Research Subsidy Spillovers, Two Ways*

Kyle Myers† Lauren Lanahan‡

Current version available here

March 29, 2020

Abstract

We study how the outputs of research spill over technological and geographic space in the context of the U.S. Small Business Innovation Research program. We infer input-output links using text analyses and identify the marginal costs of producing patents using noncompetitive grant matching policies. Due to technological spillovers, the cost of spurring patents related to specific technologies are much larger than the costs of spurring any kind of patent. Due to geographic spillovers, roughly 80% of the net patents produced by the program are from inventors that do not directly receive grants; the domestic/foreign split of output is about 75/25. The large spillovers across these two dimensions imply that the cost effectiveness of research subsidies can vary widely depending on which outputs count. Within the U.S., we identify regions likely responsible for these spillovers, which reveals a pattern that suggests the government must trade off its ability to influence either the rate or direction of invention.

JEL Codes: O31, O33, O38

*This project was conducted partly within a National Academies of Sciences (NAS) review of the Department of Energy's (DoE) Small Business Innovation Research program. 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. We appreciate helpful comments from Maryann Feldman, Scott Stern, Josh Lerner, Shane Greenstein, Joshua Graff Zivin, Douglas Hanley, Adam Jaffe, Riitta Katila, Iavor Bojinov, Jennifer Kuan, and Daniel Armanios. 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 the data. All errors are our own.

†Harvard University, Harvard Business School (corresponding author), email: kmyers@hbs.edu.

‡University of Oregon, Lundquist College of Business, email: llanahan@uoregon.edu.

1 Introduction

Research subsidies are a key mechanism that governments worldwide rely on to promote innovation. This is a response to the notion that private firms likely underinvest in research relative to the social optimum (Nelson 1959; Arrow 1962; Jones and Williams 1998). Often, targeted contracts and grants are used to direct these investments to particular areas where (one hopes) the wedge between private and public returns is largest (e.g., renewable energy technology). Policymakers also often hope these investments can spur inventions amongst a particular set of firms or individuals (e.g., domestic inventors). Understanding how these subsidies influence inventors – including those not directly receiving subsidies – is necessary for predicting how governments can manipulate the the rate and direction of invention.[1]

Depending on the specific objectives of these subsidies, evaluating their effectiveness may not just be a question of whether they spur *particular firms* to develop *any technology*. It may also be a question of whether they spur *any firm* to develop *particular technologies*. Answering these questions requires an appreciation of "spillovers" – the net effect of externalities, strategic interactions, uncertainty, agglomeration, or any other factor that amplifies or suppresses the impact of these subsidies. Because of spillovers, investments targeted to a particular point in geographic and technological space may lead to inventions elsewhere. Inventors that *do not* receive subsidies may create exactly what the government asked for. And the inventors that *do* receive subsidies may create technologies unrelated to the subsidies' objectives.

Not accounting for geographic spillovers may understate the effectiveness of these investments, while not accounting for technological spillovers may overstate their effectiveness. We do not take a normative stance in this paper. Rather, our goal is to evaluate a popular research subsidy program and provide a map of cost-effectiveness estimates that depend on whether policymakers have preferences over the types of technologies spurred by the program (e.g., due to public demand for a particular invention) or the location of inventors benefitting from the program (e.g., due to geo-political concerns).

[1] For example, macroeconomists often examine the counterfactual effect of increasing one country's "technology stock" on the GDP of other countries (e.g., Coe and Helpman 1995; Eaton and Kortum 2002). In practice, increasing this technology stock is often a goal of research subsidy programs such as the one we study. The U.S. federal government awards firms roughly $13 billion per year via research contracts and grants (excluding tax credits, which are often technology agnostic).

To capture how research spillovers play out, we compare innovative output across technological and geographic spaces that receive different levels of public investment. The program we study attempts to steer small firms toward specific technological priorities and spur inventions, which we proxy for with patents. Since we can classify both the geographic location and technological content of patents, we can estimate how the cost-effectiveness of subsidies varies depending on what inventors and what inventions "count." We take a reduced-form approach that requires minimal assumptions about how investments can lead to inventions. We then develop a systematic method of mapping inputs (R&D grants) and outputs (patents) into technology classes and estimate a technology-level production function. To handle both channels of spillovers econometrically, we use methods that are motivated by theory (Manski 1993; Toulis and Kao 2013) and are employed in the spatial economics literature (Feyrer et al. 2017; James et al. 2019).

We focus on the U.S. Department of Energy's (DoE) Small Business Innovation Research (SBIR) program. This program supports "scientific excellence and technological innovation through the investment of Federal research funds in critical American priorities." Projects related to these priorities – which range from light sensors to fuel technologies to high-performance computing tools – are solicited via Funding Opportunity Announcements (FOAs). We focus entirely on patent output since patents are the most widely used proxy for technological advance and are often cited as a metric for success by program administration. By analyzing the text of the FOAs in conjuncture with the text of patent abstracts, we map investments into patent classes, thereby connecting inputs to outputs. We identify the marginal product of public investments by leveraging plausibly exogenous variation generated by state-specific noncompetitive matching policies. Per these policies, grant recipients in certain states receive additional funds regardless of their project's quality or content (Lanahan and Feldman 2018). We present evidence consistent with our key assumption that the distribution of firms, and the technologies they are pursuing, are unrelated to these policies.

Overall, our results offer new evidence on how the visible hand of the government can influence the rate and direction of research. We highlight how the costs of spurring a patent vary widely depending on how seriously the targeted nature of these grants is appreciated. The implied costs when only counting inventions that closely align with the government's objectives can be upwards of four to seven times larger than

when all inventions are deemed relevant. We also find significant spillovers across geography which suggest that nearly 80% of the net patents generated by this SBIR program come from inventors that do not receive grants. U.S.-based inventors are responsible for roughly 75% of this net patent output. Overall, the cost-effectiveness of research subsidies varies widely, depending on whether the policymaker has strong preferences over what is invented, or who invents.

