STANFORD UNIVERSITY

THE STANFORD EMERGING TECHNOLOGY REVIEW 2023

A Report on Ten Key Technologies and Their Policy Implications

CHAIRED BY Condoleezza Rice, John B. Taylor, Jennifer Widom, and Amy Zegart
DIRECTED BY Herbert S. Lin

11

CROSSCUTTING THEMES AND COMMONALITIES

Chapters 1 through 10 addressed ten individual technology areas. This chapter pulls together some common themes that cut across them. Of course, the scientific issues at play are different because the science is different. However, there are important similarities in how people and institutions make progress that are often lost when each field is considered in isolation.

The Value and Risk of Technological Progress

Takeaway *Innovation that emerges too fast threatens the legitimate interests of those who might be negatively affected by such innovation, while innovation that moves too slowly increases the likelihood that a nation will lose first-mover advantages.*

New technologies typically bring two types of benefits. First, they can enhance or improve existing processes. Second, they can enable entirely new functions—addressing problems in novel ways and solving challenges that people did not even know they had. In the first half of the twentieth century, for example, polio afflicted thousands of people worldwide. The iron lung was invented in the 1930s to help polio victims breathe, and over the next twenty-five years, improvements were made to the iron lung.[1] But the groundbreaking Salk vaccine in 1955 brought an entirely different way to defend people against polio. Within a few years, use of the iron lung dropped to nearly zero.

Manufacturing provides another example. For decades, large-scale manufacturing has relied on the idea of an assembly line to fabricate essentially identical models of the same product. Workers were originally all human. Then robots began replacing them, performing many assembly-line tasks more rapidly and accurately while reducing production costs. In the past two decades, a complementary fabrication paradigm has emerged: custom, on-demand manufacturing of products in small quantities using 3-D printing, or what's known as additive manufacturing. This new paradigm enables production that is far more localized and customized, though it does not yield the economies of scale that mass production offers.

Technological progress also brings risks—risks of moving too fast or too slowly. Innovation that emerges too fast threatens to disrupt the often-delicate balance that has been established among many national, organizational, and personal interests. As we are seeing today with AI, the rush to deploy new capabilities may give short shrift to issues such as safety, security, employment, values, ethics, societal impact, and geopolitics. On the other hand, innovation that is too slow increases the likelihood that a nation will lose the technical, economic, and national security advantages that often accrue to first movers in a field. In both cases, policy measures are often needed to steer outcomes in a more optimal direction.

The road from scientific discovery to useful application is often rockier than expected as well, with would-be innovators finding that the realization of the benefits promised to investors and customers actually entail greater costs, deliver fewer capabilities, and take more time than anticipated. Furthermore, it may well be that only upon delivery of new products do other risks become apparent, with innovators facing issues of ethics and equity, privacy, and increased challenges to health, safety, and security—all risks that could lead to an erosion of trust in their services or capabilities.

The Central Importance of Ideas and Human Talent in Science and Technology

Takeaway *Human talent plays a central role in generating the ideas for innovation, it can be found all over the world, and it cannot be manufactured at will.*

From time to time, lone scientists working on their own achieve breakthroughs on very difficult problems. But it is far more common that successful science and technology (S&T) innovations are a result of a well-functioning collaborative effort that can bring to bear a broad range of cognitive styles and disciplinary expertise.[2]

Scientific progress obviously benefits from new ideas. New ideas are created every day by talented Americans, but Americans do not have a monopoly

FIGURE 11.1 Number of patents per year

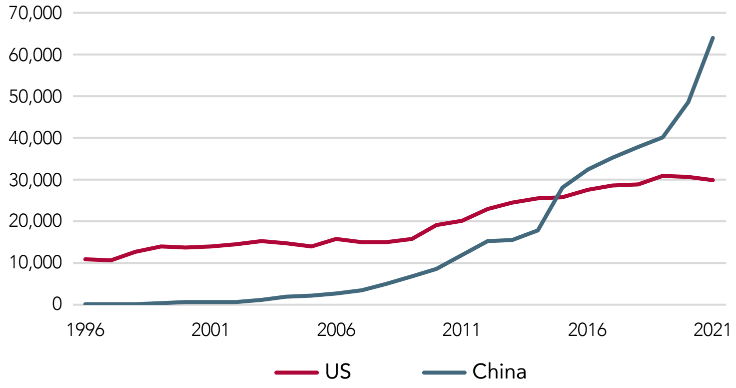

Source: World Intellectual Property Organization, "WIPO IP Statistics Data Center."

on the creation of new ideas. As one metric, consider that China's production of patents reached parity (100 percent of US patent production) around 2015 and (as of 2018) surpasses that of the United States (see figure 11.1). Even allowing for the possibility that historically, Chinese patents might be of generally lower quality than those of the United States, the trend line is clear.

Other nations are also investing heavily in research and development (R&D), which is a critical source of new ideas. According to the National Science Board, most of the world's R&D expenditures occur in a few countries, with the United States accounting for 27 percent of global R&D in 2019, followed by China (22 percent), Japan (7 percent), Germany (6 percent), and South Korea (4 percent).[3] At the same time, the concentration of R&D expenditures continues to shift from the United States and Europe to countries in East, Southeast, and South Asia. This trend is consistent with the observation that an increasing share of the world's patents are shifting over to Asia—particularly to China.[4]

How can the United States take advantage of ideas produced in other countries? One obvious way is to read the scientific and technical literature produced by scientists abroad, and that does happen in abundance. But it is well known that people are a much more effective information transfer mechanism than papers.

