupt French communications, particularly the railway system, making it very hard (or perhaps impossible) to bring up reinforcements quickly. Within a few days, German forces had reached the sea, cutting the supply lines of over 300,000 Allied troops. The British Expeditionary Force and some French troops were evacuated from Dunkirk, and France surrendered within a few weeks.

⁷ This section is an updated version of a series of policy memos co-authored with Jonathan Gruber, 2019–2024. For more details, see our book, *Jump-Starting America*. The data to identify “Top Technology Hub Candidates” are freely available on this website, <https://www.jump-startingamerica.com/>, which allows users to examine the effects of adjusting the criteria and changing the weightings used in our baseline assessment (e.g., housing prices, college graduates, working population).

digital computers. Widespread innovation boosted productivity. Household income increased faster than ever before, while inequality declined.

Since the 1970s, however, U.S. productivity growth has slowed – reflected in falling total GDP growth from 4 percent in the post-war years, to under 3 percent from the mid-1970s, and to under 2 percent since 2000. The CBO projects growth (in real potential GDP) over the next decade of only 1.9-2.2 percent.⁸

Moreover, well-paying jobs in the U.S. are now concentrated disproportionately in a small number of superstar cities. People in the rest of the country increasingly – and correctly – feel that they are being left behind. Cities and states attempt to compensate by offering large tax breaks and other subsidies to big companies, but this just enriches corporations while starving local services.

The long-term slowdown in U.S. productivity growth has multiple causes, but in part is due to congestion, restrictive zoning, and high housing prices in existing innovation mega-hubs (Seattle, San Francisco Bay Area, Los Angeles, Boston, New York City, and Washington DC). Other large U.S. cities have strengths but have not broken through as magnets for venture capital and tech business creation. Meanwhile, small towns are losing their working age populations as workers age or move away (“brain drain”).

At the same time, Asian and European countries are building their own scientific infrastructure, including technical facilities and – most importantly – expertise. The US faces a new wave of global challenges to its technological leadership and ability to sustain good jobs. (More on this in the context of thinking about potential threats from China in Section E below).

What went wrong? Policymakers forgot one of the most important lessons of the post-1945 period. Modern private enterprise is most effective when the government provides strong underlying support for science and for the commercialization of inventions.

Firms are interested in innovation only to the extent that it improves their own bottom line – and not if creating new ideas and products leads to benefits for someone else. But spillovers from discovery are incredibly important, creating both the basic scientific knowledge and the more applied ideas that help determine how fast our economy can grow. Government involvement can fill this strategic gap by enabling scientific discoveries that, while crucial for long-run progress, may not always have immediately apparent benefits to “the bottom line.”

⁸ “In CBO’s projections, real potential GDP grows at an average rate of 2.2 percent a year from 2024 to 2028— slightly higher than the average rate since the business cycle peak in 2007—and then grows at an average rate of 1.9 percent a year from 2029 to 2034.” A large part of this growth is due to expected increase in the labor force. “The productivity of the potential labor force (which equals real potential GDP divided by the size of the potential labor force) is projected to grow by an average of 1.2 percent a year from 2024 to 2028 and 1.5 percent a year from 2029 to 2034.” See p.54 in “The Budget and Economic Outlook, 2024-2034,” Congressional Budget Office, February 2024. <https://www.cbo.gov/publication/59946>

The innovation that led to rapid growth after World War II was the direct result of a fruitful partnership between the private sector, federal government, and universities that allowed us to generate and benefit from these spillovers as a country. Almost every major innovation in this era relied on federal government support, provided by both Democratic and Republican administrations. Public spending on research and development peaked at nearly 2% of our entire economy in 1964.

Public support for R&D in recent years amounts to no more than 0.7 percent of GDP (and we are now under 0.6 percent of GDP). This is equivalent to spending at least \$250 billion per year LESS than we did during the post-war boom. Lower public investment in science has contributed to the slowdown in productivity growth.

