

Funding Breakthrough Research: Promises and Challenges of the “ARPA Model”

Pierre Azoulay^{1,2}, Erica Fuchs³, Anna P. Goldstein⁴, Michael Kearney¹

¹ Massachusetts Institute of Technology, Sloan School of Management

² NBER

³ Carnegie Mellon University, College of Engineering

⁴ Carnegie Institution for Science, Department of Global Ecology

Abstract

From its 1958 origin in defense, the Advanced Research Projects Agency (ARPA) model for research funding has, in the last two decades, spread to other parts of the US federal government with the goal of developing radically new technologies. In this paper, we propose that the key elements of the ARPA model for research funding are: organizational flexibility on an administrative level, and significant authority given to program directors to design programs, select projects and actively manage projects. We identify the ARPA model’s domain as mission-oriented research on nascent S-curves within an inefficient innovation system. Finally, we describe some of the challenges to implementing the ARPA model, and we comment on the role of ARPA in the landscape of research funding approaches.

Note: All four authors were involved in the National Academies review of ARPA-E, “An Assessment of ARPA-E” (NRC 2017). Azoulay and Fuchs were Committee Members, and Goldstein and Kearney were consultants to the Committee.

Introduction

In October 1957, the Soviet Union launched Sputnik 1 into low-Earth orbit, igniting both a space race and an innovation race with the United States. The United States countered with the Apollo Program and a focus on broader innovation efforts, particularly within the Department of Defense (DOD). In 1958, the DOD launched the Advanced Research Projects Agency (ARPA) with the goal of helping the United States avoid further technological surprises. Later renamed the Defense Advanced Research Projects Agency (DARPA), the agency has grown to maintain an annual budget of \$3 billion for extramural research funding overseen by roughly 100 program managers (Gallo 2018).

DARPA is frequently heralded for its successes in funding high-risk, high-reward research. In a commissioned report to celebrate its 50th anniversary, DARPA highlighted its impact on military applications, such as missile defense and stealth technology, as well as on innovations that made their way to consumer markets, including GPS, the Internet, and autonomous navigation (DARPA 2008). Historians have attributed to DARPA a role in a number of high-profile inventions, such as the Internet (Newman 2002), the personal computer (Allan 2001), the laser (Bromberg 1991) and Microsoft Windows (Fong 2001). Piggy-backing on these perceived successes, use of the ARPA name has expanded to federal agencies outside of the DOD, to include HSARPA for Homeland Security in 2002, IARPA for the intelligence agencies in 2006, and ARPA-E in the Department of Energy (DOE) in 2009.

Given DARPA's high profile, in recent years, analysts have sought to codify the "ARPA model" of innovation (Bonvillian 2006; Van Atta 2007), but empirical evaluation of the successes and failures of the ARPA model remains elusive. One challenge for evaluating the ARPA model is the incompatibility between traditional innovation metrics and the lofty goals of ARPA organizations, e.g. "to make pivotal investments in breakthrough technologies" for DARPA,¹ or "accelerating transformational technological advances" for ARPA-E (110th Congress 2007). The returns to

¹ <https://www.darpa.mil/about-us/mission>

innovation are notoriously skew-distributed (Scherer and Harhoff 2000). For a program that makes long-term and high-risk investments, many failures can be justified by a single success. The rare nature of transformational outcomes handicaps formal program evaluation, which is geared toward measuring short-term outputs like patents and publications, or even medium-term outputs like commercialization activity. It is impossible to accurately measure the incidence of one-in-a-thousand ideas, much less one-in-a-million ideas, on a timescale relevant to political decision-making around program authorization. While DARPA has been lauded for significant technological successes ex post (e.g. the internet, GPS, etc.), ARPA-E, a relatively young agency, has yet to be associated with similar outsized wins, making the challenge of intermediate evaluation more difficult.

Nevertheless, the prospect of being able to orchestrate transformational innovation holds great appeal in the US and across the world. The question for policymakers thus remains: should the ARPA model be replicated, and if so, under what conditions? To investigate this question, we provide a summary of how the ARPA model has been implemented and what factors may be important for its success. Building on the National Academies' recent assessment of ARPA-E (NRC 2017), we focus on four defining characteristics of the ARPA model: (a) the flexible nature of its organization, (b) the identification of "technological white space" and design of programs to fill that void, (c) the discretion in selection of projects by program managers, and (d) the active management of each project using specific technical milestones and time commitments.

This paper proceeds as follows: Section 2 provides a history of (D)ARPA and its younger siblings. Section 3 outlines the elements of the ARPA model in detail. Section 4 discusses what is "ARPA-ble," to help identify technological areas where the model might bear fruit. Section 5 draws attention to several challenges for the model, and Section 6 concludes by placing ARPA-like agencies in context of the larger research-funding ecosystem.

1. History of (D)ARPA and Its Younger Siblings

Since its inception in 1958 as ARPA, DARPA has had an amoeba-like ability to evolve with new Directors in ever-changing political environments (Fuchs 2010). Defense analyst Richard Van Atta writes, “There is not and should not be a singular answer on ‘what is DARPA’—and if someone tells you that [there is], they don’t understand DARPA” (Fuchs 2010: Van Atta 2007). At various times in DARPA’s history, the agency has had many overlapping areas of focus, including broadly-defined basic research, defense applications with short-term milestones and goals, economic competitiveness and local manufacturing capabilities. Here, we provide a brief overview of this history, drawing primarily from work by Fuchs (2010), which offers a more detailed narrative.

ARPA was initially formed in response to the Soviet Union’s launch of Sputnik, a major step forward for the United States’ Cold War adversary in the space race. As such, the core mission of ARPA at inception was the prevention of technological surprises (NRC 1999). ARPA was initially tasked with overseeing America’s burgeoning space program while the National Aeronautics and Space Administration (NASA) was launched (Fuchs 2010). In 1960, with NASA up and running, ARPA shifted its focus from space to high-risk military technology that was not garnering support in the private sector, with the goal of transitioning the technologies from the lab to proof of concept (Roland 2002; Fuchs 2010). As an independent entity within the Department of Defense, the agency had to be mindful of the distribution of research and development efforts across the military services. Focusing on earlier stage research enabled ARPA to avoid overlap with more applied research and systems development.

Over the course of the 1960s, ARPA developed a flat organizational architecture, staffed with top-notch engineers and scientists from industry and academia (NRC 1999). Under Director Jack Ruina, office directors and program managers operated with significant autonomy and were encouraged to make use of available tools like “no-year money,” where the agency can reallocate unobligated funds to a subsequent year, and unsolicited proposals (NRC 1999; Fuchs 2010). Staff

were empowered to build programs independently, and the agency's direction was associated with the personalities of these staffers and the vision set forth by the Director. During the Vietnam War, however, research was redirected from high-risk, long-term research toward military applications (Flamm 1987), and ARPA rebranded as Defense ARPA (DARPA) in 1972. Under President Nixon, Congress required that all military funding have a "direct and apparent relationship to a specific military function" (NRC 1999). With these shifts, the research portfolio at DARPA took on a stricter mission orientation, which continued through the Directorship of George Heilmeier (Roland 2002).

