MELISSA FLAGG

"Getting the Questions Right Is Very, Very Important"

Policymakers must rewrite science's Cold War–era social contract by being honest about who benefits from R&D investments and about the feasibility of achieving global leadership in all scientific domains.

Since World War II, the default American goal for science, technology, and military power has been "leader." We push the frontier, we fund the basic research, we run the Manhattan Projects and the moonshots. That is an expensive role, and political and scientific decisionmakers rarely question whether it's the right one.

Maintaining a leadership position in science and technology across all domains and sectors, including basic research, technology development, and industry, made sense in the 1960s and '70s when the federal government provided an overwhelming share of domestic as well as global R&D funding and the country had a strong domestic manufacturing base to capitalize on it. That advantage showed up directly in aviation, in semiconductors, and in the things industry and the military were positioned to build. This arrangement was sometimes expressed as science's social contract: American taxpayers paid for cutting-edge research, and in exchange, they received security and a better life.

Today, however, American life occurs in a nearly unrecognizable global context. When Iranian-linked groups hacked the control systems of at least 55 municipal water systems across nine states in 2026, they weren't defeating a sophisticated defense. They were exploiting an almost infinite set of known vulnerabilities in equipment that small utilities had bought because it was what they could afford. Similar attacks had already hit New York in 2013 and Pennsylvania in 2023, but neither federal nor state governments fixed the vulnerability. Iran didn't need to out-innovate the United States; it needed to find the people who could locate weak critical points in domestic infrastructure. Having funded R&D for eight decades on the promise of security and better lives, taxpayers in affected towns found themselves disconnected from the very thing they'd been promised.

It's become common to say that science's social contract is fraying, but I don't think the science policy community has been honest about why. I argue that leadership—across every domain, all at once, was never a law of nature. It was a response, a strategy applied to a specific moment in US history, and that moment has passed. What follows is a call not to abandon ambition, but to assess what the federal government should actually be trying to do with the research dollars it still controls, what outcomes it wants, who should benefit directly, what strategies are actually feasible, and what levers can be pulled to achieve them.

Getting the questions right

The authority of this quest for technological superiority is often traced to Vannevar Bush's *Science, the Endless Frontier*, the secular scripture of American science policy. But Bush's 1945 memo to Roosevelt wasn't a timeless solution to the nation's R&D challenges. It was really a method for diagnosis and prescription: Assess the context, identify the resources and constraints, and find the feasible levers given those constraints. Bush's prescription to channel federal funding

into universities was a response to a specific diagnosis that Europe, the traditional engine of basic research, lay in ruins, while the United States faced concerns about massive postwar unemployment and genuine needs in health, security, and the public good. Funding university science was a lever that could put people to work and build a sustained research base while filling the pipeline of innovation with people and ideas.

Science's social contract also had a social context that mattered as much as the science. The postwar years saw an unprecedented expansion of the federal government. The country had just won a war and used a nuclear weapon that defied comprehension—then faced the new global strategic threat of proliferation, demanding a national response. Before this, the New Deal and its raft of government programs had pulled the country out of a depression. For the first time in the nation's history, people were willing to trust a highly centralized approach to governance.

Every one of those conditions needs to be reinterrogated today. How much leverage does federal spending on science actually offer when more than 80% of US R&D is funded by private industry and philanthropy, and many emerging areas like artificial intelligence are highly dependent on those sources? What can the federal government do in a global system where the United States and China together account for a little over half of world R&D spending, with almost $1 trillion spread across other countries? How should federal dollars be deployed in a country that no longer owns its manufacturing base, has new struggles with attracting global talent, and has lost something arguably more valuable than either—trust? If Bush were doing this diagnosis today, he would likely reach a different conclusion and prescribe something else.

I spent my early career at the Office of Naval Research and later served as the deputy assistant secretary of defense for research. More recently, as executive director of the Andrew W. Marshall Foundation, I've come to appreciate Marshall's history as an economist who became an influential postwar strategist at the RAND Corporation and leader of the Pentagon's Office of Net Assessment—often described as the Pentagon's internal think tank—from 1973 to 2015. Rather than copying the prescriptions that emerged from Cold War giants like Bush and Marshall, I believe that much can be learned by understanding and applying their methods.

