

Estimating spillovers from publicly funded R&D: Evidence from the U.S. Department of Energy

By KYLE R. MYERS AND LAUREN LANAHAN*

We quantify the magnitude of R&D spillovers created by grants to small firms from the U.S. Department of Energy. Our empirical strategy leverages variation due to state-specific matching policies, and we develop a new approach to measuring both geographic and technological spillovers that does not rely on an observable paper trail. Our estimates suggest that for every patent produced by grant recipients, three more are produced by others who benefit from spillovers. Sixty percent of these spillovers occur within the U.S., and many of them occur in technological areas substantially different from those targeted by the grants.

JEL: O31, O33, O38

The spillovers of research and development (R&D) are one of the most cited motivations for government intervention in the markets for innovation. Ideas generated by one inventor may lead other inventors to create other new ideas; this is a phenomenon that can lead to suboptimal investment in R&D (Nelson 1959; Arrow 1962). Understanding the magnitude of R&D spillovers and how they permeate geographic space (across inventors and firms) and technological space (across ideas) has been a longstanding empirical question with many theoretical and policy implications (Griliches 1992; Bloom, Van Reenen and Williams 2019; Arora, Belenzon and Sheer 2021).¹ Despite this importance, empirical evidence on the magnitude and nature of R&D spillovers has been remarkably thin.

In this paper, we provide novel evidence on the magnitude and nature of R&D spillovers by analyzing the U.S. Department of Energy (DoE) branch of the Small

* Myers: Harvard Business School, Harvard University, Boston, MA 02163, kmyers@hbs.edu. Lanahan: Lundquist College of Business, University of Oregon, Eugene, OR 97403, llanahan@uoregon.edu. This project was conducted partly within a National Academies of Sciences (NAS) review of the Department of Energy (DoE). The views expressed herein do not necessarily reflect those of the review committee, the NAS, or the DoE. 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. Data collected and created during that review (Myers and Lanahan 2020) are used throughout. We appreciate helpful comments from Maryann Feldman, Scott Stern, Josh Lerner, Shane Greenstein, Joshua Graff Zivin, David Popp, Heidi Williams, Douglas Hanley, Pierre Azoulay, Adam Jaffe, Riitta Katila, Matt Marx, Yael Hochberg, Tim Simcoe, Iavor Bojinov, Xavier Jaravel, Sharon Belenzon, Adrien Matray, Jennifer Kuan, and Daniel Armanios, as well as numerous seminar participants at the schools and conferences where we presented this work. James Bouse from the U.S. Department of Transportation and John Hardin from the North Carolina Office of Science Technology and Innovation were also very helpful in sharing and interpreting data. All errors are our own.

¹Such spillovers may be due to externalities (e.g., uncompensated “knowledge spillovers” or business stealing) or market-based transactions (e.g., licensing contracts). Externalities are important since they can distort R&D investments (Jones and Williams 1998; Bryan and Lemus 2017). All spillovers are important when examining the diffusion of ideas across geographic or technological space (Keller 2004).

Business Innovation Research (SBIR) program. The DoE SBIR program is emblematic of many governments' venture-like R&D subsidy programs: the funding agency announces priority technologies that they would like to see developed, and then awards competitive grants to the most promising small businesses that apply.² At the DoE, these priorities range from light sensors, to fuels, to high performance computing software, and are described in Funding Opportunity Announcements (FOAs). Notably, all DoE SBIR applications must respond to a specific FOA. To measure how grants awarded via FOAs ultimately lead to new ideas, we follow a long line of related work (Griliches 1998) and use patenting rates at the U.S. Patent and Trademark Office (USPTO) as our focal measure of inventive output.

We identify the magnitude of spillovers across geographic and technological space by estimating how SBIR grants ultimately lead to new patents that are (1) produced by inventors and firms located further and further away from a grant recipient, and (2) focused on technologies more and more different from the technologies targeted by FOAs. Our results indicate that, for every one patent produced by grant recipients due to a marginal increase in SBIR funding, we expect three additional patents to be produced by other inventors across the U.S. and around the world (none of whom directly received a grant). Our results imply that US-based firms and inventors are responsible for roughly 60 percent of the net patents generated by this program. Moreover, many of these new patents are topically quite distant from the technologies initially targeted by FOAs. Taken together, these results suggest that the R&D spillovers from these government grants have a wide impact across both geographic and technological space.

Our empirical approach is designed to address two of the primary challenges to estimating R&D spillovers: measurement and identification. To illustrate the measurement challenge: say that a firm in New Hampshire received an SBIR grant to develop a tool for detecting hazardous waste underground. What observable data could tell us how other inventors and firms benefitted from the R&D performed by that grant recipient?³ Prior approaches to this problem have largely relied on one of two methods. The first uses citation linkages between patents as a proxy for spillovers (Jaffe 1986; Jaffe, Trajtenberg and Henderson 1993). The second approach relies on "spillover matrices" that use other observable data (e.g., inventors' locations) to impose structure on the way in which firms' R&D influences others along dimensions of geographic, technological, or product market proximity (e.g., Bloom, Schankerman and Van Reenen 2013).

²A focus on small firms can be linked to the unique frictions they face (Holmstrom 1989; Lerner 2009) or the outsized role of young firms in the economy (Haltiwanger, Jarmin and Miranda 2013). A focus on specific technologies can be linked to distortions in how firms choose the direction of their research (Bryan and Lemus 2017).

³In this actual case, the New Hampshire-based firm, Subsurface Insights L.L.C., received this SBIR grant and ultimately expanded from a focus on hazardous waste into other areas including thermal energy storage in aquifers (suggestive of technological spillovers) and became involved in projects across the U.S. and Europe (suggestive of geographic spillovers). See <https://www.subsurfaceinsights.com/about.html> and split.to/sbirsurface for more.

We offer a new approach by drawing on methods from the spatial economics literature (Clarke 2017; Feyrer, Mansur and Sacerdote 2017; James and Smith 2020) along with natural language processing tools to capture spillovers more flexibly. Our key advantage is that we can construct a data-frame where the units of observation are areas of technological space. That is, we estimate “technology-level” regressions that relate the flow of patents to the stock of R&D grants across technologies.⁴ We identify the technological areas where spillovers are most likely to occur by using the similarity of the text in (1) patents abstracts and (2) the research objectives of the grants as described in the FOAs. On the geographic dimension, we identify spillovers by comparing the marginal product of funding when we focus on the patent output from different groups of firms and inventors (e.g., only grant recipients, all US-based firms and inventors, etc.).

Importantly, our approaches to measuring both geographic and technological spillovers require no paper trail of citation linkages and make no a priori assumptions about how “far out” in geographic or technological space spillovers from SBIR grants may permeate. For comparison, we estimate our model using the traditional approach of requiring citation linkages and the results indicate that this approach misses roughly half of the spillovers that we identify with our more flexible approach.

The identification challenge revolves around handling the endogeneity of R&D investments. In our setting, neither the technological areas targeted by FOAs nor the firms awarded the federal grants are randomly selected. There are likely unobservable differences in the supply of and demand for technologies that the DoE prioritizes. Moreover, firms of higher unobservable productivity will likely be able to obtain more SBIR funding.

We isolate plausibly exogenous variation in R&D funding by leveraging the presence of non-competitive SBIR matching policies that vary across states and time. In some U.S. states, when a firm wins a federal SBIR award, their state government automatically awards them additional funds (Lanahan and Feldman 2015). We use variation in these policies to identify “windfall” investments into technological areas that are arguably unrelated to any (unobservable) supply or demand shocks also driving R&D investments and patenting rates. To illustrate, consider two equally productive firms that are located in different states. Further, assume these firms respond to two separate FOAs focused on two different technologies, where the supply of and demand for both technologies are the same. If both firms win an SBIR award, but one firm is located in a state with a match policy, then additional funds will be invested (exogenously) in the technology pursued by that firm. We show evidence that this thought experiment often plays out in practice – states with match policies do not appear to be disproportionately populated by especially productive (or unproductive) firms pursuing especially

⁴Popp (2002) may be one of the first to estimate technology-level regressions that relate inputs and outputs within some predefined portions of technology space. Azoulay et al. (2019) also estimate a technology-level production function, though their technologies are grouped around biomedical topics.

productive (or unproductive) technologies.

We are not the first to tackle the aforementioned measurement and identification challenges associated with R&D spillovers. The two most closely related papers are Bloom, Schankerman and Van Reenen (2013) and Azoulay et al. (2019). Bloom, Schankerman and Van Reenen (2013) focus on a sample of large, publicly-traded firms, and leverage state-level R&D tax credits combined with the type of spillover matrices described above to test for both positive R&D spillovers (across technologies) and negative R&D spillovers (within product markets).⁵ Azoulay et al. (2019) are not focused on estimating R&D spillovers directly, but their empirical approach – which isolates windfalls of public R&D funding targeted to specific biomedical topics – uncovers that, in academic biomedical research, spillovers across biomedical topics account for a substantial share of the ways in which publicly funded R&D translates into realized inventions. We contribute to this line of work by offering a new methodology for capturing R&D spillovers while focusing on the increasingly important energy sector.

Our preferred estimates suggest that SBIR-funded firms are capturing anywhere between 25–50% of the net patent-based value their R&D is generating. These magnitudes are on par with a very short line of empirical studies using macro- and microeconomic methods to estimate the private and social returns to R&D (Jones and Williams 1998; Bloom, Schankerman and Van Reenen 2013; Azoulay et al. 2019; Zacchia 2020). This suggests that our best estimates of the wedge between the private and social returns to R&D are not primarily driven by any specific methodology or setting. In addition to a number of specification and robustness tests, we also provide support for our research design by benchmarking an intermediate estimate we obtain from our approach: the amount of SBIR funding necessary to spur one additional patent by (only) grant recipients. Fortunately, we have a solid benchmark with which to compare, Howell (2017), which uses a sharp regression discontinuity to identify clear evidence that these same grants enable firms to innovate.⁶ A conservative interpretation of Howell’s (2017) estimates implies a marginal cost per patent that is very close to our estimate; this gives us further confidence in our empirical design. Still, the nearly four-fold difference between this cost (of spurring one patent by grant recipients) and the cost of spurring a patent by *anyone* suggests that ignoring spillovers in program evaluations of R&D subsidies may yield very different conclusions.

We proceed as follows. Section I provides detail on the setting and primary data sources. It also outlines how we use natural language processing tools to connect the inputs (grants) and outputs (patents) of the focal production function we es-

⁵Product and revenue data is not available for the vast majority of the firms and inventors in our analyses, so we cannot fully investigate product market rivalries. However, our estimates do incorporate any business stealing effects that occur up to the patenting stage of R&D (e.g., if, after one firm successfully obtains a patent, another withdraws R&D efforts that would have otherwise led to a patent).

⁶Other studies of the SBIR program, most of which do not address spillovers, include Lerner (2000); Wallsten (2000); Audretsch (2003); Gans and Stern (2003); Link and Scott (2010); Lanahan and Feldman (2018); National Academies of Sciences, Engineering, and Medicine (2020), and Lanahan, Joshi and Johnson (2021).

timate. Section II describes the empirical model, which builds on the “knowledge production function” first introduced by Griliches (1979). We first describe how we use the noncompetitive state matching policies to isolate windfalls of funding and then how we incorporate spillovers across geographic and technological space into the model. Section III reports the main results on the size of R&D spillovers and where they occur. Section IV focuses on the value of patents spurred by these grants and some implications the results have for understanding the private and social returns to R&D. Section V concludes by connecting our findings to ongoing literatures and offering some economic and policy implications.

