

Making Technological Miracles

Mark P. Mills

Bill Gates wants a miracle. And he's willing to put his prestige and a lot of money—as much as \$2 billion—behind that pursuit. He has called for a major U.S. and global effort to stimulate research in the hope of finding a “miracle” in the science and technologies of energy production.

Why “miracle”? The motivation for Gates's campaign is his conviction that a global energy transformation of an enormous size and scope is needed to address the prospect of climate change. Hydrocarbons—oil, natural gas, and coal—currently supply close to 85 percent of global energy, and mainstream forecasts see world energy use nearly doubling, not shrinking, in the coming decades. Even many forecasts rooted in bullish expectations for growth in alternative energies concede that hydrocarbon use will still grow rather than shrink. No existing technologies can move us away from hydrocarbons on a global scale; there are no quick, easy solutions.

By “miracle,” Gates means something that seems impossible given our current technological and scientific vantage. “I've seen miracles happen before,” Gates says. “The personal computer. The Internet. The polio vaccine.” Indeed, the marvels that technology has enabled and the wonders that science has explained throughout modern history, and especially over the past century, spur confidence that comparable miracles remain undiscovered and uninvented. According to Gates, such miracles are the “result of research and development and the human capacity to innovate,” rather than “chance.” This is inherently a “very uncertain process,” Gates maintains, for which there is no “predictor function.” To help the process along, Gates has not only launched a persuasion campaign but also created the Breakthrough Energy Coalition, a group of wealthy individuals coordinating their investments and philanthropic giving to develop energy alternatives to hydrocarbons. In addition, he has recommended that government spending on energy research and development (R&D) be doubled or even tripled. And how should that government R&D money be spent? Rather than accelerating or seeking to scale up yesterday's inventions, Gates says he'd “spend it all on fundamental research.”

Mark P. Mills, a physicist, is a senior fellow at the Manhattan Institute and a faculty fellow at the McCormick School of Engineering, Northwestern University. He is also a strategic partner in a boutique energy-tech venture fund. In 2016, he was named “Energy Writer of the Year” by the American Energy Society.

SPRING 2017 ~ 37

Without regard to the climate change debate, and long before any miracles happen, though, Gates will have made a vital contribution to the public discourse by raising broader questions about science and technology. His call for miracles gives us an opportunity to think afresh and more broadly about the role that fundamental, curiosity-driven science plays in technological innovation. As we shall see, public discussions of these questions are plagued with confusion and fallacies—and even the very way we talk about these matters, using terms like “basic research” and “development,” is woefully in need of change.

Revolutionizing Energy?

In order to tackle the broader questions, it will be useful first to take a closer look at the energy domain. As noted above, hydrocarbons supply about 85 percent of the energy the world uses today. (For the United States, the figure is similar: hydrocarbons supply just over 80 percent of the energy consumed.)

For decades now—first because of pervasive fears about peak supplies of oil and natural gas, but now because so much of both are used—the United States and many other nations have pursued programs to replace hydrocarbons. The physical and economic realities inherent in that pursuit are what prompted Gates to say:

We need innovation that gives us energy that’s cheaper than today’s hydrocarbon energy, that has zero [carbon dioxide] emissions, and that’s as reliable as today’s overall energy system. And when you put all those requirements together, we need an energy miracle.

In order to appreciate Gates’s stance on transforming the world’s energy systems, one need not take a position on the urgency or severity of climate change. One need only look at the record of spending on new energy technologies. By studying the pursuit of radical technological change in the energy domain, we can better understand the challenges associated with “miraculous” technological transformations in general.

From 2009 to 2014, the U.S. government spent over \$150 billion in various forms of subsidies, grants, and research to support and advance “clean tech”—technologies aimed at replacing or reducing the use of hydrocarbon-based energy sources. Unsurprisingly, the production of energy from sources favored with so much federal largesse—solar, wind, and biofuels—has risen. Those three sources combined have grown from supplying just 2.5 percent of total U.S. energy consumed in 2009

to 5.1 percent in 2016. Meanwhile, over the same period, the growth in shale hydrocarbons—which did not enjoy subsidies or special federal programs—added 700 percent more energy to American production than did solar, wind, and biofuels combined. The explanation for this disparity: technology has advanced faster for shale than it has for solar (or wind or biofuels, for that matter). This wasn't the outcome many forecast a decade ago.

Vinod Khosla, a prominent Silicon Valley venture capitalist who invested hundreds of millions of dollars in biofuels for transportation, spoke for many when he asserted less than a decade ago: “I have no doubt that 100 percent of our gasoline use can be displaced in the next 25 years.” Thanks to tens of billions of dollars in subsidies across many years, about 40 percent of America's corn harvest is now distilled into ethanol and used for fuel. But even so, farming still supplies only about 5 percent of the energy needed for domestic (never mind global) transportation. Even if we used the entire American corn harvest to fuel cars, we would not come close to achieving what Khosla hopes for. As for biodiesel, despite subsidies and federal and private investments, this class of liquid fuel technology provided less than one-tenth of one percent of the energy consumed by the U.S. transportation sector in 2016.