Marshall is remembered for developing net assessments, a form of military strategic analysis, into a disciplined study of conditions, patterns, and dynamics to understand not only what a competitor is doing, but why. Famously, he focused on questions. "Poor, mediocre answers to good questions are more important than getting splendid answers to poor questions. That means that getting the questions right is very, very important." Marshall believed better analysis makes hard choices clearer, shows tradeoffs, and stops leaders from fooling themselves into thinking they can have something that isn't realistic.

That's the spirit behind the questions below. To rebuild the social contract for science, the scientific enterprise shouldn't just try to restore Cold War–era funding levels, patch the relationship between the administration and research universities, or force public trust back into existence. Decisionmakers also need to do the uncomfortable diagnostic work to understand what outcomes they actually want, what resources and constraints now exist, what strategies are feasible within those resources and constraints, and what levers are still attached to real power.

What outcomes do Americans want?

The United States has tied R&D to "winning" and national security for so long that the concept has become disconnected from the way that Americans define national security. National security isn't just about beating a competitor overseas, and it isn't just foreign policy that happens in embassies and multilateral discussions. It's really about securing a way of life Americans desire. For most people, that way of life has an economic basis: housing, health care, good jobs and wages, quality education, and a reasonable belief that their kids can expect to afford and access the same in the future. If these sound like domestic policy issues, it's because domestic policy and national security can no longer be pulled apart, as the Iranian hack shows.

Over the last 50 years, US leaders have tended to treat national security as an arm of foreign policy, using R&D as a geopolitical tool. This has allowed competitors to define what counts as success. Instead of asking what Americans actually want the country to be, the nation has been measuring itself against other countries' scoreboards. That drives a chronic undervaluing of domestic affairs and keeps government attention pointed outward, toward a competition that can erode the nation from within even while the United States is technically "winning" it.

Today, individual Americans face tangible, sometimes existential challenges the federal government hasn't been able to address. Americans of both parties worry about high rates of poverty, drug addiction, crime and violence, education, and the ability to recover from natural disasters. Crucial infrastructures for energy, clean water, and affordable housing are in decline. All this, along with rising income inequality, runs alongside declining trust in institutions. This is a fundamentally different environment from the one Bush was diagnosing in 1945, and it calls for a new assessment and new solutions.

Considering the scientific community specifically within this context, a clear question follows: Do American

communities actually see a return on their investment in R&D, or get to use what it produces? If people can't see the results of federally funded research showing up in their own communities, if the research isn't focused on the problems they actually have and doesn't create knowledge to solve the problem in their community, the federal government's authority to ask for that trust and the resulting investment will keep eroding.

Who benefits from R&D investments?

One issue with science's social contract is that there is no direct connection between taxpayers' priorities and the research that is done. Today, when the government funds public discovery research, it is generally assumed that the market will pick up the results, but the direct return on this investment rarely gets back to the taxpayers themselves. One standard rebuttal here is that the public already gets plenty back through innovations like GPS, the internet, and breakthrough medicines. That is true. But those benefits took decades to reach consumers, and they arrived through private commercialization that the government didn't shape, with much of the public value that is captured residing in the military. If there is a social contract with science, then the contract holders have a reasonable belief that their tax dollars will be used to address their top-priority problems. Of course, funding for exploration and basic research remains necessary, but it should not be seen as sufficient for the contract to be sustained.

One way to find and solve taxpayers' problems is to identify who is actually affected by the problem and build mechanisms connecting them to research. The military tries to do this by placing researchers with "problem owners" during times of crisis. Early in my career at the Office of Naval Research Global, the Navy sent me to sea for five days so I would understand the sailors whose problems I was working on. That lesson stayed with me for the next two decades. During operations in Afghanistan, scientists and engineers spent time in the field alongside operators before iterating in the lab and pushing solutions back out. For over a century, the US Department of Agriculture's extension system has put researchers next to farmers, exchanging real-time problems for real-time solutions in a way that drives both rapid adoption and rapid innovation. But support for extension programs has been declining since the 1980s.

In order to shift research toward taxpayer concerns, funders will have to change their outlook from being grantors and directors of research toward being partners in a heterogenous ecosystem. Right now, government tends to fund the part of the research pipeline that produces countable outputs like papers, patents, and spin-offs, then disengages once those proxy measures are counted. A meaningful partnership model would entail prioritizing a wide range of "smaller" real-world problems, staying involved through question development, pilots, procurement support, and local capacity building, and supporting implementation the community can sustain. Beyond the extension system, working examples of research that connects directly to taxpayers' needs can be seen in North Carolina's Collaboratory and Research Nova Scotia. Similar approaches could magnify the impact of US research dollars while rebuilding the trust that has been lost.