I. Setting and Data

This section describes the setting, policies, and data used in our empirical analysis. We detail: the DoE SBIR program including how they solicit applications and grant funds (Section I.A); the State Match programs that award extra non-competitive funding to federal SBIR winners (Section I.B); the U.S. Patent and Trademark Office (USPTO) patent data and the classification scheme used to categorize patents (Section I.C); how we use natural language processing to merge the SBIR and patent data (linking inputs to outputs) and create technology-level data for empirical analysis (Section I.D); and finally, summary statistics of the multiple data sets (Section I.E). Appendix B details additional data sets we use to observe or estimate geographic distances, travel costs, and other characteristics of countries and U.S. counties.

A. The U.S. SBIR Program and the Department of Energy

Since its enactment in 1982, the U.S. SBIR program has allocated over \$40 billion to support early-stage innovation by small businesses. It is emblematic of R&D subsidy programs administered by many countries worldwide.⁷ The DoE is one of eleven federal agencies participating in the SBIR program and is currently awarding roughly 250 firms a total of about \$200 million per year. Currently, the DoE SBIR program spans twelve offices,⁸ whose focus spans a broad range of energy-related initiatives including those that may not immediately come to mind when thinking of “energy,” including high-performance computing, cybersecurity, and precision measurement tools.

The DoE SBIR solicitation and award process is as follows. One to three times per year, the DoE releases a Funding Opportunity Announcement (FOA) that outlines technological areas of interest, referred to as “topics.” Each year, the FOAs include about 50 unique topics altogether. Topic descriptions are roughly

⁷Upwards of 17 countries have copied some structures of the SBIR program (see: tny.sh/sbirworldwide).

⁸The offices participating in our data include: Electric Delivery and Energy Reliability, Energy Efficiency and Renewable Energy, Environmental Management, Fossil Energy, Nuclear Energy, Defense Nuclear Nonproliferation, Advanced Scientific Computing Research, Basic Energy Sciences, Biological and Environmental Research, Fusion Energy Sciences, High Energy Physics, and Nuclear Physics.

five to ten paragraphs of text outlining the specifics of the technology that the DoE is interested in developing. Each office within the DoE is responsible for producing a set of topics contained in an FOA. To be eligible for funding, applicants must align themselves with the objectives of a particular topic in an active FOA.

Only private US-owned firms with less than 500 employees are eligible to apply for SBIR funding. Proposals are subject to an internal review to ensure there is sufficient alignment between the proposed activities and the FOA topic the firm is responding to. Then follows a competitive, external peer review process where at least three industry experts review the proposal’s technical and commercial merits. Recipients receive the first award, a Phase I grant that provides six months of support and typically \$50,000–250,000, which is intended to support development of a proof of concept. Phase I recipients are then eligible to compete for Phase II grants, which offer two years of support and typically \$750,000–1 million.⁹

The process of selecting FOA topics and recipient SBIR firms is not random. DoE program managers attempt to solicit applications for developing technologies that have a significant potential for impact. Conditional on the set of applicants, the peer review process attempts to direct funds toward firms with the most potential success. This likely leads to endogeneity issues, which we attempt to circumvent by leveraging the presence of the matching programs that are active in different states.

B. SBIR State Match Programs

Since 1984, some U.S. states have enacted policies intended to complement the federal SBIR program by awarding additional funds to firms that receive an SBIR grant. Lanahan and Feldman (2015) outline the origins of these match programs and how they operate. Following their methodology, we identify the SBIR match programs for all states up to 2018.¹⁰

Figure 1 documents the growth of these policies over time, with 15 states having a program in 2018. Most states tend to maintain the program after initial adoption. On average, 7 percent of DoE SBIR grant recipients are eligible for a state match, though with program diffusion in recent years, this increased to 17 percent in 2018. The size of these matches ranges from \$25,000–105,000 for Phase I awards and \$50,000–500,000 for Phase II awards.

These state policies are at the core of our research design. We discuss how

⁹Roughly eight percent of grants are awarded via the Small Business Technology Transfer (STTR) program, which requires a partnership between the small business applicant and a nonprofit research institution. We lack a strategy for separately identifying the effect of grants via this channel, so we aggregate all funding through both the SBIR and STTR channels into a single measure of “SBIR” funding.

¹⁰We exclude a small number of states’ programs that involve competitive elements (e.g., peer review) since our identification strategy assumes that these matching funds are automatically awarded to all federal grant winners. Despite much effort, we cannot observe the actual amount of match funds awarded to firms since states’ records are very incomplete. Still, we know how much match funding a recipient should have received, which motivates our interpretation of the empirics as an intent-to-treat analysis.

we use them in combination with some plausible, data-supported assumptions to isolate exogenous variation in R&D funding in Section II.B below.

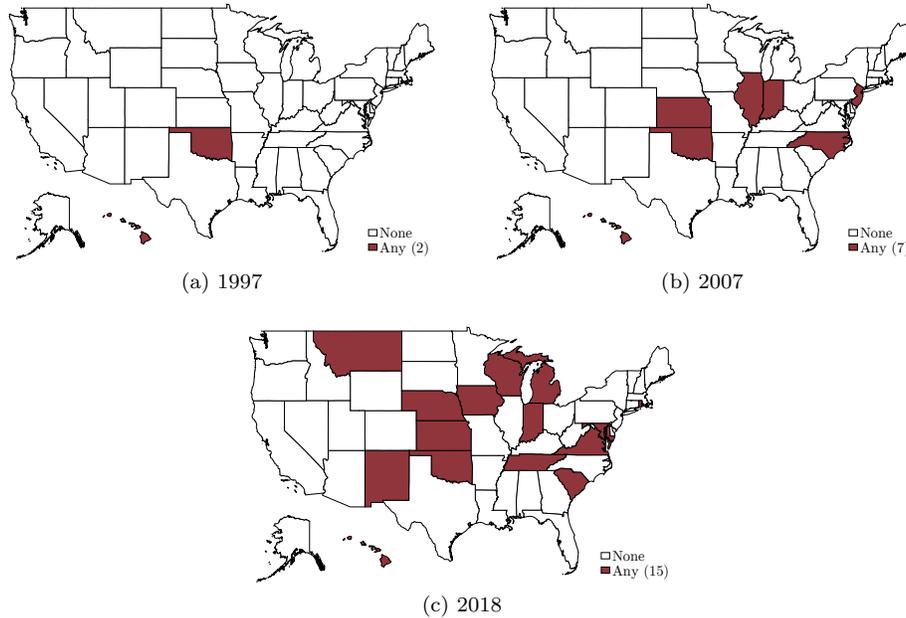


Figure 1. : State Match Programs over Time

Note: The distribution of match policies across states and over time. Most policies award an additional \$50–75,000 (103 State-Year obs.), with some smaller matches for Phase I awards (39 State-Year obs. with \$25–45,000 matches) and some larger matches for Phase II awards (48 State-Year obs. with \$250–500,000 matches).

C. Patents and Technology Indexing

We follow a long line of economic studies of invention and rely on patent activity as our measure of technological progress. We source the universe of USPTO patents from 1997 to 2018 using the USPTO’s PatentsView database (USPTO 2020). This contains patent details such as the application year, the title and abstract text, and disambiguated tables of individual inventors and firm assignees as well as their geographic locations.¹¹

Accounting for all firms in the PatentsView data and the five largest federal SBIR agencies,¹² roughly 49 percent of firms that ever receive an SBIR award ever obtain a USPTO patent between 1997 and 2018 (before or after the SBIR

¹¹Patents are often used as a proxy for innovation in policy communities, and a recent in-depth case study supports the assumption that patents are useful metrics of “real” technological progress (Igami and Subrahmanyam 2019). Furthermore, they are often referenced as a primary measure of success by the SBIR program itself (e.g., see slide 4 of the SBIR program overview slides here: tny.sh/sbirprogramslides).

¹²These agencies – DoE, Department of Defense, National Aeronautics and Space Administration,

award). However, roughly 1 percent of firms ever assigned a patent receive an SBIR award at some time.

Importantly, PatentsView also contains data on the unique Cooperative Patent Classification (CPC) terms assigned to each patent. This classification scheme is used to organize patents according to the technical features of their content. The CPC terms effectively discretize technological space into over 250,000 unique concepts that are organized into a five-level hierarchy. Of use to our design is the feature that the USPTO retroactively assigns the most recent and detailed CPC scheme (as of 2018) to all prior patents. For reasons we detail in Appendix A, we use the “main group” level of the hierarchy as the units of analysis in our regressions. At this level of the scheme there are 10,686 unique terms, which we refer to as groups for simplicity. On average, a patent is assigned five to eight CPC groups.

D. Technology-level Data on R&D Investments and Patents

Our empirical model relates the aggregate investment of SBIR funds targeted at particular technologies to the eventual aggregate flow of patents that are related to those same technologies. Our approach centers on using the CPC scheme to index technological space much like how U.S. ZIP codes index geographic space into discrete units. However, also much like ZIP codes, the CPC scheme was not developed with empirical researchers in mind. Therefore, our approach attempts to avoid relying on the CPC scheme in a manner that may produce spurious results (e.g., Thompson and Fox-Kean 2005).

As discussed in Section I.A, the DoE invests in technologies through their use of FOAs and the topics described therein. Moreover, as discussed in Section I.C all patents are assigned CPC groups that describe their technical content. Since patents are already assigned CPC groups, our goal is to infer which CPC groups most closely align with each FOA topic. With this information in hand, then, loosely speaking, if we observe \$1 million awarded through a particular FOA topic, we know in which CPC groups to “look” for new patents.

We detail our approach to linking FOA topics to CPC groups in Appendix A. What follows is a brief summary. First, we use a combination of standard optical character recognition software and hand coding to digitize DoE’s FOAs and parse each FOA into the separate topics. In the second step, we take the five to ten paragraphs of text that describe the objectives of each FOA topic, and compare the textual similarity of these descriptions with the abstracts from all USPTO patents applied for between 1997 and 2004.¹³ Using methods now commonplace

National Institutes of Health, and National Science Foundation – account for approximately 85 percent of federal SBIR funding over our study timeframe.

¹³The choice of this timeframe poses the possibility of a precision-bias tradeoff: using more years, and more recent years, in this text analysis should incorporate information that truly reflects the underlying connection between each word in the FOA topics and each CPC group. However, it also increases the risk we might incorporate information that arises endogenously. For example, when an SBIR grant recipient writes the abstract of their patent, they may use verbiage from the FOA topic related to their grant

in text analyses (term frequency–inverse document frequency weighted n-gram cosine similarity), we estimate a numerical similarity score between each FOA topic and each patent. This approach assumes that if an FOA topic and a patent abstract both use words that are very uncommon elsewhere in each corpus, then the two are likely referencing the same technologies. Throughout the paper, we refer to technological distance as the inverse of this similarity.¹⁴ Appendix A.3 contains three examples of FOA topics and the most similar CPC groups matched to those topics.