Or consider the prospect of replacing gasoline with wind-generated electricity to charge batteries in electric cars. Here, too, there are physics-based barriers to innovation. Building a single wind turbine, taller than the Statue of Liberty, costs about the same as drilling a single shale well. The wind turbine produces a barrel-equivalent of energy every hour, while the rig produces an actual barrel every two minutes. Even though the barrel-equivalent of energy from a wind turbine costs about the same as a barrel of oil, the latter is easy and cheap to store. However, storing wind-generated electricity so that it can be used to power cars or aircraft requires batteries. So while a barrel's worth of oil weighs just over 300 pounds and can be stored in a \$40 tank, to store the equivalent amount of energy in the kind of batteries used by the Tesla car company requires several tons of batteries that would cost more than several hundred thousand dollars. Even if engineers were able to double or quadruple battery efficacy, that still would not come near to closing the performance gap between energy from wind and energy from liquid hydrocarbons for transportation.

These stark facts often elicit the response that the alternative technologies will get better with time and scale. Of course they will. But there are no significant scale benefits left, since all the underlying materials

(concrete, steel, fiberglass, silicon, and corn) are already in mass production. Nor are there big gains possible in the underlying technologies given the physics we know today.

From 1980 to about 2005, wind and solar tech underwent huge gains in core efficiencies that drove costs down some tenfold. But gains since 2005 have been just a fraction of what they had been historically. Expert forecasts that show even smaller incremental gains in the future are not the result of pessimism, but a recognition that wind and solar technologies are now confronting the inevitable law of diminishing returns as they approach physical limits. Wind turbines are constrained by the Betz limit, a physical principle that shows that no more than about 60 percent of air's kinetic energy can be captured. Modern turbines can already reach 40 percent conversion efficiency. The Shockley-Queisser limit defines how much of the energy in photons can be converted into electricity by a photovoltaic cell: 34 percent. The recent announcement of a silicon cell with 26 percent efficiency shows that we are nearing that boundary too. While scientists are finding new non-silicon options for solar cells (such as the exotic-sounding perovskites), they offer incremental, not revolutionary, cost reductions, and all have similar physics boundaries.

Such considerations led Google engineers to reach a conclusion similar to Gates's regarding the state of technologies needed to replace hydrocarbons at scale. In 2007, Google launched a project aimed at developing renewable energy that would be cheaper than coal; in 2014, the company shut it down. The lead engineers explained the decision as follows: "Incremental improvements to existing [clean energy] technologies aren't enough.... We don't have the answers. Those technologies haven't been invented yet."

The anemic progress in alternative energy technologies runs counter to what many experts, policymakers, and investors expected to happen given the massive injections of private capital and generous federal and state support. Still, we hear policymakers, investors, and pundits calling for yet more spending on engineering and development, on business stimulus and subsidies. Such proposals often use the Silicon Valley buzzword "disruption," and draw analogies to prominent technological breakthroughs of yore—the Manhattan Project, the Apollo program, the rapid evolution of computing, the displacement of landline telephony by cellular technology.

The technological achievements of the space program or even Silicon Valley may indeed appear miraculous. But do those analogies make sense with regard to the prospects for achieving the fundamental, radical

transformations—the miracles—that so many seek in the energy domain, or in any other domain for that matter?

Basic Research, from Bacon to Bush

The proposition that science and technology (and hence research and development) are integral to a flourishing society is a truism endorsed by politicians of all stripes—and for good reason. As economic historian Joel Mokyr has pointed out, technological innovation offers policymakers the closest thing there is to a “free lunch.” Robert Solow was awarded a Nobel Prize in economics for his work documenting the centrality of technological progress in economic growth. Going back to early modernity, Francis Bacon was among the first thinkers to envision systematically the ways that society could be enriched by the miraculous results of material progress, in, among other places, his posthumously published story “New Atlantis” (from which this journal takes its name).

Like Francis Bacon, Bill Gates envisions a future with miracles yet to be discovered. Gates is not looking for mere economic growth—though that is vital—but a revolution in energy, a true “disruption.”