Lastly, we characterize U.S. counties based on their ability to generate spillovers and whether these spillovers are more explorative or exploitative in the sense of being more or less in line with the technological objectives stated in FOAs, respectively. In terms of fostering spillovers, it appears that travel costs are irrelevant, while private venture capital is very important. Counties with high venture capital penetration appear to be both more productive and more explorative. This suggests that depending on policymakers' objectives, they may need to tradeoff their ability to influence the rate or direction of invention.

Related Literature

This work complements a long line of empirical studies of research spillovers across firms (Jaffe 1986), how they can drive the wedge between the private and social returns to R&D (Bloom et al. 2013), and their importance for long-run growth (Jones 2005a; Aghion and Jaravel 2015). Our appreciation of the direction of innovation stems from the formal theory of Bryan and Lemus (2017), who show that even if the level of R&D in the economy is optimal, the "direction" (i.e., the level within a particular point of technology space) may not be. Related empirical work focusing on the effectiveness of using targeted subsidies to influence the direction of research has identified large directional adjustment costs facing scientists (Myers 2020).

Much of the research spillover literature focuses on how spillovers permeate geographic space (Jaffe et al. 1993; Audretsch and Feldman 1996; Kantor and Whalley 2019), can lead to agglomeration (Ellison et al. 2010), and may be mediated by travel costs in certain locations or at certain times (Agrawal et al. 2017; Kantor and Whalley 2019).[2] Most of the empirical work investigating research spillovers specifically in the sense of one firm's investment on another's output (e.g., Jaffe 1986; Bloom et al.

[2]Also of note are the long-term "environmental exposure" effects whereby children (once adults) are most likely to patent in the same technologies as the older inventors who live in their neighborhoods (Bell et al. 2019).

2013) involves estimating firm-level regressions that require assumptions about how spillovers flow through a network of firms.

An alternative approach that motivates our research design, perhaps first pursued by Popp (2002), is to focus on the network of technologies and estimate technology-level regressions that relate inputs and outputs within some predefined portions of technology space. Focusing on the energy sector, Popp (2002) relates patents to, among other things, federal R&D investments.[3] A similarly structured analysis by Azoulay et al. (2018) evaluates the product of basic research grants from the U.S. National Institutes of Health. They also estimate a technology-level production function, though their "technologies" are grouped around biomedical topics.[4] Their results imply a marginal cost of roughly $3.5 to $7 million per patent, depending on the degree to which across-topic spillovers are accounted for. We contribute to this literature by more deeply exploring how spillovers permeate both geographic and technological space.

Our analysis also contributes to the literature studying the role of public and private R&D in new energy technologies. Building on the aforementioned work by Popp (2002), Johnstone et al. (2010) show differential effects of public policies on energy-related innovation with impacts varying depending on the technological scope of the policy. Aghion et al. (2016) identify significant R&D spillovers that gives rise to path-dependent innovation within clean and dirty technologies. In their characterization of the optimal energy policy, Acemoglu et al. (2016) illustrate the importance of knowing the elasticity of innovation with respect to subsidy-induced R&D (where innovation is also proxied by patenting rates).[5]

An important related study is Howell (2017), which uses the rank-order funding protocol to evaluate a subset of the DoE's SBIR program and finds clear evidence that these grants enable firms to produce new citation-weighted patents. Extrapolating Howell's (2017) estimates suggests the DoE must award a firm an additional $150,000–$1.5 million to spur an additional patent (from that same firm). We use this estimate as a guidepost to support the plausibility of our results. We also replicate and extend

[3] Popp (2002) manually classifies patents into 11 technology classes and uses a panel data regression to estimate an implied cost per patent, in terms of federal R&D, of roughly $20 million.

[4] Economists studying patent policies have adopted units of analyses that reflect the nature of the same technology-level friction, conducting analyses at the level of genes (Williams 2013) and diseases (Budish et al. 2015).

[5] See Figure 16 in Acemoglu et al. (2016, p. 98), which illustrates the significant difference in the optimal paths of R&D subsidies and carbon taxes depending on this elasticity.

some of the results surrounding the complementary role of private venture capital in this setting (Lerner 2000; Gans and Stern 2003; Howell 2017). Furthermore, we build on Lanahan and Feldman (2018), the first to exploit the same noncompetitive state matching policy we leverage for identification. They show that these marginal dollars do push firms closer to commercialization.[6]

In summary, we advance these literatures by explicitly incorporating the stated technological objectives of targeted R&D subsidies into a large-scale economic analysis based on quasi-experimental variation. Motivated by a long line of studies of research spillovers and evaluations of R&D subsidies, we make strides in uncovering how spillovers – both technological and geographic – can amplify the value of public investments in research.

2 Setting and Data Sources

This section describes the data used to connect targeted public investments (inputs) to new inventions that are related to the objectives of those investments (outputs). We rely on multiple publicly available data sets: (1) FOA documents where the DoE outlines a number of "topics" that firms must align their proposals with to be eligible for a grant; (2) data on the DoE's SBIR grants awarded, which contain the FOA topic number and dollar amount corresponding to each award; (3) the U.S. Patent and Trademark Office (USPTO) patent record; (4) a handful of data sets that document travel costs, geographic distances, and other characteristics of U.S. counties; and (5) details of the state-specific SBIR noncompetitive matching programs, which we leverage for identification. We define the time frame of our sample from 2006 to 2017 because of data availability and the rarity of state matching policies pre-2006. For all dollar-based measures, we adjust for inflation using the 2017 Consumer Price Index. We review each data source in turn.