Across almost every dimension of technology today, America faces the imminent prospect of falling behind other nations. Around the world, including in China, government-supported research initiatives are helping to create the technologies of tomorrow, along with the associated well-paid jobs. Our competitors have studied post-1945 American history carefully – and are applying the lessons. We are at risk of falling behind and losing even more good jobs.

When Amazon announced its intention to build a HQ2 somewhere in North America, 238 cities submitted bids with generous tax incentives, infrastructure improvements, real estate deals, and other inducements. Yet the winners were two of the most economically successful places in recent decades: New York City, and Northern Virginia/Washington DC. The private sector, left to its own devices, will not close the income and opportunity gaps in America.

In contrast, placed appropriately, geographically concentrated federal investments can be truly transformative: attracting companies and helping to generate more local private sector employment. We need to jump-start the American growth engine in a way that benefits the entire country, not just a few large cities. Specifically:

1. Expand federal funding for basic science, as well as for the technology development that creates jobs through commercialization. Every incremental investment by the federal government of \$100 billion per year in this scientific infrastructure creates four million good new jobs in the near term – while boosting growth and helping to strengthen the middle class.
2. Share the benefits much more broadly across the country. In *Jump-Starting America*, Jonathan Gruber and I identified 102 urban communities that are plausible next generation tech hubs, all with large populations, highly educated workforces, and a low cost of living. These communities are home to over 80 million Americans in 36 states, across all regions of the country. For example, there are: seven potential new hub locations each in Florida, Michigan and Ohio; six each in Alabama and Indiana; five each in New York, Pennsylvania, Tennessee, and Texas; four in Georgia; and three each in Iowa, South Carolina, and Wisconsin.⁹
3. Create a rolling national competition for municipalities to become new technology hubs. Areas would bid to show that they can become an effective home for technology development. Potential hubs should also aim to increase the supply of skilled workers by: making higher education more affordable; providing appropriate practical and technical training, including for

⁹ See this website link for a map and list of the relevant Metropolitan Statistical Areas: <https://www.jump-startingamerica.com/policy-summary>.

people who begin with just a high school-level education; and linking to locally available employment opportunities that pay good wages.

4. Ensure that the benefits of technology-led growth accrue to all Americans by creating a cash Innovation Dividend. The capital appreciation (and higher rents) on government-owned real estate in new hubs would flow to all Americans through a flat annual payment – just as the Alaska Permanent Fund redistributes oil royalty revenue.

D. The Likely Workforce Impact of AI¹⁰

The world is currently flooded with optimistic forecasts regarding the growth implications of AI. [Goldman Sachs, for example](#), claims that “[AI adoption could boost productivity](#) growth by 1.5 percentage points per year over a 10-year period” and raise global GDP by 7 percent – that’s \$7 trillion of additional output, as they calculate it. People close to the industry have even [more excited estimates](#), including a supposed 10 percent chance of an “explosive growth” scenario, with global output rising more than 30 percent per year, which would be [unprecedented by almost an order of magnitude](#).

All this techno-optimism draws on a set of deep-rooted beliefs in modern economics: innovation drives productivity growth, and productivity growth is ultimately good for workers. In this long-established conventional wisdom, technological change—including various forms of automation—always has the net effect of raising wages and generating more opportunity, creating an engine that pulls everyone along and leads directly to shared prosperity.

Such optimism is at odds with the historical record and seems particularly inappropriate for the current path of “just let AI happen,” because this is focused primarily on automation (replacing people). The macroeconomic and human impacts will be much better if we first recognize that there is no singular, inevitable path of development for this (or any) new technology. Armed with that understanding, we can then address the right question: Assuming that the goal is to improve economic outcomes for more people in a sustainable manner, what policies would help put AI development on the right path, with much more focus on enhancing what *all* workers can do?

Is There a Productivity Bandwagon?