Of course, established businesses can be disrupted by new entrants to the market who are able to use existing technologies more effectively—think Uber and Amazon—with significant economic consequences for consumers and the economy as a whole. But these disruptions are not miraculous in Gates’s sense. The kind of miracles that Gates seeks are not to be found in the corporate R&D departments of Google, GM, Apple, or Monsanto, nor in the R&D budget of the U.S. Department of Energy. More efficient internal combustion engines or batteries, better photovoltaic cells or computer processors, self-navigating drones or disease-resistant crops—these are all valuable, even amazing. But they merely modify, extend, or apply existing technological frameworks; they do not leap over the miracle hurdle Gates describes.

So where are we to look? Do we get more miracles by focusing more on science or technology? Or, put in the budgetary language of government and industry, do we fund more research or more development? It is not enough to wait for miracles to emerge organically. There are often compelling practical reasons—diseases, terrorism, environmental threats—to kickstart this process. By calling for more *basic* research, Gates has put his finger on the critical issue for every domain where a miracle would be welcome, not just energy but also medicine, agriculture, transportation, and security.

Gates began pushing for a greater focus on basic research in 2015, coincidentally the seventieth anniversary of *Science, The Endless Frontier*, a government report written by Vannevar Bush. In that report, Bush, then the director of President Truman's Office of Scientific Research and Development, established the architecture for modern federal R&D policy, which remains in place today. "The Government," wrote Bush in a phrase reminiscent of Bacon, "has only begun to utilize science in the nation's welfare."

Perhaps more than ever before, any attempt to garner political and financial support for research and development must genuflect to "utility." And we spend a lot of money on research and development. Globally, the annual total of government and business spending on all forms of research and development is roughly \$1.8 trillion. In the United States—the world's biggest player—that total is more than \$450 billion per year, while some estimates put China's total R&D spending at nearly \$400 billion. Corporations and policymakers expect something back for all that money.

Accordingly, the data show that in today's research budgets there is a huge skewing in preference for the "D" in "R&D." Only around 7 percent of business R&D spending in the United States is focused on basic research; a larger portion goes to applied research (16 percent) and the lion's share goes to development (77 percent). And business accounts for 65 percent of all R&D spending.

Government R&D spending is also heavily biased toward the practical and the near-term. In fiscal year 2013, the Department of Defense accounted for about half of the total federal spending on R&D—and 91 percent of that DOD R&D spending went into development, with just 3 percent going into basic research. If you exclude DOD and look only at the remaining half of federal spending on R&D, 45 percent of it went to basic research in FY2013, with the other 55 percent going to development and applied research.

The low ratio of spending on basic research to spending on development and applied research runs counter to many scientists' views about the importance of basic research. As theoretical physicist and cosmologist Paul Davies puts it: "The reason that we do basic science is to understand how the universe works, and what our place is within the universe. It is a noble quest." There is nobility in our pursuing scientific knowledge every bit as much as literature, history, art, and music. True, scientists have long recognized that their benefactors expect utility in exchange for helping to advance basic knowledge. In the nineteenth century, the great

physicist Michael Faraday discovered that a magnetic field could induce an electric current in a wire. According to an apocryphal story, William Gladstone, then the Chancellor of the Exchequer, asked Faraday what use his discovery could possibly have. “Why sir,” Faraday is supposed to have quipped, “there is every possibility that you will soon be able to tax it!” The trick is to thread the needle so as to pursue both the noble and the useful. And there are myriad historical examples of scientists successfully threading that needle, from Archimedes—who, legend has it, put his knowledge to use inventing weapons to help defend Syracuse against the Roman siege—to Alan Turing, the mathematician who was put to work during World War II breaking the Nazi code.

Vannevar Bush was not wrong to defend the importance of curiosity-driven research, nor to emphasize the productive relationship between science and technology. The problem with his view of research and development—often called the “linear model” of technological innovation, according to which scientific research is pursued for its own sake but with an eye to eventual technological application—is that it becomes all too easy to shoehorn scientific research into a utility-driven framework, one in which research is only deemed valuable insofar as it is immediately useful to society. The legacy of the Bush model is that scientific inquiry must ultimately be steered toward a practical goal if it is to have any social utility (and thus receive public, never mind corporate, funding). The noble may be worth pursuing for its own sake, but it seems increasingly impractical in a budget-constrained world in which crises abound. Curiosity-driven projects, such as a mission to Mars or a new class of microscope, have to compete with demands to fund purpose-driven projects, such as treatments for diseases or alternative sources of energy. The temptation grows almost irresistible to cut out the curiosity-driven side of the linear model—the “useless” or, at any rate, not immediately or obviously useful part of research—leaving us only with goal-directed research.

Yet there is good evidence that the funding of undirected, noble pursuits like curiosity-driven basic research can yield as much utility as—and perhaps even more than—purpose-driven development projects, at least when it comes to producing “miracles.” Again, Bush was not wrong about the usefulness of basic research; his mistake was to cast the relationship between scientific inquiry and technological development as unidirectional and linear. But before we turn to that problem and discuss a better way to think about research, let us consider two common tropes in public debates about the nature of technological innovation—both regrettable, if unintended, legacies of Bush’s emphasis on utility.

