

Small Business Innovation Applied to National Needs

Kyle R. Myers
Harvard University
& NBER

Lauren Lanahan
University of Oregon

Evan E. Johnson
University of
North Carolina

June 13, 2025

Small businesses have long supplied a disproportionate share of major innovations in the United States. We review a centerpiece policy on this topic: the US Small Business Innovation Research (SBIR) program. We trace its legislative history and summarize program evaluations over the past four decades. Using newly matched data on SBIR awards and venture capital investments into small businesses, we show that, despite often being compared to venture-backed businesses, SBIR-backed businesses pursue very different strategies. We use simple economic theories to motivate the SBIR program as a vehicle for the government to invest in small-scale, well-defined, but risky technologies that have large externalities, and we highlight a number of case studies consistent with this framework. Because the motivating friction lies at the level of ideas, our perspective encourages future evaluations to determine how the SBIR program influences not just *who* does the inventing, but *what* gets invented. Looking forward we discuss how rising industrial concentration and the diffusion of artificial intelligence may reshape the program's comparative advantage in the innovation policy toolkit.

Myers: kmyers@hbs.edu. Lanahan: llanahan@uoregon.edu. Johnson: evanej@live.unc.edu.
There are no financial relationships or other potential conflicts of interest that apply to any of the authors. We are grateful to Maryann Feldman and Scott Stern for introducing us to each other and the many conversations about the SBIR program over the years. We also received helpful comments from Josh Lerner and David Hsu. This project was conducted partly within a National Academies of Sciences, Engineering, and Medicine (NASEM) review of the Department of Defense (DoD). The views expressed herein do not necessarily reflect those of the review committee, the NASEM, or the DoD. Portions of the text, mostly pertaining to the setting, data sources, and data construction, may appear in the documents that we assisted in preparing for the review.

1. Introduction

Virtually all empirical evidence to date indicates that businesses underinvest in innovation relative to the social optimum.¹ The key friction leading to this underinvestment is that there are many ideas (i.e., potential new technologies) for which the social value of pursuing the idea is larger than the private value of the pursuit. In turn, those ideas are less likely to be discovered.² As such, there is an essential role for the government to increase welfare by subsidizing the pursuit of new ideas.

But how exactly do we define the new ideas that should be the target of innovation subsidies? One answer lies in the US Internal Revenue Service's definition of business expenses that qualify for the research and development (R&D) tax credit. Those subsidies target any pursuit “*which is undertaken for the purpose of discovering information which is technological in nature, and the application of which is intended to be useful in the development of a new or improved business component*” (IRS Code Section 41). An alternative approach is to abandon attempts at classifying ideas and instead identify organizations that are inherently more likely to pursue ideas with large social value. This includes policies related to entrepreneurship (e.g., bankruptcy laws, investor protections) and non-profit organizations (e.g., tax exemptions, basic research grants).

An attractive feature of these broad policies is that they do not require the government to have much information about the specific ideas that should be pursued. But the vague definitions that govern these policies can introduce the possibility of moral hazard — public funds may be redirected to unintended uses. Still, they are designed to leverage the information and incentives of markets.

However, it is often the case that the government has unique information about certain ideas — ideas that have high social value and are not on the margin of business investment decisions (even in the presence of the aforementioned subsidies) because they involve high risks or low private returns. For example, the Department of Health and Human Services may be aware of a need for wearable sensors to better monitor patients in clinical trials; the National Aeronautics and Space Administration may be aware of a need for a lighter-weight robotic arm for space-bound vehicles; or the Department of Defense may be aware of a need for more accurate measurements of aircraft components. The identification of these *national needs* — specific technological ideas with a large social value relative to their private value — raises a key question: what organizations are best positioned to deliver the innovation needed? Quite often, the answer is *small businesses*.

¹For a review, see Bryan and Williams (2021).

²Formally, if a business cannot use first-degree price discrimination to fully appropriate the total surplus they would generate if they discovered a new product, then the businesses' investment in discovery will be below the socially optimal level.

One of the most important US policies at the nexus of national innovation needs and small businesses is the US Small Business Innovation Research (SBIR) program. In this chapter, we focus on the SBIR program, highlighting the critical role of small businesses in meeting national innovation needs. To do so, first, we review the history and governing legislation of the program. Second, we present new insights on SBIR-backed businesses by comparing them to small businesses backed by venture capital. Third, we review economic theory that motivates the SBIR program and highlight the types of national innovation needs that small businesses are particularly well-suited to engage with. Lastly, we conclude with two forward-looking discussions about the ongoing role of the SBIR program, and small business innovation more generally, amid increasingly concentrated private markets and the age of artificial general intelligence.

2. The Small Business Innovation Research (SBIR) Program

2.1. History and Legislative Structures

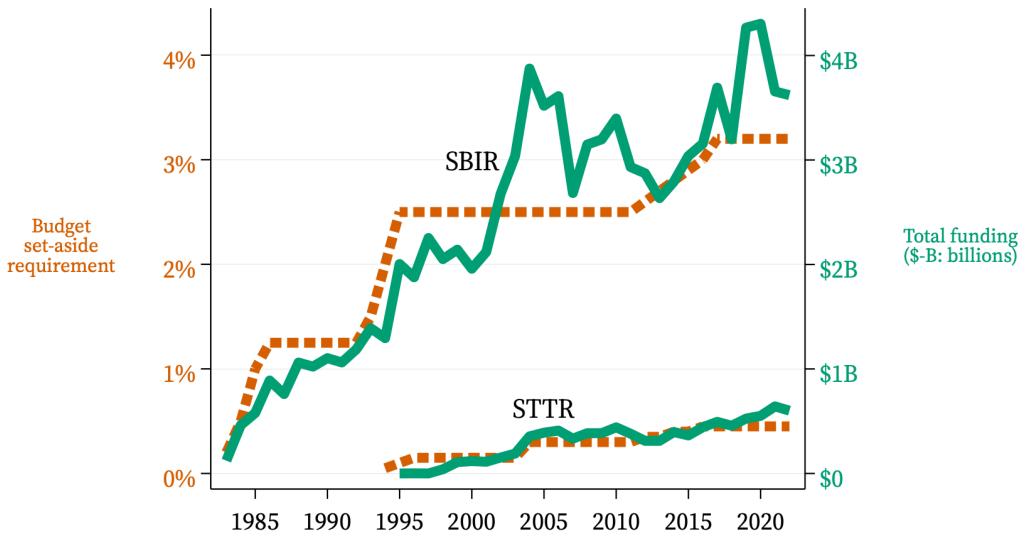
The central US policy at the intersection of small business innovation and national needs is the Small Business Innovation Research (SBIR) program. The SBIR program traces its roots to the late 1970s, when Roland Tibbets of the National Science Foundation (NSF) championed an experimental effort to fund R&D for small businesses. Tibbets believed that innovative small businesses should be given access to R&D funding that was, at the time, either explicitly or implicitly only accessible to larger federal contractors. Tibbets' NSF pilot, titled "*Small Business Innovation Applied to National Needs*", demonstrated the potential of small business innovation and helped build political support for a government-wide initiative.

In the ensuing years, lawmakers from both parties embraced the concept. This culminated in the passage of the Small Business Innovation Development Act of 1982 (Public Law 97-219), which was approved by Congress with bipartisan support and signed into law by President Ronald Reagan.

The SBIR program is not permanent. It requires ongoing congressional reauthorization, which has occurred in 1992, 2000, 2011, 2016, and 2022, with another reauthorization due for debate in 2025.³ Broadly speaking, Congress has been supportive of the program as evidenced by its persistence and the steady increases in the percentage of a federal agency's extramural research budget that must be allocated to SBIR awards (see Figure 1). In most recent years, statutory requirements dictate that any federal agency with an extramural research budget larger than \$100 million must allocate 3.2% to the SBIR program, which has annual SBIR funding approaching \$3.5 billion.

³The timing of a subsequent reauthorization is determined in the prior reauthorization.

FIGURE 1. Growth of the SBIR and STTR Programs



Note: The dashed lines (corresponding to the left y-axis) plot the statutory requirements for the percent of a federal agencies extramural research budget that must be allocated to the SBIR and STTR programs since their inception. The solid lines (corresponding to the right y-axis) plot the total spending, in USD-2025 (adjusted using the Consumer Price Index), for the SBIR and STTR programs since their inception.

In the mid-1990s, the Small Business Technology Transfer (STTR) program was developed within the umbrella of the SBIR program. In an effort to foster collaborations between businesses and academic institutions, the STTR program included an additional requirement that all applicants partner with an academic institution when competing for funding and that a minimum share of work be performed by both parties.⁴ Only federal agencies with extramural budgets in excess of \$1 billion are required to participate in the STTR program and, most recently, those participants are only required to allocate 0.45% of that budget to STTR awards. In most recent years, annual STTR funding is on the scale of \$0.5 billion (see Figure 1).

The explicit structure of the SBIR program has been relatively consistent since its inception. Businesses eligible for SBIR funding must have fewer than 500 affiliated employees and with majority-ownership by US citizens or permanent residents. There are three phases of awards. As of 2024, participating agencies may issue a Phase I award of up to roughly \$300,000. Then, agencies may award any Phase I winners a Phase II award, which may be up to roughly \$2 million.⁵ Phase III awards are unique from the first two

⁴Typically, STTR regulations have required that both the small business and the academic institution perform at least 1/3 of the work funded through any award.

⁵Agencies may submit waivers to the US Small Business Administration to exceed these limits. Also, certain agencies at certain times have been able to award “*Direct to Phase II*” awards that circumvented the requirement of a Phase I award. Although, this path is less common.

in that they have no statutory funding requirements. These awards must be funded by the agencies' non-SBIR budgets, but they afford the SBIR business some unique benefits that they would not otherwise receive if they obtained a similar contract outside of the SBIR program.⁶ Broadly speaking, these Phase III benefits are designed to insulate the SBIR business from competition in the procurement stage (i.e., after inventing some new technology), so that they stand to gain from the value they generate and their technology is not, for example, duplicated at lower cost from a competitor (which would lead to a dynamic disincentive for the SBIR business).

2.2. SBIR Processes in Practice

The founding legislation outlined four major goals, which have remained relatively intact since the program's inception, and now are written as goals to: (i) stimulate technological innovation; (ii) use small businesses to meet federal research and development needs; (iii) foster and encourage participation by emerging and undercapitalized small business concerns in technological innovation; and (iv) increase private-sector commercialization of innovations derived from federal R&D. To better understand the implicit objectives of the SBIR program, it is helpful to consider some of the key processes and regulations governing the program and discuss how these processes align with these explicit goals.