Optimism regarding shared benefits from technological progress is founded on a simple and powerful idea: the “productivity bandwagon.” This idea maintains that new machines and production methods that boost productivity will also produce higher wages. As technology progresses, the bandwagon will pull along everybody, not just entrepreneurs and owners of capital.

The theory behind the productivity bandwagon is straightforward: when businesses become more productive from new technology, they want to expand their output. For this, they need more workers, so they get busy with hiring. And when many firms attempt to do so at the same time, they collectively bid up wages.

¹⁰ This section is based on “Rebalancing AI,” by Daron Acemoglu and Simon Johnson, *Finance & Development*, International Monetary Fund, December 2023. <https://www.imf.org/en/Publications/fandd/issues/2023/12/Rebalancing-AI-Acemoglu-Johnson>

This is sometimes what happens. For example, throughout most of the twentieth century, innovations like electrical machinery and more-efficient factories boosted productivity in US industries like car manufacturing, and real wages increased in tandem. Remarkably, from the end of World War II to the mid-1970s, the wages of high school graduates and college graduates in the US and other industrial countries grew at roughly the same rate.

Unfortunately, what subsequently occurred is not consistent with the notion of an unstoppable bandwagon. How productivity benefits are shared depends on how technology changes and on the rules, norms, and expectations that govern how management treats workers. To understand this, consider the two steps that link productivity growth to higher wages.

- First, productivity growth increases the demand for workers as businesses attempt to boost profits by expanding output and hiring more people.
- Second, the demand for more workers increases the wages employers need to offer to attract and retain employees.

The Machinery Question

Contrary to popular belief, productivity growth need not translate into higher demand for workers. The standard definition of productivity is average output per worker—total output divided by total employment. The hope is that as output per worker grows, so will the willingness of businesses to hire people.

But employers are not incentivized to increase hiring based on average output per worker. Rather, what matters to companies is *marginal productivity*—the additional contribution that one more worker brings by increasing production or by serving more customers. The notion of marginal productivity is distinct from output or revenue per worker; output per worker may increase while marginal productivity remains constant or even declines.

Many new technologies, like industrial robots, expand the set of tasks performed by machines and algorithms, displacing workers who used to do these tasks. Automation raises average productivity but does not increase, and in fact may reduce, worker marginal productivity. Over the past four decades, automation has raised productivity and multiplied corporate profits, but has not created shared prosperity in industrial countries.

Replacing workers with machines is not the only way to improve economic efficiency—and history has proven this, as I describe in my recent book [*Power and Progress*](#) (co-authored with Daron Acemoglu). Rather than automating work, some innovations boost how much individuals contribute to production. For example, new software tools that aid car mechanics and enable greater precision can increase worker marginal productivity. This is completely different from installing industrial robots with the goal of replacing people.

Where Are the New Tasks?

Even more important for raising worker marginal productivity is the creation of new tasks. When new machines create new uses for human labor, this expands workers' contributions to production

and increases their marginal productivity. There was plenty of automation in car manufacturing during the momentous industry reorganization led by Henry Ford starting in the 1910s. But mass-production methods and assembly lines simultaneously introduced a range of new design, technical, machine-operation, and clerical tasks, boosting the industry's demand for workers.

New tasks were vital, not just in early US car manufacturing, but also in the growth of employment and wages over the last two centuries. And many of the fastest-expanding occupations in the last few decades—MRI radiologists, network engineers, computer-assisted machine operators, software programmers, IT security personnel, and data analysts—did not exist eighty years ago. Even people in occupations that have been around for quite a while, such as bank tellers, professors, and accountants, now work on many relatively new tasks using computers and modern communication devices. In almost all these cases, new tasks were introduced because of technological advances and have been a major driver of employment growth. These new tasks have also been integral to productivity growth—they have helped launch new products and enabled a more efficient production process.

