


The DARPA Model for Transformative Technologies

Perspectives on the U.S. Defense
Advanced Research Projects Agency

EDITED BY WILLIAM B. BONVILLIAN,
RICHARD VAN ATTA AND
PATRICK WINDHAM





<https://www.openbookpublishers.com>

© 2019 William B. Bonvillian, Richard Van Atta, and Patrick Windham. Copyright of individual chapters is maintained by the chapters' authors.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs license (CC BY-NC-ND). This license allows you to share, copy, distribute and transmit the text; to adapt the text and to make commercial use of the text providing attribution is made to the authors (but not in any way that suggests that they endorse you or your use of the work). Attribution should include the following information:

William B. Bonvillian, Richard Van Atta, and Patrick Windham (eds.), *The DARPA Model for Transformative Technologies Perspectives on the U.S.: Defense Advanced Research Projects Agency*. Cambridge, UK: Open Book Publishers, 2019, <https://doi.org/10.11647/OBP.0184>

In order to access detailed and updated information on the license, please visit, <https://doi.org/10.11647/OBP.0184#copyright>

Further details about CC BY-NC-ND licenses are available at, <https://creativecommons.org/licenses/by-nc-nd/4.0/>

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at <https://archive.org/web>

Updated digital material and resources associated with this volume are available at <https://doi.org/10.11647/OBP.0184#resources>

Every effort has been made to identify and contact copyright holders and any omission or error will be corrected if notification is made to the publisher.

ISBN Paperback: 978-1-78374-791-7

ISBN Hardback: 978-1-78374-792-4

ISBN Digital (PDF): 978-1-78374-793-1

ISBN Digital ebook (epub): 978-1-78374-794-8

ISBN Digital ebook (mobi): 978-1-78374-795-5

ISBN XML: 978-1-78374-796-2

DOI: 10.11647/OBP.0184

Cover design: Anna Gatti.

15. Does NIH need a DARPA?¹

Robert Cook-Deegan

The National Institutes of Health (NIH) recently celebrated the fiftieth anniversary of its Division of Research Grants with a symposium on peer review. NIH Director Harold Varmus introduced the theme of the day, likening competitive external peer review to democracy by invoking Churchill's quip: "the worst form of government except all the others that have been tried". This analogy expresses a belief in peer review that is widely shared among those who were in the audience. There are, however, a couple of problems with this analogy. First, it is factually incorrect. Some agencies—notably the Defense Advanced Research Projects Agency (DARPA; ARPA during some periods) and the armed services' R&D operations—have demonstrated that other methods work quite well, arguably as well as or better than those used at NIH. Second, comparing peer review to democracy implies a false dichotomy. A country cannot be at once a democracy and a dictatorship, but an agency can simultaneously use both peer review and other mechanisms to support R&D; indeed, several defense R&D agencies do just that.

The chief alternatives to competitive peer review are formula funding methods, based on political, historical, or performance factors, and what might be called the DARPA model, in which staff experts decide how to distribute research funds. Formula funding would

1 Originally published as Cook-Deegan, R. (1997). "Does NIH Need a DARPA?", *Issues in Science and Technology* 13/2, <https://issues.org/cookde/>

surely reduce transaction costs and could provide a stable flow of support to good researchers. The price of reducing transaction costs through formula funding, however, is the loss of expert judgment about innovative promise. The desire to invest in such promise, as opposed to past performance alone, is a major reason that agencies have come to rely on outside expert advice. But the DARPA approach is also a way to foster innovation.

DARPA's effectiveness depends on expert staff, clear mission, focused effort, and lean management. DARPA's main function is to quickly exploit new inventions, ideas, and concepts with potential military utility. Its eighty or so program managers distribute between \$2 billion and \$2.5 billion annually and are supervised by a half-dozen office directors, who in turn report to the DARPA director. Thus, only one management layer exists between the DARPA director and the program managers. The entire DARPA staff is roughly comparable in size to that responsible for administering extramural funds for the National Human Genome Research Institute (NHGRI) or one of the smaller NIH institutes that expends between \$100 million and \$200 million.

DARPA managers are hired for their expertise, often from industry or academia, and typically serve for four years or less. Each handle from \$10 million to \$50 million of research funding per year, of which at least 20 percent is intended for new investments. The money for new programs is a direct result of DARPA's ruthless willingness to kill programs that are not meeting expectations. Success results from a long-term strategy pursued by highly expert staff who are given great discretion to manage substantial funding commitments. Those staff members are held accountable for the results produced by the programs they fund, in quarterly reviews and detailed annual assessments by the DARPA director.