To the goal of commercializing federal R&D (Goal *iv*), one clear motivation here is the so-called “Valley of Death.” In short, federally funded basic research supported in universities may not be commercialized because neither academia nor private investors have the right mix of incentives to engage with certain developmental tasks (i.e., tasks that are not novel enough for academics to be interested, but still too risky for private investors to consider funding). The STTR program is clearly the explicit directive focused on this friction. However, the STTR program accounts for roughly one eighth of total SBIR/STTR spending (see Figure 1), which suggests this particular goal is not a leading priority. Furthermore, the data suggests that businesses may not see the STTR program as being so differentiated from the traditional SBIR program. While approximately 5% of SBIR-awardees also receive an STTR award, more than 50% of STTR-awardees also receive an SBIR awards — if the STTR program was highly differentiated, we would have expected a higher degree of separation from participants in either track of the program.

The specific language of the third goal has evolved over time, but has historically focused on small businesses owned by women, minorities, and other under-represented groups. There is clear, longstanding motivation for concern that the static and dynamic consequences of discrimination prevent inventors with great potential from accessing re-

⁶Phase III also allows for the right to sole-source the contract, provides the awardee with “*Data Rights*” which prevent the government from disclosing the SBIR business’s data to other businesses, provides an exemption from SBA size standards for a procurement, and has no limits on the dollar size.

sources necessary for their pursuits (Bell et al. 2019; Agarwal and Gaule 2020). In principle, a program could engage with these frictions with access to these explicit subsidies.⁷ The SBIR legislation does include certain reporting requirements for documenting the composition of the businesses that receive awards; however, the program has never included any processes that explicitly lower the cost (or increase the benefit) of participation for businesses owned or operated by certain groups.⁸

It is helpful to conceptualize goals *(i)* and *(ii)* as being on opposite poles of the same spectrum of objectives. At one end, the goal *(i)* of stimulating small business technological innovation writ large can be motivated by traditional financial market frictions (Lerner 2000), which inspires analogies such as “*The Government as Venture Capitalist*.⁹ What would a policy look like that truly embraced this goal? The Israeli government’s Yozma program provides one illustrative example: ten approved venture funds received matching investments from the government at a rate of 40% (up to \$8 million in total for each fund) and the funds faced a small number of limitations on which small businesses they could invest those funds in (Baygan 2003). This design leverages the information and capabilities of the private sector. The SBIR program does not have any processes as dramatic as the Yozma matching approach. Although, the original pilot of the SBIR program led by Tibbetts at the National Science Foundation did include an explicit scoring bonus for proposals that included credible commitments from private investors — a matching policy of sorts — but no such bonus has ever been included in the official program.¹⁰

In practice, the SBIR program processes reflect the goal *(ii)* of meeting federal R&D needs. Many participating federal agencies (e.g., DOD, DHS, DOT, EPA, NASA) award funds predominantly (or only) via *contracts*, which tend to include more pre-specification of a technology that the funder is interested in developing. Even in the case of the agencies that more often (or always) rely on *grant* mechanisms to make SBIR awards (e.g., NIH, NSF), applications are subject to many of the same criteria and evaluation processes that allow the agencies to influence which pursuits receive funding. Furthermore, the formal distinction between grants and contracts masks a common structure: Phase I awards offer a small investment with few explicit strings attached, but the incentive of the Phase II award encourages awardees not to veer too far from the initially stated objectives.

⁷For example, the US National Institutes of Health have a number of explicit policies aimed at assisting young researchers per their “Early-Stage Investigator” rules and processes; see: <https://grants.nih.gov/policy-and-compliance/policy-topics/early-stage-investigators>.

⁸There are, of course, rules governing participating in the SBIR program per the ownership structure of the business (e.g., nationality of owners).

⁹This is the title of Lerner’s (2000) influential analysis of the SBIR program, which sparked a long line of analyses that we review in the next Section below.

¹⁰We are also unaware of any explicit mandates that require the composition of the review panels that evaluate SBIR programs to include any minimum number of participants from the private sector in general, or the venture capital sector specifically.

Altogether, the SBIR program is designed to not just produce innovation writ large, but to steer small businesses towards technologies with out-sized public benefits. This is in contrast to the more technology-agnostic, profit-maximization behavior of venture capitalists. To dig more into the differences between the SBIR program and venture capital, and to see how much the SBIR program's implicit objectives show up in observable outcomes, the next section leverages a new dataset that tracks SBIR- and VC-backed businesses, their growth, and their innovations.

3. A New Look at the SBIR Program Compared to Venture Capital

The SBIR program is often referred to as “*America’s Seed Fund*”. In line with this moniker, the program is frequently evaluated through the same lens of success applied to the US venture capital (VC) sector. Despite an abundance of anecdotal comparisons, systematic data on small firms’ engagement with the SBIR program and/or the VC sector has been challenging to obtain. The problem is that data on the businesses’ features (i.e., their age, employment count, revenues, etc.) are typically not jointly collected with data on the businesses’ investors (i.e., VC, SBIR awards), innovative output, or engagements with the federal government more generally. So, while the empirical record consistently shows that small businesses are responsible for a disproportionate share of innovations in the US (Cohen 2010), there are still many open questions as to how best to continue to support that innovation.

Here, we highlight some statistics based on a newly assembled dataset of small US businesses that have received either (i) SBIR awards, or (ii) seed investments from VC. Our goal is to highlight the similarities and differences in how these two sets of businesses generate value in both private- and public-sector markets. The following results should be interpreted as reflecting some (unknown) combination of both *treatment effects* due to the different bundles of resources provided by the SBIR program compared to the VC industry, as well as *selection effects* due to different types of businesses (i.e., pursuing different technologies) making different choices about engaging with the SBIR program versus the VC industry.

3.1. Dataset Construction

The key inputs into this new dataset are: (i) the National Establishment Time Series dataset, which includes the name, location, age, employment, and unique identifiers for business establishments and their corporate hierarchies for 2000–2021 (Walls 2021); (ii) the Small Business Administration’s public record of SBIR awards (SBA 2024); (iii) the Crunchbase record of VC investments, which includes businesses’ names, locations, and investment amounts for 2000–2021 and for which we focus only on businesses that receive

“seed” investments; (iv) the USPTO’s PatentsView dataset, which includes businesses’ names, locations, and patent assignment information for 1976–2021 (USPTO 2024); and (v) the Federal Procurement Data System (FPDS) dataset, which includes business names, locations, and federal contracting information for 2000–2021 (BFS 2024).

There are some notable challenges to assembling this data. First, systematic comparisons of NETS data with US Census data on business establishments have shown some systematic biases due to the sampling and missing-data-imputation processes used in NETS (Barnatchez, Crane, and Decker 2017). That is, NETS appears to mischaracterize the full population of US businesses, especially for smaller organizations (Barnatchez, Crane, and Decker 2017). In light of this, we only report results that condition on a business appearing in both the NETS and either SBIR or VC datasets.¹¹

The second challenge is linking businesses across the datasets in the absence of a standardized, unique identifier common to all datasets. To overcome this challenge, we use the `linktransform` software (Arora and Dell 2023) to compute the semantic similarity between business names across all datasets. We then merge businesses across datasets under the requirements that (i) the semantic similarity is above a reasonable threshold and (ii) the location of the activity (i.e., the patent assignee’s location; the government contractor’s location) is in the same state. This still leaves a problem of potential duplication, whereby two businesses could be located in the same state and have semantically similar names. To avoid this, we restrict our sample to businesses that are uniquely matched in the NETS data from the other datasets. This reduces our sample size, but increases our confidence in the accuracy of the data construction.

A third challenge is related to interpreting the data. In the simpler case of the procurement data, we have a clear understanding of the parties involved (i.e., the US federal government and a particular contractor) and the value of that relationship (i.e., the dollar value of the contract). However, in the case of the patent data, these values are less clear. To identify relationships, we follow Fadeev (2024) and interpret citation flows from one patent to another as being indicative of a business relationship.¹² To proxy for the value of a patent, we rely on Kelly et al.’s (2021) measure of technological innovation, which has been computed for the universe of USPTO patents issued up until 2016. In short, the measure is based on a patent describing a technology that few patents had described previously (i.e., novelty) and that many patents engage with subsequently (i.e., impact). Specifically, we use their binary indicator of a breakthrough based on a patent’s 5-year performance.

¹¹The NETS data also contains a number of imputed values based on private methodologies (Barnatchez, Crane, and Decker 2017). We have no reason to assume that the imputation methodology depends on whether a business was SBIR- or VC-backed, so we are less concerned with how this may bias our relative comparisons.

¹²Fadeev (2024) finds that roughly 3/4 of citation pairs are between supplier-buyer pairs or businesses with cooperative research agreements.

Within our sample, roughly 14% of patents are classified as breakthroughs.¹³

Ultimately, our main dataset is comprised of 10,603 unique businesses (per their DUNS-based headquarters). Roughly 20% of the sample are businesses that received seed investments from VCs and no SBIR awards, with the nearly all of the remainder receiving SBIR awards and no VC seed investments.¹⁴

3.2. Initial Business Conditions

Figure 2 plots the distribution of businesses' features at the time of investment from either the SBIR program or a venture capitalist. Figure 2A highlights the significant difference in investment strategies of the two groups. For SBIR-backed firms, the structured nature of the Phase I and Phase II awards is clearly visible. The distribution of VC investments is much smoother. Furthermore, compared to the SBIR program, which awards many businesses small investments (Phase I awards), the VC distribution indicates a strategy of issuing (relatively) much fewer investments on the scale of the SBIR Phase I and much more investments on the scale of the SBIR Phase II.

The optimal distribution of investments is far from clear and depends on many factors. The SBIR distribution is consistent with traditional risk-diversification and real-options theories (Markowitz 1952; Dixit and Pindyck 1994). In contrast, the VC distribution is consistent with there being significant economies of scale and/or non-diversifiable risks (Sahlman 1990; Denes et al. 2023). An important political economy point of note here is the conjecture that the rigidity and smaller stakes of the SBIR program may be one reason the program appears relatively robust to regulatory capture (Lerner 2009).

Figures 2B–2C illustrate relatively similar age and size distributions for the two sets of businesses. The median SBIR-backed business is four years old and has five employees at the time of award,¹⁵ and the median VC-backed business is three years old with three employees at the time of investment.