For example, informal interviews with Stanford faculty across most of the technology areas yielded two important points regarding the value of direct, in-person interactions with foreign scientific colleagues. First, these interactions enable them to learn things they could not learn simply from reading papers published by the same people, as papers often do not capture vital "tacit knowledge" that enables researchers to build upon the work of others. Second, they are able to develop a much better understanding of the scope and nature of progress made and not made by their foreign colleagues—an understanding that would not result simply from reading the literature.

In fact, the same point—that ideas move much more effectively when conveyed through people than through papers—underscores how the United States actively benefits from foreign scientific input. Skilled immigrants support American innovation today. For example, immigrant college graduates receive patents at double the rate of native-born Americans.[5] In part, this is explained by a higher proportion of immigrant students pursuing STEM education in the United States. Technology companies in the United States also rely on immigration visas to bring foreign scientists and engineers to work in the United States.

Conversely, the United States suffers when foreign scientific inputs are curtailed. For instance, a Harvard Business School study found that pro-migration changes to immigration policy significantly increase innovation within a country—as measured by the production of patents—while changes that discourage immigration lead to significant declines in patent production.[6]

Human talent capable of creating ideas in S&T cannot be manufactured at will. It must be domestically nurtured or otherwise imported from abroad. Today, both paths to growing the requisite talent base to sustain and grow US innovation face serious and rising challenges. Test scores clearly show declining performance in STEM subjects in K–12 education, both in absolute terms and in comparison with other countries.[7]

Regarding US STEM education, the US Department of Defense noted in 2021 that improving the capacity and resilience of the defense industrial base requires more workers trained in the skilled trades and in STEM.[8] Yet bias against careers in the industrial trades among parents and educators has shrunk the pool of potential workers, and adverse demographic trends have led to an aging-out of a skilled workforce with irreplaceable knowledge. The Defense Department also noted the dearth of trained software engineers working on classified projects was in part because they must be US citizens.

At the same time, US immigration policies discourage or prevent foreign S&T talent from working here. Current policies are facilitating a shift of skilled immigration and associated multinational R&D investment toward other countries. Tightening immigration also can prevent companies from hiring enough skilled workers to operate their R&D facilities, increasing their incentives to relocate abroad.[9]

A significant portion of academic researchers are PhD students and professors who have immigrated to the United States to seek better educational and research opportunities. It is crucial to establish a better pathway to permanent residence upon graduation for PhD students on student visas so that the United States does not lose highly trained workers. The United States and universities invest heavily in the education of STEM graduate students, and it would be wise to find a path to allow these scientists and engineers to work and live in the country permanently. Furthermore, if ambitious goals in building up the semiconductor industry, biotechnology, or decarbonization are to be met, then increased investment is needed in the labor force. These industries hire highly trained workers who have advanced degrees. Research funding supports not only the scientific outcomes but also an essential method for training highly skilled engineers and scientists.

Finally, it is important to realize that the global talent challenge is not just about China. US allies and partners also compete for technology talent from around the world. For example, Canada has always had an immigrant-friendly policy that attracts foreign-born graduates of US universities—nearly forty thousand such individuals were recruited to Canada from 2017 to 2021.[10] More recently, Canada introduced its Tech Talent recruiting program in June 2023.[11] This program targets tech workers in the United States who hold US H-1B nonimmigrant visas, providing them more favorable terms. A 2022 survey of almost 1,500 global leaders hiring tech professionals in the United States, United Kingdom, Germany, France, the Netherlands, and Sweden found that more than one-third were searching globally.[12]

A report from Korn Ferry Consulting asserts that "human capital is the greatest value creator available to organizations,"[13] also finding that every dollar invested in human capital adds $11.39 to GDP. The same holds true for advances in science and technology more generally: the most important ingredient for S&T advances is human talent because human talent is the goose that lays the golden eggs of ideas and innovation.

None of these comments are intended to suggest that concerns about foreign appropriation of American intellectual efforts are unfounded. But the fact that American S&T efforts are deeply connected to those of the rest of the world is overall an accelerator of those efforts rather than a brake on them. Using an axe to impose blanket restrictions on engaging foreign scientists when a surgical scalpel is needed to curb only the issues that warrant serious concern is a sure way to reduce the effectiveness of US scientific efforts.

The Changing Role of Government regarding Technological Innovation

Takeaway *The US government is no longer the primary driver of technological innovation or funder of R&D.*

Many technological advancements—such as satellites and access to space, the development of jet engines, and the emergence of the semiconductor industry in Silicon Valley—have their roots in US government financial support and advocacy. But in many fields today, the US government is no longer the primary driver or funder of R&D.