[6] Other studies that examine the SBIR program specifically include: Wallsten (2000); Toole and Czarnitzki (2007); Link and Scott (2010).

2.1 The Department of Energy's Small Business Innovation Research Program

Since its enactment in 1982, the SBIR program has allocated over $40 billion to support early stage innovation by small businesses. It is emblematic of R&D subsidy programs enacted by a growing number of countries worldwide.[7] The solicitation and award process is as follows.

First, one to three times per year, the DoE releases an FOA that outlines technological areas of interest, referred to as topics. To be eligible for funding, applicants must sufficiently align themselves with the objectives of a particular topic in an active FOA.[8] The topic descriptions are roughly five to ten paragraphs of text outlining the specifics of the technology that the DoE is interested in developing. Appendix A contains three sample FOA topics that illustrate their typical breadth and depth.

Each office within the DoE is responsible for producing a set of topics contained in an FOA. Currently, the DoE SBIR program spans 12 offices/research programs within the agency: 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. As apparent from these titles, topics span a broad range of energy-related initiatives including those that may not immediately come to mind when thinking of "energy policy" (e.g., high-performance computing, cybersecurity, and measurement tools).

Once an FOA is released, small, private U.S.-owned firms are eligible to apply for SBIR funding. Their submitted proposals are subject to 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 (typically not exceeding $150,000), which is intended to support development of a proof of concept. Competitive follow-on Phase II grants, which offer two years of support (typically not to exceed $1 million), can then be

[7] It has been claimed that up to 17 countries have copied the structure of the SBIR program (see http://bit.ly/2tyjofQ).

[8] Practically, this alignment between a firm's application and a topic often involves interested firms submitting letters of intent to DoE administrators, who inform firms as to whether or not the planned use of funds is in line with a specific FOA topic.

sought by firms that received Phase I grants.

The process of choosing topics for FOAs and firms to receive awards is by no means 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. These forces likely select on features that are unobservable to us and are correlated with our inputs and outputs. We turn to the issue of handling this endogeneity in Section 4.

We obtained the full set of FOAs (in PDF form) released during our sample through a combination of internet searches and assistance from DoE staff. Award data were obtained from the Small Business Administration's public repository.[9] This reports award-level details such as the firm's name and the dollar amount of the award and, importantly, includes a unique reference to the FOA through which the grant was awarded. We use this identifier to connect SBIR funds to each topic.

2.2 Grant and Patent Record Crosswalk

We follow a long line of economic studies of invention and rely on patent activity as our measure of technological progress. This choice is also motivated by the fact that patents are often used as a proxy for innovation in policy communities, and an in-depth case study supports the assumption that patents are useful metrics of "real" technological progress (Igami and Subrahmanyam 2019). In fact, they are referenced as a primary measure of success by the SBIR program itself.[10]

We source the disambiguated USPTO patent record from PatentsView.[11] These data contain patent information, citation linkages, and disambiguated tables of inventors and firm assignees as well as their inferred locations. For simplicity, unless otherwise noted, we use the terms "inventors" and "firms" interchangeably throughout. This is because in our allocation of patents across geographic space, we give equal weight to the locations of the individual inventors and firm assignees on each patent.[12]

[9] See http://bit.ly/3afvxaF.

[10] See official SBIR program slides at http://bit.ly/30uk5TU, slide 4.

[11] See http://bit.ly/2R1Ysay for descriptions of the methodology and resulting data.

[12] For example, if a patent has one inventor located in a county in Maryland and the firm assignee is located in a county in California, then each of those two counties will receive 0.5 credit for the patent.

Importantly, PatentsView also makes available a set of tables that describe the Cooperative Patent Classification (CPC) scheme, which is a hierarchy of labels assigned to all patents. We use the CPC scheme to make the connection between the SBIR data and the patent data in a series of steps described in Section 3.2 where we further detail the CPC scheme.

To connect SBIR grant recipients to their patents, we use an approach that follows closely with the methods used by the National Bureau of Economic Research Patent Data Project.[13] In short, firm names are standardized and are used alongside their inferred locations to perform a "fuzzy match" between the SBIR award data and the disambiguated PatentsView data. Including all firms in the PatentsView data and the five largest federal SBIR agencies, roughly 30% of firms that ever receive an SBIR award between 2006–2017 ever obtain a patent (before or after the SBIR award). Roughly 5% of firms ever assigned a patent receive an SBIR award at some time.

2.3 U.S. States' Grant Matching Policies

Many U.S. states have created policies to complement the SBIR program, some as early as the 1980s. These policies include noncompetitive and some competitive award matching programs for both Phase I and Phase II recipients, in addition to some outreach services for proposal support. Lanahan and Feldman (2015) first documented these programs, providing an overview as to their motivations and interactions with the federal SBIR program. Following their methodology and extending their data, we construct a data set of the SBIR state-based matching policies for all states from 2006 to 2017.[14]

Figure 1 documents the growth of these policies over time, showing a growth from 6 states in 2006 to 16 in 2017. While most states maintain the program after initial adoption, we observe some volatility (i.e., North Carolina and New Jersey).[15] The size of these match awards range from \$25,000–\$105,000 for Phase I awards and \$50,000–\$500,000 for Phase II awards.

Ideally, we could observe how SBIR funds are used by firms. Unfortunately, access to any relevant federal data and most state data is infeasible. Fortunately,

[13] See http://bit.ly/2TxClug. Further details on our crosswalk are available upon request.

[14] We exclude policies that involve any competitive elements because one of the key assumptions we make when using these policies to generate an instrumental variable is that (some of) the funds allocated by these policies will not be correlated with firm quality.