Building an R&D strategy that acknowledges constraints

In finance, it's standard to build a portfolio that includes high-risk and low-risk investments alongside one another, sized to adapt as conditions change. The United States does not have a national research strategy, but instead funds research that successfully competes for congressional attention. In battles over budget allocation, basic research is often pitted against more applied research rather than being viewed as a system that needs to fluctuate over time.

This brings us back to the obsession with R&D "leadership." Building an outcome-based portfolio requires accepting that time and money are finite, and that different tactics are required in each field and with each competitor. Before considering strategies, decisionmakers need answers to a few blunt questions: How long would research need to be conducted in a given field to become first among competitors? How long until a competitor's advantage produces harm (military, economic, or social) that can't be reversed? Can the foundational knowledge the leader is producing be accessed and put to use by a follower? The answers to these questions display the real tradeoffs among cost, speed, and risk tolerance, informing whether the "leader by default" posture is unhelpful or potentially even harmful.

If leadership in a given area is both feasible and strategically worth it, the US enterprise can invest in leading alone or team up with allies through international collaboration. However, in specific areas where speed of adoption matters more than discovery, and leadership may not be feasible in the time or budget allowed, accepting the role of "fast follower" may be the better strategy. This allows someone else to absorb the cost of basic research and early debugging, while the follower engineers robust solutions and adopts them faster across sectors. China has built real capability, as has Singapore, by following with discipline over sustained periods and shifting into leadership where and when it was worth the cost.

The recent history of drone development demonstrates how expensive the US emphasis on leading in R&D can be. From the 1970s through the early 2000s, the United States and its allies were the dominant funders of autonomy research. But real-world implementation showed up elsewhere early. Japan allowed autonomous crop dusters by the early 1990s, and today China dominates the commercial

drone market through companies like DJI. At the same time, Iran, Turkey, and China have all integrated small armed drones into their militaries—a shift the Russia–Ukraine war has only accelerated. Those countries didn't out-invent the United States. They fast-followed straight to implementation, driven by solving concrete problems and capitalizing on clear market signals. The result is that the United States is now behind in implementation of drone technology for both civilian and military use—not because it lost a research race, but because it never adopted an implementation strategy with intention to lead all the way to the solution.

If leading or following is too expensive, there is a third option that gets little attention in strategy documents: making a technology less useful to everyone, not just the adversary. This requires policymakers to accept some cost to themselves in order to remove a competitor's advantage—investing in ways to degrade a capability's value or reliability, raise its cost to sustain, or shift the competition onto ground where the competitor's edge doesn't transfer. Rather than spending heavily to onshore chip manufacturing to avoid supply-chain risk, for instance, planners might instead push to make chip designs modular and interchangeable, removing the bottleneck the fabrication advantage depends on. There are concrete examples of this within the military, but they are rarely discussed because they run counter to the ideals of leadership without a cost to convenience, capability, or reputation.

So the original questions return: What are decisionmakers trying to secure? What outcome do they want for the nation? And who is science's social contract with? Today's policymakers need to assess the full range of levers available to them and to their competitors—as well as how much real control still sits behind each one—before deciding which to pull. In Bush's day, the value in pulling the university research lever was clear, but there are many others worth exploring today.

What are today's levers?

Today, the federal government needs to be honest about which roles it can play and realistic about which levers are still attached to real power and incentives. Sometimes that may mean assessing and sharing global advances or lessons from localized successes. Sometimes it may mean convening the range of powerful actors and the communities experiencing genuine need in the same room. Sometimes it means building a capability the market won't produce on its own or writing a regulation that changes incentive structures. But the new reality is that over the coming decades the federal government will almost never be chairing the whole system.

One place the federal government has significant unused leverage is in its ability to observe the global and domestic innovation system. Information and market intelligence is itself a powerful lever. Across federal agencies, the military, innovation offices, industry, and think tanks, staff are already compiling real knowledge of global R&D investments and developments in domestic and global academia and industry. The cost of aggregating that knowledge is relatively low, but it requires sustained investment and people with the skill to use the results. Sharing this across agencies, with the people making R&D investment decisions in and out of government, and with organizations making consequential technology purchases could transform the system's capacity to evaluate ideas, products, and vendors while strengthening research security.

> The scientific enterprise shouldn't just try to restore Cold War–era funding levels, patch the relationship between the administration and research universities, or force public trust back into existence.