An important assumption of our approach is that CPC groups describe the entirety of technological space. Motivated by Thompson and Fox-Kean (2005), we deviate from prior work (e.g., Jaffe, Trajtenberg and Henderson 1993) and do not assume that the specific hierarchical structure of the CPC scheme contains any information. In other words, we do not assume that two groups located next to each other in the hierarchy are more technologically similar than other groups.

By instead using natural language processing techniques, we offer a new methodological approach to estimating patent-based R&D spillovers that is less dependent on the structure of the patent classification system. This methodology could be applied to any setting where the researcher observes input and output data that are not classified in the same way, but where text descriptions of the units of observation can be used to generate a crosswalk.

E. Summary Statistics

Table 1 contains summary statistics of the DoE SBIR program, state match programs, and patent activity. We report annual averages for our full sample, which spans 1997–2018. The Federal DoE SBIR program releases about 50 FOA topics annually, typically awarding roughly \$1.7 million per topic. Conditional on receiving any award, firms receive about 1.8 awards per year, totaling roughly \$780,000. Approximately one-quarter of DoE SBIR awards are Phase I grants, with the remaining three-quarters awarded via Phase II grants.¹⁵

Approximately 7 percent of firms winning awards are eligible for a state-based match, with roughly 18 percent of FOA topics awarding a grant to one of these firms in a match state.¹⁶

even though those words may not actually reflect the technical content of the invention. We are not very concerned with this particular issue because (1) patent applicants face legal threats if they incorrectly describe the nature of their invention, and (2) our discussions with DoE SBIR program staff never raised such a concern. Still, we use only the first eight of our 22-year sample to estimate these similarity scores to avoid the likelihood of such problems, obtaining very similar results when we vary this window.

¹⁴This approach of using text similarity to connect units in technology-space has been shown to be an effective way to quantify economically meaningful concepts (e.g., Azoulay et al. 2019; Myers 2020).

¹⁵All dollar amounts are adjusted using the CPI (BLS 2019) into 2018-\$.

¹⁶To investigate how firms use this funding, we examined a survey administered to state match recipients from North Carolina (NC), which is the only such data we are aware of. With the caveat that the survey was based on firms receiving SBIR awards from all federal agencies (not just the DoE), firms reported that roughly 50% of the matching funds are spent on wages, which is in line with Howell and Brown (2020). About 25% of funds are spent on equipment, supplies, or facilities. It also appears that no more than 5% of these state match funds are spent on patent-related costs. As far as we can infer

Table 1—: Summary Statistics

	Type	mean	s.d.	min	p50	p99
<u>Federal DoE SBIR Program</u>						
Phase I \$, total	Dol., M	47.5	13.9	26.6	43.9	77.2
Phase II \$, total	Dol., M	152.0	45.0	100.1	132.6	240.7
Firms with awards	Count	255.6	47.0	195	257	379
Awards per firm	Count	1.86	1.89	1	1	10
\$ per firm	Dol., M	0.78	1.05	0.02	0.28	5.11
FOA Topics	Count	54.0	9.92	38	55	72
FOA Topics with awards	Count	51.4	9.90	35	52	68
\$ per FOA Topic	Dol., M	3.02	2.90	0.02	2.35	13.3
<u>State Match Programs</u>						
Awards with match \$	Count	39.5	40.4	1	27	133
Match \$, total	Dol., M	3.74	3.91	0.07	3.29	12.4
Firms with match \$	Share	0.07	0.06	0.00	0.05	0.17
FOAs with match \$	Share	0.18	0.14	0.01	0.18	0.39
<u>USPTO Firm-level Patent Flows</u>						
Ever-SBIR firms ever patent, full sample	{0, 1}	0.49				
Ever-patent firms ever SBIR, full sample	{0, 1}	0.01				
Patents per SBIR firm, post-1 st -award	Count	0.51	2.62	0	0	9
Patents by all SBIR firms, post-1 st -award	Count	487.5	167.8	146	485	727
<u>USPTO CPC group-level Patent Flows</u>						
Unique focal CPC groups, full sample	Fixed	10,686				
SBIR firms	[0, ∞)	0.06	0.49	0.00	0.00	1.23
Counties of SBIR firms	[0, ∞)	9.01	52.2	0.00	0.75	147.4
All US	[0, ∞)	10.3	56.3	0.00	1.00	161.5
Worldwide	[0, ∞)	20.8	106.7	0.00	2.13	347.4

Note: All summary statistics based on the annual averages for the years 1997 to 2018 unless otherwise indicated by “full sample,” which reports statistics for all observations.

Roughly half of all firms that ever receive an SBIR award ever obtain a patent, and roughly 1 percent of all firms in the USPTO record ever receive an SBIR award. Focusing only on years after a firm’s first SBIR award, grant recipients average 0.5 patents per year, with all patent flows showing a high degree of skewness.

Across the roughly 10,000 CPC groups we focus on, SBIR firms make up a very small percentage of the total patents within each CPC group (approx. 0.2%). However, counting all of the firms and inventors located in the same U.S. counties as SBIR grant recipients, which includes about 700 of the roughly 3,000 counties, covers roughly 90% of US-based patents.¹⁷ The domestic-foreign split of firms and inventors receiving USPTO patents over this period is roughly 50–50.

from these data and other states’ regulations, it does not appear that receiving state matching funds changes any incentives or costs surrounding patenting.

¹⁷For more on the geographic distribution of U.S. inventive activity, see Forman, Goldfarb and Greenstein (2014).

II. Research Design

Our goal is to estimate the magnitude of the R&D spillovers created by the DoE’s SBIR program. To do this, we estimate the marginal product of funding, in terms of producing patents, and compare how the marginal product depends on (1) whether the patents are produced by SBIR grant recipients, or inventors and firms some geographic distance away from the recipients, and (2) whether the patents are closely aligned with the technological objectives of an FOA topic, or the content is some technological distance away from those objectives. If it appears that SBIR funding is absolutely more productive (i.e., the SBIR-\$ cost per patent declines) as we expand the set of firms and inventors or technologies included in our analyses, we infer that to be the result of R&D spillovers.

A. Baseline Model

We focus on estimating a knowledge production function (Griliches 1979), which relates an annual flow of patents to a stock of R&D investments. Researchers have been regressing patent rates on R&D inputs for decades, but unlike most prior work where the units of analysis are firms or inventors, our units of analysis are the technological areas where firms and inventors operate, which are indexed by the CPC groups.

We use a Poisson model to describe the expected count of ultimately successful patent applications Y in each CPC group j during year t as a function of the current stock of public R&D investments K (i.e., the DoE’s SBIR grants and/or the state-based matching grants):

$$(1) \quad \mathbb{E}[Y_{jt} | K_{jt}] = \exp(\log(K_{jt})\beta + \tau_t + \omega_{jt}) ,$$

where β is the focal productivity parameter to recover and τ_t are year-specific intercepts that condition out aggregate year-to-year fluctuations and also handle the right censoring of our data (investments in year t can only lead to more patents in the current and $2018 - t$ succeeding years). ω is a parameter that captures productivity shocks that are unobservable to us (e.g., the value of patenting in each CPC group at a particular time).¹⁸

We construct the stock of public R&D investments K_{jt} using standard perpetual inventory methods. In our main specifications, we do not discount to produce conservative estimates. In Appendix E.4, we report additional results using alternative discounting assumptions to provide some suggestive evidence as to the degree of production lags. To construct the flow of patent output in each CPC group, Y_{jt} , we divide one by the number of CPC groups assigned to each patent (which typically ranges from five to eight) and sum over all patents by each CPC

¹⁸Eq. 1 does not structurally describe firm- or inventor-level production, but rather aggregates these decisions into a technology-level relationship to estimate the magnitudes of spillovers across both technological and geographic space.

group. Thus, $Y_{jt} = 1$ indicates a flow of one “patent’s worth” of output in CPC group j in year t , which ensures that our marginal product estimates are always with respect to the production of one patent (regardless of how many CPC groups it might be assigned).¹⁹

As previewed in the introduction, we face two econometric challenges for estimating Eq. 1: (1) accounting for the endogenous allocation of SBIR funds that would give rise to a correlation between K_{jt} and (the unobservable) ω_{jt} ; and (2) capturing spillovers across geographic and technological space. The remainder of this section describes how we handle these challenges and calculate the implied magnitude of R&D spillovers in this setting.

B. Identification with Windfall Funding from the Match Programs

The DoE’s SBIR investments are not made at random, but rather with some knowledge about their expected productivity (i.e., the probability they will lead to new patents and spillovers). In our model, these (unobservable) shocks are reflected in the ω_{jt} term. If these shocks are correlated with SBIR investments (i.e., $\mathbb{E}[\omega_{jt} | K_{jt}] \neq 0$), then we cannot recover an unbiased estimate of β .

To circumvent this concern, we make use of the aforementioned non-competitive State Match programs. The varied presence of these policies across states and over time generates variation in the amount of state-based R&D funding that ultimately is invested in each CPC group-year observation. We argue that a subset of this variation yields “windfalls” of R&D funding across CPC groups only because of their differential exposure to the match policies. The key identification assumption we make is that the firms located in states when an SBIR match program is in place are not abnormally productive compared to firms in other state-years, and are not responding to FOA topics that are focused on technologies with abnormal underlying supply or demand. If this is true, then we can isolate these windfalls for estimating of Eq. 1 and recover an estimate of β that is less likely to be biased by endogenous selection of R&D funding. The following three subsections overview our approach by: (1) describing how we isolate windfall investment; (2) providing evidence supporting our assumptions that these policies are plausibly exogenous; and (3) discussing the interpretation and generalizability of this approach.

The state match programs operate by multiplying endogenous federal investments. Thus, we must separate which of the state-based investments are a direct function of certain FOA topics receiving more federal dollars versus the so-called windfall dollars that arise only because of a differential exposure to certain state match policies.

We isolate these windfalls of funding by treating them as residuals in a FOA topic-level regression that predicts total state-based funding as linear, year-specific

¹⁹In Appendix E.4, we report some results from an alternative approach that of allocating patent counts to CPC groups, which yields very similar estimates of spillover magnitudes as our preferred approach.

functions of federal funding. Appendix C.1 outlines this approach in detail. In short, FOA topics and their corresponding CPC groups receive larger (or smaller) windfalls of funding if the total amount of state-based funding was above (or below) the average amount of state-based funding one would expect given the amount of federal funding observed.

Because these windfalls, which we denote with W_{jt} , are based on residuals, they take both positive and negative values. Thus, the standard logarithm transformation of the R&D stock in the knowledge production function is not possible. To circumvent this problem, we approximate the logarithm transformation by dividing R&D stocks by the sample average. Formally, our main estimating equation is now:

$$(2) \quad \mathbb{E}[Y_{jt} | W_{jt}] = \exp\left(\frac{W_{jt}}{\bar{W}}\theta + \tau_t\right),$$

where \bar{W} is the sample mean of the stock of windfalls (W_{jt}). Our approach of dividing the stocks by the sample average approximates a log transformation in terms of the coefficient estimated (i.e., $\theta \approx \beta$).²⁰ Note also that Equation 2 no longer include the unobservable ω_{jt} term because it is, by assumption, uncorrelated with the windfall funding.