[15] Anecdotal evidence suggests these switches are driven by fluctuations in state budgets.

Figure 1: State Matching Programs

(a) 2006

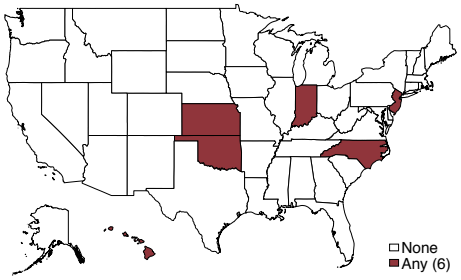

(b) 2012

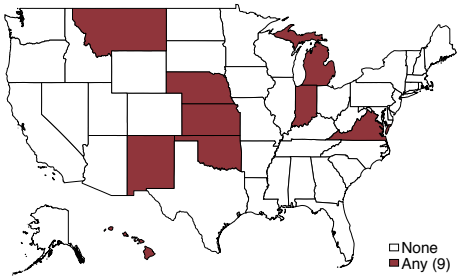

(c) 2017

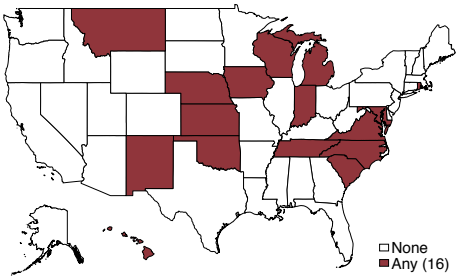

(d) Match Amounts

Phase I

Amount, $K	N State-Years
25–45	27
50–65	46
75	31
100–105	14

Phase II

Amount, $K	N State-Years
50–75	10
100	17
250	10
500	24

Notes: 1a-1c indicate which states had any matching policy in particular years. 1d tabulates the sizes of the matches awarded depending on the Phase of the award .

administrators in North Carolina (where a match policy is in place) shared results from a survey of grant recipients. With the caveat that these data are based on firms receiving SBIR awards from all federal agencies (not just the DoE), they indicate that roughly 50% of funds are spent on wages, which is in line with the notion that these firms are research-oriented and employ high-skilled labor. About 25% of funds are spent on equipment, supplies, or facilities.

One of our concerns is that the (unobservable) regulations surrounding the use of state funds might be more lax than those from federal agencies. This would be problematic for our analysis of patent outcomes if, for example, state dollars were more easily spent on patent application or on litigation fees.[16] However, at least in the case of North Carolina, and with the further caveat that we have no data on firms'

[16]Our best read of policies indicates that federal grant dollars can be spent on certain patent related costs. But a casual internet search indicates some confusion amongst SBIR participants as to the particulars of recovering these costs.

use of federal dollars, it appears that no more than 5% of these state match funds are spent on patent-related costs. As far as we can infer from these data and other states' regulations, it does not appear that receiving state matching funds changes the incentives or costs of patenting.

2.4 Travel Costs and County Characteristics

We compute travel costs within the U.S. by focusing on county-to-county pairs as a semi-dense yet computationally tractable set of regions. We draw upon archival data from the following sources: (i) Department of Transportation Airline Origin and Destination Survey; (ii) the NBER Place Distance Database; (iii) and the IRS Standard Mileage Rate program. We compute annual adjusted estimates (2006-2017) of the minimum cost of traveling between each county pair in the U.S. using the minimum cost of either driving directly, or driving to/from the nearest airports and flying. Then, we take the set of counties where DoE SBIR grants are awarded as the focal set of counties and calculate the average cost of making a round trip to these focal counties from all other counties. This allows us to draw concentric circles around the firms awarded grants as a function of these approximate travel costs. Additional detail on this data construction is in Appendix B.

To characterize each county, we rely on the collection of data sets gathered by the U.S. Cluster Mapping project (Delgado et al. 2016). These data describe a wide range of economic and socio-demographic features of states and counties.[17]

3 Empirical Model and Data Construction

3.1 Reduced-form Model

The focal aggregate production function is a negative binomial model that resembles a traditional log-transformed Cobb-Douglas function.[18] It relates the annual flow of patents to a stock of R&D investments. Researchers have been regressing patent rates on R&D stocks for decades (e.g., Griliches 1981). But unlike most prior work,

[17]The three main categories of variables are "Performance" (i.e., wage growth, establishment creation), "Business Environment" (i.e., venture capital, R&D per capita, manufacturing intensity), and "Demographic" (i.e., population, migration rate). See www.clustermapping.us/region for more.

[18]See Jones (2005b) for a derivation of the Cobb-Douglas form for aggregate production functions when the distribution of ideas for production possibilities is Pareto.

our units of analysis are not firms but rather the technological areas that those firms operate in – the "priorities" that the DoE invests in. Because we rely on the patent classification scheme to index these areas, we refer to these areas as *classes*.

The expected count of ultimately successful patent applications Y in each class j during year t depends on the current stock of R&D investments K per

$$\mathbb{E}[Y_{jt} \,|\, \mathbf{X}_{jt}] = \exp\big(\log(K_{jt})\beta + \sigma_t + \omega_{jt}\big), \tag{1}$$

where $\mathbf{X}_{jt}$ is the vector of covariates and β is the focal productivity parameter to recover. σ_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 $2017 - t$ succeeding years). ω is a state variable describing productivity shocks that are observable to all parties except for us (i.e., the unobservable value of patenting in each class at a particular time), which we address in Section 4.[19]

Eq. 1 is an aggregate and reduced-form production function in the sense that it does not structurally describe firm-level production, but rather aggregates this production into technology-level relationships. Thus, the parameters we estimate are based on the choices and capabilities of firms operating in different locations of technology space (e.g., endogenous investment responses). In other words, the model allows us to study the policymaker's return on investment but does not shed light on how exactly these investments translate to outputs.[20] Still, this return on investment is clearly relevant for policy evaluation purposes, and as we show below, we can still use this model to infer new insights about how research spillovers occur across geographic and technological space.