In the decades after the Vietnam War, industrial competitiveness and dual-use technologies were prominent in DARPA's focus (Fuchs 2010). In 1987, with U.S. semiconductor manufacturers planning to purchase the majority of their equipment from Japanese suppliers, 14 U.S. semiconductor companies came together to create a not-for-profit venture, SEMATECH, to improve domestic semiconductor manufacturing (Fuchs 2010). With semiconductor manufacturing seen as vital to defense technology, the federal government appropriated \$100M annually for five years to match the industrial funding. The funding was channeled through DARPA (NRC 1999). Between 1992 and 2000, (D)ARPA returned to having a greater focus on early-stage research. From 1993 to 1996, the agency even briefly returned to the ARPA name, before becoming DARPA once again (Fuchs 2010).

During the George W. Bush Administration, the agency's focus moved away from industrial competitiveness and dual use technologies to center on mission-driven applied research. Tony Tether took over as Director in 2001 and repositioned the agency to "bridge the gap" between fundamental discovery and military use (Fuchs 2010). During this period, the proportion of DARPA funding allocated toward academic researchers declined by 50 percent, though funding levels for DARPA as a whole remained constant (Fuchs 2010; Lazowska and Patterson 2005; Markoff 2005). In contrast to the discretion given to researchers in the 1990s, Tether increased usage of timelines for project go-no-go milestones that, if unmet, resulted in the termination of projects (Fuchs 2010).

It was in this context that the ARPA model first appeared in other Departments of the federal government. In 1998, the Central Intelligence Agency (CIA) launched the Advanced Research and Development Activity (ARDA), which became the Disruptive Technologies Office (DTO) in 2006. DTO was then combined with the newly created IARPA in 2007, with the goal of exploring transformative technical breakthroughs for intelligence activities, including data collection, analysis and computing. Similarly, the Department of Homeland Security launched HSARPA in 2002 to identify operational gaps that could be filled with technological advance within the agencies of Homeland Security: U.S. Customs and Border Protection, Secret Service, and the Coast Guard.

In its 2007 report *Rising Above the Gathering Storm*, a committee of the National Academies called for the creation of a DARPA-like agency within the Department of Energy (NRC 2007). In the same year, the America COMPETES Act authorized the creation of ARPA-E, which was tasked with “identifying and promoting revolutionary advances in fundamental science; translating scientific discoveries and cutting-edge inventions into technological innovations; and accelerating transformational technological advances in areas that industry itself is not likely to undertake because of technical and financial uncertainty” (110th Congress 2007). ARPA-E received initial funding in the American Recovery and Reinvestment Act of 2009, and has been in operation since.

There have been four agencies to use the ARPA name (DARPA, IARPA, HSARPA and ARPA-E). Others, such as the Advanced Technology Program (ATP) in the National Institute of Standards and Technology (NIST), and Biomedical Advanced Research and Development Authority (BARDA) in the Department of Health and Human Services (HHS), have been inspired by DARPA (Sen 2017). Of the ARPA agencies, DARPA has received by far the most attention, with appearances in popular media such as *Playboy Magazine* and the television show “*The West Wing*” (Fuchs 2010), Congressional testimony (Mowery 2006, Van Atta 2007), and recent books in the popular press (Belfiore 2010, Jacobsen 2015, Weinberger 2017). More recently, ARPA-E has been subject to scrutiny and empirical review by the National Academies (National Research Council 2017). In the next section, we focus on DARPA and ARPA-E as the two best-known

implementations of the ARPA model. We identify common elements across the two that, together, comprise the core of the “ARPA model.”

2. Elements of the ARPA model

Despite a great deal of commentary on DARPA, lack of access to internal archival data has hampered efforts to study it empirically. One notable exception is the work of Roland and Shiman (2002),² who offer an industrial history of DARPA’s effort to develop machine intelligence under the “Strategic Computing Initiative.” They emphasize both the agency’s positioning in the research ecosystem—carrying military ideas to proof of concept that would be otherwise neglected—as well as the program managers’ role as “connectors” in that ecosystem. Roland and Shiman are to our knowledge the only academic researchers ever to receive internal access to DARPA’s archives. Recent work by Goldstein and Kearney (2018a) on ARPA-E is to-date the only quantitative analysis using internal program data from an ARPA agency.

Others have leveraged oral histories and interviews with program managers combined with external archival data to document the practices of ARPA agencies. Fuchs (2010) describes DARPA’s function as “embedded network governance” based on the program manager’s engineering of social networks in order to influence technology directions. Supplementing interviews with invention disclosure data, Colatat (2015) finds that MIT investigators with DARPA funding were more likely to form novel collaborations than those with funding from other agencies. Sen (2017) finds that ARPA-like agencies see themselves as “islands” separated from conventional operations as well as “bridges” between R&D activities and the broader industrial context.

Meanwhile, analysts and observers have persisted in the attempt to describe the key elements of the ARPA model. Defense analyst Richard Van Atta, in his Congressional testimony on lessons from DARPA for the establishment of ARPA-E, emphasizes that DARPA is “lean and agile, with a

² For insights into this painful process, see the preface of Roland and Shiman (2002).

risk-taking culture” (2007). Bonvillian builds on Roland’s work, depicting DARPA as the heir to a WWII model of “‘connected’ science, where technology breakthroughs at the fundamental science stage were closely connected to the follow-on applied stages of development, prototyping and production” (2009). Greenstein (2011), meanwhile, portrays DARPA as a kind of “skunk works,” in which innovative projects are insulated from the routine operational concerns of the organization.

Much of the discussion about ARPA-like agencies relates to the role of the program manager. DARPA has a reputation for its flexibility and speed, partially as a result of the forward-looking and aggressive culture of program managers, often described as “100 geniuses connected by a travel agent” (Bonvillian 2009). ARPA-E has likewise been described as cultivating an “innovative startup-like culture” with energetic, highly-qualified program directors (Rohlfing, 2017).³ Van Atta summed up the importance of the program manager in this way:

