



Fund Organizations, Not Projects: Diversifying America's Innovation Ecosystem with a Portfolio of Independent Research Organizations

Ben Reinhardt

January 2022

The Day One Project Progress Studies Policy Accelerator is a joint initiative between the Federation of American Scientists & the Institute for Progress

Summary

Dominant research-funding paradigms constrain the outputs of America's innovation systems. Federal research-funding agencies like the National Institutes of Health (NIH) and the National Science Foundation (NSF) operate largely through milestone-scoped grants [that fail to incentivize](#) high-risk research, impose highly burdensome reporting requirements, and are closely managed by the government. Philanthropically-funded research organizations are an excellent mechanism to experiment with different research management approaches. However, they are perennially underfunded and rarely have a path to long-term sustainability.

A single program with two pieces can address this issue:

First, the NSF's new Technology, Innovation, and Partnership (TIP) Directorate should pilot a "organizations, not projects" program in which philanthropically matched grants fund a portfolio of independent research organizations instead of funding specific research initiatives. Partnering with philanthropies will leverage the diversity of American donors to identify a portfolio of research organizations with diverse constraints (and therefore the potential to create outlier outcomes). To have a significant impact, this pilot funding opportunity should be funded at \$100 million per year for 10 years.

Second, NSF should set aside an additional \$100 million per year to sponsor independent research organizations with impressive track records for extended periods of time. This commitment to "acquire" successful organizations will complement Part One's research-funding opportunity in two ways. First, it will encourage philanthropic participation by making philanthropies feel like their money is going *towards* something that won't die the moment they stop funding it. Additionally, it will enable the federal government to leverage the institutional knowledge created by successful experiments in research funding and management.

If successful, this two-part program can be later replicated by other federal agencies. The Administration and Congress should prioritize funding this program in recognition of three converging facts: one, that federal spending on research and development (R&D) is increasing; two, that the American innovation ecosystem is not working as well as it once did; and three, that the proliferation of new institutional structures for managing research (e.g., Focused Research Organizations, private Advanced Research Projects Agencies (ARPA), ["science angels"](#), etc. Swift action could use the increased budgets to empower new organizations to experiment with new ways of organizing R&D in order to address the current system's sclerosis!

Challenge and Opportunity

There is a growing consensus that there is a gap between the speed and efficiency of R&D projects closely managed by the government and R&D projects managed by the private sector.

Federal funding is a major part of the American R&D ecosystem. However, most federal research funding comes with a litany of constraints: earmarks that prevent researchers from spending grant money on things they think are most important (like equipment or lab automation), onerous reporting requirements, the need to get every proposal through a committee, and dozens of hours of grant writing for shockingly small amount of money. Moreover, [studies have found](#) that with a mandate to fund innovative research, federal funding decisions tend to be risk-averse.

As a result, in situations where there's a head-to-head comparison between government-managed research and technology development and privately-managed counterparts, there's little question which is more efficient.

This efficiency gap exists largely because privately-managed organizations often push control over research funds to the organization or level where the “research design” occurs. This yields powerful results. Former Defense Advanced Research Projects Agency director Arati Prabhakar [argues](#) that this mechanism, in the form of empowering program managers, is a big part of why the ARPA model works. In the business world, coupling power (money) and responsibility (research design) is simply common sense. In the research world, the benefits of [“embedded autonomy”](#) are straightforward. Autonomy enables an organization or individual to react quickly to unexpected circumstances. Research is highly uncertain by nature. Coupling embedded autonomy with research design means that funding will be spent in the most useful way possible at a given moment based on knowledge gained as experimentation progresses — not in the way that a researcher *thought* would be most useful at the time they submitted their grant proposal.

Recognizing the power of embedded autonomy to enable powerful, diverse research, there is currently an explosion of experiments in non-academic research organizations. Many are too new to have clear results, but non-academic research organizations — including HHMI Janelia, Dynamicland, Willow Garage, and early SpaceX — have created new fields, won Nobel prizes, and changed the paradigms of entire industries. But even the most successful research organizations struggle to raise money unless there is a clear business case, which leaves public-goods oriented research in the lurch. Philanthropists are strongly motivated by legacy, so they want to fund things that will last. As a result, Private funders often hesitate to fund research organizations that produce public-good R&D.

Understanding this problem suggests a potent new way of deploying the federal government's R&D budget: partnering with philanthropists to build a diverse portfolio

of research organizations with autonomy over their own budgets, and then providing long-term support to the most effective of those organizations.