The two key econometric challenges of estimating Eq. 1 are (1) handling ω, the unobservable productivity shock, and (2) extending the model to identify spillovers. We turn to these issues in Sections 4 and 5, focusing the remainder of this section on how we construct the underlying data and estimate the model.

[19] We make the standard assumption about the dispersion of the negative binomial model: it equals $1 + \alpha\mathbb{E}[Y_{jt} \,|\, \mathbf{X}_{jt}]$, where α is the overdispersion parameter to be estimated.

[20] Without product market data, we cannot examine a likely source of negative spillovers: product competition. Still, our reduced-form results encapsulate any strategic interactions at the R&D stage of patent production.

3.2 Input-output Connections

We need to classify both R&D stocks K and patent flows Y into the same categorization scheme. In particular, we are interested in using j to index "technology space," much like how U.S. counties index U.S. geographical space. To do so, we leverage the CPC scheme, a hierarchy of categories used to organize the patent record. Much like counties, the CPC provides us with a discretized version of the space we are investigating. However, also much like counties, the CPC was not developed with empirical researchers in mind. Therefore our approach attempts to avoid relying on the CPC scheme in a way that may produce spurious results (Thompson and Fox-Kean 2005).

Figure 2: Mapping Inputs from Funding Announcements to Patent Classes

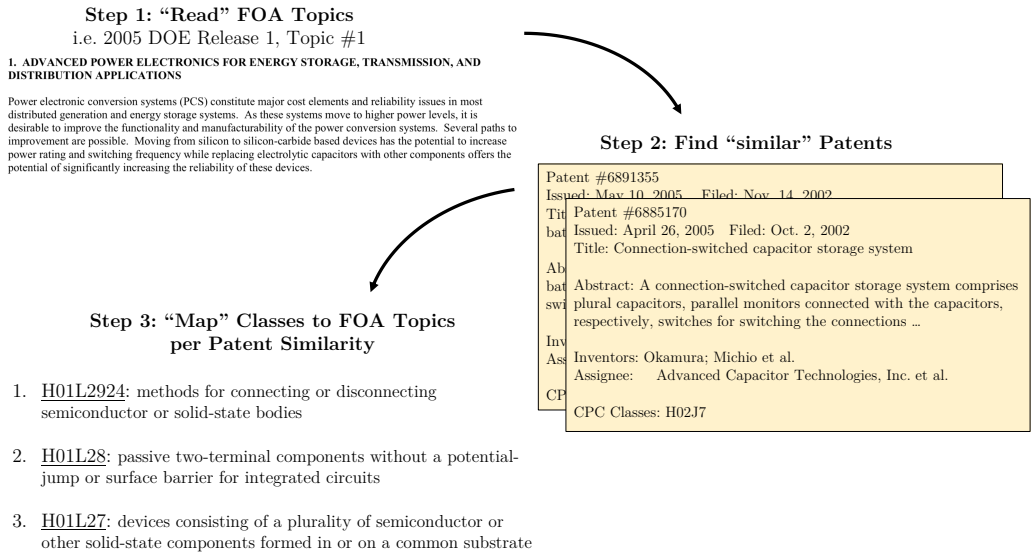

Notes: See Appendix B for more details.

Figure 2 provides an overview of our process. First, we use standard optical character recognition software to digitize the 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 topic, and we compare the textual similarity of these descriptions with the abstract of every USPTO patent awarded between 2001 and 2004 (the four years preceding our sample of FOAs).[21] Using methods now

[21] Using these presample patents for the similarity calculation is an attempt to avoid any textual similarity that arises from the DoE's in-sample choice of priorities.

commonplace 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. Loosely speaking, 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.

Finally, we leverage the fact that all patents have pre-assigned CPC classes and assume that these CPC classes describe the entirety of technological space. However, we do not need to assume – like most prior work – that the specific hierarchical structure of the patent classification scheme reflects that classes are more or less technologically similar to each other. Instead, we can collapse the FOA-patent-level similarity scores to FOA-class-level similarity scores and build FOA-specific rankings of class similarity.[22] For each FOA, the resulting data describe what are the most technologically similar CPC classes.

Using these FOA class rankings and the FOA award data, we calculate the value of grants awarded via FOAs that were similar to each class j – our measure of investment flows, I_{jt}. Furthermore, we calculate separate bins of investment from FOAs with different degrees of similarity to class j. In practice, we must make assumptions about the structure of these bins (especially with respect to the minimum threshold of similarity between an FOA and a class to warrant the flow of *any* funds). We discuss need for and determination of these thresholds in Section 5.

Calculating the flow of patent output in each class, Y_j, is more straightforward. Most patents are labeled with multiple CPC classes. For simplicity, we divide the patent evenly between each of these classes. Thus, for all calculations and regressions, a value of $Y_j = 1$ indicates one "patent's worth" of technology output in that CPC class j.

For more on the construction of the main data set, see Appendix B. For other notes related to the creation of supplementary data and the estimation of the regression model, see Appendix C.

[22] In this collapsing process, outlined in detail in Appendix B, we take steps to remove spurious correlations in text usage including using Bayesian shrinkage factors and demeaning scores at the FOA topic level.

4 Identification and Inference

Figure 3 illustrates the causal pathways that we are concerned with. To summarize, the DoE's SBIR investments facilitate R&D performed by grant recipients. The state match programs moderate the total amount invested in particular firms. The R&D performed by these firms may lead to patents and other outputs that may in turn (either directly or indirectly via patent disclosure) influence the R&D performed by other firms in the economy.