3.3. Market Outcomes

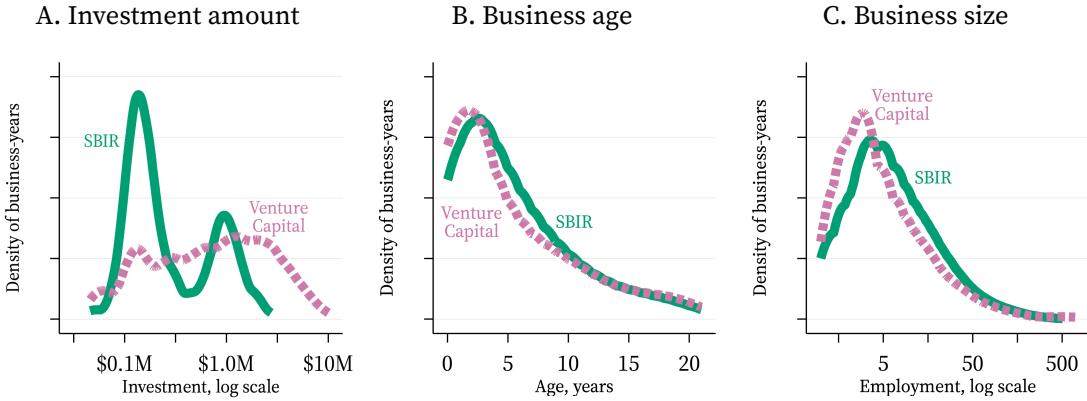
Here, we report a series of estimates from regressions of business-year-level outcomes (i.e., survival, employment, revenues, patents, federal contracts) on a binary variable indicating that a business previously received an SBIR award. We include a vector of year-specific indicators to control for any secular trends correlated with the growth of VC over this period. While we can estimate regressions that include only business-year

¹³In the full Kelly et al. (2021) data, over our sample period, the breakthrough rate is roughly 12%.

¹⁴Only 1% of the sample appears in both the VC and SBIR data.

¹⁵As a benchmark, the SBA award data reports an SBIR awardees' number of employees for roughly 60% of business-year observations. The median of the non-zero values reported is 7, which is consistent with our estimate of 5 per the NETS data.

FIGURE 2. SBIR- vs. VC-backed: Features at the Time of Investment



Note: Plots the distribution of business-year observations, where the business receives either an SBIR award or a series round of investment from venture capitalists (Panel A.), the age of the business in that year (Panel B.), and the number of employees at the business that year (Panel C.).

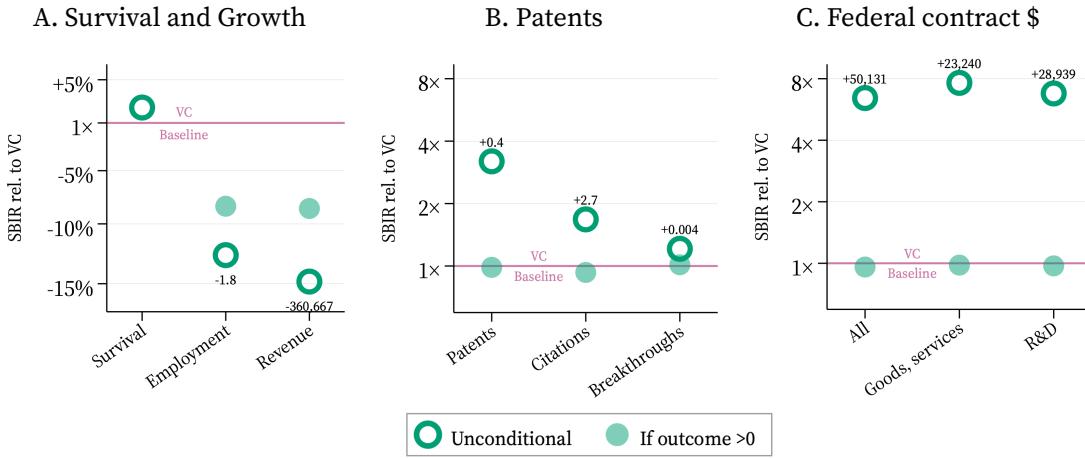
observations where the outcome is non-zero, given the structure of the data we also observe business-years whether or not the business has produced any of the outcomes of interest. This allows us to run what we term *unconditional* regressions, which include zeroes for any business-year observation that has not survived or did not produce any of the focal outcome. Given the inclusion of these zeroes and the non-negative nature for the outcomes, we estimate Poisson regressions in all cases, which yields an estimate of the percentage difference between SBIR- and VC-backed businesses with respect to each outcome (Chen and Roth 2024).

The unconditional regressions are useful in that they incorporate all treatment and selection effects (including survival differences) that might give rise to any difference in outcomes. The regressions where we condition on a non-zero outcome remove some, but certainly not all selection effects — they remove the survival selection effect, and they also remove the extensive margin selection effect (i.e., some businesses simply choosing not to pursue any patents, or not to pursue any federal contracts).¹⁶

Figure 3A reports the results from the unconditional and conditional regressions, focusing on several traditional private market outcomes of business growth. First, we consider survival, for which the unconditional and conditional regressions are the same: we find that, on a year-to-year basis, SBIR-backed businesses are 1.4% more likely to survive (on a base survival rate of 80%). We also find that SBIR-backed businesses have annual employment and revenues that are roughly 15% lower than VC-backed businesses. This amounts to roughly 2 fewer employees and \$360,000 less revenue each year.

¹⁶However, since we cannot directly observe businesses' choices *per se*, we are also conditioning on some amount of the differential treatment effect since, for example, one funding source may improve a businesses' odds of getting any patents conditional on filing a patent application.

FIGURE 3. SBIR- vs. VC-backed: Private and Public Market Outcomes



Note: Reports the relative differences in outcomes for SBIR-backed businesses compared to VC-backed businesses along multiple dimensions based on unconditional and conditional (outcome > 0) regressions. Numbers reported next to the unconditional estimates convert the relative difference into absolute levels per the sample averages.

Looking at patent-based metrics (Figure 3B), the unconditional results show large relative differences in patent outcomes. SBIR-backed firms produce roughly 3 times more patents, 2 times more citations, and 20% more breakthrough patents. But it should be noted that these large relative differences are based on small absolute differences (i.e., +0.4 more patents, +2.7 citations, and +0.004 breakthrough patents). The rarity of these events in the unconditional sample yields large relative differences.

Figure 3C reports the results from the unconditional and conditional regressions, focusing on outcomes based on the federal procurement data. The unconditional regressions illustrate a dramatic relative difference in the dollar value of contracts that these two sets of businesses obtain. Whether looking at contracts related to the procurement of goods and services, or those based on some research and development activity, SBIR-backed businesses have nearly 8 times as much engagement with the federal procurement system per the unconditional regressions. Still, just as we saw in Figure 3B, these differences are explained by the extensive margin difference. Conditional on having any federal contracts, the dollar value of these contracts is virtually identical regardless if they involve an SBIR- or VC-backed business in our sample.

The results illustrated in Figure 3 indicate that SBIR- and VC-backed businesses engage in very different strategies. VC-backed businesses often focus on growing their business in ways that do not involve the patent record or federal contracts. The opposite appears true for SBIR-backed businesses. Overall, this is consistent with the SBIR program not simply trying to mimic venture capitalists, but rather, generating some combination of selection and treatment effects that pull or push participants towards the development of

technologies that are patentable and needed by federal agencies. In theory, the program should be pulling and pushing these businesses towards technologies that have high social value but a low private value (and, therefore, would go unpursued in the absence of the SBIR program; i.e., Bryan and Williams 2021).¹⁷ In the next section, we use economic theory to highlight the features of such national needs that the SBIR program is well-suited to target.

4. What National Needs can Small Business Innovations Solve?

4.1. Theoretical Framework

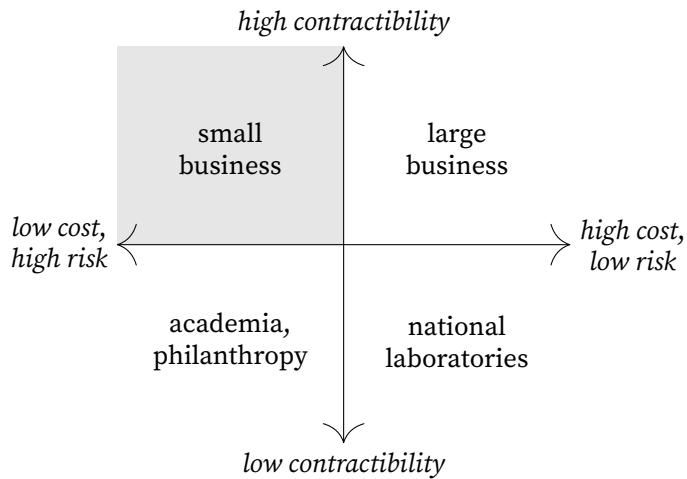
Economists have long been interested in the comparative advantages of different organizational forms vis-à-vis innovation, with the unique capabilities of small businesses often highlighted. Here, we are less concerned with the ability of different organizations to innovate writ large. Rather, we are concerned with the ability of different organizations to pursue different types of national innovation needs. Thus, for the purposes of this discussion, it is helpful to assume that the wedge between social and private returns to pursuing the innovations in question (which determines the optimal subsidy rate) is the same.

First, it is helpful to classify innovations along three dimensions: (i) contractibility, (ii) scale, and (iii) risk. *Contractibility* refers to the degree to which the innovation can be legally specified ex-ante. For example, if a new technology must be able to achieve very well-defined performance metrics, then it is highly contractible. *Cost* refers to the expected financial cost of resources necessary to pursue the innovation. *Risk* refers to the uncertainty in the expected value of pursuing the innovation. Figure 4 uses these three dimensions (converted into two axes) to illustrate the organizational types best suited for different types of innovations as suggested by prevailing economic theories.

Broadly speaking, the contractibility of a potential innovation — the degree to which the government can legally specify the requirements of the technology needed — is the chief determinant as to whether the technology should be sourced from the private market. In general, there is a tradeoff between the high-powered incentives of the for-profit market and the costs of writing and negotiating contracts with businesses. Most economic theories yield a prediction that, as the transaction costs of interacting with businesses decrease (e.g., the government can more easily and clearly specify the technological need), the more efficient it is to source the innovation from businesses (Hart, Shleifer, and Vishny

¹⁷On this point, Narain (2025) provides compelling evidence that VC-backed businesses pursue technologies that are more likely to pay off sooner, compared to SBIR businesses, who tend to pursue technologies with a longer time horizon of development.