In other words, the federal government should experiment with funding organizations rather than projects.

Such an approach would position the federal government to act like a limited partner (LP) in multiple venture capital funds. In this capacity, the federal government would avoid setting overly specific requirements around how a particular grant is spent. The government would instead set very high-level priorities (e.g., “create new manufacturing paradigms” or simply “do impactful research”), give funded organizations the autonomy to figure out how to best achieve this goal, and then evaluate success after the fact.

The time is right to invest in creative federal research-funding approaches. There is bipartisan support for large increases to federally funded R&D. But pushing huge amounts of money through outdated R&D funding structures is like slamming on the accelerator of a car that needs an engine repair: incredibly inefficient and with the potential to backfire. By contrast, embedding autonomy in a diverse portfolio of organizations could unlock the sort of unexpected, game-changing inventions and discoveries that have driven the American economy: electricity, airplanes, the internet, the transistor, cryptography, and more.

Plan of Action

The current Administration should launch a two-part program at NSF to test a research-funding system that prioritizes organizations over projects.

As Part One of this program, the NSF’s TIP Directorate should pilot a research-funding opportunity in which philanthropically matched grants fund a portfolio of independent research organizations instead of funding specific research initiatives. This pilot funding opportunity should be funded at \$100 million per year for 10 years. The Directorate should target funding between 5 and 15 organizations this way, [quadratically matching](#) philanthropic funds at values between 100% and 1000% depending on the number of participating philanthropic donors.

As Part Two of this program, NSF should set aside an additional \$100 million per year to sponsor independent research organizations with impressive track records for extended periods of time. The Directorate should set a goal of identifying two organizations during the ten-year pilot that would be good candidates for this long-term funding, funding each at \$50 million per year.

More detail on each of these program components is provided below.

Part One: Philanthropically matched grants

Partnering with private donors is key to the success of the proposed organization-focused funding opportunity. By funding only organizations that have already raised philanthropic dollars, the federal government will leverage philanthropists' due diligence on screening applicants to ensure high-potential awardees. Similarly, the funding opportunity should employ [quadratic matching funding](#) to use donors' confidence as an indicator of how much money to give each organization and to reduce bias favoring organizations that are able to raise a large amount of money from a small number of donors.

Leveraging philanthropic opinion in this way does come with the risk of biasing awards towards organizations working on particularly popular areas or that are particularly good at sales or marketing. The organization-focused funding opportunity could address this risk by establishing a parallel funding pathway whereby a large number of researchers can file a petition for an organization to be selected for funding. The TIP Directorate obviously must impose additional criteria beyond the endorsements of the philanthropic and research communities. It will be tempting for the Directorate to prioritize funding organizations working on specific, high-interest technology areas or themes. But the goal of this program is to advance the long term health of the American innovation ecosystem. Often, tomorrow's high-priority area is one that doesn't even exist today. To that end, the Directorate should evaluate potential grantee organizations on their "counterfactual impact": i.e., their capacity to do work that is disincentivized in other institutional structures.

The question of how best to evaluate success of the funding opportunity is a challenging one. It is notoriously hard to evaluate long-term research output. The whole point of this proposal is to move away from short-term metrics and rigid plans, but at the same time the government needs to be responsible to taxpayers. Metrics are the most straightforward way to evaluate outcomes. However, metrics are potentially counterproductive ways to evaluate new and experimental processes because existing metrics presume a specific way of organizing research. We therefore recommend that the TIP Directorate create a Notice of Funding Opportunity to hire an independent, nonpartisan, and nonprofit board whose job is to holistically evaluate funded organizations. The board should include people working in academia, industrial research, government research, and independent research organizations, as well as some "wildcards". The board should collectively have deep experience performing and guiding high-uncertainty, long-term research and development.

The board would regularly (but not over-frequently) solicit opinions on output and impacts of funded organizations from the program's philanthropic partners, members of the government, people working with the organizations, unaffiliated researchers, and members of the organizations themselves. At the end of each year, the board should give each organization an evaluation "report card" containing a holistic letter grade and an explanation for that grade. Organizations that receive an

F should immediately be expelled from the funding program, as should organizations that receive a D for three years in a row.

Part Two: Invest deeply in demonstrated success

In the same way that governments took over funding libraries once they were started by Gilded Age philanthropists, the government should take over funding immensely successful research organizations today.