Private companies have taken up much of the slack. For example, while the US government once used its own rockets to launch satellites, it now often does so by contracting with companies that provide access to space as a service. These companies, however, may be under the jurisdiction of nations or controlled by senior executives whose interests are not aligned with those of the users of their services. For example, for a several-year period in the mid-2010s, the United States was entirely dependent on Russia to transport American astronauts to the International Space Station. More recently, the Starlink satellite communications network has been an essential part of Ukrainian battlefield communications; however, the CEO of Starlink curtailed Ukrainian access on a number of occasions in ways that affected Ukrainian battlefield strategy.[14] Such concerns are most serious when there is only one or a small number of private-sector providers of the services in question.

Many US officials recognize the growing role of a handful of private actors behind influential innovations in technology. They believe that supporting closer public-private sector cooperation, as well as informed government regulation of emerging technologies, is a pressing imperative. Even if the government cannot rely on its own capabilities to remain at the forefront of technological innovation, it still has an important role to play in funding and promoting R&D, facilitating the broad adoption of key innovations and standards, and convening coalitions of like-minded actors both domestically and internationally.

A Trend toward Increasing Access to New Technologies Worldwide

Takeaway *National monopolies on technology are increasingly difficult to maintain. Even innovations that are exclusively American born (an increasingly rare occurrence) are unlikely to remain in the exclusive control of American actors for long periods.*

Access to technologies such as synthetic biology, robotics, space, and blockchain often spread from rich nation states and large corporations to less wealthy nations, smaller corporations, universities, and individuals. Even innovations that are American born—an increasingly rare occurrence—are unlikely to remain in the exclusive control of American actors for long. Many emerging technologies exhibit a long-term trend of riding a declining cost-curve over time, making them accessible to an ever-larger set of individual actors. Export controls may delay this spread in some cases for a limited time, but the overall trend is toward decentralized access.

This trend has several implications:

- **Greater policy complexity results from more actors.** Other actors, both state and nonstate, will have capabilities to challenge US interests that they did not have before.
- **Technological advantages will diminish.** The United States may have the most technologically advanced capabilities, but even the more rudimentary instantiations of these capabilities available to other actors eliminate monopolies and narrow the relative advantages once enjoyed by the United States.
- **Winning isn't winning anymore.** The old paradigm of "winning" a technological race to achieve gains that last decades and are hard to replicate—traveling to the moon, developing the atomic bomb—will be replaced by the paradigm of constant competition.
- **More diversity in bureaucracy and ethics has consequences.** Actors less subject to bureaucratic and ethical constraints will be able to exploit technology more nimbly and adapt more rapidly to conditions on the ground.

On the other hand, for physical technologies whose effectiveness depends on deployments in large quantities, geography still plays a role. Natural resources such as rare-earth metals are geographically constrained, and production facilities for physical artifacts still matter.

It may be possible to extend periods of American monopoly on certain technologies, but not indefinitely. Such extensions can have valuable short-term benefits, not least because they buy time for US policymakers to better anticipate a world of democratization. But all too often buying time becomes an end unto itself, and actions to craft a better policy—such as improving targeted immigration reform and sustainability—are not taken.

To be sure, there are probably exceptions to this democratization trend. One is the first appearance of an emerging technology. At such a point in time, the democratization process has not yet begun, at least not in full force, and it may indeed be that the technology in question will be characterized by the dominance of a few key actors.

For example, the training of large language models (LLMs) from scratch is a new capability still only belonging to a few large companies, and it is completely out of reach even for large coalitions of top research universities. On the other hand, research is already underway to build applications based on foundational models at much lower cost. In many cases, the approach taken is to fine-tune a foundational model that has already been trained from scratch, thus building on previous efforts. Therefore, it may well be that this particular exception is apparent only because of the quirk of timing, and LLMs will also be further democratized as time goes on.

In other cases, the impact of emerging technologies depends on the scale of deployment. Universities may be able to develop sustainable energy technologies such as better batteries, but they lack the infrastructure to manufacture them at scale, an enterprise that requires enormous investments from the private sector.

Synergies between Different Technologies

Takeaway ***The synergies between different technologies are large and growing as advances in one technology often support advances in other technologies.***

The technologies described in this report span a broad range, but most have in common a synergistic relationship to other technologies. Improvements in technology A can be used to improve the performance of technology B, while improvements in B help C— then C and B together can help improve A.

For example:

- AI contributes to advances in synthetic biology in addressing the protein-folding problem, predicting protein shape from the DNA sequence of base pairs.
- AI helps to screen many candidate compounds to predict the ones most likely to exhibit desirable properties for materials science.
- Materials science is central to the identification of new semiconductors that may be useful in developing more energy-efficient chips, which in turn can reduce the cost of training AI models.
- New materials are important in space research for the construction of spacecraft and satellites.
- New materials are needed to enable the development of neural probes that can send and receive electrical signals in neural tissue.
- Energy technologies help to improve the performance of robotics and spacecraft.
- Synthetic biology can build organisms that produce certain specialized materials.
- Cheaper semiconductors have driven down the cost curve of DNA sequencing, which itself is a fundamental technology for synthetic biology.[15]

Some of the S&T areas in this review—AI, synthetic biology, materials science, and energy—have a foundational flavor impacting a variety of problem or technology domains.[16] Others are better characterized as technology applications—space and robotics, for example—which focus on solving specific problems through an artful blend of a number of technologies. But these are differences of degree, not of kind, and even more foundational technologies can benefit from advances in different technology applications.