FIGURE 4. Organizational and Innovation Type Alignment



Note: Plots the organizational types best suited for innovations with certain features as suggested by prevailing economic theories.

1997; Levin and Tadelis 2010).¹⁸

The cost and risk of the project at hand are important determinants of the scale of the organization best suited to pursue the project. Cost is relatively straightforward. Assuming that there are some economies of scale in the organizations' operations, larger organizations should be matched with more costly projects. For highly contractible projects, this would be large businesses. For less contractible projects, this would be the large national laboratories or other intramural research programs.

Risk is complicated. There is some intuitive appeal to the notion that larger organizations may be able to diversify the risk from any particular innovation project they undertake; however, we are less concerned with different organizations' abilities to handle risks that are forced upon them. We are more concerned with the organizations' appetite for choosing to take risks.

Overall, there are many theoretical arguments with empirical support that smaller organizations will have inherently larger appetites for risks. There are two broad motivations for this result. First, smaller organizations will have fewer potential tasks to be undertaken, which will generate less tension between more measurable, routine tasks and unmeasurable, risky tasks. This allows for more resources in the organization to be allocated to novel, innovative activities (Sah and Stiglitz 1986; Holmstrom 1989). Second, smaller organizations will have fewer opportunities in their choice set. The ownership

¹⁸The sourcing distinction can also be conceptualized as a continuum of control with outsourcing on one end and full in-house provision on the other end. Through this lens Lerner and Malmendier (2010) and Bruce, de Figueiredo, and Silverman (2019) highlight the role of termination options and cooperative agreements (i.e., more control to the funder) in less contractible settings.

over the potential gains from pursuing those opportunities will be divided between fewer people. This creates less pressure from external and internal capital markets to choose safe, short-term opportunities and motivates more innovative effort (Holmstrom 1989; Aghion and Tirole 1994; Anton and Yao 1995; Stein 1997).

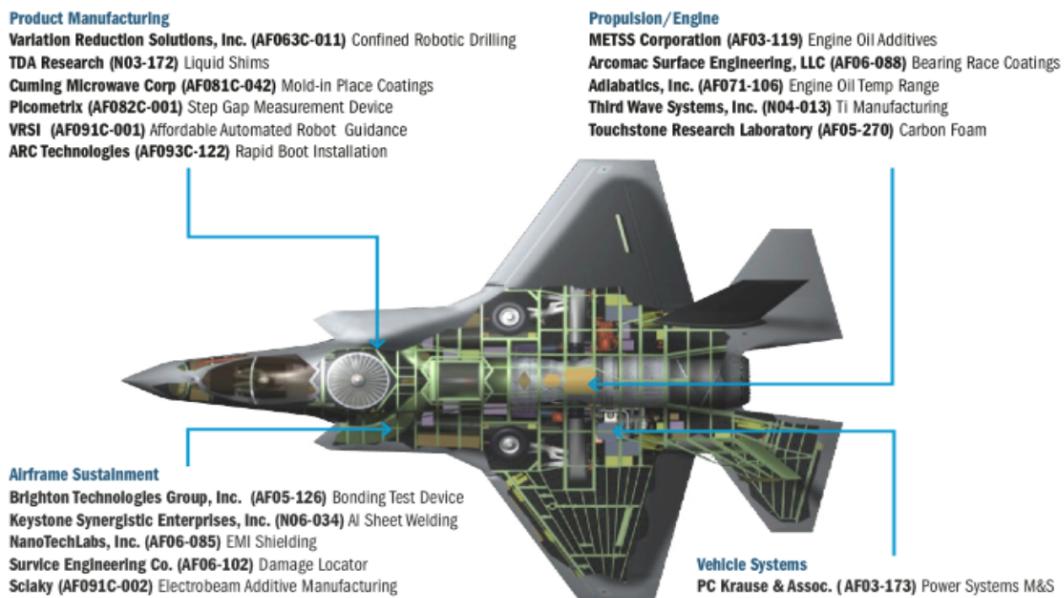
To summarize, small businesses tend to be best suited for developing innovations that solve small-scale, uncertain, well-defined problems. As we argue below, there are an increasing number of national needs that fit these criteria.

4.2. Small Parts of Big Networks

While small businesses tend to be best suited for solving small-scale problems (i.e., scale of costs), it is certainly not the case that they can only generate solutions with large scale benefits. As the world increasingly revolves around networks of technologies with inter-dependencies (Katz and Shapiro 1994), an ever-growing body of theoretical arguments and empirical evidence continues to emphasize that there can be points in those networks that are small when viewed in isolation, but are very important to the functioning of the network as a whole. The “network” in question may revolve around a single platform product that is a network of technological components (Baldwin and Clark 2000; Baldwin, Woodard et al. 2009; Kretschmer et al. 2022), industries connected together via a network of businesses (Milgrom and Roberts 1990; Kremer 1993; Jones 2011; Acemoglu et al. 2012; Elliott, Golub, and Jackson 2014; Bigio and La'o 2020), or even a network of countries connected through trade and migration (Kremer 1993; Jones 2011). Here, we provide illustrative examples of how SBIR awardees have made significant technological contributions of each. These examples highlight the presence of small technologies addressing large externalities (both positive and negative) within each type of network.

Network of Components. The federal government is increasingly involved in the procurement of extremely complex products that involve a large network of individual components. In these networks of components, there are often technologies that are, by themselves, of relatively little value, challenging to develop, but relatively easy to specify. These narrow technological challenges are a prime example of a national need well-suited for a small, innovative business. In line with this rationale, it is easy to find many examples of SBIR-backed businesses contributing small, but crucial components to some of the most complicated products that the federal government procures. See Figures 5 and 6 for illustrations of the components of the F-35 jet and the Mars Perseverance rover developed by SBIR awardees. A good example involves Picometrix, Inc., which developed a measurement tool to inspect the seams of the panels on the F-35 and ensure proper tolerance. The firm’s nondestructive testing system minimizes manufacturing costs for F-35s, filling a small but important niche. Similarly, Motiv Space Systems, Inc. engineered the mechan-

FIGURE 5. SBIR Technology in the F-35



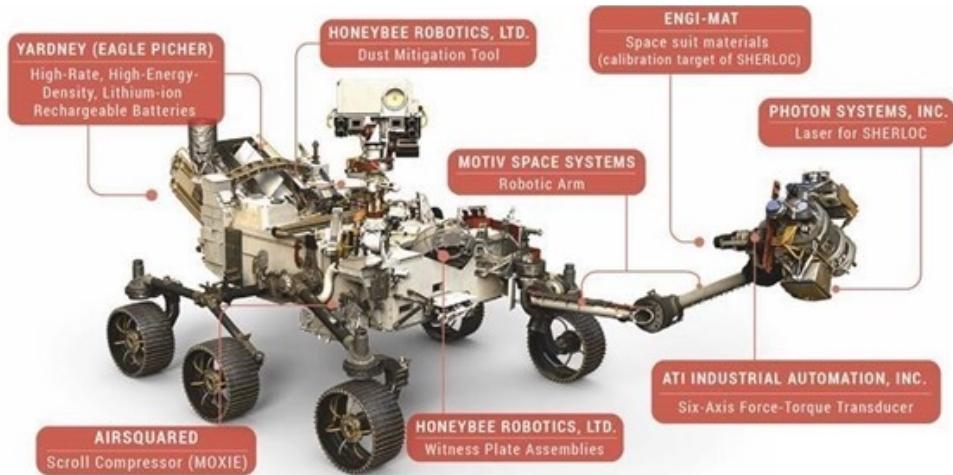
Credit: SBIR at the Department of Defense (2014) National Academies of Sciences, Engineering, and Medicine. Washington, DC: The National Academies Press. *Note:* Shows the SBIR-funded companies, and their respective technologies, that were integrated into the F-35 Lightning II.

ical arm on the Perseverance Rover, which facilitated rock sampling and collection on Mars. This technology is not only central to the essential aims of the US mission, but the business has also applied this technology in additional (terrestrial) markets including disaster relief, oil and gas extraction, bomb disposal tasks, chemical manufacturing, power generation, and subterranean activity. These companies were integral to the development of very specific technologies for national needs that may not have had immediately clear commercial applications.

Network of Businesses. When industrial sectors and supply chains are viewed as networks, it is easy to find instances where small innovations yield large benefits for the network as a whole. Often these innovations take the form of measurement devices or tools. These innovations tend to be under-delivered by the private sector because they generate information. Though, information is generally challenging to monetize (Arrow 1962). So, it is not surprising that there are many success stories of SBIR-backed businesses developing new measurement technologies that help generate value for multiple businesses within a sector. In the healthcare sector, these technologies often revolve around monitoring patients' health or improving our diagnostic capabilities.¹⁹ In the energy and

¹⁹Exemplar SBIR awardees include: BioSensics, LLC, which has made substantial contributions in the arena of wearable sensors with enhanced digital technologies; QT Medical, Inc., which developed an at-home electrocardiogram (ECG) for children.

FIGURE 6. SBIR Technology in the Mars Perseverance Rover



Credit: Jet Propulsion Laboratory; NASA. *Note:* Shows the SBIR-funded companies, and their respective technologies, that were integrated into the 2020 Perseverance Rover.

transportation sector, these technologies often focus on monitoring the quality of air or infrastructure.²⁰

Network of Countries. Networks now, more often than not, span the entire globe, which leads to the possibility of there being national needs based in other nations. A classic example of this sort of (negative) international externalities is pollution. For a developing country, the value of using dirty energy technologies (e.g., to grow capital-intensive sectors of their economy) is well beyond the direct cost of the pollution within the boundaries of their country. Of course, other countries bear additional costs due to this pollution. Would an innovative small business based in the US expect large gains from improving the environmental impact of the energy technologies used by such a developing country? Perhaps not, because that business does not have much to gain privately (because the willingness of businesses in that developing country to pay for that technology is low). This is despite the fact that the US may have much to gain from the reduction in pollution. However, with support from the SBIR program companies may be incentivized to engage with these sorts of national needs.²¹

²⁰Exemplar SBIR awardees include: 2B Technologies, which developed a novel air pollution monitoring system that can more easily attached to drones, kites, balloons, and aircraft, for expansive monitoring of air quality; and Fuchs Consulting, Inc., which developed a non-destructive method to inspect bridge maintenance without the need of traffic control or disruption.