Two Fallacies of Innovation

At first blush, analogies between energy innovation and iconic technological achievements like the Manhattan Project or the Apollo program might appear not just defensible but also quite apt. The United States built the first atomic bomb, won the space race, and became the world's indisputable powerhouse in science and technology. American scientific and technological prowess may be judged by the facts that the United States is home to the majority of the world's leading research universities and its residents or citizens have been awarded about half of all Nobel Prizes in science, medicine, and economics.

The problem with these analogies is that they paint an incorrect, though attractive, picture of how innovation works: The government selects a practical goal and provides funding for the research, with a straight path from idea to insight, then to innovation and industrial application. Following this model, its proponents tell us, the government and its orbit of researchers and advisors not only built the atomic bomb and put a man on the moon, but also invented the Internet—from which emerged many new products and whole new industries. Why not do that again?

But this picture is flawed, in at least two ways. First, it suffers from what we might call the *moonshot fallacy*, which goes like this: “If we can put a man on the moon, surely we can [fill in the blank with any aspirational goal].” It is true that engineers have achieved amazing feats when tasked with particular, practical goals. But not all goals are equally achievable. Bill Gates's hoped-for miracle of transforming the global energy economy is not like putting a few people on the moon a few times; it is more like putting *everybody* on the moon—permanently. The former was a one-time engineering feat; the latter would require an array of new technologies to be invented and then integrated into the world economy at every level. Most of our present challenges, from curing disease to feeding or transporting billions of people, are also fundamentally unlike the discrete engineering challenge posed by the race to the moon. Engineers are capable of performing remarkable feats once or a few times, especially if (as was more or less the case with the Manhattan Project and the Apollo program) cost is no object. But scales matter—especially in physics. It is not just more challenging but qualitatively different to engineer devices or systems that are both effective and affordable at a global scale.

Moonshot enthusiasm began while the Apollo program was still underway. In the aftermath of the 1973 Arab oil crisis, editors at *U.S. News & World Report* published a book titled *1994: The World of Tomorrow*

summarizing expert forecasts. The editors began the book observing: “So staggering is today’s rate of change that nothing stays as it is for long. Man is progressing so rapidly that things no longer move gradually. They leap. They soar. They bring the future tumbling upon us.” Among the futurist pundits included in that collection were the earliest participants in the now-decades-long tradition of predicting a cataclysmic energy crisis that urgently demanded technological solutions: “Experts predict that unless a massive effort to solve the [energy] problem is launched immediately, Americans face a doomsday future.” The experts of 1973 were sure that, in the wake of Apollo, the internal combustion engine would “become a thing of the past” before the twenty-first century began. That didn’t happen, of course. Instead, the number of petroleum-burning automobiles in the world rose dramatically (motor vehicles are being produced today at double the rate of the early 1970s), as did rates of global air travel; oil production not only increased to match demand but overshot to such an extent that we now have a global glut that has collapsed oil prices.

The linear model of innovation also suffers from something Vannevar Bush could not have anticipated. Call it the *Moore’s Law fallacy*. First proposed in 1965 by Gordon Moore, who would later go on to co-found Intel, Moore’s Law is a prediction that the number of transistors fabricated on a single silicon microchip doubles every two years (commensurately dragging down the cost of computing). It has come to epitomize the relentless and astonishing gains in computing power (and cost-effectiveness) characteristic of modern information technology. More and more information can be stored and transported at ever-smaller scales, using profoundly fewer atoms and less energy per unit of data. On top of this, software engineers use clever mathematical codes—themselves enabled by increasingly powerful computing—to parse, slice, and shrink information itself, compressing it without loss of integrity. The combination is profound. Compared to the dawn of modern computing, today’s information hardware consumes over 100 million times less energy per logic operation, while working in a physical space more than one million times smaller. A single smartphone is thousands of times more powerful than a room-sized IBM mainframe from the 1970s.

Many tech entrepreneurs, along with the politicians and pundits who are in awe of them, seem to believe that such disruptive Silicon Valley innovation is imminently achievable across nearly every domain—including and especially energy. They point to the way the Internet has disrupted big-box retail, newspapers, taxis, hotels, and more—the ‘Uberization’

of everything. When it comes to energy, digital disruptors believe that machines like car engines can follow the same tantalizing tech trajectory as computer chips.

The Moore's Law fallacy was already in play at the end of the last century, during the era of "irrational exuberance." In 1996, the editors of *Wired* published a portfolio of futurist predictions, and while many of the forecasts about information technology would prove accurate, the forecasts about the world of atoms and energy completely missed the mark. For example, had the futurists been right, supersonic commercial air travel would now be the norm. The problem with their predictions is encapsulated by Peter Thiel's observation that "we've had enormous progress in the world of bits, but not as much in the world of atoms."