Nonlinear Paths from Research to Useful Application

Takeaway *The traditional "linear" model of R&D, in which basic research leads to applied research, which then leads to development and prototyping, which finally leads to novel and useful products or services, is only one model for how societies obtain value for investments in technology innovation.*

The traditional "linear" model of R&D in science and technology starts with basic research leading to applied research, which then leads to development and prototyping, which finally leads to marketable products.

- Basic research is activity aimed at fundamental scientific understanding without any particular applications in mind.
- Applied research is activity to deepen this scientific understanding with an application area or specific problem in mind.
- Development is activity that builds on scientific understanding to construct engineering prototypes and proofs-of-concept.

This model of scientific development has a long history, but it is by no means the only path. Many analysts argue this model is so unrepresentative of how scientific development actually proceeds that it can be harmful to the scientific enterprise.

Other models are less linear in nature; they acknowledge and even exalt the need for feedback between the various activities. For example, some problems or application areas are so challenging that they entail obtaining a deeper fundamental scientific understanding of a real-world phenomenon—what some call "use-inspired basic research." One of the most famous examples is the work of the French chemist Louis Pasteur on milk spoilage, an applied problem that required advancements in fundamental biological science regarding bacterial processes.

In other cases, deep fundamental scientific knowledge may be necessary in technology areas adjacent to the primary problem of interest. New drugs that are effective against cancer tumors do no good if they cannot be delivered in lethal concentration to the tumor. Therefore, research on drug delivery mechanisms is as important here as the development of new anticancer compounds.

The Relationship of Political Regime Type to Technological Progress

Takeaway *Democracies provide greater freedom for scientific exploration, while authoritarian regimes can direct sustained funding and focus to technologies they believe are most important.*

Technological innovation occurs in both democracies and autocracies, but different regime types face different advantages and challenges. True democracies enjoy the rule of law and a free flow of ideas and people, as well as the ability of individuals to pursue research goals of their own choosing. Perhaps most important, because failure in a democracy does not lead to persecution or necessarily result in professional ostracism, individuals are freer to experiment and explore.

Authoritarian regimes are more aptly characterized by the rule of the state, or the whim of a single "supreme leader." This can lead to the constrained flow of ideas, coercion to limit individual freedom of action and thought, and top-down direction to explore only topics of interest to the state. In this environment, failure may carry very high consequences for individuals. Under such circumstances, it would be understandable if individual scientists limit themselves to studying "safe topics" for which failure to make progress is unlikely.

Yet it must be noted that authoritarian regimes have the advantage of being able to direct funding and public attention to problems that they believe are important. They can sustain that focus for long periods of time more independently of short-term considerations such as profit or politics. To the extent that technology-based solutions are known, authoritarian leaders can exploit that knowledge to implement those solutions, regardless of any downsides.

Successful innovation requires both an exploration of the relevant space of possible solutions, to eliminate paths that will not lead to viable solutions, and an exploitation of viable solutions to focus resources on specific problems deemed important to the regime.[17] Competitors such as China have taken advantage of US scientific exploration in many domains through means both legal and illegal and have gone on to exploit that knowledge through a variety of commercial and military efforts.

Attempts to obtain some of the benefits of a more centralized direction to the technology policy efforts of the United States have been described as steps in the direction of adopting an industrial policy. Critics often argue that such efforts unduly interfere in a free market and that picking "winners and losers" leads to inefficiencies. Advocates argue that only through such action will the United States be able to offset some of the advantages that authoritarian nations would otherwise enjoy over it. The public policy problem is acknowledging some truth in both perspectives and seeking an appropriate balance consistent with both American values and economic competitiveness.

Punctuated Technological Progress

Takeaway *Technology often progresses in fits and starts, long periods of incremental results are followed by sudden breakthroughs, and the speed of change is hard even for leading researchers to anticipate.*

Taken as a whole, technological progress exhibits a variety of patterns. Some technologies have demonstrated consistent progress for extended periods. For example, semiconductor technology is characterized by Moore's law, an exponential reduction in the cost of semiconductors over time. Solar cells and LED efficiency have followed similar cost reduction curves.

Other technologies have demonstrated much more uneven rates of progress. These technologies see long periods of incremental development and refinement that are punctuated by short bursts of radical innovation. In some cases, the bursts are the result of some particular breakthrough—examples might be the emergence of the personal computer in the 1980s or the World Wide Web in the early 1990s. In other cases, the bursts are due to the simultaneous availability and maturity of several key technologies that are required to make significant progress in some other technological domain. Here, an example might be electric cars, where battery technology, lightweight materials, sensors, and computing power have come together to make such cars more economically feasible.

When punctuated progress characterizes a technology, forward projections of progress based on past rates may well be misleading. Successful forecasting depends on familiarity with a wide variety of technologies precisely because it is hard to predict which specific technologies will prove critical. Indeed, even experts in a given field can be surprised by the rapidity of progress, as has happened in the last year with artificial intelligence and applications such as ChatGPT and large language models. Geoffrey Hinton, a 2018 Turing Award winner for his work on artificial intelligence, said, "I have suddenly switched my views on whether these things are going to be more intelligent than us. I think they're very close to it now and they will be much more intelligent than us in the future."[18] This sentiment is shared by other experts in LLMs, who have told us they too have been stunned by the speed of advances in their own field.