Our identification approach relies on the fact that the CPC groups that the DoE invests are differentially exposed to the state matching policies. Specifically, the key assumption is that these policies are not more or less prevalent amongst states with either: (1) particularly more or less productive firms; and/or (2) firms operating in CPC groups that are particularly more or less productive.

To explore these assumptions, we perform two empirical tests. First, we estimate the association between the presence of a matching policy in a state and the amount of federal SBIR funds per capita awarded to firms in that state. The results from these regressions are in Appendix C.2. Across all specifications, we cannot reject a null of no differences. Whether we examine across- or within-state variation, there does not appear to be a meaningful difference in the flow of total SBIR funds between states with and without the matches, nor between states with different match rates.

Second, we assess whether the program solicits selection concerns among firm behavior. Because these match policies technically increase the expected size of an SBIR award, firms may move to states that enacted these policies. However, with match amounts ranging from \$25,000–500,000, and federal-level success rates reported at roughly 15%–20%, this is likely not a meaningful change in the expected benefit of moving across state lines, especially compared to the costs of relocation.

Still, we explore this possibility of firm movement being correlated with these

²⁰A one unit increase in $\frac{W_{jt}}{\bar{W}}$ describes a 100% increase in investments from the mean. See Appendix C.2 for more on this approximation.

policies by constructing a panel data set of the locations of all SBIR grant recipients from the major federal agencies.²¹ We test whether, over the course of their early lifespans, firms are more or less likely to move into states after a matching program is enacted. If the enactment of these policies significantly altered the value of SBIR program participation, or were perhaps correlated with meaningful economic shocks, we would expect to identify an association between the beginning of the policy and the movement of SBIR grant recipients into (or out of) the state. The results of this analysis are also presented in Appendix C.2 and show no meaningful evidence that grant recipients are more or less likely to move into states after these policies are enacted. This is consistent with the findings of Lanahan and Feldman (2018), the first to use these policies to identify plausibly exogenous public R&D funding, as well as work on the location choices of new firms more generally (e.g., Stam 2007; Li et al. 2016).

While our identification approach gives us confidence that we are focusing on useful (and plausibly exogenous) variation in R&D funding, we acknowledge that it is not without some important caveats and limitations. As we outline below, we believe that these limitations generally push us towards more conservative productivity estimates when relying on the state match windfalls.

The windfall funding amounts are much smaller than total funding levels, typically 1–10% of what the DoE directly invests in each CPC group. On the one hand, the small relative size of these investments raises some concerns about measurement error in the independent variable, which would bias our estimates toward zero (yielding conservative estimates). On the other hand, from the perspective of the firms, these windfalls are a 50–100% increase in funding, which is certainly economically meaningful.

Two other factors may lead us to overly conservative estimates. First, we cannot observe the actual amount of state-based match funding each grant recipient receives due to that data not being systematically collected. This frames our approach as an “intent to treat” analyses. Second, while we assume constant returns to scale in our empirical model (which aggregates output across many firms), it would be very reasonable to assume that each firm has decreasing returns to scale. In this case, the marginal funding from state matches would be less productive than the inframarginal funds from the DoE.

More broadly, few policies are truly random, and in their description of these match programs, Lanahan and Feldman (2015) note that they are associated with a range of factors. They find that, among other things, states with greater budget slack, those with below average high-tech employment, and those with neighboring state adopters of the program are more likely to have a program in place.

Additionally, many of the exceptionally inventive states, including California,

²¹This includes SBIR grants from the Department of Defense, National Institutes of Health, National Science Foundation, and National Aeronautics Space Administration. Firm locations were found via match (with an 80% match rate) to the National Establishment Time Series.

Massachusetts, and New York, do not have matching programs in our sample. This confluence of factors is perhaps why the analyses above do not suggest that states with programs are especially unique in terms of total SBIR funding or small firm movement.²² Still, we acknowledge that our design is based on policies present in a relatively modest subset of the U.S. R&D enterprise as illustrated in Figure 1. And given the highly skewed nature of innovative inputs (much of which are flowing into states that do not have match policies), this may lead us to underestimate the extent to which a matching program in these states might generate spillovers.

C. Incorporating Spillovers into the Empirical Model

This section describes how we modify our baseline model to allow for the possibility of R&D spillovers across geographic and technological space. In overview, first we allow for spillovers across geographies by estimating multiple regressions that include the patent output from regions where it is more and more costly to travel to grant recipients following Feyrer, Mansur and Sacerdote (2017). Second, we allow for spillovers over technologies (within each geographic regression) by incorporating the results of our natural language processing of the FOA topics with methods for estimating the treatment effect spillovers following Clarke (2017) and James and Smith (2020). We discuss these approaches in turn.

For a spillover to occur, information needs to be transmitted. Geographic distances have led to sizable costs of this transmission. So, it is no surprise that a large line of research, spurred in large part by Jaffe, Trajtenberg and Henderson (1993), has documented evidence consistent with geographic distances being a constraint of R&D spillovers.

For the purposes of capturing geographic spillovers across firms and inventors, we take the straightforward approach of estimating separate regressions that increase the size of concentric “circles” that successively include patents from firms and inventors further and further away from DoE SBIR grant recipients in the dependent variable (Feyrer, Mansur and Sacerdote 2017).²³ The first, exclusive regression corresponds to counting only patents awarded to firms that receive DoE SBIR grants, which effectively replicates a traditional firm-level analysis of the program (e.g., Howell 2017). For the most inclusive regression, we count all USPTO patents, regardless of where the firm or inventor is located.

²²Anecdotally, discussions with State Match administrators indicated that firms were unaware of the policy in their state until the firms received notification from the state that they would be receiving additional matching funds.

²³James and Smith (2020) raise an important issue with this approach that arises when the dimension on which these concentric circles are drawn is correlated with variation in the independent variable. In Feyrer, Mansur and Sacerdote (2017), the units of analysis and the dimension of spillovers are both geographic, and their independent variable is not random across geographies. However, the unit of analysis in our study is technologies, which combined with the fact that we are leveraging the state match policies for identification, leaves us unconcerned with James and Smith’s (2020) critique. However, their critique is very relevant for our approach to capturing technological spillovers, which is why we employ a version of their solution when tackling spillovers in that dimension.

To explore how spillovers play out in between these extremes, we rely on concentric circles based on weighted travel costs for U.S. activity. This offers a more precise construct to assess geographic spillovers given that the costs of human travel, rather than geodesic distances, constrain the flow of ideas (Agrawal, Galasso and Oettl 2017; Catalini, Fons-Rosen and Gaulé 2020). The relevant policy levers here are related to subsidizing (or taxing) travel.

We compute travel costs within the U.S. by focusing on county-to-county pairs as a rich yet computationally tractable set of regions. We estimate the cost of making a round-trip from every county in the U.S. to all counties where an SBIR grant recipient is located. For this cost, we use the minimum cost of either driving directly, or driving to/from the nearest airports and flying based on average fares from the Department of Transportation (BTS 2020) and the distances between regions (??). We detail the computation of this metric in Appendix B.

Motivated by gravity models of trade, we construct our final measure of weighted travel costs by dividing these costs by the product of the total amount of SBIR funding that flows into the destination county and the populations of the origination and destination counties (?). Just as a gravity trade model would predict a large exchange of goods between two large, neighboring countries, so too do we predict a large exchange of ideas between two large countries with cheap travel connections. Appendix B documents this process in further detail.

Returning to our focal knowledge production function, instead of a single dependent variable describing total patent flows Y_{jt} , we examine multiple patent flows denoted by Y_{jt}^d , where the superscript d denotes different thresholds for patents to be included in the dependent variable based on the location of the inventors or assignee firms. Instead of a single regression, we estimate a series of regressions indexed by d :

$$(3) \quad \mathbb{E}[Y_{jt}^d | W_{jt}] = \exp\left(\frac{W_{jt}}{\bar{W}}\theta^d + \tau_t^d\right),$$

where the vector of θ^d estimates describes how the output elasticity changes as the base of firms and inventors whose patents are included in the dependent variable count is expanded. We estimate versions of this where we include USPTO patents from: only grant recipients, all firms and inventors in the same counties as grant recipients, all US-based firms and inventors within each decile of the weighted travel cost metric, and lastly, all firms and inventors worldwide.

In our setup, identifying technological spillovers amounts to estimating peer effects. Our approach is motivated by Clarke (2017) and James and Smith (2020), who include regressors that indicate the treatment status of nearby units.²⁴ In those papers, “nearby” refers to geographic distances. Here, we use the similarity scores generated from our mapping of FOA topics to CPC groups. Just as travel

²⁴See also the statistical literature for more on estimating causal effects in the presence of interference (e.g., Toulis and Kao 2013).

costs constrain the flow of ideas over geographies, we expect the cumulative nature of R&D to constrain the flow of ideas over technologies.

To incorporate this into our empirical model, first recall that we construct investment flows and corresponding stocks (W_{jt}) that describe the dollar amount of SBIR funds allocated to CPC group j in time t . Even more specifically, we can construct bins of investments flowing into each CPC group as a function of how similar that group is to the FOA topic from which the investment originates. Let $b \in \{1 - 5, 6 - 10, \dots, 96 - 100\}$ index these bins, giving W_{jtb} such that, for instance, $W_{jt,b=6-10}$ describes the stock of investments in group j at time t originating from FOA topics where group j is in the 6th to 10th percentile of technological distance.

Substituting these bins of investments into the main production function, and allowing their effect on patent rates to vary, yields:

$$(4) \quad \mathbb{E}[Y_{jt}^d | W_{jtb}] = \exp\left(\sum_{b \in \mathcal{B}} \frac{W_{jtb}}{\bar{W}} \theta_b^d + \tau_t^d\right),$$

where \mathcal{B} is a set of technological distance bins.

To summarize, Equation 4 represents multiple regressions indexed by the superscript d , one for each concentric set of firms and inventors whose patents are included in Y . Within each regression, we obtain multiple output elasticity estimates indexed by the subscript b , one for each bin of investments. We will use our estimates of θ_b^d from *the same regression*, holding d fixed, to investigate spillovers across technological space, and we will use our estimates of θ_b^d from *different regressions*, holding b fixed, to investigate spillovers across geographic space.

This approach is more flexible (albeit more demanding) than the spillover matrix approach to estimating R&D spillovers, since those methods impose structure on the magnitudes of spillovers across spatial dimensions. Our approach, however, estimates the structure of spillovers from the data.

As noted by Manski (1993), identification here can be obtained only with assumptions about connections of units in the network. In our setting, this amounts to identifying some technological distance beyond which we must assume no spillovers occur. In other words, \mathcal{B} cannot include all technological distance bins. Practically speaking, we must identify some boundary of similarity scores between an FOA topic and CPC group beyond which it is reasonable to assume that investments originating in the FOA topic have no effect on output in that CPC group.

We search for this boundary using the data-driven procedure proposed by Clarke (2017). This frames the issue as an optimal bandwidth problem, and uses cross-validation to identify the boundary of spillovers that best fits the data. Appendix D reports further details and the results of this procedure. In short, the results suggest a boundary of the 60th percentile of technological distance

when focusing on grant recipients, and a boundary of the 40th percentile when focusing on all other groupings of firms and inventors. After making these assumptions, our estimates of θ_b can then be used to infer the marginal product of SBIR funding, when the focal output is patents on technologies that are some distance from an FOA topic described by b .