Automation Also Rises

Automation in an industry can also boost employment—in that sector or in the economy broadly—if it substantially increases productivity. In this case, new jobs may come either from nonautomated tasks in the same industry or from the expansion of activities in related industries. In the first half of the twentieth century, the rapid increase in car manufacturing stimulated massive expansion of the oil, steel, and chemical industries (think gasoline, car bodies, and tires). Vehicle output at mass scale also revolutionized the possibilities for transportation, enabling the rise of new retail, entertainment, and service activities.

The productivity bandwagon is not activated, however, when the productivity gains from automation are small—what we call “so-so automation.” For example, self-checkout kiosks in grocery stores bring limited productivity benefits because they merely shift the work of scanning items from employees to customers. When stores introduce self-checkout kiosks, fewer cashiers are employed, but there is no major productivity boost to stimulate the creation of new jobs elsewhere. Groceries do not become much cheaper, there is no expansion in food production, and shoppers do not live differently.

Even nontrivial productivity gains from automation can be offset when they are not accompanied by new tasks. For example, in the industrial heartland of the American Midwest, the rapid adoption of industrial robots has contributed to mass layoffs in the manufacturing sector and ultimately prolonged regional decline.

The situation is similarly troubling for workers when new technologies focus on surveillance. Increased monitoring of workers may lead to some small improvements in productivity, but its main function is to extract more effort from workers.

All this underscores perhaps the most important aspect of technology: *choice*. There are often myriad ways of using our collective knowledge to improve production and even more ways to

direct innovation. Will we invent and implement digital tools for surveillance? For automation? Or to empower workers by creating new productive tasks?

When the productivity bandwagon is weak and there are no self-correcting mechanisms to ensure shared benefits, these choices become more consequential—and the few tech leaders who make them become more powerful, both economically and politically.

The Human Complementary Path

New technology may complement workers by enabling them to work more efficiently, perform higher quality work, or accomplish new tasks. For example, even as mechanization gradually pushed more than half of the US labor force out of agriculture, a range of new blue-collar and clerical tasks in factories and newly emerging service industries generated significant demand for skilled labor between (roughly) the years 1870 and 1970. This work was not only better paid but also less dangerous and less physically exhausting.

This virtuous combination—automation of traditional work alongside creation of new tasks—proceeded in relative balance for much of the twentieth century. But sometime after approximately 1970, this balance was lost. While automation has maintained its pace or even accelerated over the ensuing five decades due to computerization, the offsetting force of new task creation has slowed, particularly for workers without four-year college degrees. As a result, these workers are increasingly found in low-paid services such as cleaning, security, food service, recreation, and entertainment. These jobs are socially valuable, but they require little specialized education, training, or expertise, and hence pay poorly.

The critical question we face in the new era of AI is whether this technology will primarily accelerate the existing trend of automation without the offsetting force of good job creation—particularly for non-college workers—or whether it will instead enable the introduction of new labor-complementary tasks for workers with diverse skillsets and a wide range of educational backgrounds.

It is inevitable that AI systems will be used for some automation. A major barrier to automation of many service and production tasks has been that they require flexibility, judgment, and common sense—things that are notably absent from pre-AI forms of automation. Artificial intelligence, especially generative AI, can potentially master such tasks. It is unclear how much this type of automation will contribute to aggregate productivity growth while these technologies are immature, but they could contribute to sizable productivity gains as costs fall and reliability improves.

The dominant intellectual paradigm in today’s digital tech sector—among both business leaders and academic researchers—also favors the automation path. A major focus of AI research is to attain human parity in a vast range of cognitive tasks and, more generally, to achieve “artificial general intelligence” that fully mimics and then surpasses capabilities of the human mind. This intellectual focus encourages automation rather than the development of human-complementary technologies.

However, there is a case for qualified optimism: AI offers an opportunity to complement worker skill and expertise, if we direct its development accordingly.

Human productivity is often hampered by lack of specific knowledge or expertise, which could be readily supplemented by next-generation technology. AI tools could assist in such cases by boosting human expertise, supporting workers in unfamiliar situations, providing on-the-spot training, and improving all forms of information translation. AI holds great potential for training and retraining expert workers, such as educators, medical personnel, software developers, and other workers with modern “crafts” (such as electricians and plumbers). AI could also create new demands for human expertise and judgment in overseeing these processes, communicating with customers, and enabling more sophisticated services that leverage these tools.