What prevents the use of this lever is that today the federal government is positioned to command the system rather than support, convene, or conduct it. In fact, sharing this kind of information with other organizations can be viewed as creating an unfair competitive environment.

Another lever is the government's power to influence the structure of the market through industrial policy. One example is the 1956 consent decree against AT&T's Bell Systems, which forced the monopoly to license thousands of patents royalty-free and is credited with helping create the commercial semiconductor industry. That logic hasn't been applied consistently to industrial policy. Federal procurement, in fact, often rewards market concentration; between 2020 and 2024, the Department of Defense sent nearly a third of its dollars to just five firms. If the goal is a wider, more resilient base of suppliers and ideas, procurement itself is a lever the government already holds—but too rarely uses.

Finally, the government could leverage the role of convener to support the compilation and sharing of lessons learned. This country's problems are diverse; even where they are similar, they are not identical. They cannot be solved by a system that propagates a single, one-size-fits-all solution.

The government should use its convening power to build knowledge across contexts so that results scale across many "small" problems in different domains and geographies. Understanding what leads local efforts to succeed or fail is critical but deeply undervalued, as is supporting the people and institutions who build capacity to implement sustainable solutions.

There are certainly other levers still undiscussed, but they won't be found by admiring yesterday's top-down strategies or China's apparent centralized efficiency. Envy of China's top-down control misses the fact that efficient systems are inherently brittle. A central plan has no redundant path, no competing program, no dissenting lab that survives to offer an alternative. Tidiness is a vulnerability.

By contrast, the US system is a glorious mess. Dozens of agencies, companies, philanthropies, universities, and private labs fund overlapping research. Jurisdictions collide, programs duplicate, and no one is fully in charge. In peacetime, that looks like waste. In conflict or fierce economic competition, redundancy is resilience and, through a military lens, a form of deterrence in its own right. The programs that look duplicative in a budget hearing are the ones still functioning when the "priority" program fails. If an adversary knows a single "priority" capability is the only way to solve a problem, that capability becomes a target. If competitors know there are multiple paths through a challenge, that knowledge itself deters action. The US "mess" isn't a design flaw competent management could, or should, fix.

Rather than trying to abolish redundancy, policymakers should treat it as another lever of power, a deliberate tool. Buy a diversity of approaches on purpose, write them into standards and contracts, map where these capabilities already sit across the ecosystem, and build support for states, localities, and critical sectors to plan for it instead of hoping it shows up by happy accident after a crisis.

This is where resilience and the social contract meet. A social contract built around "leadership" promises Americans a single, one-size-fits-all answer to every problem: the fastest chip, the most advanced weapon, the definitive cure. A social contract built around resilience promises something humbler and, I'd argue, more honest. It suggests that when the best answer fails, as it eventually will, there will be another path, closer to home, that someone had the foresight to keep alive. That's a harder story to fund and a harder one to campaign on. It is also the one that could have kept those water utilities from being a target. If R&D is to create security for future generations of Americans, decisionmakers must work to build resilience across the entire country's geography—in K–12 classrooms, hospitals, jobs, and cultural life.

I am not arguing against government funding of basic and applied research. Nor am I arguing that the United States should lose its ambition. But when the government sees money as its primary tool for leverage and leadership as its only possible goal, it overlooks the potential of other levers and other goals. This is reflected in the very way science agencies are often described—as funders—rather than coordinators, conveners, partners, strategists, enablers, evaluators, or even leaders. It is also reflected in the way researchers and administrators view their work, in which the political rationale for increased federal R&D investment has long been leadership and national security. Along the way, leadership has also become central to the identity and pride of the federal science and university communities.

> If there is a social contract with science, then the contract holders have a reasonable belief that their tax dollars will be used to address their top-priority problems.

If science's social contract is to be rebuilt, the science enterprise must reconsider its own values and assumptions.

It's worth returning to where this started. The municipal water systems Iran found and exploited didn't fail to benefit from 80 years of American scientific leadership because nobody was building anything. They failed to benefit because US leadership was never asking the questions that mattered to those towns: Are your basic needs met? Do you feel safe? Can you access and implement the knowledge created by the government? Can you afford the technologies that industry created on the backs of that knowledge? A social contract worth rebuilding starts by answering those questions first.

Melissa Flagg *is the executive director of the Andrew W. Marshall Foundation and a former deputy assistant secretary of defense for research.*

This article is part of an exploration of science's social contract that is generously supported by The Kavli Foundation, the Rita Allen Foundation, and the Dana Foundation.