While there is broad recognition that we may be at the cusp of a moment of radical technological change across a number of fields (AI, synthetic biology, nuclear fusion), the precise contours, speed, and implications of this moment are much harder to ascertain.

Nontechnological Influences on Technological Innovation

Takeaway *Technology applications in society require scientific proof-of-concept, engineering feasibility, economic viability, and societal acceptability.*

Technology plays a key role in supporting and advancing national interests, and headlines in the news often tout scientific breakthroughs that offer opportunities to solve societal problems and to improve the quality of life. But to play a valuable role, any given technology application must demonstrate not only technical feasibility but also economic viability. It must also be acceptable to the relevant constituencies, including the public at large. People and organizations must be able to adapt to its use, despite the disruption it may cause. The requirements and burdens imposed by law, policy, and regulation must be compatible with widespread adoption and use of the application.

There is often a large gap between a demonstration of scientific feasibility—let's call it Q—and a product or a service based on Q that is useful to society. Press reports of scientific breakthroughs often give the impression that useful exploitations of these breakthroughs are just around the corner. That is almost never true. Scientific feasibility is a necessary prerequisite, but it may well be that engineering or economic feasibility does not follow. After scientific proof-of-concept is achieved, engineering feasibility must be demonstrated, which includes considerations of cost and ease of use.

Scientific proof-of-concept is only the first step. Engineering feasibility must also be demonstrated, which includes considerations of cost and ease of use. Take the case of the technical success of early attempts to build supercomputers out of superconducting components that required liquid helium to cool them. Technical feasibility was demonstrated, but the liquid helium requirement would make these computers difficult to deploy and use in practice; alternative computing technologies appeared to offer comparable performance at lower cost.

In other cases, engineering feasibility can be demonstrated but cost considerations must first be resolved. For example, when carbon fibers were first being investigated in the laboratory, they cost $10 million per pound—clearly infeasible for large-scale use.[19] A substantial amount of work was required to reduce the cost by what is today a factor of a million.

In still other cases, it may prove too difficult to develop a manufacturing process to build a device based on Q, or the materials used to demonstrate Q are too expensive or rare to support large-scale production. Less expensive or cumbersome alternatives to devices based on Q may be available, thus reducing the marketplace viability of Q-based devices.

Societal acceptability matters as well. In Europe, though much less so in the United States, genetically modified organisms as food are highly controversial, and concerns over their safety have prevented the uptake of GMO foods consumed widely in the United States. The psychology of individuals and cultural practices and beliefs of a community or society also contribute to the adoption and use of any given technology application. The essential point here is that technology in society is not just about the technology.

Lastly, given that some technological demonstrations of scientific feasibility do not advance to the marketplace and become "orphaned," an important public policy question is how to manage them.

For example, a start-up company may be established to commercialize Q. If the company fails in the marketplace for economic reasons, a competitor or another nation with a different cost structure may be able to make it economically viable—and that competitor's interests may not align with those of the United States. What if the competitor is a bad actor and simply buys the now-defunct start-up, thereby acquiring the rights to the intellectual property underlying Q?

The Role of Universities in Tech Innovation

Takeaway *US universities play a pivotal role in the innovation ecosystem that is increasingly at risk.*

The US infrastructure for funding and conducting R&D (to which innovation is closely related) is broad and deep. For example, the private sector is the second-largest supporter of R&D in the United States (the first is the federal government).[20] Entities such as Bell Laboratories, IBM's Thomas J. Watson Research Center, and Xerox PARC once performed substantial amounts of basic scientific research. Today, their present-day equivalents focus most of their R&D efforts on process and product development closely related to the bottom line of their parent companies. This focus has two important implications. First, companies tend to focus their efforts on research with foreseeable commercial application, not frontier or fundamental research where the connection between breakthroughs and application may not be apparent and where it may take years, if not decades, for a technology to mature. Yet companies depend on nonindustrial research. For example, more than 80 percent of the algorithms used today (not just in AI but in all kinds of information technology) originated from sources other than industrial research.[21] Second, corporate R&D outputs tend to be restricted and proprietary to preserve any market-competitive edge that they may afford to the company that paid for them.

The federal government also operates a large number of laboratories and federally funded research and development centers. For example, the Department of Energy operates seventeen national laboratories for conducting research and development that serve the department's core missions in energy, science, national security, and environmental stewardship.[22] These labs specialize in particularly difficult problems that fall beyond the capabilities of private industry or individual universities. (The fusion breakthrough described in chapter 6 was conducted at the Lawrence Livermore National Laboratory.) Nor is the Department of Energy unique in having mission-driven laboratories that actually conduct research and development work; the Department of Defense, NASA, the National Institutes of Health, and the Department of Commerce are just a few of the departments with their own mission-driven labs.

But the role of universities is unique and pivotal in the innovation ecosystem, and this role is often underappreciated. For example, in contrast to research done in mission-driven federal laboratories, the scope and breadth of research conducted at major research universities tends to be much larger than that conducted at federal labs simply because the various foci of university research are not constrained by particular missions.