²¹Exemplar SBIR awardees include: ASAT, which developed an integrated stove that significantly reduces biomass emissions; and Porifera, which has developed technologies to efficiently extract water from a wide range of (possibly contaminated) solutions.

5. SBIR Program Evaluations

5.1. Literature Review

Here, we review the academic literature on the SBIR program for a more systematic understanding of the program. Lerner (2000) undertook one of the first empirical analysis of the SBIR program, sparking a number of subsequent studies. Using matching techniques to compare SBIR awardees to observationally equivalent firms, Lerner (2000) found that SBIR awardees grew significantly faster over a ten year period, when focusing on outcomes such as sales and employment. Furthermore, there was significant heterogeneity in this growth, with the positive returns confined mostly to regions with significant venture capital activity. In the decade following Lerner (2000), a number of studies used similar methodologies and documented similarly positive returns. These studies explored a range of other metrics for success (e.g., patents, commercialization of technologies, strategic alliances) and identified a number of mechanisms that moderated and mediated the effect of the SBIR program.²²

In more recent years, scholars have identified new data sets and natural experiments that have facilitated more quasi-experimental evaluations of the SBIR program. Howell (2017) leveraged administrative data that facilitated a regression discontinuity analysis of the program. This analysis showed a clear positive difference in SBIR award winners propensity to patent and receive follow-on venture capital investments (compared to marginal SBIR applicants that did not win awards). Another set of studies made use of state-specific matching programs, whereby SBIR award winners in certain states would receive additional investments simply due to their location. A number of studies incorporated these state matching policies into their research designs to study the effect of the SBIR program on access to private financing (Lanahan and Armanios 2018), employment (Lanahan, Joshi, and Johnson 2021), business certification and signaling (Lanahan and Armanios 2018; Lanahan, Armanios, and Joshi 2022) and to study the broader R&D spillovers of the program as well (Myers and Lanahan 2022). Notably, the quasi-experimental evidence generally failed to find a significant effect of the program on businesses' employment levels (Lanahan, Joshi, and Johnson 2021). However, these studies continued to find significant effects on the awardees innovative efforts (Lanahan and Armanios 2018) and, perhaps more importantly, on the innovative efforts of other businesses not directly engaged with the program (Myers and Lanahan 2022). To this latter point, Myers and Lanahan (2022) find that for every patent produced by SBIR awardees, roughly three more are produced by others who benefit from R&D spillovers.

A particular point of contention in SBIR evaluations is the role of multiple award

²²For example, see: Wallsten (2000); Audretsch, Weigand, and Weigand (2002); Audretsch (2003); Gans and Stern (2003); Toole and Czarnitzki (2007); Link and Ruhm (2009); Link and Scott (2010).

winners or “SBIR mills”. Of course, a high concentration of any government subsidy among a small set of recipients will raise concerns of regulatory capture. Part of this concern stems from the fact that some of the aforementioned studies find the effect of an SBIR award on a particular business’s outcomes is negatively correlated with the number of prior awards that business has received. There are (at least) two scenarios that could give rise to this result: (i) multiple award winners do effectively achieve regulatory capture and have developed an ability to win SBIR awards despite being less productive than alternatives; and (ii) the multiple award winners are pursuing systematically different technologies than other awardees, and these technologies are more costly and difficult to pursue. Separating these two scenarios is very difficult. Feldman et al. (2022) offers a combination of quantitative and qualitative analyses indicating that multiple award winners are engaged in unique technological pursuits compared to other awardees. In a similar vein, Link and Swann (2024) show that the correlation between an awardees commercialization outcomes and their number of prior SBIR awards is *negative* if all prior awards are counted but *positive* if technologically relevant prior awards are counted. These results are consistent with multiple award winners filling a distinct niche within the SBIR program’s landscape. Furthermore, if regulatory capture was occurring, an intuitive hypothesis would be that the concentration of SBIR awards among awardees would be increasing over time. However, our analyses indicate this is not the case — the concentration of SBIR awards among awardees has remained effectively flat since the programs’ inception.²³ However, it is still unclear how much some firms are able to, for example, write SBIR applications that win awards but do not produce valuable outcomes.

The National Academies of Science, Engineering, and Medicine (NASEM) has also conducted several reviews of the program across the leading participating agencies (i.e., DOD, HHS, NSF, DOE, and NASA). Beyond the traditional academic analyses, these NASEM reviews also conduct important quantitative and qualitative reviews of the operations and functioning of the SBIR program (i.e., application and award procedures, applicant outreach). Since 2020, NASEM has published reports for DOE (National Academies of Sciences, Engineering, and Medicine 2001, 2020), NSF (National Academies of Sciences, Engineering, and Medicine 2023), and HHS (National Academies of Sciences, Engineering, and Medicine 2022) with reviews of the SBIR programs at DOD and NASA actively underway as of 2025. Across the published reports, there is broad evidence of positive returns to the program when focusing on a range of innovative outcomes. However, many suggestions

²³For example, the gini coefficient, which measures how unequal the distribution is, of business-level SBIR dollars has been essentially flat since the program began. Since the second year of data availability, 1984, the gini coefficient has been 0.6 when calculated using the crude business names in the SBA data. When we restrict our analysis to the years 2000-2020, for which we have the more rigorous business disambiguation, we estimate a gini coefficient of 0.5.

for improvement to the program are made, often revolving around outreach efforts to potential applicants, a desire to increase the speed of review and award, and a need for more transparent data systems. On the last point, the reports typically note that administrative data on the program's operations are rarely, if ever, made available to external researchers, which limits evaluation efforts.

Outside of the NASEM reports, evaluation of the *design* of the SBIR program has proven challenging since there have not been many empirically useful natural experiments in the features of the program. However, one study of particular note is Bhattacharya (2021), which develops a structural model of the program and businesses' decisions to engage with it. The results, based on data from the Navy's SBIR program, suggest that the program could be improved by: inviting more businesses to apply; increasing the size of awards; combining Phases I and II into a single Phase; or requiring Phase I awardees to share intermediate progress with each other during Phase II pursuits. Although, the results also show that the benefits of incentivizing more efficient R&D efforts by the businesses can come at a significant cost to the agency funding the program (Bhattacharya 2021). More work like this, focused on the design specifics of the SBIR program, would certainly be valuable.

5.2. New View on Average Returns

A challenge that all analyses of the SBIR program grapple with is choosing what metric to use as the outcome of interest. The theoretical concept is clear: we want to know the net impact of the SBIR program on social welfare, which could be quantified, for instance, in terms of the marginal value of public funds committed to the program (Hendren and Sprung-Keyser 2020). But while the costs of the program are relatively clear (i.e., the total dollar value of awards disbursed plus administrative costs), quantifying the benefits is a more challenging exercise.

With a few exceptions (c.f., Howell et al. 2025), evaluations of the SBIR program have not been able to connect SBIR investments in a given business with the flow of federal contracts (from non-SBIR-related work) to that same business. This has been unfortunate given the explicit goal of the program to meet federal needs, and the results we documented in Figure 3C. Thankfully our dataset allows us to dig deeper into the flows of SBIR funding and non-SBIR contracts to SBIR-awardees.²⁴

Since our data has linked SBIR awardees to the FPDS data, we can observe the dollar amounts of the SBIR awardees' contracts with the government, importantly, including all non-SBIR contracts. Furthermore, if we make three (admittedly large) assumptions,

²⁴See Belenzon and Cioaca (2022) for an important, broader effort to match businesses in the Federal Procurement Data System to details in other databases.

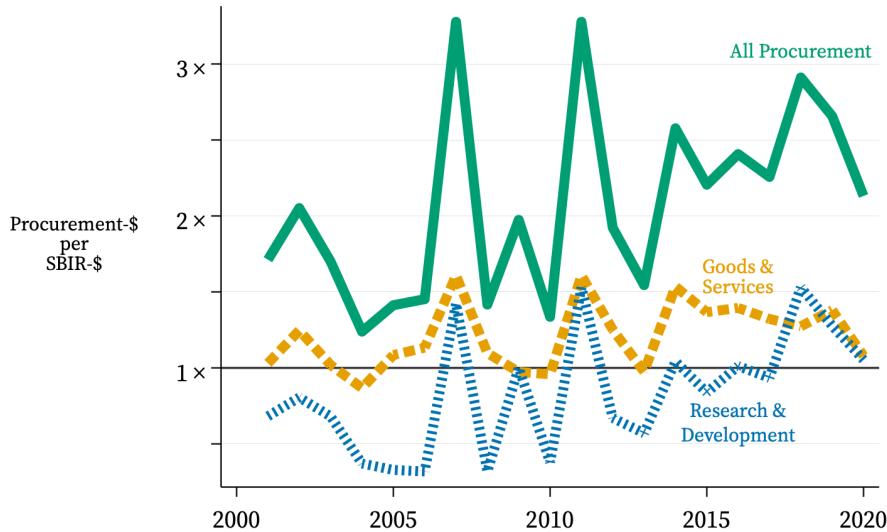
we can provide a new view of the *average* returns to the SBIR program. The assumptions are as follows: (i) the government's willingness to pay an SBIR business for a contract is a proxy for the public value of the contract; (ii) the SBIR business would not have obtained any contracts with the government if not for their participation in the SBIR program; (iii) the government would not have awarded the contract with an SBIR business to a non-SBIR business in the absence of the SBIR business.

Given these assumptions, we can construct a simple measure of the average flow of returns to the SBIR program by dividing total non-SBIR spending by total SBIR awards. This provides a simplistic rate of return metric that is greater than 1 (i.e., indicating a positive return) if SBIR awardees are obtaining more non-SBIR-contract-dollars than SBIR-award-dollars. For the years where we have data on federal contracts, 2000–2021, this ratio is equal to:

$$\frac{\$ \text{ value of all non-SBIR-contracts awarded to SBIR-businesses}}{\$ \text{ value of all SBIR awards}} = 2.7,$$

which indicates that, in a given year between 2000–2021, for every \$1 awarded to SBIR businesses, prior SBIR awardees received an average \$2.7 in non-SBIR contracts with the federal government.