Redirecting technological change is not easy, but it can be done. Specifically, there are five steps that can help put AI development onto a human-complementary, rather than human-displacing, path. Governments everywhere should consider these steps, but the impact will be strongest in places where the technology is under active development—particularly in the U.S.

What Can Be Done?

Reform business models: The dominant developers of AI easily expropriate consumer data without compensation, and their reliance on digital advertising incentivizes grabbing consumers’ attention through whatever means possible. To force this business model to change, we need to establish clear ownership rights for all consumers over their data, and to tax digital ads. Enabling a more diverse range of business models is likely essential for the development of AI technology that is helpful to all humans.

Tax system: The current tax code places a heavier burden on firms that hire labor than on those that invest in algorithms to automate work, and the same is true in many other countries. Policymakers should aim to create a more symmetric tax structure, equalizing marginal tax rates for hiring (and training) labor and for investing in equipment/software. This would shift incentives toward human-complementary technological choices by reducing the bias of the tax code toward physical capital over human capital.

Labor voice: The direction of AI will have profound consequences for all workers. Creating an institutional framework through which workers have a voice in its development would be constructive—and there is an important role for civil society in pressing for this to happen. At the same time, government policy should restrict deployment of untested (or insufficiently tested) AI for applications that could put workers at risk, for example in high-stakes personnel decision-making tasks (including hiring and termination) or in workplace monitoring and surveillance. Health and safety rules need to be updated accordingly. The same policies are relevant for other countries to consider.

Funding for more human-complementary research: Additional support for the research and development of human-complementary AI technologies could have significant impact. With this field in its infancy, the US federal government should foster competition and investment in technology that pairs AI tools with human expertise, aiming to improve work in vital social sectors

like education and healthcare. Other countries could pursue research efforts, either by themselves or (more efficaciously) by working together on shared problems. Once enough progress has been made, governments can encourage further investment by advising on whether purported human-complementary technology is appropriate to adopt in publicly funded education and healthcare programs.

AI expertise within government: AI will touch every area of government investment, regulation, and oversight, including (but not limited to): transportation, energy production, labor conditions, healthcare, education, environmental protection, public safety, and military capabilities. Developing a consultative AI division within government that can support the many agencies and regulators tackling these challenges will support more timely, effective decision-making.

The Potential Impact on the Macroeconomy

AI could increase global GDP in the next five years, although not as substantially as the enthusiasts claim. It might even modestly raise GDP growth in the medium-term. However, on our current trajectory, the first-order impact is likely to be increased inequality within industrial countries. Middle-income countries and many lower-income countries also have much to fear from the existing path. A new capital-intensive technology is arriving and will soon be applied in offices and factories everywhere. There is no guarantee that, on its current path, AI will generate more jobs than it destroys.

If we can redirect AI onto a more human-complementary path, while using it to address pressing social problems, all parts of the planet can benefit. But if the “just-automate” approach prevails, shared prosperity will become even harder to achieve.

E. A Better Response to China¹¹

China’s strategy for economic growth is no secret. In the short-term, the authorities endeavor to keep the Renminbi exchange rate at a level that is consistent with robust manufacturing exports, while over the longer run they build the financial capital, physical infrastructure, and – most important – human expertise necessary to become more innovative.

Such a strategy has compelling antecedents, with versions finding remarkable success over the past 70 years in Asia (Japan, Taiwan, Singapore, and South Korea) and in Europe (Germany but also France, Austria, Sweden, Finland and others). If a country can advance its innovative capacity sufficiently, it can create a lot of good jobs at the forefront of designing and building what comes next.