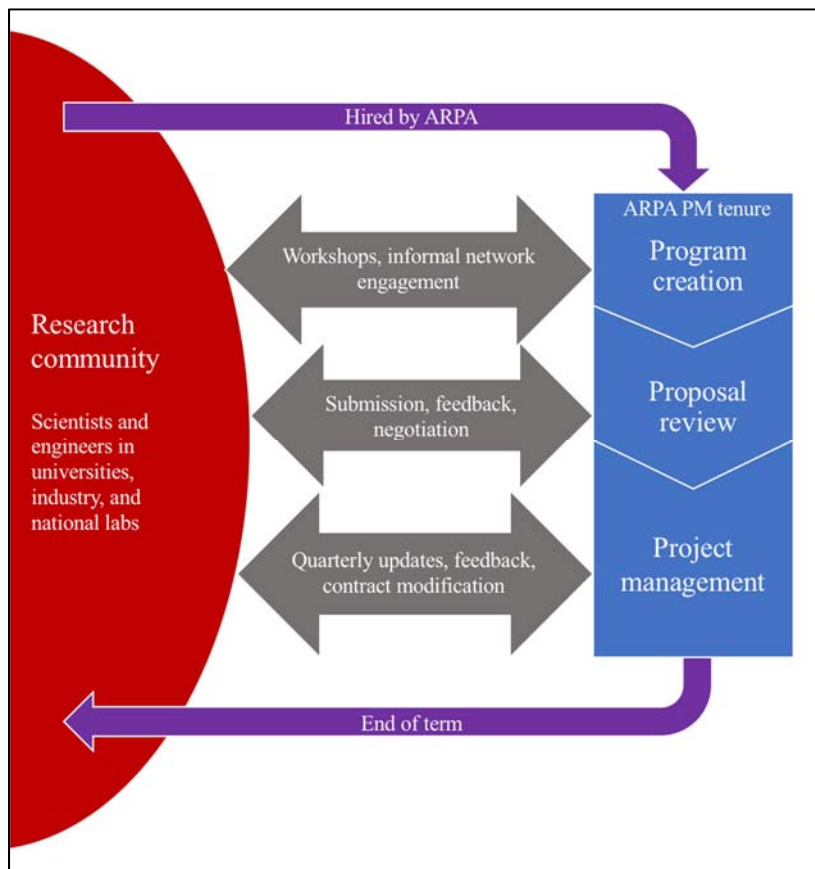
The DARPA program manager is, in fact, the key. He or she is the technical champion who conceives and owns the program. He is not told what to do, though he does have to have approval from his office director, and from the DARPA Director. Once he starts that program, it is his, and he makes it happen, and he has to make the choices involved in that. So, in essence, they are risk-taking, idea-driven entrepreneurs heading up their own practice. (Van Atta, 2007)

DARPA and ARPA-E program managers leave their positions in academia, industry, or elsewhere in government to join the agency. They remain continually engaged with the research community throughout the planning and execution of a technical program (Figure 1), and then return to work outside the agency after a short three to five year term. Serving as a DARPA program manager has been a launching pad for prominent careers in science and technology management across the country: Pradeep Khosla (former Dean of Engineering at Carnegie Mellon University and current Chancellor of UCSD), Zachary Lemnos (former Chief Technology Officer

³ Program Directors at ARPA-E are roughly equivalent to the position of Program Manager at DARPA. We use each term when referring to the program staff of the relevant agency, and we use “program manager” to refer to the position in a generic sense.

of MIT Lincoln Laboratory and former U.S. Assistant Secretary of Defense for Research and Engineering), Arati Prabhakar (former Director of the U.S. National Institute of Standards and Technology and former Director of DARPA), and Robert Taylor (founder of Xerox PARC's Computer Science Laboratory) were all DARPA program managers earlier in their careers.

Figure 1. The Lifecycle of a Program Manager at an ARPA Agency



In this section, we structure our discussion of the ARPA model by highlighting four common institutional features of DARPA and ARPA-E. Although each of these may be adhered to by an ARPA-like agency to varying degrees at any given time, we argue that together they form the core of the ARPA model. They fit into four buckets: (1) general organizational flexibility, (2) bottom-up program design, (3) discretion in project selection, and (4) active project management.

Importantly, the last three elements highlight the critical responsibility of program managers within ARPA-model agencies, and each element corresponds to one stage of an ARPA program:

General Organizational Flexibility. Contrary to most agencies in the federal bureaucracy, the ARPA model entails significant organizational flexibility, allowing agencies to respond more quickly to changing technological conditions. Between DARPA and ARPA-E, there are a number of policies and practices that provide this flexibility:

Independence. ARPA-E and DARPA both operate independently from other branches of DOE and DOD respectively. The Director of ARPA-E reports directly to the Secretary of Energy, which allows ARPA-E to operate outside the many levels of bureaucracy encumbering other research organizations within DOE. DARPA is also known for reporting “to the top” with minimal bureaucracy (Van Atta 2007). Indeed, part of the motivation for DARPA’s founding was the argument that rivalry between the services led to the Soviets being first to space; after early criticism, ARPA avoided hiring program managers out of the services (Roland 2002).

Flat organization. Within the agencies, ARPA-E and DARPA both have a flat internal structure. All ARPA-E program directors report to the agency’s Deputy Director for Technology. With an order of magnitude larger budget than ARPA-E, DARPA has only one additional degree of hierarchy with its six program offices. The fact that DARPA only funds extramural research allows it to be less subdivided than other Department of Defense agencies such as the Army Research Lab, which is organized into directorates each consisting of multiple divisions.

Hiring outside of civil service regulations. DARPA and ARPA-E can hire program managers without being confined by the typical requirements for hiring federal employees or contractors. This allows them to hire top talent quickly and with more competitive salaries. In the context of DARPA, existing office directors and program managers often “tap” young star researchers they meet during their own work at DARPA and encourage them to consider becoming a DARPA program manager.

Fixed-term employment. Program managers are hired from technical positions in academia, industry, and government, and they typically serve one term of 3-5 years (Fuchs 2010, NRC 2017). This organizational turnover is likely advantageous both for the organization, which receives a regular infusion of new ideas, and the program directors themselves, who can use this prestigious appointment as a stepping stone to top-tier science and engineering leadership positions. Rotation between industry, academic, and government positions is common for leaders in Japan, many of whom believe that the “bounded vision” of for-profit firms and government can be overcome by bringing the two together (Fransman 1993).

Flexible contracting. Performers at ARPA-E and DARPA may be funded through common mechanisms (grants, contracts, cooperative agreements), but both DOD and DOE also provide agencies “other transaction authority” or OTA, which they may use in cases where the terms of traditional agreements are not suitable. Examples include transactions for prototyping, or technology investment agreements (TIA) with customized terms related to intellectual property (Institute of Medicine, 2012). This flexible contracting enables DARPA program managers to fund investigators with millions of dollars practically overnight (Fuchs 2010).

“I have described the role of DARPA PM like that of a conductor of an orchestra. The conductor does not play any instrument, but he must understand every instrument and the music. He brings all the harmony together.” – DARPA Program Manager

Bottom-Up Program Design. Using the flexibility of the organization as described above, program managers at ARPA-like agencies are able to pursue technical areas based on the real-time changing need for innovation and opportunities for progress. These short-lived technical programs contribute to the reputation for agility at ARPA-like agencies, as they can redirect resources with relative speed. This bottom-up approach to designing technical programs is the second core element of the ARPA model.

At the start of an ARPA-E program director’s three-year term, they are tasked with designing a technical program—a specific challenge under which a set of projects will be funded, e.g. \$30

million for a program that may award 10 teams with \$3 million each. Program directors engage with their target research communities and host stakeholder gatherings to refine their idea, and then they pitch their program to ARPA-E leadership, who either accept the idea or encourage further exploration. This process can take up to 18 months of a 36-month contract, so each program director typically creates one program during their tenure. With 15 or so program directors at a time at ARPA-E, and each program consisting of projects that are typically 3 years long, there have already been dozens of programs completed in the short lifespan of the agency.