FIGURE 7. Relative Government Spending on SBIR Firms



Note: Plots the ratio of (*numerator*) the total amount of non-SBIR contract dollars awarded to prior SBIR award-winning businesses per (*denominator*) the total amount of SBIR dollars awarded. The ratio is reported annually for three different *numerator* values: (i) all procurement contract dollars, (ii) R&D contract dollars only, and (iii) all non-R&D contract dollars for “Goods & Services”. Note the log scale. For example, if the government awarded \$1 billion in total SBIR funding in a given year, then a ratio of 1 indicates that, in that same year, the government awarded prior SBIR award-winners \$1 billion in procurement contracts.

Figure 7 plots this average return metric for SBIR awardees from 2000–2020. It also splits total non-SBIR contract dollars into spending due to R&D contracts (e.g., to continue developing SBIR-originated technologies) and spending due to non-R&D contracts for goods and services (e.g., to procure fully-developed SBIR-originated technologies). This ratio has been trending upwards since the late 2000s, leveling off in most recent years.

In one sense, this ratio is a lower bound on the average returns of the program since there is a host of socially valuable outcomes from the SBIR program that do not directly result in federal contracts. On the other hand, interpreting the ratio as a metric of average returns does require strong assumptions about the causal effect of the SBIR program. And while the evidence referenced thus far does generally support the view that the marginal dollar invested into the program does spur innovation, the aggregate impact of the program being applied to a national need are still unclear.

5.3. Evaluations of the future

The common view of the SBIR program as a quasi-venture-capital fund has led evaluators to ask questions centered around individual businesses: did an SBIR awardee hire more people, raise more private money, or survive longer than an observationally similar business that did not receive an SBIR award? Impressive data collection efforts and clever research designs (i.e., regression discontinuities around agency score cut-offs) generally conclude that the marginal award boosts business-level performance.²⁵ But the theoretical motivation and practical implementation of the SBIR program suggests a “*SBIR as venture capitalist*” lens misses an important point: venture capitalists are focused on generating private returns based on the development of a specific business, regardless of which ideas they pursue; however, the SBIR program is focused on generating social returns based on the development of a specific idea, regardless of which business is responsible.

Looking forward, the empirical challenge will be to understand *idea-level* counterfactuals: how much more likely is it that some technology is invented, refined, and deployed because of SBIR-backed businesses? Regression-discontinuity designs will remain useful for many important dimensions of this questions. For example, Howell et al. (2025) shows that open-topic solicitations pull new entrants into the defense base, which surely is relevant for a healthy churn of innovative ideas. But many traditional business-level analyses will not tell us whether society ends up with different innovations, or just a different set of winners. Myers and Lanahan (2022) provides a hint of a way forward, in that they are able to estimate technology-level regressions to compare progress surrounding ideas that receive more or less SBIR investment. However, Myers and Lanahan (2022)

²⁵With the notable exception of the null result for traditional SBIR awards identified by the regression discontinuity design in Howell et al. (2025).

rely entirely on the patent record for measures of value, which is still far from the more ideal notion of the marginal value of public funds that policymakers could use compare across programs. Credible future evaluations will need richer outcome measures that convert complex benefits of technologies (e.g., downstream emissions reductions, defense readiness gains, health improvements) into dollar values, and research designs that can identify plausibly-exogenous investments at the idea level, not merely at the firm level (although, aggregations of business-level variation may prove useful in identifying idea-level variation in investments). These efforts will be essential for improving the nation's portfolio of innovation policies.

6. The Future of Small Business Innovation and National Needs

Here, we focus on two major trends relevant to small businesses' role in the US innovation ecosystem: market concentration, and the emergence of artificial intelligence (AI). The innovation policies of the future will need to grapple with both issues.

6.1. Case Study of Consolidation: Defense Innovation Base

The rise of “superstar” or “mega” firms has been one of the most important evolutions of the US economy in recent years (Autor et al. 2020). One sector full of national technological needs where market concentration has become quite pronounced is the national defense industry — see Figure 8A for an illustration. The consolidation of the Department of Defense’s (DoD) major prime contractors continues to receive a significant amount of attention, especially in terms of how it has changed the role of small businesses in the defense innovation base. As highlighted by a recent DoD report: “*Insufficient competition may leave gaps in filling [the DoD’s] needs, remove pressures to innovate ... result in higher costs to taxpayers ... and raise barriers for new entrants.*” (pg. 1, US Department of Defense 2022). Anecdotes of unintended consequences and misaligned incentives are abundant (Shah and Kirchhoff 2024). However, quantitatively diagnosing how this consolidation has played out, especially with respect to small businesses’ role, is quite challenging.

For example, Carril and Duggan (2020) find these consolidations to have led to a significant increase in the rate of “*cost-plus*” contracts. On one hand, that is concerning because, compared to the alternative “*fixed-price*” contracts, the cost-plus contracts increase the potential scope for inefficiencies (i.e., moral hazard). However, Carril and Duggan (2020) find no significant changes to acquisition costs. This may be because the government is still able to exercise significant monopsony power as the only buyer of many goods and services the primes produce (Carril and Duggan 2020). Furthermore, economic theory suggests that the use of cost-plus contracts is more efficient when procuring more complicated goods and services (Bajari and Tadelis 2001), and, by most accounts, the complexity

of defense technologies continues to increase over time.²⁶

In theory, how might this consolidation have affected the rate of defense innovation and the role of small businesses in the defense innovation base? Unfortunately, economic theories yield a wide range of possibilities. Consolidation at this point of a supply chain could plausibly increase or decrease prime contractors' incentives and capabilities to innovate, which in turn has an ambiguous effect on the small businesses that supply primes with innovation (Hart et al. 1990).²⁷

While the trend illustrated in Figure 8A has been widely circulated, it has proven difficult to observe how the rest of the defense innovation base has evolved.²⁸ To provide a new view of the defense innovation base during this period of consolidation among the mega prime contractors, we again look to the patent record.

We identify businesses that are plausibly active suppliers of innovation to the mega defense prime contractors by identifying corporate assignees that meet two criteria: (i) one of the businesses' patents is cited by a patent of one of the mega primes in a given year, and (ii) the business obtains a patent within five years that the citation was received. The first criterion is based on the result due to Fadeev (2024), that most patent citations reflect business arrangements, and the second criterion helps ensure that the business cited by the mega prime is still in operation (as evidence by them recently obtaining a patent). In this case study, since we want to observe the industry trends prior to the shift in consolidation, we also infer business size directly from the patent record based on the number of unique inventors assigned to a businesses' patents in a given year.²⁹

Figure 8B illustrates the change in the number of small businesses that are connected to mega defense primes in the patent record over the same period of consolidation. From 1980 to 2000, there was roughly a 4–5 fold increase in the number of these defense innovation suppliers, and this increase was roughly equal to the relative decrease in the number of mega primes over the same period. However, from 2000 onward, there has been only very modest growth in the number of small and large businesses that are innovation suppliers to the mega primes. Now, while there is 1/16 as many mega primes compared

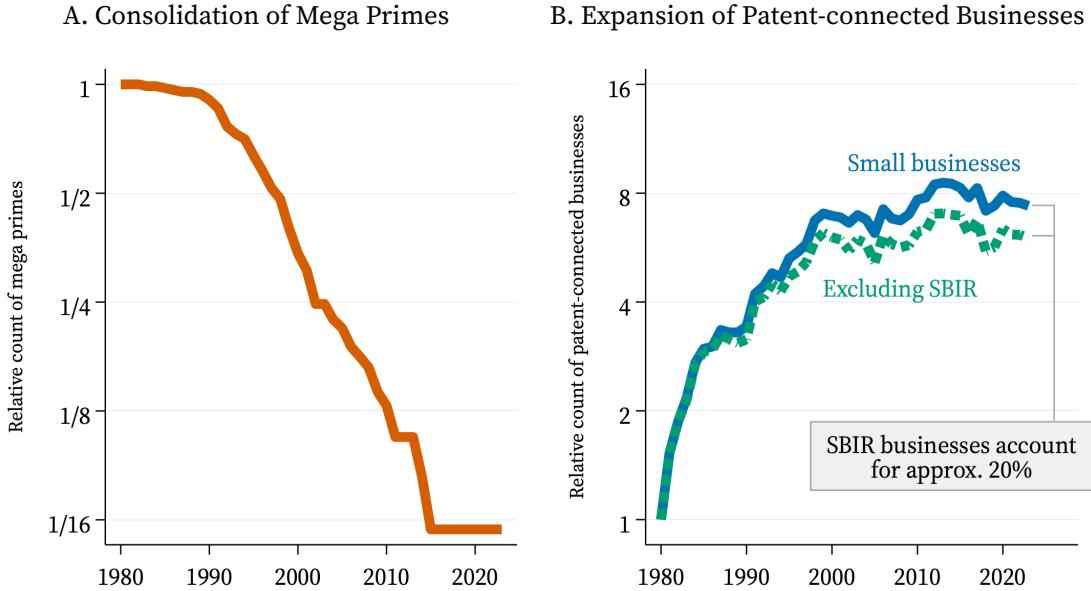
²⁶For example, the number of source-lines-of-code (a measure of software complexity) in the DoD's leading fighter jets have increased by roughly an order of magnitude with each new generation (West and Blackburn 2017).

²⁷Prime contractors may use their increased monopoly power to squeeze rents out of their suppliers, which could shrink the defense base; or, the primes may use that power to foster increased competition and innovation among their suppliers.

²⁸One commonly reported estimate is that the number of small businesses in the defense *industrial* base (which presumably includes non-innovating businesses) shrunk by 40% during the 2010s (Cronk 2021); however, it is unclear how this estimate was constructed.

²⁹Our NETS data does not cover businesses prior to 1990. We label businesses with fewer than 50 unique inventors listed on their patents in a given year, since, for the years we can jointly observe patents and actual business size, very few small businesses (per total employee counts) have more than 50 unique inventors appear in the patent record in any given year.

FIGURE 8. The Evolving Defense Innovation Landscape



Note: Panel (A.) plots the relative count of mega prime defense contractors per Figure 2 of US Department of Defense (2022). Panel (B.) plots the relative count of patent assignee businesses that, in a given year, (i) are cited by a mega prime and (ii) obtain a new patent within the following five years. Small businesses are proxied by having fewer than 50 unique inventors listed in the patent record.

to 1980, there is only roughly 8 more small business innovation suppliers compared to 1980.

Figure 8B also illustrates the extent to which SBIR-backed firms comprised this set of innovation suppliers. In most recent years, we estimate that roughly 20% of these patent-connected small businesses were SBIR program participants.