Underlying the Moore's Law fallacy is a category error. The software challenge of how to store the most information in the smallest possible physical space is distinct from the hardware challenge of how to move physical objects using as little energy as possible. Different laws of physics come into play. In the world of people, cars, planes, trucks, and large-scale industrial systems—as opposed to the world of algorithms and bits—hardware tends to expand, not shrink, along with speed and carrying capacity. The energy needed to move a ton of people, heat a ton of steel or silicon, or grow a ton of food is determined by properties of nature whose boundaries are set by laws of gravity, inertia, friction, mass, and thermodynamics. If energy technology had followed a Moore's Law trajectory, today's car engine would have shrunk to the size of an ant while producing a thousandfold more horsepower. While it is true that engineers can build ant-sized engines, such engines produce roughly 100 billion times less power than, say, a Subaru. No amount of money or Silicon Valley magic will cause a car engine's power, or its equivalent, to disappear into your pocket.

Moore's Law-like improvements in energy are not just unlikely, they cannot happen given the physics we know today. This is not to say that Silicon Valley and information technology will not dramatically affect the production of energy and physical goods. On the contrary, there is enormous opportunity with information and analytics to wring far more efficiency out of our physical and energy systems. But wringing efficiency out of existing infrastructures—as valuable as that is—is akin to 'Uberizing' solar panels and shale rigs. It will get us more efficiency but not the kind of disruptions—the miracles—that would be analogous to discovering petroleum or nuclear fission, or the invention of the photovoltaic cell.

The Pursuit of Disruption

When genuinely miraculous technologies do emerge, they frequently come not from extensions of known science and technology but from foundational conceptual revolutions. Consider modern chemistry, which took us from alchemy all the way to combustion, weapons of war, and modern medicine; or quantum physics, including the discovery of the photoelectric effect, which earned Einstein his Nobel Prize and was important in the history of television, among other things; or electromagnetism, pioneered by Faraday and Maxwell, without which we would have no electric motors; or Turing's mathematics, which made possible modern computer operating systems; or von Neumann's and Morgenstern's mathematics of game theory so widely used in business, economics, and political science.

Of course, such "miraculous" ideas do not occur in a vacuum, nor do they generate devices, products, tools, services, and companies overnight or in a straightforward, linear fashion. There is a dynamic interaction between scientific insights and the technologies, financing, engineering, as well as the standards, regulations, and policies that complement, enable, and develop them. But the point is that such foundational ideas are the stuff of true disruptions—rather than linear extensions of existing capabilities and knowledge, or the result of goal-directed research.

When he accepted a shared Nobel Prize in 2013 for his discoveries related to the molecular "traffic" within living cells, Randy Schekman said that the prizes

reflect the value of curiosity-driven inquiry, unfettered by top-down management of goals and methods.... And yet we find a growing tendency for government to want to manage discovery with expansive so-called strategic science initiatives at the expense of the individual creative exercise we celebrate today.

Policymakers are understandably impatient for research to yield practical solutions to seemingly intractable challenges, from finding better food and fuel to curing diseases, fighting terrorism, and improving cybersecurity. Accordingly, as we have seen, federal spending on research and development tends to echo corporate behavior, directed toward specific, practical problems or projects—and sometimes even to specific *products* such as solar panels or batteries. In other words, federal R&D spending resembles—and therefore competes with or directs—business R&D.

Unsurprisingly, most business R&D is focused almost exclusively on developments that can reasonably be expected to improve competitiveness or profits in the near term. As noted above, only about 7 percent of overall R&D spending by U.S. businesses goes toward basic science. Even many corporations lauded as innovators do not really support *scientific* research, choosing instead to concentrate on engineering projects. Google's Advanced Technology and Projects group, for example, reportedly gives projects two-year deadlines. Pharmaceutical companies work on longer timelines, but they are still driven by utility (and constrained by time and money). Such a focus may be sensible in business, but it bears no resemblance to the kind of curiosity-driven inquiry Schekman lauded.

So when the denizens of Silicon Valley search for the next disruption, it seems they are looking in the wrong place. The world awaits the first Nobel Prize to be awarded to a researcher from Google, Apple, Amazon, Microsoft, or one of our other modern high-tech firms. These corporations are a very long way from business models that foster Nobel laureates.