In contrast to the private sector, university research is almost always open, enabling would-be innovators to take advantage of it. Open research promotes transparency, accountability, collaboration, reduced duplication, and wider impact. By sharing study details, data, and results openly, researchers allow others to verify, replicate, and build on their work more easily. This accelerates discovery as more minds can work on problems in a collaborative way, with reduced redundant efforts. Published open research also reaches more people, magnifying its educational and societal benefits.

The role of universities in building the national economy has been recognized since 1862 with the passage of the Morrill Act establishing land-grant universities. Government-supported university research made a critical difference in World War II on technologies such as radar, proximity fusing, computers—and the atomic bomb. University research has since generated knowledge whose exploitation creates new industries and jobs, spurs economic growth, and supports a high standard of living, while achieving national goals for defense, health, and energy.[23] University research has been a rich source of new ideas, particularly for the longer term, and universities are the primary source of graduates with advanced S&T skills.

Universities have the mission of pursuing high-risk research that may not pay off in commercial or societal applications for a long time, if ever.[24] For example, research in number theory—a branch of pure mathematics—was undertaken for decades before it became foundational to modern cryptography. In the 1960s, academic research on perceptrons sought to develop a computational basis for understanding the activity of the human brain. Although this line of research was abandoned after a decade or so, it ultimately gave rise to the work in AI on deep learning several decades later. The term "mRNA vaccines" entered the public lexicon in 2021 when COVID-19 vaccines were released. Yet development of these vaccines was built on university research with a thirty-year history.[25] Magnetic resonance imaging (MRI) was first discovered in university studies in the 1940s, but it took another three decades of research, much of it university based, for the first medical MRI imagers to emerge.

Finally, the increasingly blurred distinction between fundamental research and export-controlled research is creating challenges in academia in fostering international collaboration, particularly in fields such as the semiconductor industry, nanotechnology, AI, and neuroscience. Some researchers are concerned that fundamental research could now be considered export controlled and may steer clear of foreign collaborations out of an abundance of caution. This policy ambiguity can deter collaborations and create obstacles for non-US researchers wishing to contribute to work in the United States. This is particularly concerning as international cooperation could expedite progress in emerging fields like nanomaterials, where countries like Korea are making significant strides, especially in biomedical applications. These policy issues, widely recognized among the research community, underscore the urgent need for clarification and reform to advance research and promote effective international collaborations.

The Structure of Research and Development Funding and the Valley of Death

Takeaway *Sustaining American innovation requires long-term government R&D investments with clear strategies and sustained priorities, not wild swings from year to year, which is increasingly common.*

Budget is one obvious aspect of government funding for R&D. But three other aspects deserve as least as much attention. First, government has an important role in funding important research with long horizons, as industry is not generally structured to support long-term R&D efforts. Such government-funded research should generally be regarded as precompetitive in nature.

Second, wide swings in funding from year to year—increasingly common in government funding—are antithetical to a systematic R&D effort. In a free market economy, talented scientists can choose where and in what domains to work, and they have a natural aversion to work environments that do not provide stability. Therefore, wide swings in funding have the effect of driving away the scientific talent that can best find employment in that field elsewhere.

Third, the so-called valley of death remains a significant problem. This refers to the period after research has demonstrated the engineering feasibility of a particular innovation—a step beyond scientific feasibility—but before the innovation achieves adoption on a scale large enough to establish the viability of a business model based on the innovation.

When an innovation is first offered to customers, it is expected to provide new functional capability. Its cost matters as well, as the new functional capability may or may not be worth the price of adopting it. At one cost, a potential customer might choose not to acquire it, while at a lower cost, that customer may well do so.

If the initial cost is too high, customers will be scarce and the firm producing or providing the innovation is likely to fail commercially if it does not receive funding from a source not related to production to stay afloat. But it often happens that the per-unit production cost will decrease as the total cumulative volume of production increases. Known as the learning curve in manufacturing, this phenomenon is primarily due to the efficiencies and knowledge gained from repetitive production processes. Such cost reduction is particularly important when, as is true for energy production, significant societal benefits accrue and new technologies are deployed at scale.

Research funding generally disappears after feasibility has been demonstrated. For a firm then to get through the valley of death, it must either demonstrate the viability of its business model to investors who believe in the promise of the innovation or attract enough customers on its own to sustain it. True commercial viability is unlikely to start until the per-unit cost has dropped to levels affordable by most would-be customers.

While in the valley of death, it is typical that no party is willing to invest the minimum level for product or manufacturing refinement to continue, and projects often have to stop or progress much more slowly than before. In some cases, the innovation never scales beyond the initial stages, regardless of its technical sophistication or desirability.

These points suggest bridge funding may help in establishing commercial viability. The sticking point, however, is the difficulty of distinguishing between real innovations that would be truly valuable if only they could get through the valley of death and ersatz innovation look-alikes for which valley-of-death concerns are merely a smokescreen to cover up their genuine inadequacies and problems in the face of market realities.