At the five-year midpoint and ten-year endpoint of the pilot funding program, the evaluation board should identify any funded organizations that have produced outstanding output. The TIP Directorate should then select up to two of these candidates to receive indefinite government support, at a funding level of \$50 million per organization per year. These indefinitely funded organizations would become a line item in the TIP's budget, to be renewed every year except in extreme circumstances. The possibility of indefinite federal support as an “exit strategy” for philanthropic funders will encourage participation of additional philanthropic partners by providing (i) philanthropically funded organizations a pathway for becoming self-sustaining, and (ii) philanthropies with a clear opportunity to establish a legacy.

What qualifies as “outstanding output”? Like evaluating success, it's a challenging question. We recommend using the same board-based grading scheme outlined above. Any organization that receives an A grade in two of the past five years or an A+ in any one of the past five years should be eligible for indefinite support. This approach will require grading to be very strict: for instance, an A+ should only be given to an organization that enables Nobel-prize-quality work.

Conclusion

Building portfolios of independent research organizations is an incredibly effective way of spending government research money. The total federal research budget is almost \$160 billion per year. Less than 1% of that could make a massive difference for independent research organizations, most of which have budgets in the \$10 million range. Funding especially promising independent research organizations with an additional \$10 million or more per year would have a huge effect, empowering organizations that are already doing outstanding work to take their contributions to the next level.

Even the highest-performing private research organizations in the world — like Google DeepMind and HHMI Janelia Farm — have budgets in the range of \$200 million per year. Sponsoring a select number of especially high-performing research organizations with an additional \$100 million per year would hence have similarly transformative impacts. These large indefinite grants would also provide the major incentives needed to bring the world's leading philanthropies to the table and to

encourage the most cutting-edge independent research organizations to dedicate their talents to the public sector. The sum total of achieving these outcomes would still account for only a tiny fraction of the overall federal R&D budget.

Finally, we emphasize that the goal of this pilot program is not solely to establish an independent research organization portfolio in the TIP Directorate. It is also an opportunity to test a novel research-funding mechanism that could be replicated at numerous other federal agencies.

Frequently Asked Questions

1. What is the relationship between a portfolio of independent research organizations and the existing university ecosystem?

Independent research organizations are not entirely separate from the university ecosystem; you could imagine them funding work in a university lab, collaborating with professors, or taking the next steps on work that originated in a university. However, to achieve the diversity of incentives that is core to this proposal independent research organizations do need to be independently managed. That is, they can't just be a subsidiary of a university or existing company; at the end of the day a subsidiary is under the same constraints as the parent organization.

2. Is there historical precedent or context that exists that this proposal builds on or refutes to be aware of?

The National Aeronautics and Space Agency (NASA)'s [Commercial Orbital Transportation Services \(COTS\)](#) program is one of the most successful examples to date of the federal government delegating R&D management. The government's guaranteed purchase of COVID-19 vaccines as part of project Warp Speed similarly embedded authority with vaccine developers and researchers.

3. What is the first step to get this proposal off the ground?

The first step is for the federal government — perhaps through NSF — to partner with philanthropists to award small matching grants to existing independent research organizations. It is low risk but will encourage people to see independent research organizations as a real possibility, begin building relationships with both philanthropic organizations and independent research organizations themselves, and start surfacing inevitable frictions while the stakes are low.

4. What would a less ambitious version of this proposal look like?

A less ambitious version of this proposal would moderately decrease funding levels for Part One of the program and consider delaying Part Two until there is more confidence in the workability of an organization-first funding structure. We caution that if grant amounts are reduced too much, it will be difficult to attract the top tier of

philanthropic partners and independent research organizations to participate. We recommend reducing the number of grants awarded before significantly reducing award levels.

5. Why can't the NSF (or any other federal research-funding agency) simply diversify its research-management approach internally instead of funding independent research organizations?

In the process of experimenting with different approaches to research management, there will inevitably be failures. Approaches need to be able to be ramped up and shut down quickly. But organizational practices in government tend to take a while to implement and stick around for a long time. Additionally, culture matters. Different approaches to research management will require very different institutional cultures. It's hard to create an entirely new culture inside of an existing institution. Studies have shown that even when the NSF sets aside money for "high risk" research, the work it funds is not that high risk. For all of these reasons, funding independent research organizations is a more promising way to achieve more effective research management in practice than is seeking to reform research management within NSF (or any other federal agency).