DARPA uses similar processes for program creation. DARPA program managers identify technology directions by simultaneously engaging with military liaisons to understand their mission-related needs, talking to and convening scientists to brainstorm new technology directions to meet those needs (Fuchs, 2010). Identifying technology directions occurs both formally through working groups such as the Information Science and Technology (ISAT) group and informally, through one-on-one conversations and flying out a group of scientists to brainstorm ideas together (Fuchs 2010).

When choosing an area of “white space” in which to operate, program managers consult with industry and other experts, and they seek to address bottlenecks in the innovation system (NRC 2017). They craft programs to tackle areas where a need exists and where a targeted investment could make a significant difference. The agency’s administration may seed this process by hiring program managers from a certain technical background, e.g. ARPA-E hiring someone with experience in geology in hopes that they will explore advanced geothermal energy, but the program directors are then allowed to explore with relative freedom.

ARPA programs are often designed with specific technical targets. For example, the ARPA-E Grid-Scale Rampable Intermittent Dispatchable Storage (GRIDS) program was created in 2010 to support novel energy storage or battery concepts, but a pre-requisite for funding was a cost projection of \$100 per kilowatt-hour—a step-change in storage costs that would fundamentally transform the electricity industry. Interestingly, the cost threshold set in this program has become

the *de facto* benchmark for industry costs going forward, as well.⁴ The specificity of these targets at the program-level translates to concrete benchmarks for the success of each project, which is critical for active project management, discussed later in this section. These targets also ensure that each program is directly aligned with some component of the agency's mission.

The structure of bottom-up program creation is key to the ARPA model, in as much as it enables ideas to meet mission challenges to be emergent and program managers to deliberately connect people who are otherwise unconnected to shape the direction of research. As one DARPA program manager described this process, "You get communities together that don't naturally talk and you give them some latitude and some life, and you push them forward and see what comes out of it" (Fuchs 2010). Beginning with pre-program workshops and continuing through quarterly presentations to each other (in the case of DARPA) or annual meetings of ARPA performers (in the case of ARPA-E), a community of people is built around a particular goal, and this function can have lasting impact beyond the specific outcomes of any particular project.

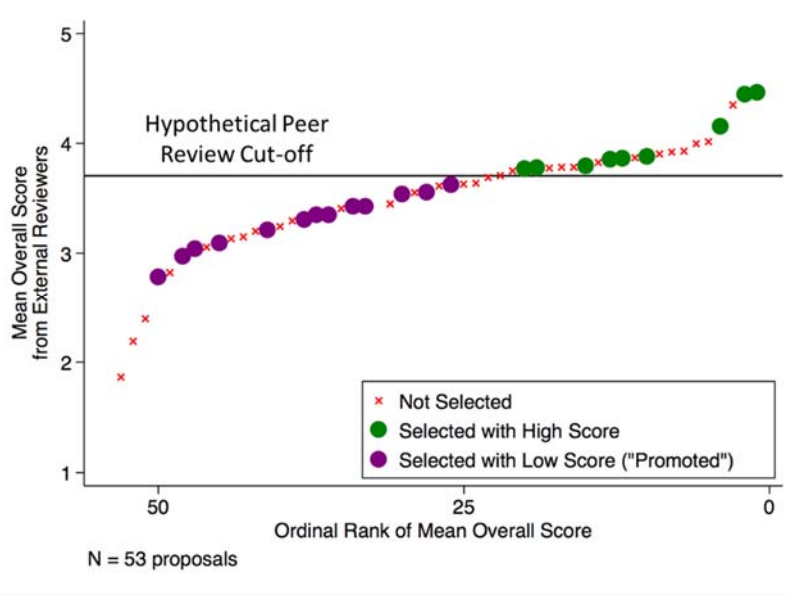
Discretion in Project Selection. The third important feature of the ARPA model is the use of discretion by program directors to decide how to allocate funds within a program. Any research funding program must tackle fundamental questions: Which researchers should we support? Who can offer the greatest return on investment? What technologies fit best within the agency's mission? These questions are extremely difficult to answer due to the uncertainty of research outcomes (Rosenberg 1996). Research creates new knowledge or information, and so by definition, it is impossible to predict the result of a particular project. Furthermore, even if a project meets its technical goals, there is great uncertainty in how the research will be used and what the impact of that application will be. Regardless of which selection criteria a funding program uses, e.g. scientific merit or applicability (Gans and Murray 2012), the inherent uncertainty of research presents serious difficulty for evaluating proposed projects.

⁴ See for example, "Eos Energy Storage drives down costs on battery systems to below \$100 per kWh": https://www.renewableenergymagazine.com/energy_saving/eos-energy-storage-drives-down-costs-on-20170419

Given this inherent uncertainty, the most common method of coping with technical uncertainty is peer review, i.e. to solicit opinions of the technical proposals from a panel of expert reviewers, who are presumably the most informed and least burdened by uncertainty. Of course, subject matter experts may disagree, so programs typically combine multiple opinions in order to arrive at a minimally controversial set of proposals. For example, proposals may be sorted on the basis of mean numeric ratings received from the review panel. There exists a vast literature on the deficiencies of peer review for research proposals, such as its biases based on demographics, seniority and, importantly in this context, novelty (Boudreau et al. 2016). This bias against novelty is particularly concerning for research organizations that are committed to pursuing transformative research, e.g. ARPA-style programs.

The alternative strategy employed in the ARPA model allows program directors to use their own expert judgment in deciding which proposals to fund. Program managers at DARPA have great latitude in selecting performers and distributing funds without soliciting external peer review. ARPA-E tends to be under greater scrutiny than DARPA—clean energy being more politically vulnerable than defense—and so it has adopted a more moderate strategy of collecting external peer reviews as informational inputs. Goldstein and Kearney (2018a) describes how ARPA-E program directors select projects from across the distribution of review scores, i.e. average review scores are only slightly correlated with an increased probability of selection. The authors construct a counterfactual cut-off line for review scores, below which an application would be rejected by a selection method based on ranking average scores. In the example program shown in Figure 2 below, 22 proposals were selected for funding. Goldstein and Kearney (2018a) creates a hypothetical peer-review score cut-off line to represent the cut-off line used if peer-review scores completely determined selection. Instead, over half of the projects selected were selected despite being below the hypothetical peer-review cutoff line. Across the entire ARPA-E portfolio 2009-2015, approximately half of all projects were similarly “promoted” (Goldstein and Kearney 2018a).

Figure 2: Review Score Distribution and Project Selection for an ARPA-E Technical Program



Rather than indiscriminately following the advice of peer reviewers, some program directors report selecting projects based on the review comments instead (Goldstein and Kearney 2018a). Program directors at ARPA-E and DARPA often balance programs with complementary, and potentially competitive, technologies, in order to create a diverse portfolio (Fuchs 2010). To the extent that expert disagreement is useful in measuring the degree of technical uncertainty within a research portfolio, the use of individual discretion in the case of ARPA-E results in a portfolio of projects that carry a higher degree of uncertainty, as defined by disagreement among reviewers. Moreover, under the premise that uncertainty is a corollary to novelty, individual discretion is an antidote to novelty bias in peer review.