With estimates from Eq. 4 in hand, we calculate the implied magnitude of spillovers across the two dimensions we study. For example, if we estimate that an additional \$1 million would spur one additional patent from SBIR grant recipients, but the same amount would spur three patents from all firms and inventors (including grant recipients), then we would say that SBIR grant recipients are responsible for only 33% of the net patent output of the program. In this hypothetical case, we would deem geographic spillovers to be twofold, since for every one patent produced by the grant recipients, two more are produced by others.

Still, we need to estimate the marginal product of SBIR investments in a way that reflects the realities of the DoE’s investment choices – they invest in FOA topics, not CPC groups. Thus, we use our elasticity estimates to construct FOA topic-level marginal products (in terms of additional patents) based on the distribution of observed FOA topics and which CPC groups are targeted within. Appendix B.3 details our approach, which also accounts for the flow nature of our dependent variable. We report the average of these topic-level estimates (along with the 5th and 95th percentiles).

III. Main Results on R&D Spillovers

A. Results at Geographic Edge Cases

We begin by reporting the results based on the patent flows from four geographic edge cases: grant recipients, all firms and inventors in the same U.S. counties as grant recipients, and then all US-based and finally all inventors and firms worldwide. We use the empirically derived boundaries of technological spillovers, and start with a simple case where we use a single stock of investments that spans the entire technological distance over which we expect spillovers; that is, in the notation of our main equation (Eq. 4), we use a single bin b in the set \mathcal{B} .²⁵

Table 2 reports the regression estimates of the output elasticities for the four edge cases. Appendix E.1 plots the raw data underlying these estimates, which support our assumption of a constant elasticity. We also report our conversion of these estimates into average marginal products (in terms of patents per \$1 million – the approximate size of a Phase II SBIR grant), which are intended to represent the average return on additional investment in an FOA topic. In all cases, we cluster our standard errors at the CPC group level.²⁶

²⁵The estimating equation for each of the four d is: $\mathbb{E}[Y_{jt}^d | W_{jtb}] = \exp\left(\frac{W_{jtb}}{W} \theta_b^d + \tau_t^d\right)$, where $b = \mathcal{B}$ and is defined as the “technology boundary” in Table 2.

²⁶While we have the universe of data on DoE SBIR grants and can observe all USPTO patents, we are still concerned with design-based uncertainty (Abadie et al. 2020). When a CPC group receives

Table 2—: Results at Geographic Edge Cases

	Grant recipients		Recipients' counties	US-wide	Worldwide	
	(1)	(2)	(3)	(4)	(5)	(6)
Total \$	2.274 (0.196)					
State Match \$		1.014 (0.094)				
Windfall \$			0.134 (0.021)	0.125 (0.016)	0.123 (0.015)	0.130 (0.014)
$\frac{\partial \text{patent}}{\partial \$1\text{M}}$	9.28 [9.0–10.2]	4.14 [4.0–4.5]	0.55 [0.5–0.6]	1.40 [1.2–1.6]	1.73 [1.5–1.9]	2.96 [2.6–3.3]
<i>N</i> obs.	235,406	235,406	235,406	235,384	235,384	235,384
Tech. boundary	<i>p</i> 60	<i>p</i> 60	<i>p</i> 60	<i>p</i> 40	<i>p</i> 40	<i>p</i> 40
Year F.E.	Y	Y	Y	Y	Y	Y

Note: Reports the output elasticity estimates from regressions using the “simple” model that aggregates all technological spillovers into a single bin. Standard errors clustered at the CPC group level are reported in parentheses. $\frac{\partial \text{patent}}{\partial \$1\text{M}}$ reports the additional patents expected from a \$1 million dollar investment, averaging across all Funding Opportunity Announcements, with the fifth and ninety-fifth percentiles reported in brackets. Tech. boundary indicates the boundary of technological spillovers in terms of the percentile of technological distance between the CPC group and the Funding Opportunity Announcement topic.

Focusing first just on the patents produced by DoE SBIR grant recipients, columns 1–3 of Table 2 illustrate the different estimates we obtain when focusing on total funding (federal plus state match), just the match funding, and then just the windfall portion of the match funding. We estimate an elasticity of roughly 2.3 when including total SBIR funding in the model (col. 1), but hypothesize that a significant portion of this variation is driven by the selection bias discussed previously. Focusing just on the state match funding yields an elasticity close to 1.0, which corresponds to 4 patents per \$1 million (col. 2). But we are still concerned with the possibility of selection bias here, since match funding, on average, is a direct effect of federal funding.

Column 3 reports the results from our preferred specification. Here, we focus just on the windfall funding generated by the match programs, which yields an elasticity estimate of 0.134. This magnitude corresponds to about 0.5 patents per \$1 million. Acknowledging the caveats of our focus on the match-based windfall funding, it appears that not accounting for the endogenous allocation of funding overstates the productivity of these grants by an order of magnitude.

Columns 4–6 of Table 2 shift the focus geographically outward from grant recipients. We continue to estimate average elasticities between 0.12 and 0.13. But because the base of patent flows grows quickly with each successive grouping (while the elasticity does not decline), we estimate much larger marginal products on a patent-per-dollar basis: 1.4, 1.7, and nearly 3 patents per \$1 million for the

additional investment in a given year, that flow of variation persists in the stock of investments in that CPC group for all future years.

same-county, US-wide, and worldwide sets of inventors and firms, respectively. These magnitudes suggest that grant recipients might be responsible for only one sixth of the net patent output created by these grants – the R&D spillovers in this setting could be as large as five-fold. In the next section, we relax our assumption of a single bin of investments (by introducing multiple bins b in the set \mathcal{B}) to more flexibly estimate spillovers over technology-space.

B. Net Spillovers Across Geographies and Technologies

The results in Table 2 suggest the potential for geographic spillovers is large. But how closely do these new patents align with the DoE’s objectives as expressed in the FOA topics? And how concentrated are those geographic spillovers within the US? Here we report the results from our full and most flexible model described by Eq. 4.

Figure 2 reports the elasticity estimates based on Eq. 4 using the same four geographic edge cases as before. Each panel is based on a separate regression, and we plot the estimates of $\theta^{d,b}$, where d indicates the geographic set of patents used in the regression and b indicates the technological distance bin. As before, and still based on our data-driven search for the boundary of technological spillovers, we estimate b -specific elasticities up to the 60th percentile of technology distance when focusing only on patents from grant recipients and up to the 40th percentile for all other cases.

In each case, we observe a clear and intuitive trend: the productivity of SBIR funding is the highest within the CPC groups that are closest in technology space to the objectives of the FOA topic through which that funding flows. The elasticities decrease roughly monotonically until we estimate a series of null effects. This pattern of estimating relatively precise zero elasticities at the same point where the data-driven choice for the threshold of spillovers gives us further confidence in our approach.

We use the elasticities reported in Figure 2 (as well as further unreported regressions based on different sets of within-US firms and inventors) to generate our full set of marginal product estimates. Figure 3 plots these estimates, focusing on spillovers across just technological space (Fig. 3a), across just geographic space within the U.S. (Fig. 3b), and within the U.S. and abroad (Fig. 3c).

Focusing first on technological space, Figure 3a illustrates how spillovers along this dimension are not exactly monotonic. Up to the 25th percentile of technological distance, we estimate roughly equal output per dollar.²⁷

Figures 3b and 3c help illustrate spillovers across geographic space. Within the U.S., grant recipients themselves are responsible for nearly half of all US-based

²⁷This pattern does not align with the more monotonic elasticity estimates shown in Figure 2 because the CPC groups that receive the largest amount of “high match” funding (as in, lower technology distance percentiles) also tend to have patent flow rates that are below average. Although we cannot investigate this result further with our research design, this is consistent with the notion that the government is directing funding towards technologies where it is more difficult to patent, which is in line with the theoretical motivation for these subsidies.

patents. The estimates suggest that grant recipients generate roughly 0.75 patents for every \$1 million they receive. In Appendix E.2, we walk through some simple math that indicates this estimate is well within a reasonable range one would expect based on Howell’s (2017) sharp regression discontinuity evaluation of this program – our confidence intervals do not reject Howell’s (2017) results.

The remainder of US-based patents are almost entirely due to firms and inventors in the closest 10% of U.S. counties (per travel costs), with output dropping essentially to zero for the vast majority of counties. This concentration is consistent with prior work on the geographic distribution of inventive activity in the U.S. (Forman, Goldfarb and Greenstein 2014).

Comparing domestic and foreign patenters, Figure 3c indicates that US-based firms and inventors are responsible for about 60% of net patent output, with the remaining 40% due to foreign firms and inventors. We are unaware of previous attempts to directly estimate this sort of domestic-foreign split of the returns to public R&D investments (in terms of patent production), which indicates that international R&D spillovers can be large.

Figure 3d plots the joint distribution of output across both geographic and technological dimensions. Within the U.S., grant recipients are responsible for most of the patents that are technologically the closest to the DoE’s objectives (1st–10th percentile of technology distance), with nearby firms and inventors responsible for the vast remainder of the within-US spillovers. Foreign firms and inventors tend to focus on relatively close technologies, although we find meaningful effects out to the 25th percentile of technological distance.

Table 3 summarizes all of these results succinctly. We report the average marginal products and costs and decompose the net patent output of these grants along both the geographic and technological dimensions. We find that grant recipients are responsible for only about one quarter of the net patent output, which indicates geographic spillovers are on the scale of $3\times$ – for every patent they produce, we expect three more from other firms and inventors around the world. Similarly, we find that the nearly two-thirds of net patent output is related to technologies that are not very similar to the original objectives of the DoE’s FOA topics – technological spillovers are on the scale of twofold.

One notable difference between geographic and technological spillovers is that it appears spillovers permeate geographic space much further than technological space. On the geographic dimension, we find that countries all around the world benefit from these grants (a result we explore further in Appendix G). However on the technological dimension, it appears that spillovers permeate only about 40% of technology-space. In other words, it appears the costs of taking an idea around the world are much lower than the costs of using an idea related to one technology to advance a completely different technology. This is consistent with related work that suggests the adjustment costs of changing the direction of research are large (Myers 2020). We further discuss the implications of these magnitudes and patterns in Section V.

Table 3—: Summary of Outputs and Costs

	% of net patents	patents \$1M	\$ patent
Counting all USPTO patents and...			
...only grant recipients	26%	0.75	\$ 1,330,000
...only non-recipient firms & inventors nearby recipients	20%	0.59	\$ 1,684,000
...only remainder of US non-recipients	14%	0.40	\$ 2,476,000
...all US firms & inventors	59%	1.75	\$ 572,000
...all foreign firms & inventors	41%	1.19	\$ 839,000
Counting all firms & inventors, only USPTO patents that are...			
...very similar to grants' tech. objectives	37%	1.10	\$ 909,000
...somewhat similar to grants' tech. objectives	40%	1.17	\$ 853,000
...least similar to grants' tech. objectives	23%	0.67	\$ 1,496,000
Counting all USPTO patents, all firms & inventors	100%	2.94	\$ 340,000

Note: Reports average marginal products and costs when focusing on a particular set of patents or firms and inventors. The bottom row defines output and costs when all patents are considered, so “% of net patents” is 100% by construction. “patents / \$1M” reports the net number of patents expected from a marginal investment (awarded only to grant recipients) of \$1 million. “\$ / patent” reports the marginal cost expected to produce one additional patent.