6. Why route funding through an independent research organization instead of giving unrestricted funding directly to researchers?

Research management matters. Giving money directly to researchers with very low friction is a strong research-management strategy, as demonstrated by [Fast Grants](#) and others. However, direct funding is just one strategy among many. Some degree of top-down coordination and direction-shaping (as has historically been done by DARPA and various corporate R&D labs) can lead to better results than loosely coordinated researchers alone. In addition, giving money to a portfolio of independent research organizations is a useful way to test different research-management strategies (some of which could ultimately be adopted by federal agencies directly).

7. What are some examples of head-to-head comparisons between government-managed research and privately managed research?

Select examples include:

- The development of SpaceX's starship vs. NASA's Space Launch System (SLS).
- The [ITER](#) fusion program vs. fusion companies like Commonwealth Fusion and Helion.
- NIH grants for COVID-19 research vs. privately funded "fast grants" for the same.
- NIH grants for neuroscience research vs. HHMI Janelia.

In each of these cases, the privately funded research effort outperformed the federally funded and managed effort. SpaceX managed to test (unsuccessfully and then successfully) multiple iterations on Starship in the time it took contractors to test-fire (unsuccessfully) SLS' main engines once. ITER is taking so long that companies like Commonwealth and Helion are able to build systems architectures around technologies that didn't even exist when ITER's plans were locked in. Privately funded "fast grants" for COVID-19 research were yielding results when NIH was just starting to issue requests for proposals? How many other technologies are stagnating because money isn't getting to high-efficiency privately funded research organizations?

8. What are some examples of independent research organizations?

Select examples include:

- **Covariant Research** spins up Focused Research Organizations (FROs): temporary organizations (funded at \$50 million over five years) aimed at scaling execution that wouldn't occur in other institutions.
- **NewScience** funds non-professor researchers directly to work on things that professors are not incentivized to do.
- **Activate** gives scientists time, money, and lab space to focus on translating their research into a product.
- **Dynamicland** is building an entirely different computing paradigm based on interaction with a physical space.

For a longer list of independent research organizations, see [The Overedge Catalog](#).

9. How does a portfolio of autonomous research organizations compare to other mechanisms the government uses to fund external organizations?

For an answer to this question, refer to the summary table below.

	Portfolio of independent research organizations	Grants	Cooperative R&D Agreement (CRADA)	SBIR/STTR	Space Act-style commercial partnership
Clear deliverables	No	Yes	Yes	No	Yes
Constraints on how	No	Yes	Yes	No	No

money can be spent					
Heavy reporting requirements	No	Yes	No	Yes	No
Research organized by a for-profit	Sometimes (independent research organizations are sometimes B corps or similar)	No	No	No	Yes
Research organized by a university	No	Yes	No	No	No
Long-term agreement with mid-term checks	Yes	No	No	No	Yes
Approximate amount of money involved per "unit" per year	\$1-100 million	\$100,000-\$1 million	No data	\$500,000-5 million	\$10M-\$1B
Operates in pre-commercial phase	Yes	Yes	Sometimes	No	No
Enables coordinated and systems-level work	Yes	No	Yes	Yes	Yes
Likelihood of success judged	No	Yes	Yes	No	No

before funding					
Can generate unexpected upside	Yes	Sometimes	No	Sometimes	No

About the Author



At the end of the day, **Ben Reinhardt** is trying to enable more awesome sci-fi stuff (that ideally empowers people and leads to human flourishing). The best way to do that is to create new institutional structures that enable different sets of activities than the current set. Specifically he is trying to create a private ARPA to go after things that are too research-y for startups, too engineering heavy for academia, and too weird for the government. His official training is a B.S. in history and mechanical engineering, and a Ph.D in space robotics. Professional history includes: plumbing/HVAC, NASA, working at and attempting to create startups, and venture capital.

About the Day One Project



The Federation of American Scientists' Day One Project is dedicated to democratizing the policymaking process by working with new and expert voices across the science and technology community, helping to develop actionable policies that can improve the lives of all Americans. For more about the Day One Project, visit dayoneproject.org.

The Day One Project offers a platform for ideas that represent a broad range of perspectives across S&T disciplines. The views and opinions expressed in this proposal are those of the author(s) and do not reflect the views and opinions of the Day One Project or its S&T Leadership Council.