In addition to using discretion to select performers, ARPA program staff take an additional step of actually engineering the teams that submit proposed ideas. At ARPA-E, program directors work to create connections between different researchers with complementary expertise, such as introducing academics to potential commercial partners. At DARPA, evidence suggests that program managers are even more involved in the engineering of research teams (Fuchs 2010, Colatat 2015). Berlin (2017) documents the case of ARPANET, where the program manager,

Robert Taylor, who went on to become the founder and manager of Xerox PARC, decided to award the contract to an assembled team of performers including Bolt, Beranek and Newman, rather than to Raytheon. Ultimately, the program manager believed the assembled team to be a better “cultural fit.”


Active Project Management. The fourth core element of the ARPA model is active project management, which entails empowering program staff to make decisions related to capital, tasks, milestones and technical goals throughout the project. The practice of active project management is consistent with the literature on R&D project management within the private sector. For example, Huchzermeier and Loch (2001) detail the importance of maintaining flexibility for a firm through the use of “real options” for managing a research project, such as delay, abandon, contract, expand and/or switch. Similarly, Trigeorgis (1997) finds that flexibility improves the potential upside of a project while limiting downside losses. Dixit and Pindyck (1994) find that as uncertainty in project payoffs increase, additional measures should be taken by managers to maintain flexibility.

ARPA agencies generally stand out among other public grant-making organizations for retaining control rights over the research after allocating funds. The conventional approach is to allow researchers to manage their own research process independently. This approach is especially well-supported for academic researchers, who value having creative control over their research directions (Aghion et al. 2008; Stern 2004). An often-cited model of investigator independence is the Howard Hughes Medical Institute (HHMI), which funds “people, not projects” and gives its performers authority to choose their own research direction. Azoulay et al. (2011) found that HHMI performers, who are free to choose and modify their own research directions, publish more high-impact papers compared to similarly qualified scientists funded by the National Institutes of Health (NIH) with pre-defined project aims.

In this context, reducing the level of authority and flexibility held by funded researchers might be expected to hinder exploratory research (Manso 2011). However, at ARPA agencies, this loss coincides with the empowerment of ARPA program managers. In the ARPA model, the nexus of exploration occurs first and foremost at the level of the program director, through the program creation and project selection processes described above. The program directors themselves have tolerance for failure as they make decisions, i.e. setting the technical goals, selecting a preferred set of projects, and managing those projects along the way. For this reason, hiring talented program staff with a penchant for exploration is pivotal to the success of ARPA-like programs.

Over its history, DARPA’s project management practices have varied between more and less active; broad area announcements and betting on the right people were hallmarks of DARPA in the 1960s and 1990s, while milestones and “go/no-go” reviews were emphasized under Directors Heilmeier (1975-1977) and Tether (2001-2008) (Fuchs 2010). Throughout these periods DARPA performers have been required to report quarterly on progress. At ARPA-E, program directors also remain closely engaged with researchers over the course of a project, receiving quarterly progress updates and giving feedback on the same schedule (Goldstein and Kearney 2018b). ARPA-E program directors frequently revise project milestones, budget, and timeline for deliverables; in some cases, they cancel projects that are not able to meet technical milestones.

Figure 3: Budget and Project Length Modifications at ARPA-E



As seen in Figure 3, budgets and project lengths were more likely to be expanded than to be cut at ARPA-E. These adjustments correlate significantly with overall quarterly status ratings; projects rated well are likely to see their budgets increased, and projects that miss major milestones are far more likely to be terminated (Goldstein and Kearney 2018b). By cutting projects short in response to poor performance, the agency is able to mitigate any increased risk in their portfolio caused by funding more controversial projects.

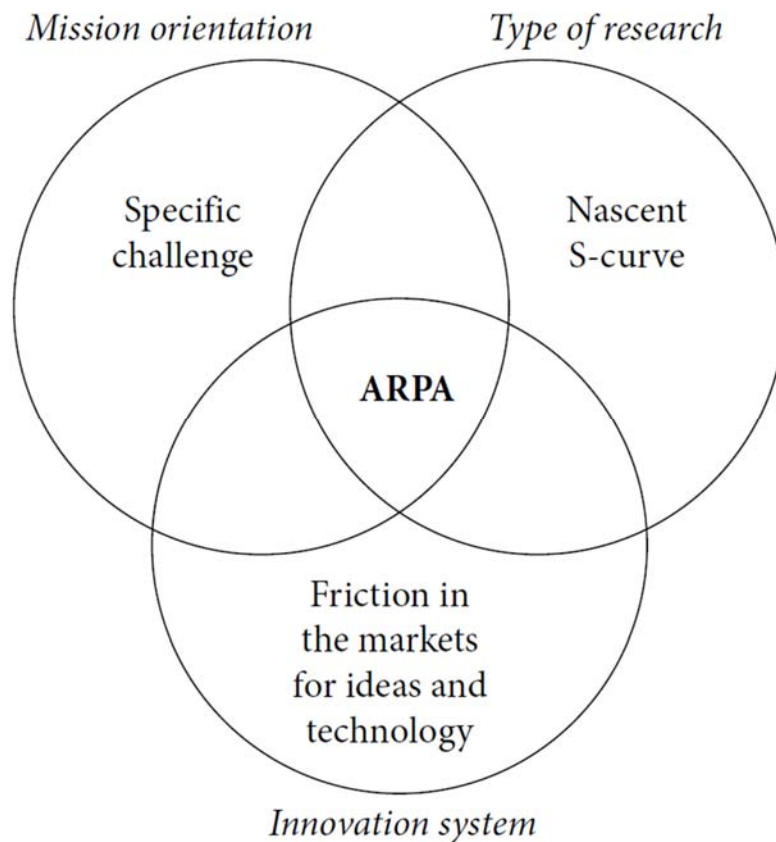
Complementarities. So far in this section, we have outlined the four practices related to organizational structure, program design, project selection and project management that are the hallmarks of the ARPA model. Together these choices lodge much more control and decision-making power with the funding organization than is typical for public research funders. We propose that this is no coincidence—the elements of the ARPA model may comprise a bundle of complementary practices. Milgrom and Roberts (1990) articulate both the benefits of complementary elements within a bundle of organizational practices and the drawbacks to selecting elements of a bundle that do not cohere; many examples of complementary business practices have been identified (Brynjolfsson and Milgrom 2013). The parameters of innovation management discussed in this section may function the same way: adopting one feature increases the return to adopting another.

Each of the four elements of the ARPA model relies on a highly talented, independent, and empowered program staff. These individuals, having been given great autonomy in some aspects of their position, are better prepared (and may be expecting) to exercise autonomy in other aspects as well. For this reason, mixing and matching among these elements would be less effective than implementing them all together. A program that uses individual discretion to select funding recipients may underperform in the absence of active project management. These interactions among policy choices are as yet unexplored for R&D programs and are an important area for future research.