In DARPA culture, managers are self-avowed scientific and technological fanatics. Their base skill is recognizing talent that is relevant to defense needs and providing funds for its expression. The institutional ethos is described as "80 decision makers linked by a travel office", which emphasizes its highly interactive (and, at times, intrusive) style. It is ironic that within one of the world's most notorious bureaucracies, the Department of Defense, resides a tribe of rambunctious technological entrepreneurs.

Created by the Eisenhower administration in the wake of the Soviet launch of Sputnik, DARPA played a crucial early role in the development of computer time-sharing, interactive computing, space launch vehicles, satellite surveillance, lasers, stealth technology, and many other technological innovations. Its twenty-five-year-old Information Processing Techniques Office (IPTO) is DARPA's best-known program outside defense technologies. IPTO spawned the first departments of computer science, bolstered an academic base for large-scale integrated chip design at a time when that foundation was eroding perilously, and created the prototype for today's Internet. It is safe to say that many computing activities we take for granted in the 1990s, such as e-mail, computer graphics, interactive computing, alternative chip architectures, and networking, can be traced to DARPA funding decisions made in the 1960s and 1970s.

Biomedical Success

This period has also been a time of remarkable progress in biomedical research, and NIH has played a central role. NIH funding accounts for almost 30 percent of the world's biomedical research literature, compared to about 40 percent from other U.S. sources and about 30 percent from all foreign sources. The volume and excellence of U.S. biomedical research, as well as the innovative power of industries dependent on such research (such as pharmaceuticals, medical devices, and biotechnology), can largely be attributed to NIH and its system of peer review.

But is peer review the only way to achieve success in this field? In materials science, telecommunications, space, lasers, and microelectronics—other fields in which the United States is the world leader—the nation's advantages in R&D arguably derive as much from mission-oriented agency-directed research and technology development as from peer-reviewed science. In many fields of engineering, mathematics, and physical sciences, the National Science Foundation's (NSF) base of peer-reviewed grants is complemented by other agencies' dynamic portfolio of mission-related science and technology, much of which is funded outside of peer review.

Many of these fields do seem more like engineering than pure science, and some people assume that DARPA's funding procedures are suited to technology with definite aims but not to science. Experience suggests otherwise, however. Packet switching for electronic communication, computer time-sharing, integrated large-scale chip design, and networking were as conceptually "basic" when DARPA was funding them as most molecular biological experiments are today. Nothing was there except a notion that computers could be made to do things they had never done before. When NSF and NIH both frowned upon funding work on neural networks, Leon Cooper received funding thanks to the judgment of a program manager at the Office of Naval Research (ONR), which uses a mix of peer review and DARPA-like funding mechanisms. ONR also led the way toward single-atom chemistry, "squeezed" states of light, and acoustics—all fields with a heavy dose of basic science.

Another reason to consider the DARPA approach is its lower transaction costs. Administrative review costs at NIH or NSF rise arithmetically with the number of applications. External costs, however, rise much faster as the percentage of proposals that are funded falls. If half of all proposals result in funding, which was the case at NIH several decades ago, one unfunded grant proposal is prepared for each one funded. When success rates fall to one in five or six, as they have in several areas, four or five proposals are wasted for everyone funded. Preparing a grant proposal is a substantial effort, and the total external costs for all applicants may approach or even exceed the amount awarded to the successful one. Physicist Leo Szilard once noted that, at some point in a competitive grant system, applying for grants would consume all of a scientist's time, leaving none for research. With 15- to 20-percent success rates, a "Szilard point" (where waste exceeds benefit) is no longer a frivolous speculation, but a real possibility. Whereas NIH extramural administrators spend most of their time crafting rules for competition and then selecting among applicants, DARPA staff spend most of their time keeping abreast of their fields and camping in sparsely populated outposts along the technological and scientific frontiers.

Many scientists and engineers fear that grant competition has pushed peer review well past its power to distinguish the truly outstanding from the merely excellent. The least painful solution to this problem,

at least for the scientists and engineers seeking funds, is more money for grants, so that more are funded, the success rate rises, the relative external costs fall, and reviewers need only separate the good from the excellent. To relieve the tension in the peer review system would require at least a doubling of federal research support in combination with a “birth control” policy to stem the growth of the applicant pool. Although NIH enjoys stalwart bipartisan support, a budget increase of this magnitude is unlikely; and even if budgets grow, the applicant pool may well grow faster, if history is any guide.