Compare today's tech leaders to those of yesteryear. As of this writing, Google is nearly nineteen years old. Bell Labs was just a dozen years old when one of its scientists was awarded the 1937 Nobel Prize in physics along with George Thomson for confirming experimentally Louis de Broglie's wave-particle theory of matter. Researchers at Bell Labs would go on to receive a total of eight Nobel Prizes. Five Nobels have also gone to researchers at IBM's labs, which were established in 1945 to pursue what IBM itself calls "pure science." Bell and IBM were not alone in those days; support for pure science or basic research was common at many other storied corporate labs, including those of Xerox, Kodak, DuPont, and even Exxon.

The problem is not that today's American tech sector lacks the money. Apple's futuristic-looking new headquarters, nicknamed "the spaceship," cost \$5 billion. (The Pentagon, completed in 1943, cost roughly \$1 billion in inflation-adjusted dollars.) Facebook recently moved into its own new mega headquarters, and Google is planning a similar upgrade. The collective market value of the top 100 tech companies is measured in trillions of dollars; only fifteen countries in the world have a GDP of over one trillion dollars. And the collective revenues of America's Fortune 500 equals two-thirds of the entire U.S. GDP.

But there is no evidence that any corporation, much less any corporation in the high-tech sector, is interested in returning to anything like the Bell Labs model. It is notable that Eric Schmidt, the executive chairman of

Google's parent company, is a prominent supporter of the 70-20-10 model for how employees, or at least technical employees, should spend their time: 70 percent on core business activities, 20 percent on related projects, and 10 percent on projects unrelated to core business. The engineers who led Google's energy initiative wrote that the last 10 percent should be dedicated to "strange new ideas that have the potential to be truly disruptive." But while the allotment of employee working hours might seem like a good proxy for a tech company's overall R&D spending, there is no reason to believe that the 10 percent figure corresponds to anything like basic research. Nor is there reason to believe that any of the other major tech firms is interested in basic research at anything approaching the levels of the big technology companies of the twentieth century. It remains to be seen whether Bill Gates's Breakthrough Energy Coalition will end up funding any curiosity-driven research or will simply emulate the model of goal-driven and directed R&D. Similarly, many of the large philanthropic research grants to universities and research institutes are mission-driven.

As it stands today, over half of *all* spending on *basic* research in the United States (59 percent) comes from the federal government and universities. (And the overall federal share is likely greater than reported since industry self-identifies research activity categories, leaving "basic" subject to definitional exaggeration.) While the United States remains the world's foremost supporter of research, the nation is now rapidly losing its lead. At stake is not mere prestige but the erosion of the foundation of innovation, and thus the long-term weakening of the economy, not to mention the capacity to find new "miracles" in any field.

The Structure of Technological Revolutions

The terms "basic research," "applied research," and "development" have appeared throughout this essay. This taxonomy—which echoes the ideas offered by Vannevar Bush in *Science, The Endless Frontier*—is used by the federal government to classify R&D funding and work. Here, for example, is how the Office of Management and Budget defined each of the terms in a 2013 budget document:

- Basic research: "study directed toward fuller knowledge or understanding of the fundamental aspects of phenomena and of observable facts without specific applications towards processes or products in mind"—although it "may include activities with broad applications in mind."

- Applied research: “study to gain knowledge or understanding necessary to determine the means by which a recognized and specific need may be met.”
- Development: “application of knowledge or understanding, directed toward the production of useful materials, devices, and systems or methods, including design, development, and improvement of prototypes and new processes to meet specific requirements.”

The political scientist Donald Stokes criticized Bush’s formulation of basic research for providing “too narrow an account of the actual sources of technological innovation” and “too narrow an account of the motives that inspire such work.” In 1997, in an effort to improve upon this R&D taxonomy, Stokes introduced what he considered a more “interactive” way of thinking about work in science and technology. Stokes’s alternative taxonomy was based on two binary questions about knowledge and utility, resulting in four quadrants. (See Figure 1.) Stokes illustrated the nature of the quadrants by picking an iconic name for three of them: the physicist Niels Bohr, whose “quest of a model atomic structure was a pure voyage of discovery” without consideration of utility; the entrepreneur

Fig. 1. The Stokes Model

		Considerations of Use?	
		No	Yes
Quest for Fundamental Understanding?	Yes	Pure Basic (Bohr)	Use-Inspired (Pasteur)
	No	Golden Fleece (Proxmire)	Pure Applied (Edison)

Adapted from Donald E. Stokes, *Pasteur’s Quadrant* (1997)

Thomas Edison, who “exemplified the applied investigator wholly uninterested in the deeper scientific implications of his discoveries”; and the chemist and microbiologist Louis Pasteur, representing a mixed quadrant, as someone whose scientific research was fueled by practical motives such as curing disease.