3. What is ARPA-ble?

In the previous section, we described four elements of the ARPA model and proposed that these four are complements with respect to empowerment of the program staff. The obvious next question is: when should this particular set of research management practices be chosen over another? In short, what is ARPA-ble? We draw attention here to three important characteristics of a research area that determine whether it is appropriate for an ARPA-like agency: the presence of a clear mission, the type of research being conducted, and the status of the broader innovation system within which technology is researched, developed and deployed (Figure 4).

Figure 4: What is ARPA-ble?



Mission Orientation. The first requirement for an ARPA-ble challenge is that it must be possible to organize the domain of research around a technology-related mission or a set of overarching goals. According to Sen (2013), transformational innovation programs are targeted specifically toward desired social or economic outcomes, rather than intangible goals or undefined technical development. The ARPA model is not appropriate for so-called basic research, which “is performed without thought of practical ends” (Bush 1945) and is far upstream from any commercial activities, e.g. high energy physics.

DARPA and ARPA-E both have specific missions.. In the case of DARPA, the agency’s mission is to advance technologies that could be critical to national security. In the case of ARPA-E, the agency’s mission, as codified in statute, is to enhance energy security through the reduction of energy imports, the reduction of greenhouse gas emissions and increasing energy efficiency throughout the economy. An ARPA agency’s mission is the foundation of the program design process described in Section 3, in which a program manager seeks a particular innovation or technological change. He or she then uses active project management to guide the pursuit of breakthroughs to enable this change (Carleton 2010). Bonvillian (2009) calls this type of challenge-based model “right-to-left,” in that it moves opposite the archetypal innovation pipeline, starting with a desired technology and working backwards to determine the research needs.

Mowery (2006) argues that in addition to a “clear mission” DARPA having the military as its “single customer” is essential to the agency’s success. However, in contrast to the early years of DARPA, the military is no longer the primary customer for many of the technologies it needs, such as computers or microprocessors. As such, DARPA itself has had to find ways to influence technology directions in technologies critical to the military, but in which the military is only a very small fraction of overall demand. The military’s clear mission is therefore more important to DARPA’s function than the military’s role as a customer.

It is important to note that not every technical area united by a specific mission or challenge is a good fit for the ARPA model. The mission must be associated with quantifiable goals and sub-

goals with trackable progress metrics. Technical pursuits like the NIH BRAIN Initiative, with the goal of “revolutionizing our understanding of the human brain”,⁵ may not be a good fit for an ARPA model organization, as the degree to which scientists understand the human brain is not practically measurable. Similarly, agencies such as the Department of Energy’s Office of Science that pursue “blue sky” or basic research, for which sub-goals are often hard to define, are ill-suited for the ARPA model. We turn to this topic in our next section.

Nascent S-Curve. From the discussion of mission orientation, it is apparent that ARPA-ble research is distinct from basic research—and yet it is also distinct from so-called applied research, which is geared toward performance improvement of existing technologies. The ARPA model’s emphasis on risk-taking and novelty make it a poor fit for a program focused on extending an existing technological paradigm. But if ARPA agencies do neither “basic” nor “applied” research, then how can we describe their function?

Recent research by Goldstein and Narayanamurti (2018) suggests that ARPA-E occupies a productive middle ground between basic and applied research. ARPA-E awards have a higher rate of patenting and the same or higher rate of publishing than similar awards from other branches of the Department of Energy, which are sharply divided between “basic” and “applied”. ARPA-E awards are also more likely to produce *both* a patent and a publication, indicating an openness to both invention and discovery. Relatedly, soon after its founding, DARPA began to emphasize the transition of technology out of the laboratory and into the hands of users or producers who would bring it to full adoption (Roland and Shiman 2002). This role of ARPA-like agencies in bridging the gap between basic and applied research is our second distinguishing characteristic of ARPA-ble challenges.

One way to define the ARPA domain takes advantage of the concept of technology S-curves by Foster (1986), where technical performance tends to improve as a sigmoidal function of research

⁵ <https://www.braininitiative.nih.gov>

effort. Individual technologies progress along an individual S-curve, with occasional discontinuities resulting from a radical breakthrough and creation of a new S-curve. The beginning of an S-curve represents a novel concept that has not yet been the focus of much effort. Here, the returns to effort are small. The middle portion of the S-curve includes the inflection point where the returns to effort are highest. The final segment of the S-curve entails diminishing returns to effort and slow improvements to technical performance.

It seems that the ARPA model is optimized for technical areas that reside in nascent S-curves – the technology exists, is relatively unexplored, and has great potential for improvement. The practice of program design identifies those areas; the practice of project selection finds uncertain but promising research efforts in pursuit of the specific technical challenge; and the practice of project management focuses effort as more information becomes known and some ideas begin to show more potential than others. Organizational flexibility gives the agency the ability to quickly respond to opportunities with minimal administrative friction.

We can also consider the ARPA agencies in the taxonomy introduced by Stokes (1997), in which research activities are divided across two dimensions: whether the use is considered, and whether fundamental understanding is the goal. In this framework, research is either purely basic (Bohr’s quadrant), purely applied (Edison’s quadrant), or use-inspired basic research (Pasteur’s quadrant). The ARPA model certainly puts “use”, i.e. technology, at the forefront—the potential applicability of research is the foundation of the ARPA-like agency’s strategy—but what is the role of science in the ARPA model?

At the early stage of an S-curve, there may still be scientific challenges to be met before the technology can progress. Indeed, ARPA agencies are open to advancing the scientific frontier. And yet, because fundamental understanding is not *in itself* a goal of the ARPA-like agency, it is not properly identified as “use-inspired basic research” in Pasteur’s quadrant. Instead, the ARPA model occupies a portion of Edison’s quadrant wherein the “use” under consideration is still novel and immature. If scientific discovery is necessary for a given research project, then that project may journey incidentally into Pasteur’s quadrant as it furthers fundamental understanding.

Frictions in the Markets for Ideas and Technology. One final important element to consider for use of the ARPA model is the state of the broader innovation system in a given technical area. The economic justification of government funding for research is that private firms would not conduct the same work. The private sector chronically under invests in R&D because firms cannot accrue all of the benefit of an innovation. This justification is commonly applied to early-stage, upstream research, and yet it can also apply to later-stage R&D if private markets are not able to successfully move technologies to commercialization. If there are frictions in the pathways through development, demonstration and broad-scale deployment, the ARPA model for intervention is suitable. The defense and energy industries are two examples of industries where such frictions curtail innovative effort.

In the energy industry, the path from idea generation to impact for a new technology can be extraordinarily arduous. Many innovations require large amounts of capital for demonstration and scale-up, and innovators generally have to compete with legacy providers immediately upon market entry. With no clear path to commercialization and significant regulatory constraints, financing the development of new energy technologies is often extremely challenging. Howell (2017) documents the presence of financial constraints for technologies in the energy sector that are particularly acute for early-stage and/or hardware firms. The situation is very different in the pharmaceutical industry, where there is a well-established commercialization path for new drugs where a start-up will develop a drug through early stages of clinical trials, and then more established firms will buy out the start-up for further commercialization (Goldstein et al. 2017).