Although important, budget constraints and administrative inefficiency are not the most compelling reasons to experiment with DARPA-like funding mechanisms. The most serious threat to science under the peer review system is conservatism—the safe squeezing out the novel. A look at the history of NIH involvement in DNA sequencing illustrates how a DARPA-like mechanism might prove more effective than external, prospective peer review. In 1981, Leroy Hood and his colleagues at Caltech applied for NIH (and NSF) funding to support their efforts to automate DNA sequencing. They were turned down. Fortunately, the Weingart Institute supported the initial work that became the foundation for what is now the dominant DNA sequencing instrument on the market. By 1984, progress was sufficient to garner NSF funds that led to a prototype instrument two years later. In 1989, the newly created National Human Genome Research Institute (NHGRI) at NIH held a peer-reviewed competition for large-scale DNA sequencing. It took roughly a year to frame and announce this effort and another year to review the proposals and make final funding decisions, which is a long time in a fast-moving field. NHGRI wound up funding a proposal to use decade-old technology and an army of graduate students but rejected proposals by J. Craig Venter and Leroy Hood to do automated sequencing. Venter went on to found the privately funded Institute for Genomic Research, which has successfully sequenced the entire genomes of three microorganisms and has conducted many other successful sequencing efforts; Hood’s groups, first at Caltech and then at the University of Washington, went on to sequence the T cell receptor region, which is among the largest contiguously sequenced expanses of human DNA. Meanwhile, the army of graduate students has yet to complete its sequencing of the bacterium *Escherichia coli*. The point is

not that the study section bet wrong—any research funding must be fault-tolerant and take risks—but that it bet on old technology over new.

NIH and NSF have long struggled with the tendency toward conservatism in peer review. NSF has set aside small grants for exploratory research that is subject only to expeditious staff review. With NSF's tradition of grant managers rotating into and out of their fields in academia, this is similar in spirit to DARPA, although the dollar amounts are generally too small to fund more than pilot projects. NSF has a good idea, but there is no reason to believe that innovative projects are always small. Besides, requiring that innovation prove itself early with small grants may lead to premature declarations of failure and force investigators to write a follow-up grant at the same time as they have only a few months' funding to do the pilot work. At NIH, some study sections set aside specific grants or are given the option of selecting one or a few especially novel proposals for special consideration. But this does not avoid the inefficiencies of the group process and of grant proposal preparation, and it ultimately amounts to a few groups doing sporadically what individual experts might do better.

A Small Dose of DARPA

A DARPA-like funding mechanism cannot cover the same breadth of science and technology as NIH or NSF. Even if a pilot test of a DARPA-like program is a success, it still should be considered as an alternative for a few select programs only. Much of the most important work supported by NIH and NSF is conducted through tens of thousands of relatively small grants. Innovation bubbles up in unexpected places thanks to the flexibility of the grant mechanism, which leaves funds largely under the control of investigators. NIH handles 45,000 grant applications per year. It would be folly to adopt DARPA's methods for so many small projects covering enormous areas of science. The DARPA system cannot scale up easily, because its effectiveness depends on a flat bureaucracy and strong direct accountability from manager to agency director. The DARPA process is best suited to force scientific and technical progress in critical areas and to accomplish tasks when a new technology is promising but not yet proven. It is not suited to sustaining the bulk of scientific research.

DARPA-like pilot projects might be tried first by one or a few NIH institutes or center directors working with their respective councils to foster specific fields or to develop needed technical capacities. If NIH were to experiment with a DARPA-like mechanism, it should focus on areas that are ripe for such experimentation, such as:

An emerging technological capacity that would be widely beneficial if successfully developed,

- An advance promising a major leap, not an incremental improvement,
- A capacity whose development requires substantial sustained funding,
- A field or technique that is unlikely to be developed by ongoing academic efforts or within industrial firms,
- An emerging scientific field or technical area that lacks a natural disciplinary base, or
- A promising new field populated by only a few individuals.

NIH has amply demonstrated its agility and excellence, maintaining scientific quality and administering a credible and effective process for allocating funds. That solid base of peer-reviewed science should be not be chipped and fragmented. The edifice could benefit from a new wing, however, that poses little danger to its foundations. One or two institute directors could hire some rising stars and make them responsible for moving their fields ahead rapidly. After four or five years, the results of NIH's "DARPA corps" could be compared to the record of peer review groups in similar areas.

Testing a DARPA mechanism within NIH is not a call to end peer review as we know it, or even a substantial fraction of it. But neither is the generally excellent track record of NIH and NSF any proof that a DARPA-like mechanism can't improve the system. In the 1960s, C. Jackson Grayson wrote a classic work on oil drilling that demonstrated why a long-term diversified strategy is important for success when confronting uncertainty. Peer review is best regarded as a way to contend with moderate uncertainty, but it is not a good way to decide where to wildcat. DARPA's methods seem better suited to that, and some wildcatting is a good idea.