Stokes left the fourth (bottom-left) category unnamed, since it represents research that advances neither knowledge nor utility. I’ve labeled it for Senator William Proxmire, Democrat from Wisconsin, who created the Golden Fleece Award in 1975, which he announced regularly to pillory examples of wasteful government spending. First “awarded” to the National Science Foundation for studying why people fall in love, the run for the Golden Fleece Award ended with Proxmire’s retirement in 1988.

Stokes improves upon, though he does not move beyond, Bush’s original taxonomy. Indeed, neither Stokes’s taxonomy nor the Bush-inspired linear classifications used by the federal government does justice to the non-linearity and productive serendipity characteristic of how science and technology advance and, in particular, how “miracles” happen.

Stokes is right, of course, that scientists *are* often aware, with varying degrees of certitude, of potential applications of their research. And he is right that scientific research can be motivated by practical pressures and enabled by technological developments. But it is far too often impossible to know in advance whether to categorize research as “useful” or not, or as “practical” or not.

Consider, for instance, that Randy Schekman’s Nobel Prize-winning discovery emerged from work he did on how molecules in yeast proteins operate. Clearly, such knowledge could have some utility; yeasts have long been used to make bread and beer, and can be used to produce ethanol. One might have even thought that research on yeast proteins would be the proper domain of brewers and biofuel companies, not university professors. But in his Nobel banquet speech, Schekman said he and his team had “no notion of any practical application” and were driven entirely by curiosity. Yet their research ended up providing a roadmap for the biotechnology industry to manufacture “commercially useful quantities of human proteins”—an especially useful result when one considers that “one-third of the world’s supply of recombinant human insulin is produced in yeast.”

Schekman’s example would seem to corroborate Bush’s notion that pure, scientific inquiry can yield practical results. And so it does. But more importantly, it illustrates how the fruits of research—be they technological or scientific—cannot be predicted in advance. There is no “predictor function,” as Gates puts it, for either utility or knowledge.

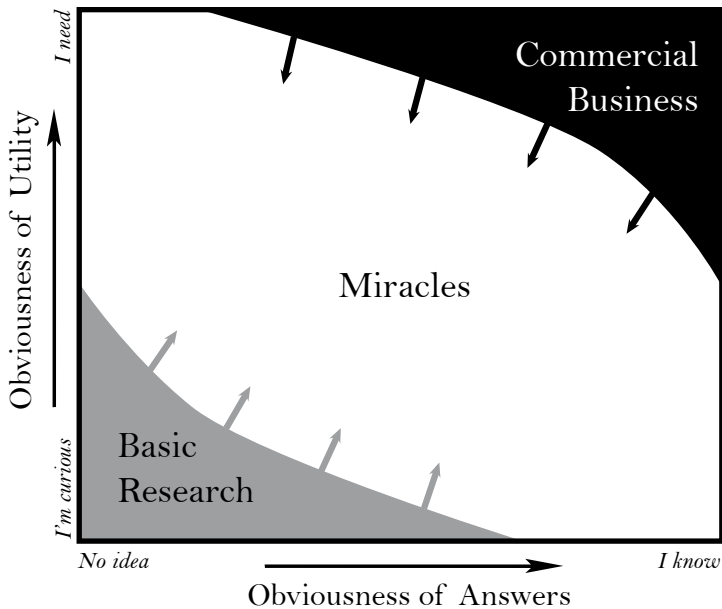
There are many more examples of the productive serendipity associated with curiosity-driven research. When Watson and Crick in 1953 identified the structure of DNA, they were not seeking to improve the criminal justice system. Similarly, as Schekman pointed out in his speech, undirected research in the molecular basis of neurotransmitters resulted in practical uses of the toxins produced by botulinum bacteria, both for treating neuromuscular diseases and paralysis and for cosmetic use in Botox. Princeton chemist Edward Taylor's curiosity-driven study of butterfly wings led, eventually, to the development of cancer therapeutics. This research was commercialized in partnership with Eli Lilly & Co., and the royalties ended up funding the construction of a new chemistry building at Princeton—another practical outcome, which, itself, was aimed at advancing scientific understanding.

This is not to say, of course, that all basic research leads to useful or practical outcomes, much less Nobel Prizes or “miracles.” Most does not. Nor does the relationship between scientific research and technological innovation run in one direction; often some new technology will lead to an improved (but not necessarily useful) understanding of nature. The steam engine was not born out of an attempt to apply principles of thermodynamics to a practical problem, but rather the reverse: the field of thermodynamics was born out of Carnot's curiosity about the steam engine.