In the defense industry, there are a separate set of frictions that impede the commercialization process. National security is a public good with an associated collective action problem whereby there will be chronic underinvestment in mission-critical innovation. In addition, if a firm does produce an innovative product for the defense department, that product could be subject to secrecy

requirements, whereby the firm may be unable to repurpose the product in other industries or abroad. While DARPA performers themselves do not typically have military clearances, DARPA program managers can require sensitive information to connect their funding decisions to specific military needs.

4. Challenges for ARPA-like Agencies

Even in areas where the ARPA model is a good fit, there remain challenges that must be overcome to maximize the effectiveness of the model:

Measuring long-term transformation. The mission of an ARPA-like agency often involves transformation of some industry or market to address a societal need. This transformation invariably occurs over the long-term, and yet, the ARPA model requires that individual projects be held to strict timelines; projects often only run for three years. In the case of ARPA-E, for example, a project that aims to create a new design for an energy storage device may be a tremendous success, and still the impact in the market and/or the environment will not be seen for decades. This limitation must be kept in mind during short-term efforts by Congress or the agency itself to measure the agency's progress toward its mission.

One size does not fit all. ARPA model organizations fund a diverse set of performers, including large firms, academics, start-ups and national labs. They generally have one set of practices that apply equally to all research teams, despite the very different incentives and organizational contexts underlying participation across these different types of performers. Successfully implementing the ARPA model means navigating these differences to avoid conflicting motivations for researchers.

The dark side of discretion. In the ARPA model, program directors have significant freedom to craft programs and select and manage projects, and the Director also enjoys freedoms afforded by the organization's flexibility. As more freedom is given to both the Director and the program staff, the quality of the agency's activities are increasingly dependent on the talent of these individuals. Any research agency requires a Director with strong leadership ability and staff members with high

technical acumen, and yet the organizational flexibility and program staff empowerment in the ARPA model raises the stakes on this requirement. The core elements of the ARPA model cannot exist under micromanagement by the Director or with under-qualified program staff.

Tension between autonomy and accountability. Another risk faced by the ARPA model is that, with fewer checks and balances in place, empowered staff members are less accountable for the impact of their decisions. The Director is accountable to Congress and must defend the organization's budget, and the program directors are accountable to the Director, who decides the funding amount allocated to each program. However, the short tenure of program directors means that each program is an isolated endeavor, and success or failure of the program does not necessarily affect the program director's next career move. Some accountability is levied informally by the network of program directors within the technical community, in that program directors who are perceived as having good performance can use ARPA as a stepping stone to subsequent high-level R&D management opportunities.

Building trust. The program design and proposal solicitation phases of the ARPA model depend heavily on the willingness of the researcher to share their ideas and their ongoing, and potentially confidential, work. If program managers or the agency director do not treat these ideas with appropriate care, the agency can lose its reputation as a trustworthy partner (Fuchs, 2010). The agency must cultivate a healthy relationship with the research community in order to maintain a steady influx of ideas.

Nurturing culture with high turnover. An energetic and enterprising culture is a hallmark of the ARPA model. However, culture is a byproduct of the individuals that make up the organization, and in ARPA-like agencies, the program directors work on three-year contracts. It is a challenge to maintain and nurture an organization's culture with such high turnover. Nonetheless, there are a few mechanisms for perpetuating a culture of risk-taking and exploration: first, in the program creation process when new program directors pitch their ideas to existing program directors and the agency's administration, and second, when the program directors meet again with the

administration to justify their selections of which proposals to fund. The alumni networks of program directors also help to maintain the reputation and culture of the organization.

Transitioning to market. The focus of the ARPA model is on making progress along technical S-curves, but innovation also requires diffusion of new technologies. Overseeing the transition from research to development is a major challenge for ARPA-like agencies, especially when there is no public customer for the innovation. ARPA-E has a “tech-to-market” team dedicated to working with all performers on their strategy for transitioning technologies out of the lab, and yet a recent assessment by a committee of the National Academies found that this program should be “reconceptualized” (NRC 2017).

5. Conclusion

In this paper, we review the key elements of the ARPA model for research funding: organizational flexibility on an administrative level, alongside significant authority vested with program directors to design programs, select projects and actively manage projects. We identify the ARPA model’s domain as mission-motivated research on nascent S-curves within an inefficient innovation system. We also describe some of the challenges to implementing the ARPA model.

It is important to note that the ARPA model is but one part a landscape of research funding programs, each with different goals and accordingly different approaches. In Figure 5 below, we depict two dimensions of research management: the source of idea generation, and the locus of control for project execution. The bottom right quadrant represents the ARPA model, based on two of the four features that we describe above: ideas are submitted in response to mission-inspired programs created by the funder, and projects are then actively managed by the funder.

At the national or regional level, innovation ecosystems could embrace a diversified mix of funding models, including some programs that adhere to the ARPA model, while others fill out the rest of the landscape, producing different types of outcomes and associated benefits. The top left quadrant represents programs where researchers generate ideas and then pursue them with

little oversight and often over long time horizons. HHMI and some programs within NIH or NSF fit this category. Programs in the top right quadrant solicit proposals for a particular mission-inspired area and then allow the researcher to freely explore, as with the Gates Foundation or the Chan-Zuckerberg Initiative. In the bottom left quadrant, researchers propose ideas, but program staff are deeply involved in the execution of the project, as is common in venture capital.

Figure 5: Innovation management strategies

		<u>Idea Generation</u>	
		Investigator initiation:	Mission-inspired solicitation:
Investigator Freedom:	<u>Project Execution</u>	HHMI NIH NSF	Gates Foundation Chan-Zuckerberg Initiative
Empowered Program Staff:		Venture Capital Google X	DARPA ARPA-E

One of the challenges of the ARPA model is the difficulty of evaluating its effectiveness toward long-term goals, and yet the task of evaluation is critically important nonetheless. Following on the recent assessment of ARPA-E, research on the ARPA model should continue to probe the effectiveness of specific features. Within the existing ARPA agencies, there is some low-hanging fruit for evaluation: varying the tenure of program directors, frequency of reviews, duration of projects, and funding levels.

The available evidence on R&D management indicates that different practices yield different types of innovation-related outcomes. The ARPA model addressed here is one set of practices geared toward mission-oriented research on nascent technologies. It fits within a broader innovation landscape that includes support for early-stage “blue sky” research and later-stage incremental development. Importantly, ARPA-like organizations are not a substitute for other

sources of R&D support, but instead serve as a complementary piece of a diverse innovation system.