A firm's failure to pass through the valley of death may also have competitive implications internationally. Such a firm is ripe for acquisition by foreign competitors with deeper pockets who may be willing to invest in innovative products that have not yet reached market viability. Chinese investors, for example, were successful in acquiring Atop Tech, a firm with an automated designer capable of producing high-end microchips, after it went bankrupt in 2017. This transaction failed to elicit any reaction from the Committee on Foreign Investment in the United States (CFIUS), despite its mandate to review and, if necessary, block certain transactions involving foreign investment that may impact US national security. The Foreign Investment Risk Review Modernization Act of 2018 was enacted in part to improve the ability of CFIUS to review just such transactions.

A new funding model, known as focused research organizations (FROs), seeks to fill the gap inherent in the valley of death. The FRO provides funding to assemble scientists and engineers with the required expertise to rapidly prototype and test materials and technologies for their applications. One initiative to support FROs was launched in 2021, Convergent Research, a nonprofit organization with the mission of incubating and funding new FROs. In March 2023, it received $50 million in philanthropic donations to launch two new FROs.[26]

Cybersecurity

Takeaway *Researchers working in highly competitive environments who neglect cybersecurity place their research progress at risk.*

Cybersecurity refers to technologies, processes, and policies that help to protect computer systems, networks, and the information contained therein from malicious activities undertaken by adversaries or unscrupulous competitors. It is often believed that cybersecurity is an issue that primarily affects private-sector businesses and government, but the world of academic R&D faces a variety of cyber threats as well.

One important cybersecurity interest is in ensuring the integrity of data. Scientific experiments produce data, and if important data are deleted or destroyed, scientific progress can be significantly retarded. A possibly more worrisome scenario is that the data are altered in hard-to-detect ways that subtly and invisibly skew the subsequent work based on that data, possibly putting scientific investigators on the wrong track and wasting significant effort. Adversaries or competitors seeking to delay scientific work have significant incentives to engage in activities that could compromise data in this manner.

Similar comments apply to the computer programs used to analyze data. If a computer program is maliciously altered in a subtle way, it may be a long time before the alteration is noticed. Once such an alteration is noticed, all previous analyses performed using that program are inevitably called into question.

A second cybersecurity interest is in ensuring the confidentiality of various work products, such as datasets and working papers. Datasets may have been collected under promises of confidentiality or nondisclosure agreements, and unauthorized access to such datasets quite possibly violates such promises or agreements. Premature disclosure of working papers can compromise claims of priority, an important currency in which academic R&D trades. Additionally, draft working papers are often incomplete, inconsistent, or downright wrong and are not in any sense defensible—premature disclosure of such papers as though they did in fact reflect completed work is a nightmare of any scientist.

Many laboratories rely on computers to control or supervise data collection from various instruments. Compromising these computers through a cyberattack could cripple data collection efforts or corrupt the data being collected. The instruments in question could also be damaged by hacking the controlling computers.

Technical safeguards are available for most cybersecurity problems, of which the above are just a sample. But especially in academic laboratories, maintaining and operating such safeguards consistently calls for a serious management effort to impose the necessary discipline on all those working in those labs. Such discipline often conflicts with informal laboratory cultures that stress collegiality, openness, and flexibility.

A second cyber-related threat to the R&D enterprise is selective targeting of key personnel working on important research projects. It is a matter of public record that a number of Iranian nuclear scientists have been killed since 2007, reportedly because they were associated with Iran's nuclear program.[27] But assassination is not the only form of targeting. Much less violent forms of targeting could involve what might be broadly termed harassment, which often originates in or is perpetrated through cyberspace. For example, compromising the personal life of a principal investigator (e.g., draining bank accounts, interfering with the investigator's personal finances, threatening the investigator's family) can all be accomplished through the internet. Dealing with such matters will inevitably reduce the work effectiveness of the individuals targeted in this manner. Calling into question the investigator's professional conduct and ethics is another approach that could have comparable effectiveness.

NOTES

1. Science Museum, "The Iron Lung," October 14, 2018, https://www.sciencemuseum.org.uk/objects-and-stories/medicine/iron-lung.

2. Keith Sawyer, *Group Genius: The Creative Power of Collaboration* (New York: Basic Books, 2007).

3. Amy Burke, Abigail Okrent, and Katherine Hale, *The State of U.S. Science and Engineering 2022* (Arlington, VA: National Science Board, National Science Foundation, 2022), https://ncses.nsf.gov/pubs/nsb20221.

4. World Intellectual Property Organization, "World Intellectual Property Indicators: Filings for Patents, Trademarks, Industrial Designs Reach Record Heights in 2018," October 16, 2019, https://www.wipo.int/pressroom/en/articles/2019/article_0012.html.

5. Jennifer Hunt and Marjolaine Gauthier-Loiselle, "How Much Does Immigration Boost Innovation?," *American Economic Journal: Macroeconomics* 2, no. 2 (April 2010): 31–56, https://doi.org/10.1257/mac.2.2.31.

6. Dany Bahar et al., *Talent Flows and the Geography of Knowledge Production: Causal Evidence from Multinational Firms*, Harvard Business School, Technology and Operations Management Unit, Working Paper no. 22-047, last modified December 23, 2022, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4005693.

7. Amy Burke, Abigail Okrent, and Katherine Hale, "U.S. and Global STEM Education and Labor Force," in *The State of U.S. Science and Engineering 2022* (Arlington, VA: National Science Board, National Science Foundation, 2022), https://ncses.nsf.gov/pubs/nsb20221/u-s-and-global-stem-education-and-labor-force.