The economic success of perceived rivals has led to trade friction with the United States in the past, including for Germany in the early 1970s and Japan in the 1980s, but the challenge from China is viewed as more existential for three main reasons. First, China’s economy is already

¹¹ This section is an updated version of “To Counter China, Out-Invent It,” by Jonathan Gruber and Simon Johnson, *Foreign Affairs*, September 12, 2019. <https://www.foreignaffairs.com/articles/china/2019-09-12/counter-china-out-invent-it>

big (more than half the size of the U.S.) and could easily become the largest in the world. Second, China has acquired a reputation for not always playing by existing trade and exchange rate rules, with significant negative implications for employment in the United States. Third, China is clearly projecting its influence around the world, sharing technology and potentially creating its own international sub-system.

The real China issue is not an undervalued exchange rate (that was the problem 15-20 years ago), or weak protection for intellectual property (to encourage domestic innovation, China is already moving toward stronger protections), or even the much-discussed U.S.-China trade deficit (a completely misleading guide to sensible economic policy.) The main threat from China is both much more serious and considerably easier for the U.S. to address than any of the current headline concerns – it is their increasingly effective government-led investments in Research and Development (R&D). China’s commitment to this area seems likely to make it the global innovation leader of the near future – displacing the U.S. from a position it has held for 70 years. Quite apart from the direct national security consequences of no longer leading the way, such a shift would reduce the number of good jobs in the U.S. that flow from new breakthrough industries.

Some changes in Chinese policy could help, but there is actually no reason to force or persuade China to do anything different. What the U.S. needs to do is change its own policy drastically to favor innovation and spread opportunity more broadly around the country.

The Real Threat from China

Chinese higher education was decimated as a result of the Cultural Revolution from 1966-1976, but during the 1990s and 2000s the country took a major step forward in science and engineering. From 1990 to 2010, Chinese enrollment in higher education rose eightfold, rising from 6 percent to 17 percent of total world higher educational enrollment. Over the same time period, the number of college graduates rose from 300,000 to nearly 3 million per year. In 2017, [7 million students graduated from Chinese universities](#). (The US equivalent number is around 3 million new graduates).

The quality of scientific education in China can be debated, but the numbers are not in question. In 1990, China graduated only one-twentieth as many science and engineering Ph.Ds. as the U.S., whereas by 2010 China (28,000 Ph.Ds.) had surpassed the U.S. (24,500). Chinese universities remain of a lower quality than their U.S. counterparts, but the gap is closing. In 2003, China had only 10 universities ranked in the top 500 in the world; [today they have 45](#). (This assessment is run by a Chinese group, which sees 8 out of the top 10 universities as American.)

More scientists and more engineers add up, over time, to more innovation – when combined with adequate funding. [Total worldwide spending on Research and Development](#) currently runs around \$2 trillion per year (latest data are for 2015). About a quarter of this spending is by the US, down from a 37 percent share in 2000. [Up from \\$33 billion \(roughly the same as France; one-tenth of the US level\) in 2000](#), China now ranks second, with \$408.8 billion – 21 percent of world R&D spending.

Technology is Power – and Jobs

Why should we care which country makes new discoveries? Ultimately, technological development benefits the world as a whole as ideas spread. A look back at American history, however, shows that technological leadership matters for both national security and the kind of jobs that are created in a country.

In 1938, on the eve of World War II, US federal and state governments spent a combined 0.076 percent of national income on scientific research, a miniscule amount. By 1944, the US government was spending nearly 0.5 percent of national income on science – a sevenfold increase that was used to fund the creation of radar systems, the development of penicillin, and the invention of the atomic bomb. The effects of this unprecedented surge were simply incredible and, for America's enemies, ultimately devastating.

The lesson for peacetime (and the Cold War) was straightforward: use public resources to create new science and encourage the private sector to bring applications of these ideas to market. From 1940 to 1964, federal funding for research and development increased twentyfold. At its peak in the mid-1960s, this spending amount was around 2 percent of annual gross domestic product.