The pathway of interest is from public investments to patents, as operationalized by Eq. 1. However, Figure 3 makes clear that when estimating Eq. 1 we face the challenge of simultaneity. Investments are not made at random but rather with some knowledge about their productivity (or the demand for the downstream good). In our model, this is reflected by the ω term. If there are forces that jointly influence both public investments and the productivity with which those investments are used to generate patents, we cannot estimate how changes in public investments directly lead to new patents (as captured by β).

Figure 3: Model of Investment and Production

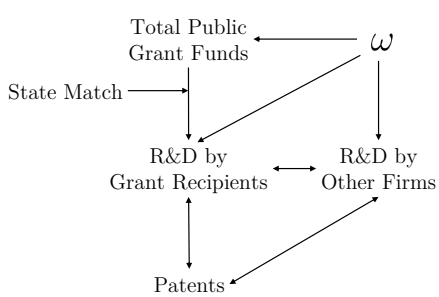

Notes: ω is the unobservable state variable (e.g., productivity or demand shocks).

Practically speaking, there are three channels through which not accounting for ω might lead to biased productivity estimates. The first channel is the selection of topics included in FOAs. We would overestimate productivity if the DoE solicits proposals related to technologies where patenting is the easiest or the most valuable.[23] Second is the selection of topics by applicants. Conditional on the set of topics an-

[23]Still, requests from Congress, input from National Laboratories, path dependence, and other pressures to diversify have been posed as additional forces at play. See Chapters 2 and 4 in National Academies of Sciences, Engineering, and Medicine (2001) for details.

nounced, potential applicants will likely pursue topics that have the largest potential for generating and appropriating returns. Lastly is the selection of applicants to receive grants. The DoE's policies are intentionally designed to award more funds to more productive firms. These forces would likely lead to upwardly biased productivity estimates, especially because patents are an objective output for the DoE.

To circumvent these issues, we make use of the aforementioned state-specific matching policies. Doing this allows us to isolate variation in public investments that is plausibly orthogonal to these unobservable variables and allows us to speak to a clear policy question – holding fixed the number of SBIR awards and without changing any other policies (i.e., the review process), what would happen if the size of SBIR awards increases?

The following three subsections describe (1) evidence that these policies are plausibly exogenous with respect to our analyses, (2) the construction and interpretation of the windfall investments, and (3) our inference approaches.

4.1 Evidence of Match Policy Exogeneity

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

To explore these assumptions, we perform two 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 of these regressions are in Appendix D. Across all specifications, we cannot reject a null of no effect. Whether we examine across- or within-state (across time) 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. It is worth noting that the notoriously productive states of California and Massachusetts do not have matching programs in our sample.

Because these state policies increase the expected size of an SBIR award, it may have been that as firms became aware of these policies, they moved into states that enacted the matches. Or, perhaps these policies were enacted in anticipation of

an influx of businesses into the state. However, with match amounts ranging from \$25,000 to \$500,000, and federal-level success rates reported at roughly 15%–20%, this is likely not a tremendous change in the expected benefit of moving across state lines, especially considering it does not change the extensive margin value of winning. Still, we explore this possibility in our second test by constructing a panel data set of the locations of SBIR grant recipients.[24] Using these data 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 were 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 presented in Appendix D and show no meaningful evidence that grant recipients are more or less likely to move into states after these policies are enacted.

4.2 Estimating and Interpreting the Windfall Investments

The prior analyses suggest the particular firms, and thus the particular technology classes, that are more or less exposed to the match may receive investments plausibly orthogonal to our worrisome ω term. However, the fact that these state-based policies operate by multiplying (endogenous) federal investments poses a challenge. We must separate which of the state-based investments are a function of receiving more federal dollars versus being windfall investments that arise only because of each technology classes' differential exposure to certain match policies. In short, we do this by predicting the average amount of state funds that each class would receive based on the flow of federal funds and the average matching policy at the time (across all states) and then use deviations from this prediction – the state match windfall – as an instrument. For the formal derivation of the instrument, see Appendix E.

However, because they are effectively residual investments, these windfalls take both positive and negative values. The logarithm transformation is no longer applicable. Instead, we demean investments by dividing by the sample average. Formally,

[24]This 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% march rate) to the National Establishment Time Series.

we rewrite the main production function to now be

$$\mathbb{E}[Y_{jt} \mid \mathbf{X}_{jt}] = \exp\Big(\frac{K_{jt}}{\overline{K}}\theta + \sigma_t\Big), \tag{2}$$

where $\overline{K}$ is the sample mean of whichever investment flows we plug into K (e.g., total public investment, or just state match windfalls). Thus, a one unit increase in $\frac{K_{jt}}{\overline{K}}$ describes a 100% increase in investments from the mean. As we show in Appendix C, this transformation approximates a log transformation in terms of the coefficient estimated (i.e., $\theta \approx \beta$).[25]

Some other caveats and limitations of this instrumental variable are worth mentioning. First, due to the recency of many state matching programs, the magnitudes of the resulting windfall investments are much smaller than the total federal investments, typically on the order of 1%–10% of what the DoE directly invests in each class. On the one hand, the small relative size of these investments raises some concerns about measurement error, which would bias our estimates toward zero (yielding conservative cost-effectiveness estimates). On the other hand, from the perspective of the firms receiving these windfalls, a 50% or 100% increase in investments is certainly economically meaningful. Additionally, we assume constant returns to scale and, because we cannot observe the actual flow of state investments, we interpret the instrumental variable results as an intent-to-treat effect. We further discuss the estimation and interpretation of these windfall investments in Appendix E.