Or, to choose more recent examples, it was curiosity that drove Felix Bloch and Edward Mills Purcell in the 1940s to explore the odd phenomenon of nuclear magnetic resonance that laid the foundation for the invention of the MRI scanner, an astonishingly useful medical tool. Inversely, tinkering to achieve something useful, Arno Penzias and Robert Wilson of Bell Labs were trying to improve commercial radio antennas when they inadvertently discovered the cosmic microwave background radiation. This discovery, for which Penzias and Wilson shared the Nobel Prize in 1978, was taken by many scientists as ending the debate over whether the universe began with the Big Bang—a non-utilitarian debate if ever there was one. In these and many other similar cases, the common denominator is neither research divorced from all technological and practical concerns, nor the crucible of mere practical necessity, but rather unpredictability—the serendipity of curiosity-driven inquiry.

But it is precisely the dynamic unpredictability of curiosity-driven research that vexes linear-thinking policymakers and outcome-obsessed pundits. Consequently, it may be easier to justify the funding of basic research if we can offer a new taxonomy of research: a less linear, non-binary way of characterizing the work that leads to discovery and invention.

Fig. 2. The Structure of Technological Revolutions



Such a taxonomy—which I propose here (see Figure 2)—should reflect the dynamic interaction between domains of science and technology, between foundational research and commercial research. It should also eschew the binaries of “utility” vs. “the pursuit of knowledge” or “basic” vs. “applied” research, instead using sliding scales for the *obviousness* of utility and knowability. On one axis, we can plot utility in terms of the degree to which the pursuit is curiosity- or necessity-driven (from “I’m curious” to “I need”). Thus, the question “How do proteins behave in yeasts?” is importantly different from the question “How can I manufacture human insulin?” On the other axis, we can plot the degree to which we think an answer can be found.

In this taxonomy, the search for extraterrestrial life using the next generation of mega-telescopes certainly belongs in the extreme corner of purely “I’m curious” and have “no idea” how easily the answer can be obtained. And yet, these new radio and optical telescopes may generate such a staggering quantity of data that the unprecedented demand for new forms of storage and analytics may lead astronomers to create, inadvertently, tools useful for big data in industry and health care. A similar story could be told about the new class of super-microscopes intended

to advance our understanding of proteins for drug discovery (“I need”); in ways we cannot yet know for sure, they will doubtless prove useful beyond the laboratory, perhaps by discovering molecular magic that enables useful ways to store electricity, or by discovering ways to deliver therapies inside cells.

Harking back to some earlier examples of where serendipity belongs in this taxonomy, Carnot was probably halfway along the x-axis of “obviousness” but only a little way up the y-axis of “obviousness of utility”—yet the field of thermodynamics was deeply foundational, useful, and truly miraculous. On the upper side of the map we could place the work done by General Electric in pursuit of more efficient or powerful aviation turbines, which the company and its customers clearly “need” (top of the y-axis), but which involves relatively well understood, if challenging, physical principles (the right side of the x-axis, deep in the “Commercial Business” zone with minimal likelihood of surprises). Penzias and Wilson began on the edge of the “Commercial Business” zone, well up the y-axis of “I need” (to get a better antenna), but in their case they were further to the left, since they likely had less idea of how to achieve their goals.

The dynamic flow of knowledge and experience between the two opposing corners of the structural taxonomy in Figure 2 is where the miracles happen. Some miraculous outcomes may seem possible at the outset, since scientists often do have a sense that they are in a kind of Holy Grail territory (that is, the territory marked out in the center of the map). Other outcomes are total surprises emerging out of both ‘left’ field (Penzias and Wilson) and further to the ‘right’ (Carnot). Put another way, utility-driven research on difficult questions can produce miracles, just as curiosity-driven research can deliver miracles in domains that are considered “obvious.” In both cases, the key ingredient is the freedom to be curious.

In a 2011 lecture, then-Federal Reserve chairman Ben Bernanke noted that “we know less than we would like about which [R&D] policies work best.” The challenge for federal funding of research is to strike the right balance between the extremes of curiosity and utility, and between the extremes of ignorance and certainty. There are utility-driven questions like: “How do we cure cancer?” or “How can we store electricity cheaply?” or “How can we resupply the International Space Station?” or “How can we detect hidden explosives?” Then there are curiosity-driven questions like: “How do proteins operate inside a cell?” or “How can we model quantum electrochemical behavior in biological systems?” or “Is there life on other planets?” or “Do gravity waves exist?” While it is usually obvious

which questions are utility-driven and which are curiosity-driven, it is *not* obvious, as history shows, which may lead to foundational discoveries—producing not just utility but perhaps also “miracles.”

The energy miracle Bill Gates seeks may well be found as a consequence of new science or new technology arising from research programs studying questions that seem distant from energy—and indeed, distant from any practical use. It is often said that necessity is the mother of invention. But curiosity is the mother of miracles. We need both. And a preoccupation with the former will not produce more of the latter. Perhaps Bill Gates’s enduring legacy will be the revitalization of federal and philanthropic support for the curious pursuits of brilliant minds.