The result was modern pharmaceuticals, microelectronics, digital computers, jet aircraft, satellites, GPS, the internet, and much more. The number of university places for studying science, engineering, and all their applications greatly increased. The United States became, without question, the best country in the world to study, develop, and commercialize new technology.

Median family income in the United States doubled from 1947 through 1970. This increase in prosperity was shared throughout the country, with growth not only on the coasts but in the industrial Midwest and the newly dynamic South. Not all of this shared growth was due to the new prominence of Research and Development, but public investment in new sectors and activities played a catalytic role. New ideas based on government-funded R&D, both through universities and private sector contracts, were the basis for the growth of some of our most iconic companies and best employers – companies like IBM, AT&T, and Xerox.

From the late 1960s, the U.S. government commitment to Research and Development declined, but in the early 1980s the U.S. still spent 1.2 percent of GDP on publicly funded science. Total R&D spending, including private sector efforts, amounted to 2.3% of GDP. Few other countries came even close.

This situation changed dramatically over the next twenty-five years, as U.S. public sector R&D slipped below 0.7% of GDP. Today, nine countries spend a higher share of GDP on public support for Research and Development. Seven countries now spend a higher share of their GDP on total R&D (public plus private) than does the U.S. China, the rising challenger, currently spends less than the U.S. currently (2.07 percent vs. 2.74 percent, for total R&D as a share of GDP) but plans to change that in the coming years.

The private sector in the U.S. continues to innovate, but mostly in the kind of software projects favored by modern venture capitalists. Outside of life sciences, the private sector does not spend much on fundamental breakthroughs – new knowledge is great for an economy but seldom pays off for the investors who fund the work.

Corporations have moved away from basic research, with a 60 percent decline in publications by corporate scientists in recent decades. In the critical area of clean energy, for example, the type of long-term innovative investments that lead to future jobs are not the current priority of US companies and financiers – the amount of capital involved is too large and the time horizon for returns is regarded as too long.

Other promising industries of the future, such as synthetic biology and hydrogen power, were hatched in the U.S. but now seem likely to develop further in China and elsewhere – based on where governments are providing more support for basic research and the translation of new science into production.

At the same time, many local US innovation ecosystems – which used to be vibrant creators of good factory jobs – have fallen into disrepair. In their analysis of what they call “[The China Trade Shock](#)”, David Autor, David Dorn, and Gordon Hanson identified a form of hollowing out of the middle class created by relatively cheap (and not high tech) Chinese imports in the early 2000s. In any important set of geographies, across the industrial Midwest and in smaller towns, low-cost Chinese imports wiped out existing employers or induced firms to move jobs to lower labor cost locations.

American consumers were able to buy less expensive goods but American workers, in some geographies, lost out rather badly. The U.S. economy has proved good ultimately at generating replacement jobs, but most of these do not pay high wages or provide the same level of benefits. Now China is poised to move further up the technology ladder, into more sophisticated products – this is the point of all that scientific investment. They have already demonstrated substantial progress in sectors such as telecommunications equipment, payments systems, and the use of artificial intelligence in supervising social media. One industry view is that China will soon overtake the U.S. as the world R&D leader in pharmaceuticals.

The most pressing question for the U.S. role in the world economy now becomes: Do Americans want to buy the next generation of machines, medicine, and software from China – falling back from the forefront of global technology leadership just like Britain did after World War II? Or do Americans (and their leaders) want to increase scientific investment and reassert leadership through the creation of opportunity – which was the U.S. response after the Soviet Union launched Sputnik in 1957?

Any tariff-based confrontation with China today or in the near future is unlikely to help boost productivity or create well-paid jobs for Americans of all abilities and education levels. We need a different set of policies.

F. The National Security Benefits of Funding Science¹²

Congress should fund the science and economic development programs authorized in the “CHIPS and Science Act” in the FY 24 appropriations bills, including by fully funding the Commerce Department’s regional technology hubs program.