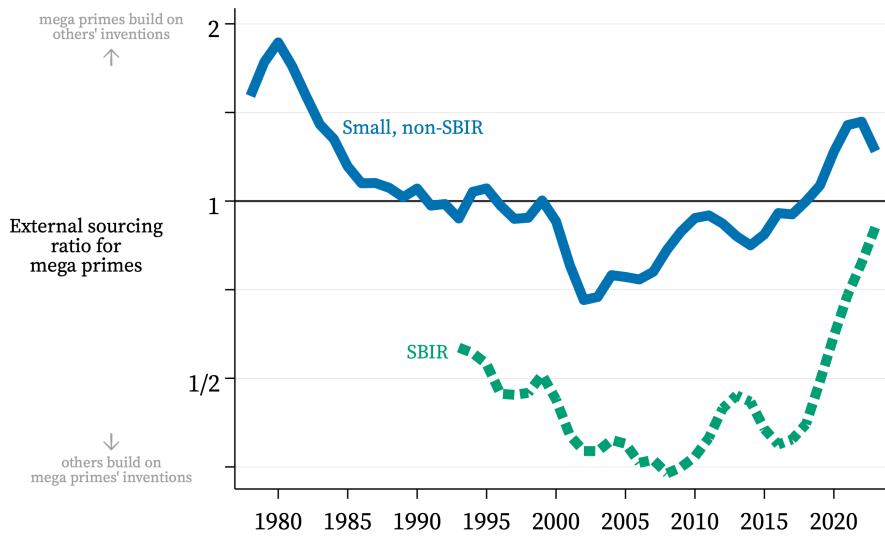
Another way the patent record can provide some insight here is by using it to identify whether innovations have become more *internally sourced*, with mega primes dictating the direction of innovation, or more *externally sourced*, with other businesses developing innovations that mega primes integrate into their own technological systems. Again, we use patent citation flows to proxy for the sourcing of innovation based on whether (i) a patent citation flows from a non-mega prime to a mega prime, which we label as internally sourced, or (ii) a patent citation flows from a mega prime to a non-mega prime, which we label as externally sourced. This allows us to construct a metric of external innovation sourcing that is the ratio of those two citation flows:

$$\text{external sourcing ratio} = \frac{\text{cites to other businesses from mega defense primes}}{\text{cites to mega defense primes from other businesses}}.$$

We do not take a position on the optimal sourcing ratio, but instead use the ratio to understand shifts over time in where the new ideas for defense innovations are sourced from. Figure 9 plots the external sourcing ratio over the past 40 years for non-SBIR small businesses and SBIR businesses.³⁰

The onset of consolidation among the mega primes in the 1990s is apparent. We find a marked shift towards internal sourcing from 1980 to 2000 — it was more common for outsiders to build on the mega primes' inventions than vice versa. Another pattern illustrated by Figure 9 is that SBIR businesses are much more likely to be building inventions on top of mega primes' own inventions. This is consistent with the DoD's historic use of the SBIR program. Awards are targeted at pre-specified needs of the military branches, which will very often revolve around improving the components of the systems developed by the mega primes. This also indicates that the SBIR program has not historically been a vector for new technologies (i.e., inventions outside the scope of the mega primes' systems) to enter the defense base.

FIGURE 9. External Sourcing of Innovation and Mega Defense Primes



Note: Plots the external sourcing ratio, as defined in the main text. Values above 1 indicate that mega primes cite other businesses' patents more than those businesses cite mega primes, and vice versa. Small businesses are proxied by having fewer than 50 unique inventors listed in the patent record.

A final pattern that emerges from Figure 9 is the dramatic increase in external sourcing in recent years. Part of the jump is plausibly pandemic-related; Covid-19 may have created demand for technologies that mega primes had not historically specialized (e.g., medical

³⁰We exclude the initial decade of SBIR data due to the small sample size likely generating a significant amount of noise in the proxy.

logistics, secure tele-work). Institutional reforms within the Department of Defense also may have played a role in this increase. The Defense Innovation Unit (DIU) has shown great success at incorporating new ideas into the defense innovation base with efforts involving novel contract designs (i.e., Other Transaction Authorities) and ensuring geographic proximity to leading technology hubs (i.e., offices in Silicon Valley, Boston, and Austin).³¹ Similarly, AFWERX, within the Air Force, pioneered a more expansive use of “open” SBIR funding announcements that invite businesses to propose dual-use technologies of their choosing (Howell et al. 2025). Likewise, the Army’s xTechSearch prize competition now offers cash awards for promising prototypes and channels winners into follow-on SBIR or procurement contracts.

Taken together, the evidence from the patent record is consistent with anecdotal concerns that incumbent consolidation could make the industrial base stale, but recent efforts may change that. Still, even these metrics we present here are limited in their ability to convey the social value of the technologies at hand. As we noted in the prior section, continued work to develop datasets that quantify and track the social value of technologies will be crucial for being able to monitor the defense industry as it evolves.

6.2. Small Business Innovation and National Needs in the Age of AI

The rapid emergence of generative artificial intelligence (AI) presents some unique challenges and opportunities for the role of small businesses in the innovation economy. An increasingly likely possibility is that AI tools will allow for businesses to simultaneously be both *small*, as defined by the number of humans employed, and *enormous*, as defined by the scale of their business operations. Small businesses with this extreme digital leverage (e.g., 10 humans managing 10,000 agentic AIs) could be able to perform orders-of-magnitude more operations than current businesses, so long as those operations can be digitized.

How might the emergence of small businesses with high digital leverage shape the innovation landscape? First, it is helpful to note that many of the comparative advantages of small businesses vis-à-vis innovation could remain. Although the digital workforce of these firms could be much larger than their human workforce, the division of equity is likely to remain small — there will still only be a few humans in the firm who have much to gain from their innovations. Thus, the incentive misalignment concerns highlighted in Section 4 could still favor small businesses as the ideal setting for small, high-risk, contractible R&D. However, many other important factors could change.

Opportunity costs. With access to such a large digital labor force, small businesses’ opportunity costs may change dramatically. Traditionally, a small business with specialized

³¹See Shah and Kirchhoff (2024) for an excellent review of the DIU.

capabilities would face a narrow set of feasible market opportunities — they would face small opportunity costs if they choose to engage with a narrow national need instead of pursuing a large commercial projects. But as AI capabilities improve, small businesses' ability to pivot to more lucrative commercial markets will increase, and so too will the opportunity cost of engaging with the government. Such a shift would decrease the government's ability to direct small businesses towards the most pressing national needs. As emphasized by Acemoglu (2011), if the costs of switching between technological pursuits are too low, there may be a suboptimal amount of diversity in the innovation base. Monitoring the diversity of small business innovators amid the rise of AI will certainly be important.

Digital versus physical. Another interesting dimension to monitor will be how national innovation needs evolve in the age of AI. As the AI workforce rapidly overcomes many digital bottlenecks, it will expose a new set of technological challenges that cannot be solved by software alone — challenges that will rely on physical skills, analog capabilities, and tacit knowledge. Thus, the canonical image of small business innovation involving software intensive operations may not persist when it comes to national needs. Furthermore, if the bottlenecks to progress become less amenable to digitization, then national needs will be less visible in the data, which implies that data-based management and evaluation may become less possible.

Pacing and sourcing. Historically, international pressures have created national needs that often emerge at the technological frontier. Studies of innovation in the wake of World War II and the Space Race suggest that governments have been responsible for major technological developments (Gross and Sampat 2023; Kantor and Whalley 2025). That is to say, national needs have frequently set the pace of innovation. However, the age of AI may push the pace of progress in consumer-oriented technologies beyond what current government structures can manage. This will require new institutions that ensure that government managers are kept up-to-date on the capabilities of the private sector. Efforts such as the Defense Innovation Unit and the Air Force's AFWERX (discussed in Section 6.1) are promising steps in that direction. Ultimately, success will depend on continued innovations in the process of innovation itself — testing new ways to find and support small businesses, connecting public and private efforts, and shortening the path from prototype to practical use.

6.3. Conclusion

Each group in this innovation ecosystem has a distinct comparative advantages. Small businesses excel at tackling small-scale, well-defined, highly uncertain technical problems. Venture capitalists excel at spotting businesses with large private upside and supplying

the capital and managerial capabilities to scale them rapidly. Government agencies excel at spotting technological externalities (i.e., when there is a large gap between the social and private returns from pursuing technologies that are in national need).

The SBIR program is a unique innovation policy at the intersection of these three groups, and there may be ways to further leverage the comparative advantages of each group. Efforts by organizations such as the Defense Innovation Unit have demonstrated the value of providing small businesses not just with raw capital, but with connections and channels for commercialization. Perhaps the SBIR program of the future could include more explicit resources or guarantees along the commercialization journey so that the small businesses can focus on what they do best: inventing novel and valuable technologies. The perspective of venture capitalists is partially incorporated into the program via participation on review panels, but more is possible. Tibbet's original NSF pilot provides an example of what a more explicit incorporation of the VC perspective could look like: it awarded a bonus in the review stage for businesses with pre-commitments from private investors. And while there is no doubt that government officials work hard to surface national needs (i.e., problems in need of solutions), there may still be gains from making sure officials have the most up-to-date information about small businesses' capabilities at the technology frontier (i.e., solutions in need of problems). Stitching these groups together, while keeping an eye on the value created by the entire SBIR program, not just the success of any specific business, can help this ecosystem continue to deliver outsized benefits to the nation.

References

Acemoglu, Daron. 2011. "Diversity and Technological Progress." In *The Rate and Direction of Inventive Activity Revisited*, 319–356: University of Chicago Press.

Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. "The network origins of aggregate fluctuations." *Econometrica* 80 (5): 1977–2016.

Agarwal, Ruchir, and Patrick Gaule. 2020. "Invisible geniuses: Could the knowledge frontier advance faster?" *American Economic Review: Insights* 2 (4): 409–424.

Aghion, Philippe, and Jean Tirole. 1994. "The management of innovation." *The Quarterly Journal of Economics* 109 (4): 1185–1209.

Anton, James J, and Dennis A Yao. 1995. "Start-ups, spin-offs, and internal projects." *The Journal of Law, Economics, and Organization* 11 (2): 362–378.

Arora, Abhishek, and Melissa Dell. 2023. "Linktransformer: A unified package for record linkage with transformer language models." Mimeo.

Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 609–626: Princeton University Press.

Audretsch, David B. 2003. "Standing on the shoulders of midgets: The US Small Business Innovation Research program (SBIR)." *Small Business Economics* 20 (2): 129–135.