4.3 Inference

The identifying variation in windfall investments arises because grant recipients located in states with different match policies are pursuing inventions in different technology classes. Thus, each technology class is exposed to windfall investments to the extent that inventors pursuing that class are located in states with any/large match rates. When it comes to estimating standard errors, we are concerned about both spatial and temporal correlations (e.g., agglomerative forces will likely lead different firms in the same state to pursue similar technology classes; specialization will likely

[25] The strict interpretation of the θ coefficient is not the traditional average elasticity across the sample values of the dependent and independent variables, as given by β. It is the average elasticity at the mean value of the independent variable across the sample values of the dependent variable. We assume that this difference is negligible.

lead the same firm to pursue similar technology classes over time). To account for these correlations, we implement the score bootstrap (Kline and Santos 2012). Bootstrap sampling occurs using a higher level of CPC class aggregation compared to what determines our units of observations, and two-way clustering occurs across both these aggregate classes and years.[26]

Still, these standard errors are based on traditional asymptotics and are intended to guide inference about effects in an infinitely sized population. In cases where researchers observe a practically finite population and are interested in estimating internally valid standard errors – as we do and are – randomization inference tests provide an alternative approach. However, the match policies are not a true experiment in the sense that we do not know the underlying (pseudo-)random process that generated the pattern of observed policies. Thus, motivated by Bertrand et al. (2004), we conduct an approximate permutation test where we construct a null distribution using random permutations of the observed match policies. Then, again following Bertrand et al. (2004), and in the spirit of Fisher's exact test, we construct p-values as the share of estimates from the null distribution that are less extreme than the estimate identified using the observed match policies. See Appendix G for more.

5 Estimating Spillovers

Our overall approach to estimating the magnitude of spillovers in a policy-relevant way is as follows: first, we recover multiple elasticity estimates when counting different sets of patents; second, we use these elasticities to estimate the implied cost (to the DoE) of spurring a marginal patent amongst these different sets; finally, to approximate the degree of spillovers across different dimensions, we compare these marginal costs when different sets of patents are counted. For example, if we estimate that it would cost \$1 million to spur one additional patent to be generated by SBIR grant recipients, and only \$500,000 to spur one patent out of all inventors (including grant recipients), then we would say that SBIR grant recipients are only responsible for 50% of the net patent output of the program – geographic spillovers would be twofold in this hypothetical case.

[26] We use the level 3 CPC codes for the clustering of the standard errors. There are about 120 of these level 3 classes compared to the roughly 8,700 level 4 classes that define our units. In Stata, this is implemented using the `boottest` package (Roodman et al. 2019).

5.1 Geographic Spillovers

For the purposes of capturing geographic (across firm) spillovers, we take the straightforward approach of estimating separate regressions that increase the size of concentric "circles" of which patents are counted in the dependent variable (Feyrer et al. 2017). We draw these circles first at the level of the focal firms that receive DoE SBIR grants, which effectively replicates a traditional firm-level analysis of the program. Then, looking within the U.S., we draw a series of circles based on percentiles of approximated travel costs to include patents from inventors and firms that are further and further away (in terms of travel costs) from the focal recipient firms. Finally, because we do not have international travel cost data, we include all patents from inventors and firms not located in the U.S. Thus, now instead of a single dependent variable describing total patent flows Y_{jt}, we can examine patent flows denoted by Y_{jt}^d, where d describes the threshold of the geographic threshold for patents to be included in the dependent variable.

Now using this approach, instead of a single production function, we estimate a series of models of the form:

$$\mathbb{E}[Y_{jt}^d \,|\, \mathbf{X}_{jt}] = \exp\Big(\frac{K_{jt}}{\overline{K}}\theta^d + \sigma_t^d\Big), \tag{3}$$

where the vector of θ^d coefficients describe how the output elasticity changes as the base of possibly affected firms and inventors expands.

5.2 Technological Spillovers

Identifying across-technology spillovers is not as straightforward. As noted by Manski (1993), identification here can only be obtained with assumptions about connections of units in the network. Following Toulis and Kao (2013), we make an informed assumption (using our instrumental variable) about how far in technology space spillovers can occur.[27]

In our setting, this amounts to identifying some level of similarity between each FOA and the full set of patent classes beyond which we assume no spillovers will occur. In other words, we must search for a similarity score between an FOA and

[27] In the context of Toulis and Kao (2013), the assumption is about which units are "neighbors" and which are not. It is assumed that only neighbors can have nonzero spillover effects on each other.

CPC class beyond which it is reasonable to assume that investments originating in the FOA will have no effect on output in that CPC class.

Recall that we construct investment flows I_{jt} and corresponding stocks K_{jt} that describe the dollar amount of SBIR funds allocated to class j in time t. Even more specifically, as outlined in Appendix B, we can construct bins of investments flowing into each class as a function of how similar that class is to the FOA topic from which the investment originates. Let $b \in \{100\text{–}95,\ 96\text{–}90,\ \ldots\ ,\ 10\text{–}6,\ 5\text{–}1\}$ index these bins, giving $K_{jt,b}$ such that $K_{jt,b=96\text{-}90}$ describes the stock of investments in class j at time t originating from FOA topics where class j is in the 96th to 90th percentile of similarity scores. Substituting these bins of investments into the main production function, and allowing their effect on patent rates to vary, yields a more flexible version of the model:

$$\mathbb{E}[Y_{jt}^{d} \,|\, \mathbf{X}_{jt}] = \exp\Big(\sum_{b \in \mathcal{B}} \frac{K_{jtb}}{\overline{K}} \theta_b^d + \sigma_t^d\Big), \tag{4}$$

where $\mathcal{B}$ is the set of similarity bins within the assumed boundary of technological spillovers. Per Toulis and Kao (2013), $\mathcal{B}$ cannot include all possible similarity bins.