The “CHIPS and Science Act” provided appropriations for the CHIPS activities in the bill but only set authorization levels for the important science programs in the bill, which are designed create a growing base of new research and innovation around the country. The FY 23 appropriations bills started funding the programs but at levels significantly below what was authorized. The FY 23 appropriation for the National Science Foundation, for example, was about \$2 billion below the authorized level, and the House and Senate numbers for FY 24 are about \$6 billion below the authorization. The FY 23 appropriation for the regional technology hubs at the Department of Commerce was \$500 million – just a quarter of what was authorized.

Appropriating the money for the science portions of “CHIPS and Science” would represent a significant commitment towards ending the dramatic almost sixty-year slide that we have seen in government support of science in the U.S. In the mid-1960s, the U.S. spent almost 2% of our GDP supporting public research and development, which was far more than any nation on earth. We now spend less than 0.6% of GDP – placing us 14th in the world and, critically, well behind China.

But the argument for funding science is not just about history or international competition – it is about jobs and long-run economic growth. Evidence clearly shows that higher public funding of research and development [creates economic opportunity](#): *Investing \$10 billion a year in government science funding would create 400,000 new jobs.* And the long-term growth prospects are even more impressive, with estimates suggesting a rate of return on public R&D spending nearly over 50%.

Importantly, these benefits are not just national – localized investments in science have outsized returns for the localities themselves. Recognizing this, the Commerce Department regional technology hubs program would reward areas with the best potential to be next generation technology hubs. This program has the potential to transform areas around the country that have operated in the shadow of the existing “superstar” cities on the coasts but which themselves can become the superstar cities of the future.

There is a long historical precedent for such transformational government investments. From the decades of economic growth that followed from government investments in public universities, to the jobs created by ongoing public funding of university research, government science funding has been a local growth catalyst. Consider east Orange County, Florida, part of the Orlando metro area, which due to a forward-looking collaboration between the University of Central Florida and the Navy, along with decades of military research funding, has become the world hub of the computer simulation industry and created 100,000 jobs over the past thirty years.

¹² This section reproduces a policy memo by Jonathan Gruber and Simon Johnson, circulated on February 1, 2024.

A common criticism of government-led efforts at economic development is that political considerations will interfere with allocating resources to the most promising economic opportunities. Fortunately, we now have two strong counterexamples to this contention. The first is the \$1 billion Build Back Better Regional Challenge run by the Economic Development Administration. This was essentially a “test case” for the tech hubs competition, with similar criterion for success but at a smaller funding scale. [The results of this competition](#) were 21 grants of \$25 to \$65 million that covered 24 states from all over the country, covering every region of the nation, rural and urban projects, including blue and red states alike. The second is the first round of the regional tech hubs competition. [EDA designated 31 new Tech Hubs](#), once again spanning the entire country and a wide range of potential technologies. There has been no evidence that these processes did anything other than what they were designed to do: select the places with the best potential to develop into new technology hubs through innovative growth.

But this is just a start. Becoming a next generation tech hub requires competing with the existing coastal superstar cities that are the location of [75% of all venture capital funding](#) and where [90% of the nation’s innovation sector growth was generated](#). This requires a massive commitment to “jump start” a region. The federal funding requires a local contribution, so the federal investment will generate the needed public and private investment at the local level.

For the 31 areas on the initial list of tech hubs, the argument for more funding for the tech hub program is clear. The \$500 million initially appropriated will not be enough to jump start more than a handful of these locations. Yet all of the places on this list have clearly shown the potential to benefit from transformational government spending. Moreover, [another 18 areas](#) were designated as Strategic Development Grant recipients, receiving small grants to allow them to develop their case for tech hub status further. That is 49 places around our country that are already positioned to benefit from more funding.

For this program to reach its potential, Congress must fully fund the Biden Administration’s requests for science funding. Only by investing fully in science can the U.S. reach its growth potential and return to the golden era when we led the world both in innovation and economic growth. And only by recognizing and rewarding those places with the most potential to grow can we spread the benefits of innovation-led growth around the country.