C. Robustness and Alternative Specifications

We estimate a number of additional regressions to explore (1) the robustness of our key identification assumptions and (2) the sensitivity of our results to any of the choices made in our data construction. Below, we highlight some of these findings, with the full set of results included in Appendix E. Overall, we obtain very similar results across the range of reasonable alternatives we explore.

Our identification strategy relies on an assumption that each CPC group’s exposure to windfall funding from the state-based match programs is orthogonal to any underlying shocks in the supply of or demand for the technologies related to that group. We discuss this assumption and evidence supporting it in detail in Section II.B. But as a robustness test, we also estimate two models that include a series of time-varying fixed effects at aggregated levels of the CPC scheme. These fixed effects remove time-varying variation in patent flows and SBIR funding that is common to aggregations of CPC groups, some of which may be due to (unobservable) supply or demand shocks.

Appendix E.3 reports and discusses the results of these regressions in detail. Overall, we obtain relatively similar, though sometimes more conservative, estimates when including these controls for technology-specific time trends. Almost all of the estimates are within the confidence intervals of our preferred specification. There is some attenuation, but some of this is to be expected given that these sorts of fixed effects tend to bias production function estimates towards zero (Griliches and Mairesse 1995).

Appendix E.4 reports results from seven different regressions that focus on al-

ternative specification choices either in the data construction process (e.g., related to the similarity score calculations or the technological distance threshold) or in the regression model itself (e.g., whether zero-valued observations are included, the range of sample years, or the use of a negative binomial instead of a Poisson model). All of these alternative specification choices yield very similar results, which suggests that no single decision in our empirical approach is driving our main estimates.

We also investigate our assumptions about the discount rate on the stock of funding, which is relevant for getting a sense of any production lags. As shown in Appendix E.4, we estimate a similar pattern of point estimates when we introduce non-zero discount rates, with larger discounting depressing the estimates towards zero. The magnitude of this compression towards zero is suggestive of production lags on the scale of five to seven years. We interpret these results cautiously, since the stock-and-flow empirical model we use is not well suited to investigate these lags in depth.

D. Additional Results

Here, we compare our approach to the traditional front-page citation approach of capturing spillovers.²⁸ To identify what share of spillovers identified by our approach are reflected in front-page citations, we first estimate the model using only patents connected via citations to patents of grant recipients. Then we compare the net output expected per those models to our preferred model that does not require any citation linkage. We explore two approaches to making the citation linkages. First, we take the popular approach of using what we term “one-degree” citations, as in direct citations, to grant recipients’ patents as links. Second, we use what we term an “all-degree” approach and include any patent that can be connected by any degrees of citations to a grant recipient’s patent.

Appendix E.5 reports the full results of these tests. We find that requiring citation linkages results in significantly lower spillover estimates. The one-degree approach yields spillover estimates that are 55% smaller than our preferred model, and the all-degree approach yields estimates that are 43% smaller.

This finding suggests that not much more than half of the R&D spillovers we identify are reflected in front-page citation trails. Despite the many aforementioned critiques of front-page citations, this result is, to our knowledge, the first to quantify how much may be “missed” when forced to rely on paper trails. Given the pervasive use of these citation linkages as proxies for R&D spillovers, we hope these results highlight the value of our alternative approach. However, one large upside of relying on citation data is the ability to investigate the specific lag struc-

²⁸Although survey evidence suggests citations can reflect meaningful knowledge flows (Jaffe, Trajtenberg and Fogarty 2000), econometric analyses suggest a significant portion of citations are strategic and/or are in response to unobservable shocks (Alcácer, Gittelman and Sampat 2009; Arora, Belenzon and Lee 2018). See Bryan, Ozcan and Sampat (2020) for recent work exploiting the, likely more informative, “in-text” citations as evidence of spillovers.

tures of knowledge production, which can be long and complex (Azoulay et al. 2019).

Our main specification puts equal weight on individual inventors and firms (assignees) in terms of where patents are produced in geographic space. At the boundaries of our geographical sets – focusing just on patents from grant recipients and focusing on all USPTO patents – this distinction is irrelevant. However, for the intermediate cases, attributing patents entirely to inventors or to firms can shed some light on which entity appears more important for mediating spillovers over geographic space. An important caveat is that inventor and firm locations are fairly highly correlated.²⁹

Appendix E.5 reports the results from these alternative approaches, and in general we do not estimate substantially different results whether we assign locations entirely to inventors or firms. This suggests there may not be dramatic differences in their relative importance as conduits of spillovers. There is some evidence that firms may be more important for facilitating geographic spillovers over very short distances and that individual inventors may be more important for facilitating international spillovers.³⁰ But again, we detect no major differences.

In Appendix G, we report the results of a purely descriptive search for the correlates of R&D spillovers. We cannot make any causal statements, but our setting and data provide a unique opportunity to (1) estimate how specific regions (e.g., U.S. counties, foreign countries) are more or less likely to benefit from spillovers, and then (2) regress these estimates on features of each region to explore what is more or less correlated with the level of spillovers into that region.

In addition to features motivated by prior work, we focus on the degree to which a region appears to exploit knowledge produced by SBIR firms and focus on technologically similar inventions versus using that knowledge to explore technological space and focus on technologically distant inventions. We term this feature simply the “exploit-explore index” of a region. The notions of exploitation and exploration are pervasive in innovation economics and typically posed as a tradeoff (e.g., Manso 2011). But the importance of these alternative strategies, and whether such a tradeoff exists, has not received much attention in the context of R&D spillovers.

The results indicate that the size of spillovers across regions varies closely with features that reflect the supply of and demand for energy technologies in those regions. Interestingly, few features consistently explain more variation in a region’s ability to capitalize on spillovers than their exploit-explore index. Even after conditioning on large vectors of other relevant controls, regions that appear more willing to focus on the same technologies that the SBIR firms pursued are more likely to ultimately create more patents. In the case of domestic spillovers,

²⁹83% of all inventor-assignee pairs are from the same country. Amongst pairs where one is located in the U.S., 50% are located in the same state and 31% are located in the same county.

³⁰The importance of individuals for international spillovers is consistent with Griffith, Harrison and Van Reenen (2006), who find that UK firms with large inventor bases in the U.S. appear to differentially benefit from spillovers.

the exploit-explore index is the fourth most important feature (out of 50 possible), with only industry clusters in IT, production technologies, and oil and gas appearing more important. At the international level, the index is found to be the most important feature (out of 83 possible). Altogether, the results suggest a need for continued work on why different groups of innovators are more likely to exploit or explore technology space, especially after they learn new knowledge from others.

IV. Private Versus Social Value

An important question remaining is what share of these spillovers is due to externalities – what is the wedge between the private and social returns to R&D? We cannot answer this question in general, but we can answer a closely related question: how much of the total patent-based value stimulated by these R&D grants is ultimately captured by the grant recipients themselves?

Our analyses thus far have implicitly treated all of the patents spurred by these grants as equally valuable to the different sets of firms and inventors that obtain them. That is, one could assume that the marginal patents produced by grant recipients are equally as valuable as the spillover-based patents produced by non-grant recipients, on average. If this assumption is true, then it implies that the share of patent-based value captured by any group of firms or inventors is simply equal to their share of the net patent output produced by these grants. As outlined in Table 3, we estimate grant recipients are responsible for about 25% of the net patent output. This would imply they capture roughly 25% of the patent-based value their own R&D ultimately generates.

To relax this assumption about the average value of all new patents being equal, we turn our focus to evaluating whether the SBIR grants have any effect on the quality of new patents. To do so, we estimate a series of regressions using the main specification, instead focusing on the number of forward citations per patent as the dependent variable. While we noted in detail above the drawbacks of using patent citations to proxy for spillovers, they are currently still the most accessible proxy for value. Many studies have shown that forward citations are strongly correlated with the private value of patents to firms (e.g., Harhoff et al. 1999; Hall, Jaffe and Trajtenberg 2005; Kogan et al. 2017). Thus, this regression allows us to investigate whether there is any important difference in the average citation-proxied value of patents spurred by SBIR funding.

Appendix F reports the results of these regressions, which show a clear pattern that the marginal patents produced by grant recipients (due to the SBIR funding) are cited much more often than average. This result alone is interesting in that it suggests the SBIR funding is not merely enabling these firms to patent an idea of marginally lower quality, but rather it is enabling R&D that is increasing the quality of the firms' inventions in terms of their private value, but also (to the extent these citations are reflecting externalities) their social value.

When we expand the set of firms and inventors in the regression though, there

are only very minor increases in the average citation-based quality of patents. In other words, the effect of SBIR funding on the quality of patents appears important only for grant recipients themselves, with the spillover-based patents appearing to be very close to average quality.

This pattern suggests that our initial estimate – that grant recipients capture 25% of the net patent value they generate – is likely an underestimate. How much of an underestimate might this be? This depends on many factors, but for simplicity and given our data constraints, we focus on one key factor: the relative value associated with a marginal patent versus the value associated with a marginal citation to a patent. That is, we assume that the only difference in expected value across patents is based on the number of forward citations each patent receives. With this assumption, if we know how valuable a marginal patent is compared to a marginal citation, we can estimate the share of patent-based value each group of firms and inventors captures, while incorporating this quality-margin effect.

Figure 4 plots the distribution of patent-based value captured by the grant recipients over a range of assumptions about the relative value of citations versus patents. This range spans from zero to one third and is based on estimates from prior work to value patents and citations (Harhoff et al. 1999; Hall, Jaffe and Trajtenberg 2005; Kogan et al. 2017). Obviously, if no value is assigned to citations, we revert back to the quantity-only based 25% estimate. Using the upper end of the relative value of citations suggests that grant recipients may be capturing upwards of 50% of the net patent-based value stimulated by these grants.³¹

Figure 4 also plots the share of patent value captured by all US-based firms and inventors. Our range of estimates spans from about 60% (when all patents are treated equal on average) up to roughly 75%. This suggests that international spillovers might allow foreign countries to capture anywhere from 25–40% of the patent-based value generated by the DoE’s SBIR grants. To our knowledge, this is one of the first direct estimates of the domestic share of returns to public R&D investment. Macroeconomists have used equilibrium models and aggregate data to estimate how much foreign countries benefit from domestic technological advances (Eaton and Kortum 2002; Coe, Helpman and Hoffmaister 2009). We cannot directly map our estimates to theirs for comparison, but continued work that can link these micro- and macroeconomic estimates is worthwhile.³²

If we take our estimates of patent value capture as reflecting the wedge between the private and social returns to R&D, our results suggest that the marginal social

³¹In Appendix F.2, we take an alternative approach and estimate the effect of funding on both the quantity and quality margins at the same time using citation-weighted patent counts as dependent variables. As we discuss there, these approaches put significantly more weight on citations than patents, which yields a larger range of estimates of the private value captured by the grant recipients (between 38–75%).