Audretsch, David B, Juergen Weigand, and Claudia Weigand. 2002. "The impact of the SBIR on creating entrepreneurial behavior." *Economic Development Quarterly* 16 (1): 32–38.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. 2020. "The fall of the labor share and the rise of superstar firms." *The Quarterly Journal of Economics* 135 (2): 645–709.

Bajari, Patrick, and Steven Tadelis. 2001. "Incentives versus transaction costs: A theory of procurement contracts." *RAND Journal of Economics*: 387–407.

Baldwin, Carliss Y, and Kim B Clark. 2000. *Design Rules, Volume 1: The Power of Modularity.*: MIT press.

Baldwin, Carliss Y, C Jason Woodard et al. 2009. "The architecture of platforms: A unified view." *Platforms, Markets and Innovation* 32: 19–44.

Barnatchez, Keith, Leland D Crane, and Ryan A Decker. 2017. "An assessment of the National Establishment Time Series (NETS) database." Finance and Economics Discussion Series 2017-110. Washington: Board of Governors of the Federal Reserve System.

Baygan, Günseli. 2003. "Venture capital policies in Israel.", OECD Publishing.

Belenzon, Sharon, and Larisa C Cioaca. 2022. "Guaranteed Public Demand and Corporate Scientific Research." Mimeo.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen. 2019. "Who becomes an inventor in America? The importance of exposure to innovation." *The Quarterly Journal of Economics* 134 (2): 647–713.

BFS. 2024. "Federal Procurement Data System." <https://www.usaspending.gov>.

Bhattacharya, Vivek. 2021. "An empirical model of R&D procurement contests: An analysis of the DOD SBIR program." *Econometrica* 89 (5): 2189–2224.

Bigio, Saki, and Jennifer La'o. 2020. "Distortions in production networks." *The Quarterly Journal of Economics* 135 (4): 2187–2253.

Bruce, Joshua R, John M de Figueiredo, and Brian S Silverman. 2019. "Public contracting for private innovation: Government capabilities, decision rights, and performance outcomes." *Strategic Management Journal* 40 (4): 533–555.

Bryan, Kevin A, and Heidi L Williams. 2021. "Innovation: Market failures and public policies." In *Handbook of Industrial Organization*, vol. 5, 281–388: Elsevier.

Carril, Rodrigo, and Mark Duggan. 2020. "The impact of industry consolidation on government procurement: Evidence from Department of Defense contracting." *Journal of Public Economics* 184: 104141.

Chen, Jiafeng, and Jonathan Roth. 2024. "Logs with zeros? Some problems and solutions." *The Quarterly Journal of Economics* 139 (2): 891–936.

Cohen, Wesley M. 2010. "Fifty years of empirical studies of innovative activity and performance." *Handbook of the Economics of Innovation* 1: 129–213.

Cronk, Terri M. 2021. "Deputy Secretary: How Industry Can Continue Supporting DOD's National Security Priorities." *DOD News*.

Denes, Matthew, Sabrina T Howell, Filippo Mezzanotti, Xinxin Wang, and Ting Xu. 2023. "Investor tax credits and entrepreneurship: Evidence from us states." *The Journal of Finance* 78 (5): 2621–2671.

Dixit, Avinash K, and Robert S Pindyck. 1994. *Investment Under Uncertainty*.: Princeton University Press.

Elliott, Matthew, Benjamin Golub, and Matthew O Jackson. 2014. "Financial networks and contagion." *American Economic Review* 104 (10): 3115–3153.

Fadeev, Evgenii. 2024. "Creative construction: Knowledge sharing and cooperation between firms." Mimeo.

Feldman, Maryann, Evan E Johnson, Remi Bellefleur, Savannah Dowden, and Eshika Talukder. 2022. "Evaluating the tail of the distribution: the economic contributions of frequently awarded government R&D recipients." *Research Policy* 51 (7): 104539.

Gans, Joshua, and Scott Stern. 2003. "When does funding research by smaller firms bear fruit?: Evidence from the SBIR program." *Economics of Innovation and New Technology* 12 (4): 361–384.

Gross, Daniel P, and Bhaven N Sampat. 2023. "America, jump-started: World War II R&D and the takeoff of the US innovation system." *American Economic Review* 113 (12): 3323–3356.

Hart, Oliver, Andrei Shleifer, and Robert W Vishny. 1997. "The proper scope of government: Theory and an application to prisons." *The Quarterly Journal of Economics* 112 (4): 1127–1161.

Hart, Oliver, Jean Tirole, Dennis W Carlton, and Oliver E Williamson. 1990. "Vertical integration and market foreclosure." *Brookings Papers on Economic Activity* 1990: 205–286.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A unified welfare analysis of government policies." *The Quarterly Journal of Economics* 135 (3): 1209–1318.

Holmstrom, Bengt. 1989. "Agency costs and innovation." *Journal of Economic Behavior & Organization* 12 (3): 305–327.

Howell, Sabrina. 2017. "Financing innovation: Evidence from R&D grants." *American Economic Review* 107 (4): 1136–64.

Howell, Sabrina T, Jason Rathje, John Van Reenen, and Jun Wong. 2025. "Opening up Military Innovation: Causal Effects of Reforms to US Defense Research." *Journal of Political Economy* Forthcoming.

Jones, Charles I. 2011. "Intermediate goods and weak links in the theory of economic development." *American Economic Journal: Macroeconomics* 3 (2): 1–28.

Kantor, Shawn, and Alexander T Whalley. 2025. "Moonshot: Public R&D and growth.", Mimeo.

Katz, Michael L, and Carl Shapiro. 1994. "Systems competition and network effects." *Journal of Economic Perspectives* 8 (2): 93–115.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. "Measuring technological innovation over the long run." *American Economic Review: Insights* 3 (3): 303–320.

Kremer, Michael. 1993. "The O-ring theory of economic development." *The Quarterly Journal of Economics* 108 (3): 551–575.

Kretschmer, Tobias, Aija Leiponen, Melissa Schilling, and Gurneeta Vasudeva. 2022. "Platform ecosystems as meta-organizations: Implications for platform strategies." *Strategic Management Journal* 43 (3): 405–424.

Lanahan, Lauren, and Daniel Armanios. 2018. "Does more certification always benefit a venture?" *Organization Science* 29 (5): 931–947.

Lanahan, Lauren, Daniel Erian Armanios, and Amol M Joshi. 2022. "Inappropriateness penalty, desirability premium: What do more certifications actually signal?" *Organization Science* 33 (2): 854–871.

Lanahan, Lauren, Amol Joshi, and Evan Johnson. 2021. "Do Public R&D Subsidies Produce Jobs? Evidence from the SBIR/STTR Program." *Research Policy* 50 (7): 104286.

Lerner, Josh. 2000. "The government as venture capitalist: The long-run impact of the SBIR program." *Journal of Private Equity* 3 (2): 55–78.

Lerner, Josh. 2009. *Boulevard of Broken Dreams: Why Public Efforts to Boost Entrepreneurship and Venture Capital Have Failed and What to Do About It.*: Princeton University Press.

Lerner, Josh, and Ulrike Malmendier. 2010. "Contractibility and the design of research agreements." *American Economic Review* 100 (1): 214–246.

Levin, Jonathan, and Steven Tadelis. 2010. "Contracting for government services: Theory and evidence from US cities." *Journal of Industrial Economics* 58 (3): 507–541.

Link, Albert N, and Christopher J Ruhm. 2009. "Bringing science to market: Commercializing from NIH SBIR awards." *Economics of Innovation and New Technology* 18 (4): 381–402.

Link, Albert N, and Christopher A Swann. 2024. "SBIR mills and the US Department of Defense." *The Journal of Technology Transfer*: 1–30.

Link, Albert, and John Scott. 2010. "Government as entrepreneur: Evaluating the commercialization success of SBIR projects." *Research Policy* 39 (5): 589–601.

Markowitz, Harry. 1952. "The utility of wealth." *Journal of Political Economy* 60 (2): 151–158.

Milgrom, Paul, and John Roberts. 1990. "Rationalizability, learning, and equilibrium in games with strategic complementarities." *Econometrica: Journal of the Econometric Society*: 1255–1277.

Myers, Kyle R, and Lauren Lanahan. 2022. "Estimating spillovers from publicly funded R&D: Evidence from the US Department of Energy." *American Economic Review* 112 (7): 2393–2423.

Narain, Namrata. 2025. "How Patient is Venture Capital?" Mimeo.

National Academies of Sciences, Engineering, and Medicine. 2001. *SBIR/STTR at the Department of Energy.*: National Academies Press.

National Academies of Sciences, Engineering, and Medicine. 2020. *Review of the SBIR and STTR Programs at the Department of Energy.*: National Academies Press.

National Academies of Sciences, Engineering, and Medicine. 2022. *Review of the SBIR and STTR Programs at the NIH.*.: National Academies Press.

National Academies of Sciences, Engineering, and Medicine. 2023. *Review of the SBIR and STTR Programs at the National Science Foundation.*.: National Academies Press.

Sah, Raaj Kumar, and Joseph E Stiglitz. 1986. “The Architecture of Economic Systems: Hierarchies and Polyarchies.” *American Economic Review*: 716–727.

Sahlman, William A. 1990. “The structure and governance of venture-capital organizations.” *Journal of Financial Economics* 27 (2): 473–521.

SBA. 2024. “Award Data.” <https://www.sbir.gov/awards>.

Shah, Raj M, and Christopher Kirchhoff. 2024. *Unit X: How the Pentagon and Silicon Valley Are Transforming the Future of War.*.: Simon and Schuster.

Stein, Jeremy C. 1997. “Internal capital markets and the competition for corporate resources.” *Journal of Finance* 52 (1): 111–133.

Toole, Andrew A, and Dirk Czarnitzki. 2007. “Biomedical academic entrepreneurship through the SBIR program.” *Journal of Economic Behavior & Organization* 63 (4): 716–738.

US Department of Defense. 2022. “State of Competition within the Defense Industrial Base.” Department of Defense Report.

USPTO. 2024. “PatentsView.”

Walls, D. 2021. “National Establishment Time-Series (NETS) Database.”

Wallsten, Scott. 2000. “The effects of government-industry R&D programs on private R&D: The case of the Small Business Innovation Research program.” *RAND Journal of Economics* 31 (1): 82–100.

West, Timothy D, and Mark Blackburn. 2017. “Is digital thread/digital twin affordable? A systemic assessment of the cost of DoD’s latest Manhattan project.” *Procedia Computer Science* 114: 47–56.