8. US Department of Defense, *Industrial Capabilities Report to Congress: 2020 Annual Report*, January 2021, https://www.businessdefense.gov/docs/resources/USA002573-20_ICR_2020_Web.pdf.

9. Britta Glennon and David R. Dollar, "What's Behind the Globalization of R&D?," April 26, 2021, in *Dollar & Sense*, podcast produced by Brookings Institution, MP3 Audio, 25:02, https://www.brookings.edu/articles/whats-behind-the-globalization-of-rd.

10. Cecilia Esterline, "Previously Unreported Data: The U.S. Lost 45,000 College Grads to Canada's High-Skill Visa from 2017 to 2021," Niskanen Center, March 14, 2023, https://www.niskanencenter.org/previously-unreported-data-the-u-s-lost-45000-college-grads-to-canadas-high-skill-visa-from-2017-to-2021.

11. Immigration, Refugees, and Citizenship Canada, "Canada's Tech Talent Strategy," Government of Canada, last modified June 27, 2023, https://www.canada.ca/en/immigration-refugees-citizenship/news/2023/06/canadas-tech-talent-strategy.html.

12. Clare McDonald, "Quarter of UK Businesses Looking Abroad for Tech Talent," Computer Weekly, December 14, 2022, https://www.computerweekly.com/news/252528353/Quarter-of-UK-businesses-looking-abroad-for-tech-talent.

13. Michael Distefano et al., *The Trillion-Dollar Difference*, Korn Ferry Institute, 2016, https://www.kornferry.com/content/dam/kornferry/docs/article-migration/Korn-Ferry-Institute_The-trillion-dollar-difference.pdf.

14. Adam Satariano et al., "Elon Musk's Unmatched Power in the Stars," *New York Times*, July 28, 2023, https://www.nytimes.com/interactive/2023/07/28/business/starlink.html.

15. James M. Heather and Benjamin Chain, "The Sequence of Sequencers: The History of Sequencing DNA," *Genomics* 107, no. 1 (January 2016): 1–8, https://doi.org/10.1016/j.ygeno.2015.11.003.

16. Jean-Lou Chameau, William F. Ballhaus, and Herbert Lin, eds., *Emerging and Readily Available Technologies and National Security: A Framework for Addressing Ethical, Legal, and Societal Issues* (Washington, DC: National Academies Press, 2014), https://doi.org/10.17226/18512.

17. James G. March, "Exploration and Exploitation in Organizational Learning," *Organization Science* 2, no. 1 (1991): 71–87, http://www.jstor.org/stable/2634940.

18. Will Douglas Heaven, "Geoffrey Hinton Tells Us Why He's Now Scared of the Tech He Helped Build," *MIT Technology Review*, May 2, 2023, https://www.technologyreview.com/2023/05/02/1072528/geoffrey-hinton-google-why-scared-ai.

19. American Chemical, "High Performance Carbon Fibers: National Historic Chemical Landmark," September 17, 2003, https://www.acs.org/education/whatischemistry/landmarks/carbonfibers.html. The original publication is by Roger Bacon, "Growth, Structure, and Properties of Graphite Whiskers," *Journal of Applied Physics* 31, no. 2 (February 1960): 283–90, https://doi.org/10.1063/1.1735559.

20. Congressional Research Service, *Federal Research and Development (R&D) Funding: FY2022*, January 2022, https://crsreports.congress.gov/product/pdf/R/R46869.

21. Neil C. Thompson, Shuning Ge, and Yash M. Sherry, "Building the Algorithm Commons: Who Discovered the Algorithms That Underpin Computing in the Modern Enterprise?," *Global Strategy Journal* 11, no. 1 (2021): 17–33, https://doi.org/10.1002/gsj.1393.

22. National Laboratories, "What We Do," accessed September 19, 2023, https://nationallabs.org/our-labs/what-we-do.

23. National Research Council, *Research Universities and the Future of America: Ten Breakthrough Actions Vital to Our Nation's Prosperity and Security* (Washington, DC: National Academies Press, 2012), https://doi.org/10.17226/13396.

24. National Academies of Sciences, Engineering, and Medicine, *Information Technology Innovation: Resurgence, Confluence, and Continuing Impact* (Washington, DC: National Academies Press, 2020), https://doi.org/10.17226/25961.

25. Elie Dolgin, "The Tangled History of mRNA Vaccines," *Nature*, October 22, 2021, https://www.nature.com/articles/d41586-021-02483-w.

26. Alex Knapp, "Why Billionaires Ken Griffing and Eric Schmidt Are Spending $50 Million on a New Kind of Scientific Research," *Forbes*, March 17, 2023, https://www.forbes.com/sites/alexknapp/2023/03/17/why-billionaires-ken-griffin-and-eric-schmidt-are-spending-50-million-on-a-new-kind-of-scientific-research.

27. Mehdi Jedinia, "History of Assassinations of Iran's Top Nuclear Scientists," Voice of America, December 03, 2020, https://www.voanews.com/a/extremism-watch_history-assassinations-irans-top-nuclear-scientists/6199135.html.