³²For example, Eaton and Kortum (2002) estimate that when the state of US-based technology increases, the relative magnitude of welfare gains for each major foreign country (due to spillovers) is often between 10% to 20% of the US’s welfare gains.

returns to R&D in this setting are anywhere from about 100–300% of the marginal private returns (since grant recipients only obtain somewhere between half and one-third of the patents their R&D ultimately spurs). While seemingly large, this order of magnitude is very much in line with the few firm-level studies that have focused on large public firms (Bloom, Schankerman and Van Reenen 2013; Zacchia 2020) and macroeconomic studies (Jones and Williams 1998) that have estimated this same wedge.

We would emphasize that this conclusion should be interpreted with caution as it is based on a number of large assumptions that generally would imply we are considering an upper bound of this wedge.³³ Continued empirical work to more precisely estimate this wedge will certainly be important.

V. Discussion

The presence of R&D spillovers across geographic and technological space are a key motivation for government intervention in markets for innovation. Here, we provide another set of results in the still very short line of causal estimates of the magnitude these spillovers (Bloom, Schankerman and Van Reenen 2013; Azoulay et al. 2019; Zacchia 2020). In a new setting with a new methodology, we again find that R&D spillovers are very large. The firms that receive SBIR funding from the U.S. Department of Energy obtain only 25% of the net patents that their R&D ultimately stimulates. For every one patent they produce, three more patents are produced by others who benefited from their R&D. Even after accounting for marginal differences in patent quality, our estimates suggest that, at most, SBIR firms may appropriate only half of the patent-based value their R&D creates. We provide another piece of evidence that the marginal social returns to R&D may be orders of magnitudes larger than the marginal private returns.

Our focus on smaller firms is important since many governments worldwide have R&D policies that target subsidies to small firms. Our focus on public subsidies in the energy sector suggests that this increasingly important sector is fertile ground for the diffusion of new inventions. And while our findings provide another important data point suggesting a large wedge between the private and social returns to R&D, continued work that disentangles externalities and spillovers is essential since it has direct implications for the optimal level of R&D subsidies.

Our results also have implications for program evaluations of R&D subsidies. Recent work has continued to improve the methodologies used to study public R&D investments (e.g., Einiö 2014; Howell 2017), but rarely are spillovers accounted for. In our setting, even if we ignore international spillovers and focus only on patents from US-based firms and inventors, we would have missed more

³³This includes (at least) the following assumptions: (1) product market rivalry spillovers after the patenting stage are negligible; (2) the marginal private costs of producing these patents is equal for all firms and inventors; and (3) none of the spillovers are due to market transactions (e.g., via patent licenses). On last point, Arque-Castells and Spulber (2019) show that ignoring markets for ideas can lead to overestimating the social returns to R&D by a significant amount.

than half of the net patent output from the SBIR program.

On the methodological front, our use of text analyses to connect the units of inputs (the funding opportunities) to outputs (patent groups) could be applied in a variety of other contexts. In our case, it frees us from relying on any citation linkages between patents, which much of the related literature to date has been forced to do. Our results suggest that relying only on these paper trails might overlook a substantial amount of patent-based R&D spillovers. However, our abandoning of the paper trail does limit our ability to investigate the lag between R&D investment and outputs which has been shown to be long and complex (Azoulay et al. 2019). Outside of our context, it is commonplace that the inputs and outputs of a production function are not indexed at the same unit of analysis, but our approach could be used to generate a useful data frame that directly incorporates spillovers across the units of observation.

As noted throughout the text, there are a number of important caveats to our research design and the interpretation of our results. An important limitation is that our focus on the state match policies mechanically narrowed our focus to firms and inventors in a subset of the U.S. that excludes many of the more historically productive regions (e.g., Silicon Valley, Seattle, New York City), where the effect of these grants may be larger. Likewise, since the match policies are awarding marginal funds on top of the federal grants, our estimates will be lower than the causal effect of the federal grants if there are decreasing returns to scale at the firm level.

Another tradeoff we make in focusing on these match policies is that we are confident that the estimates do not reflect selection effects in terms of firms applying to the SBIR program. This helps with the internal validity, but likely narrows our external validity. The firm location analysis described in Appendix C.2, along with discussions with managers of the state match programs, indicate that it is highly unlikely that firms are making any decisions based on the presence of the match policies. However, if the entire SBIR program were to announce that all grants would be larger, we expect this could lead to marked changes in the composition and number of applicants. How this might ultimately affect the productivity of these public R&D investments would be a very interesting counterfactual to pursue.

Ultimately, we do not take a normative stance on which patents produced by this program should or should not “count.” Instead we highlight how the costs of spurring a patent vary widely depending on how seriously spillovers are appreciated. In our setting, counting all patents as relevant would lead to a large overestimate of the program’s effectiveness if specific technologies are of interest. Conversely, focusing only on grant recipients would lead to a large underestimate of the program’s effectiveness if the output of other firms and inventors is of interest.

Overall, our results offer new evidence on how the visible hand of the government can influence the rate and direction of R&D. We have shown that the

output of publicly funded R&D can vary widely on the extent to which spillovers across technological and geographic space are appreciated. We hope future work continues to investigate the magnitudes and determinants of these spillovers to better inform both economic theories and innovation policies.

REFERENCES

- Abadie, Alberto, Susan Athey, Guido Imbens, and Jeffrey Wooldridge.** 2020. "Sampling-Based versus Design-Based Uncertainty in Regression Analysis." *Econometrica*, 88(1): 265–296.
- Agrawal, Ajay, Alberto Galasso, and Alexander Oettl.** 2017. "Roads and innovation." *Review of Economics and Statistics*, 99(3): 417–434.
- Alcácer, Juan, Michelle Gittelman, and Bhaven Sampat.** 2009. "Applicant and examiner citations in US patents: An overview and analysis." *Research Policy*, 38(2): 415–427.
- Arora, Ashish, Sharon Belenzon, and Honggi Lee.** 2018. "Reversed citations and the localization of knowledge spillovers." *Journal of Economic Geography*, 18(3): 495–521.
- Arora, Ashish, Sharon Belenzon, and Lia Sheer.** 2021. "Knowledge spillovers and corporate investment in scientific research." *American Economic Review*, 111(3): 871–98.
- Arque-Castells, Pere, and Daniel Spulber.** 2019. "Measuring the Private and Social Returns to R&D: Unintended Spillovers versus Technology Markets." *Mimeo*.
- Arrow, Kenneth.** 1962. "The Rate and Direction of Inventive Activity: Economic and Social Factors." Chapter Economic Welfare and the Allocation of Resources for Invention, 609–626. Princeton University Press.
- Audretsch, David.** 2003. "Standing on the shoulders of midgets: The US Small Business Innovation Research program (SBIR)." *Small Business Economics*, 20(2): 129–135.
- Azoulay, Pierre, Joshua Graff Zivin, Danielle Li, and Bhaven Sampat.** 2019. "Public R&D investments and private-sector patenting: Evidence from NIH funding rules." *Review of Economic Studies*, 86(1): 117–152.
- Bloom, Nicholas, John Van Reenen, and Heidi Williams.** 2019. "A toolkit of policies to promote innovation." *Journal of Economic Perspectives*, 33(3): 163–84.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen.** 2013. "Identifying technology spillovers and product market rivalry." *Econometrica*, 81(4): 1347–1393.

- BLS.** 2019. *Consumer Price Index for All Urban Consumers*.
- Bryan, Kevin, and Jorge Lemus.** 2017. “The direction of innovation.” *Journal of Economic Theory*, 172: 247–272.
- Bryan, Kevin, Yasin Ozcan, and Bhaven Sampat.** 2020. “In-Text Patent Citations: A User’s Guide.” *Research Policy*, 49(4): 103946.
- BTS.** 2020. *Airport fares and distances*.
- Catalini, Christian, Christian Fons-Rosen, and Patrick Gaulé.** 2020. “How do travel costs shape collaboration?” *Management Science*, 66(8): 3340–3360.
- Census.** 2010. *2010 Census of Population and Housing*.
- Clarke, Damian.** 2017. “Estimating Difference-in-Differences in the Presence of Spillovers.” *Mimeo*.
- Coe, David, Elhanan Helpman, and Alexander Hoffmaister.** 2009. “International R&D spillovers and institutions.” *European Economic Review*, 53(7): 723–741.
- Eaton, Jonathan, and Samuel Kortum.** 2002. “Technology, geography, and trade.” *Econometrica*, 70(5): 1741–1779.
- Einiö, Elias.** 2014. “R&D subsidies and company performance: Evidence from geographic variation in government funding based on the ERDF population-density rule.” *Review of Economics and Statistics*, 96(4): 710–728.
- Feyrer, James, Erin Mansur, and Bruce Sacerdote.** 2017. “Geographic dispersion of economic shocks: Evidence from the fracking revolution.” *American Economic Review*, 107(4): 1313–34.
- Forman, Chris, Avi Goldfarb, and Shane Greenstein.** 2014. “Information technology and the distribution of inventive activity.” In *The Changing Frontier: Rethinking Science and Innovation Policy*. 169–196. University of Chicago Press.
- 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.
- Griffith, Rachel, Rupert Harrison, and John Van Reenen.** 2006. “How special is the special relationship? Using the impact of US R&D spillovers on UK firms as a test of technology sourcing.” *American Economic Review*, 96(5): 1859–1875.
- Griliches, Zvi.** 1979. “Issues in assessing the contribution of research and development to productivity growth.” *Bell Journal of Economics*, 92–116.

- Griliches, Zvi.** 1992. “The search for R&D spillovers.” *Scandinavian Journal of Economics*, S29–S47.
- Griliches, Zvi.** 1998. “Patent statistics as economic indicators: A survey.” In *R&D and productivity: The econometric evidence*. 287–343. University of Chicago Press.
- Griliches, Zvi, and Jacques Mairesse.** 1995. “Production functions: The search for identification.” *Mimeo*.
- Hall, Bronwyn H, Adam Jaffe, and Manuel Trajtenberg.** 2005. “Market value and patent citations.” *RAND Journal of Economics*, 36(1): 16–38.
- Haltiwanger, John, Ron Jarmin, and Javier Miranda.** 2013. “Who creates jobs? Small versus large versus young.” *Review of Economics and Statistics*, 95(2): 347–361.
- Harhoff, Dietmar, Francis Narin, Frederic M Scherer, and Katrin Vopel.** 1999. “Citation frequency and the value of patented inventions.” *Review of Economics and statistics*, 81(3): 511–515.
- 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, and J Brown.** 2020. “Do cash windfalls affect wages? Evidence from R&D grants to small firms.” *Mimeo*.
- Igami, Mitsuru, and Jai Subrahmanyam.** 2019. “Patent Statistics as an Innovation Indicator? Evidence from the Hard Disk Drive Industry.” *Japanese Economic Review*, 70(3): 308–330.
- Jaffe, Adam.** 1986. “Technological Opportunity and Spillovers of R&D: Evidence from Firms’ Patents, Profits, and Market Value.” *American Economic Review*, 76(5): 984–1001.
- Jaffe, Adam, Manuel Trajtenberg, and Michael Fogarty.** 2000. “Knowledge spillovers and patent citations: Evidence from a survey of inventors.” *American Economic Review*, 90(2): 215–218.
- Jaffe, Adam, Manuel Trajtenberg, and Rebecca Henderson.** 1993. “Geographic localization of knowledge spillovers as evidenced by patent citations.” *Quarterly Journal of Economics*, 108(3): 577–598.
- James, Alexander, and Brock Smith.** 2020. “Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution: Comment.” *American Economic Review*, 110(6): 1905–1913.