References

- 110th Congress, 2007. America COMPETES Act. United States.
- Aghion, P., Dewatripont, M., Stein, J., 2008. Academic freedom, private sector focus and the process of innovation. *RAND J. Econ.* 39, 617–635.
- Allan, R., 2001. *A History of the Personal Computer*. Allan Publishing, Ontario, CA.
- Azoulay, P., Graff Zivin, J.S., Manso, G., 2011. Incentives and Creativity: Evidence from the Academic Life Sciences. *RAND J. Econ.* 42, 527–554.
- Berlin, L., 2017. *Troublemakers: how a generation of silicon valley upstarts invented the future*. Simon & Schuster UK.
- Belfiore, M.P., 2010. *The department of mad scientists: how DARPA is remaking our world, from the internet to artificial limbs*. Harper.
- Bonvillian, W.B., 2009. The Connected Science Model for Innovation—The DARPA Role, in: *21st Century Innovation Systems for Japan and the United States: Lessons from a Decade of Change*. National Academies Press.
- Boudreau, K.J., Guinan, E., Lakhani, K.R., Riedl, C., 2016. Looking Across and Looking Beyond the Knowledge Frontier: Intellectual Distance and Resource Allocation in Science. *Manage. Sci.*
- Bromberg, J., 1991. *The Laser in America, 1950–1970*. MIT Press, Cambridge, MA.
- Brynjolfsson, E., Milgrom, P., 2013. Complementarity in Organizations, in: *The Handbook for Organization Economics*. pp. 11–55.
- Bush, V., 1945. *Science - The Endless Frontier*.
- Carleton, T.L., 2010. *The Value of Vision in Radical Technological Innovation*. Stanford University.
- Colatat, P., 2015. An organizational perspective to funding science: Collaborator novelty at DARPA. *Res. Policy* 44, 874–887.
- DARPA, 2008. *DARPA: 50 Years of Bridging the Gap* [WWW Document]. Faircount Media Gr. URL <https://issuu.com/faircountmedia/docs/darpa50> (accessed 4.1.18).
- Dixit, A.K., Pindyck, R.S., 1994. *Investment under Uncertainty*. Princeton University Press.

- Flamm, K., 1987. Targeting the computer: government support and international competition. Brookings Institution.
- Fong, G., 2001. ARPA does windows: the defense underpinnings of the PC revolution. *Business and Politics* 3 (3), 1469–3569.
- Foster, R.N., 1986. The S-Curve: A New Forecasting Tool, in: *Innovation: The Attacker's Advantage*.
- Fuchs, E.R.H., 2010. Rethinking the role of the state in technology development: DARPA and the case for embedded network governance. *Res. Policy* 39, 1133–1147.
- Gallo, M.E., 2018. Defense Advanced Research Projects Agency: Overview and Issues for Congress. Congressional Research Service.
- Gans, J.S., Murray, F.E., 2012. Funding Scientific Knowledge: Selection, Disclosure, and the Public-Private Portfolio, in: Lerner, J., Stern, S. (Eds.), *The Rate and Direction of Inventive Activity Revisited*. University of Chicago Press, pp. 51–103.
- Goldstein, A.P., Azoulay, P., Zivin, J.G., and Bulovic, V., 2017. “Promoting Energy Innovation with Lessons from Drug Development.” *The Hamilton Project*.
- Goldstein, A.P., Kearney, M., 2018a. Uncertainty and Individual Discretion in Allocating Research Funds. Manuscript under review.
- Goldstein, A.P., Kearney, M., 2018b. Active Project Management at ARPA-E. Manuscript in preparation.
- Goldstein, A.P., Narayanamurti, V., 2018. Simultaneous Pursuit of Discovery and Invention in the US Department of Energy. *Res. Policy*. In press. <https://doi.org/10.1016/j.respol.2018.05.005>
- Greenstein, S. (2011) “Nurturing the Accumulation of Innovations: Lessons from the Internet” in Henderson and Newell (2011) *Accelerating Energy Innovation, Insights from Multiple Sectors*, The University of Chicago Press, Chicago, IL.
- Huchzermeier, A., Loch, C.H., 2001. Project Management Under Risk: Using the Real Options Approach to Evaluate Flexibility in R & D. *Manage. Sci.* 47, 85–101.
- Institute of Medicine, 2012. *Accelerating the Development of New Drugs and Diagnostics: Maximizing the Impact of the Cures Acceleration Network: Workshop Summary*.
- Jacobsen, A., 2015. *The Pentagon's brain: an uncensored history of DARPA, America's top secret military research agency*. Back Bay Books.
- Lazowska, E.D., Patterson, D.A., 2005. An endless frontier postponed. *Science* 308, 757.
- Markoff, J., 2005. Pentagon Redirects its Research Dollars. *New York Times*.
- Manso, G., 2011. Motivating Innovation. *J. Finance* 66, 1823–1860.

- Milgrom, P., Roberts, J., 1990. The Economics of Modern Manufacturing: Technology, Strategy, and Organization. *Am. Econ. Rev.* 80, 511–528.
- Mowery, D., 2006. Testimony before the Hearing on “Should Congress Establish ‘ARPA-E’, the Advanced Research Projects Agency-Energy?” Committee on Science and Technology, United States House of Representatives, Washington DC.
- National Research Council, 1999. *Funding a Revolution: Government Support for Computing Research.* Computer Science and Telecommunications Board, Commission on Physical Sciences, Mathematics, and Applications, Washington, DC.
- National Research Council, 2007. *Rising Above the Gathering Storm: Energizing and Employing America for a Brighter Economic Future.* Committee on Prospering in the Global Economy of the 21st Century: An Agenda for American Science and Technology and Committee on Science, Engineering, and Public Policy, Washington, DC.
- National Research Council, 2017. *An Assessment of ARPA-E.* Board on Science, Technology and Economic Policy, Policy and Global Affairs, and Board on Energy and Environmental Systems, Division on Engineering and Physical Sciences, Washington, DC.
- Newman, N., 2002. *Net Loss: Internet Prophets, Private Profits, and the Costs to Community.* Pennsylvania State University Press, University Park, PA.
- Rohlfing, E., 2017. Overview and Update on the Advanced Research Projects Agency-Energy (ARPA-E) [WWW Document].
- Roland, A., Shiman, P., 2002. *Strategic Computing: DARPA and the Quest for Machine Intelligence 1983-1993,* MIT Press, Cambridge, MA.
- Rosenberg, N., 1996. Uncertainty and Technological Change, in: *The Mosaic of Economic Growth.* pp. 334–353.
- Scherer, F.M., Harhoff, D., 2000. Technology policy for a world of skew-distributed outcomes. *Res. Policy* 29, 559–566.
- Sen, A., 2013. Totally radical: From transformative research to transformative innovation. *Sci. Public Policy* 41, 344–358.
- Sen, A., 2017. Island + Bridge: how transformative innovation is organized in the federal government. *Sci. Public Policy* 44, 707–721.
- Stern, S., 2004. Do Scientists Pay to Be Scientists? *Manage. Sci.* 50, 835–853.
- Stokes, D.E., 1997. *Pasteur’s Quadrant.* Brookings Institution Press, Washington, DC.
- Trigeorgis, L., 1996. *Real Options: Managerial Flexibility and Strategy in Resource Allocation.* MIT Press.

Van Atta, R. H., 2007. Testimony before the Hearing on “Establishing the Advanced Research Projects Agency-Energy (ARPA-E),” Subcommittee on Energy and Environment, Committee on Science and Technology, United States House of Representatives, Washington DC.

Weinberger, S., 2017. The imagineers of war: the untold history of DARPA, the Pentagon agency that changed the world. Knopf.