- Jones, Charles, and John Williams.** 1998. “Measuring the social return to R&D.” *Quarterly Journal of Economics*, 113(4): 1119–1135.
- Keller, Wolfgang.** 2004. “International technology diffusion.” *Journal of Economic Literature*, 42(3): 752–782.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman.** 2017. “Technological Innovation, Resource Allocation, and Growth.” *Quarterly Journal of Economics*, 132(2): 665–712.
- 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.
- Lanahan, Lauren, and Maryann Feldman.** 2015. “Multilevel innovation policy mix: A closer look at state policies that augment the federal SBIR program.” *Research Policy*, 44(7): 1387–1402.
- Lanahan, Lauren, and Maryann P Feldman.** 2018. “Approximating exogenous variation in R&D: Evidence from the Kentucky and North Carolina SBIR State Match programs.” *Review of Economics and Statistics*, 100(4): 740–752.
- 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.
- Li, Minghao, Stephan J Goetz, Mark Partridge, and David A Fleming.** 2016. “Location determinants of high-growth firms.” *Entrepreneurship & Regional Development*, 28(1-2): 97–125.
- Link, Albert, and John Scott.** 2010. “Government as entrepreneur: Evaluating the commercialization success of SBIR projects.” *Research Policy*, 39(5): 589–601.
- Manski, Charles.** 1993. “Identification of endogenous social effects: The reflection problem.” *Review of Economic Studies*, 60(3): 531–542.
- Manso, Gustavo.** 2011. “Motivating innovation.” *Journal of Finance*, 66(5): 1823–1860.
- MCDC.** 2014. *Geographic Correspondence Engine*.
- Myers, KM.** 2020. “The Elasticity of Science.” *American Economic Journal: Applied Economics*, 12(4): 103–134.

- Myers, Kyle M, and Lauren Lanahan.** 2020. “Data Assembled and Created for Review of the Small Business Innovation Research and Small Business Technology Transfer Programs at the Department of Energy.” <https://doi.org/10.7910/DVN/YVSV5P>, 1.
- National Academies of Sciences, Engineering, and Medicine.** 2020. *Review of the SBIR and STTR Programs at the Department of Energy*. National Academies Press.
- NCI.** 2019. *nci2019seer*.
- Nelson, Richard.** 1959. “The simple economics of basic scientific research.” *Journal of Political Economy*, 67(3): 297–306.
- Popp, David.** 2002. “Induced innovation and energy prices.” *American Economic Review*, 92(1): 160–180.
- Stam, Erik.** 2007. “Why butterflies don’t leave: Locational behavior of entrepreneurial firms.” *Economic Geography*, 83(1): 27–50.
- Thompson, Peter, and Melanie Fox-Kean.** 2005. “Patent citations and the geography of knowledge spillovers: A reassessment.” *American Economic Review*, 95(1): 450–460.
- Toulis, Panos, and Edward Kao.** 2013. “Estimation of causal peer influence effects.” *International Conference on Machine Learning*, 1489–1497.
- USPTO.** 2020. *PatentsView*, <https://patentsview.org>.
- 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.
- Zacchia, Paolo.** 2020. “Knowledge spillovers through networks of scientists.” *Review of Economic Studies*, 87(4): 1989–2018.

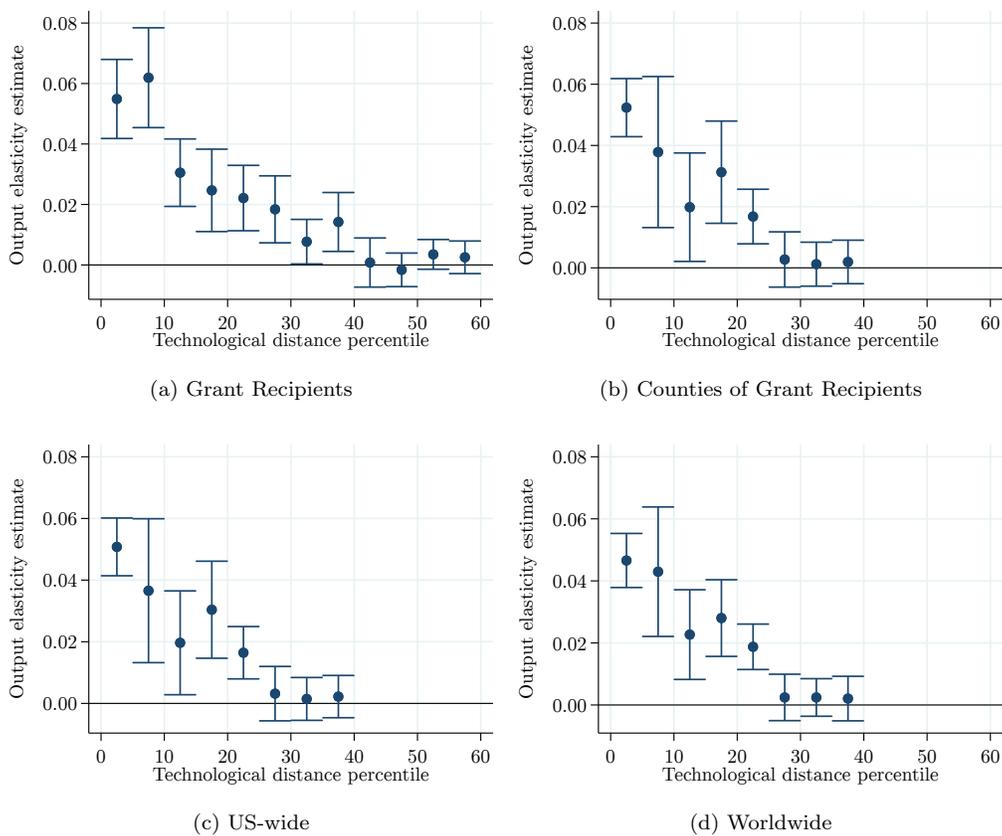


Figure 2. : Patent Output Elasticity Estimates

Note: Plots the estimates of θ_b^d from Eq. 4, where d corresponds to geographies and varies across the panels (with different panels corresponding to different regressions), and b corresponds to technology distances and varies within each panel (with estimates within each panel from the same regression). 95% confidence intervals based on standard errors clustered at the CPC group level are shown with brackets.

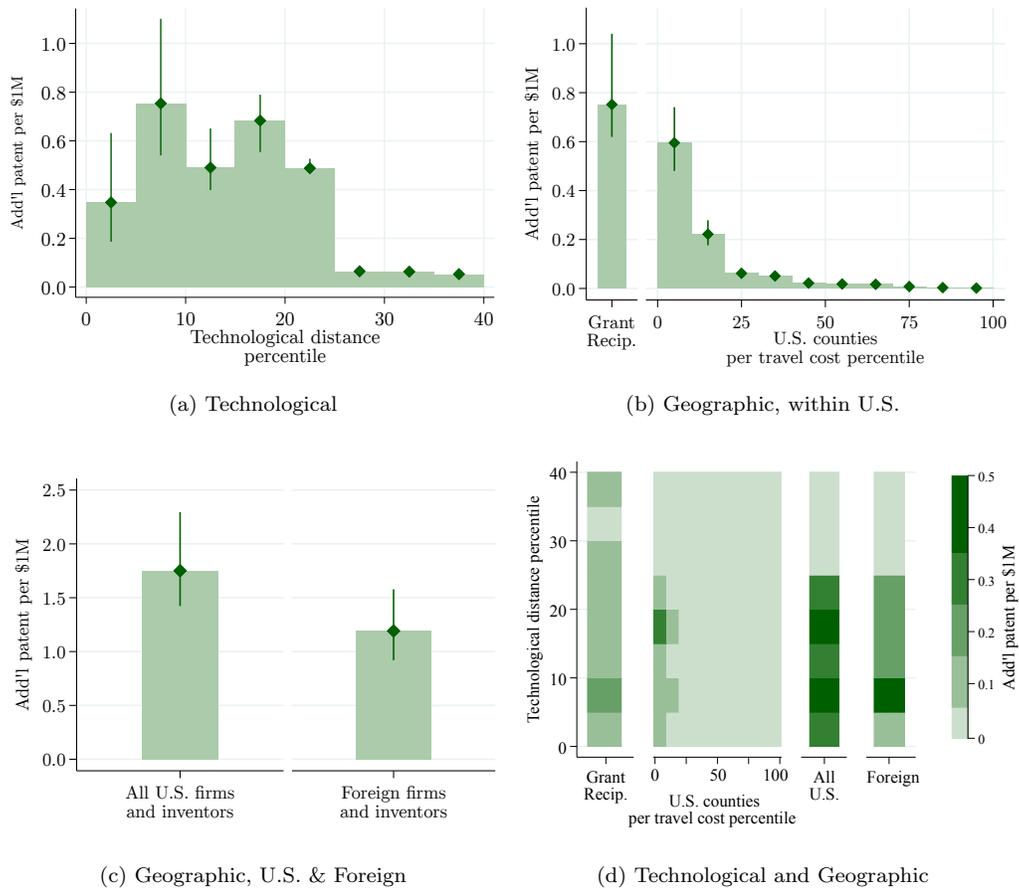


Figure 3. : Marginal Products across Geographic and Technological Dimensions

Note: Plots the the net patents from specific subsets of firms and inventors to be expected from a \$1 million dollar investment, averaging across all Funding Opportunity Announcements, with the fifth and ninety-fifth percentiles of the FOA topic-level statistics shown by the vertical lines.

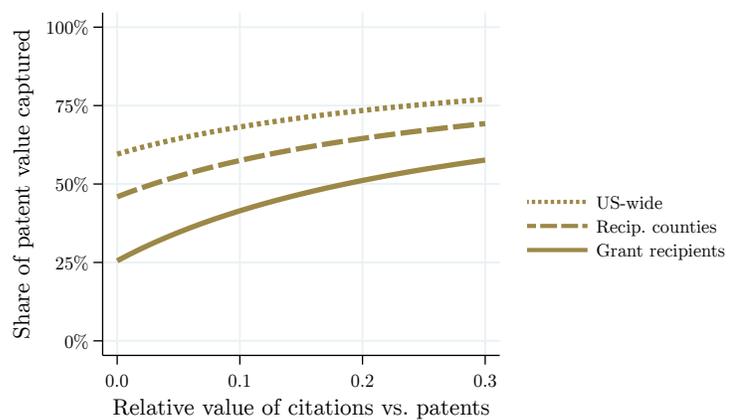


Figure 4. : Share of Net Patent Value Captured by Different Firms and Inventors

Note: Plots the share of net patent value captured by successive aggregations of firms and inventors (y -axis) based on a range of assumptions (based on the existing literature) about the relative average value of an additional patent compared to the average value associated with an additional forward citation to a patent (x -axis).