To identify a reasonable threshold for which similarity bins to include in $\mathcal{B}$, we use the state match windfall investments to search for a point where we begin to estimate null effects. We begin by only allowing the highest levels of similarity to be relevant for spillovers and then systematically relax this constraint.

Figure 4 plots the vector of θ_b^d coefficients for $\mathcal{B} = \{100\text{–}95,\ \ldots\ ,\ 45\text{–}40\}$ and $d =$ SBIR grant recipients, U.S. inventors, and worldwide inventors.[28] We find that, beginning around the bins that include the 55th to 50th percentile of similarity scores, we start estimating effect sizes that cannot (or marginally) reject the null. In other words, this pattern of coefficients suggests that investments originating from a particular FOA topic will not have a meaningful effect on patenting rates in CPC classes that are lower than the 50th percentile of similarity with that FOA.

Thus, for all main regressions, we set $\mathcal{B} = \{<100\text{–}95,\ \ldots\ ,\ 60\text{–}55\}$. In Appendix G, we report replications of our main results where we vary the boundary plus or minus two percentile bins, and the implied marginal costs per patent are similar.

To summarize our approach: we set up a technology-level production function to examine the impact of public investment on targeted technologies. We map federal

[28] In Appendix F we present the set of regressions that led us to this stopping point. It illustrates a negative selection effect whereby not conditioning on investments from "nearby" CPC classes identifies significantly smaller output coefficients.

Figure 4: Technological Spillover Boundary Search

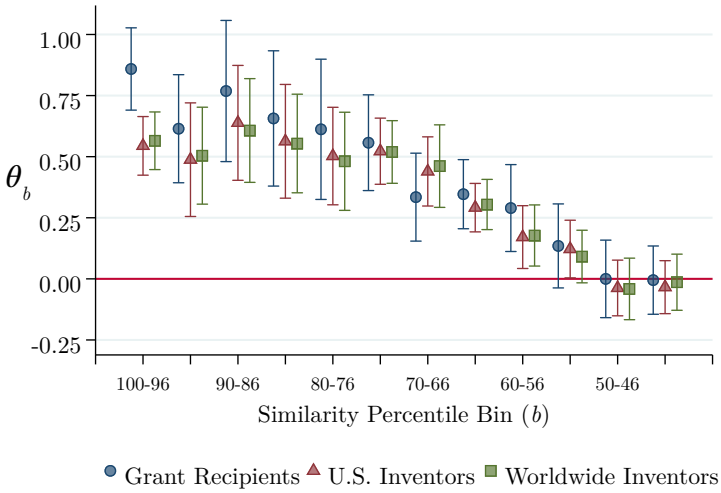

Notes: This figure plots the vector of θ_b^d coefficients from Eq. 4 for d = ("grant recipients", "U.S. inventors","worldwide Inventors") with 95% confidence intervals bracketing the estimates.

funding opportunities onto patent classes and identify the marginal costs using the noncompetitive state SBIR match program. We then estimate the production function using a negative binomial regression and rely on the state windfall investment as a measure of plausibly exogenous inputs. We trace two forms of spillovers – geographic and technological – to estimate the cost of public funding on patent activity. For the former, we compute travel costs from each SBIR grant recipient across U.S. counties. For the latter, we draw upon the FOA-to-CPC similarity score to trace across technology space.

6 Main Results

??

6.1 Summary Statistics

Appendix A contains tables that describe the summary statistics of the final technology-year input-output data set used in the analyses as well as the top technology classes that receive investments. Annual patent flows are highly skewed, as is typically observed. There appears to be a small set of classes that receive a very disproportionate

amount of funding from "high-match" FOAs, which is unsurprising when considering the set of technologies the DoE might pursue relative to the entirety of (patentable) technology space.

Appendix A also displays the top classes per federal investment. The DoE invests largely in energy-centric patent classes but also in chemistry, computing, and communication classes. The top classes per state match windfall investment are also shown. We see that firms in states with the matching policies are not always pursuing technological areas with the same rank-order preferences as the DoE as a whole. We take this pattern as further evidence that the firms in states with matching policies are not pursuing technological areas that are particularly more or less attractive relative to those in states without these policies.

6.2 Results at Boundary Cases: Grant Recipients, U.S., and Worldwide Inventors

Before unpacking technological spillovers, we begin by examining USPTO patent flows using the three boundaries of the geographical sets of inventors: SBIR grant recipients, all U.S. inventors, and inventors worldwide. In these analyses, we lump all investments coming from FOAs with similarity scores at or above the 55th percentile into a single bin. Figure 5 plots the raw data for these three cases, using either the total amount of federal and state investments or only the state match windfall investments. As seen by the differing slopes of the fitted lines, these graphs suggest that ignoring the likely endogeneity of federal investment could overstate productivity.

Table 1 presents estimates of the output elasticities and the implied average marginal costs of producing one additional patent. Focusing only on the patents produced by SBIR grant recipients (Panel A), we estimate an elasticity of about 1.1 per federal investments, which implies an average marginal cost of about $330,000. Introducing the state-level investments identifies a slightly smaller elasticity. And when identifying the elasticity only using the state match windfall, we identify a significantly smaller elasticity of roughly 0.49, which translates into costs of about $790,000 per additional patent. Column (4) weights grant recipients' patents by their citation counts and finds implied costs of roughly $1 million. This finding is in line with the magnitude of the citation-weighted cost estimates implied by Howell (2017), which gives us confidence in our instrumental variable approach.

Figure 5: Input-Output Relationships for Boundary Cases, Demeaned and Residualized

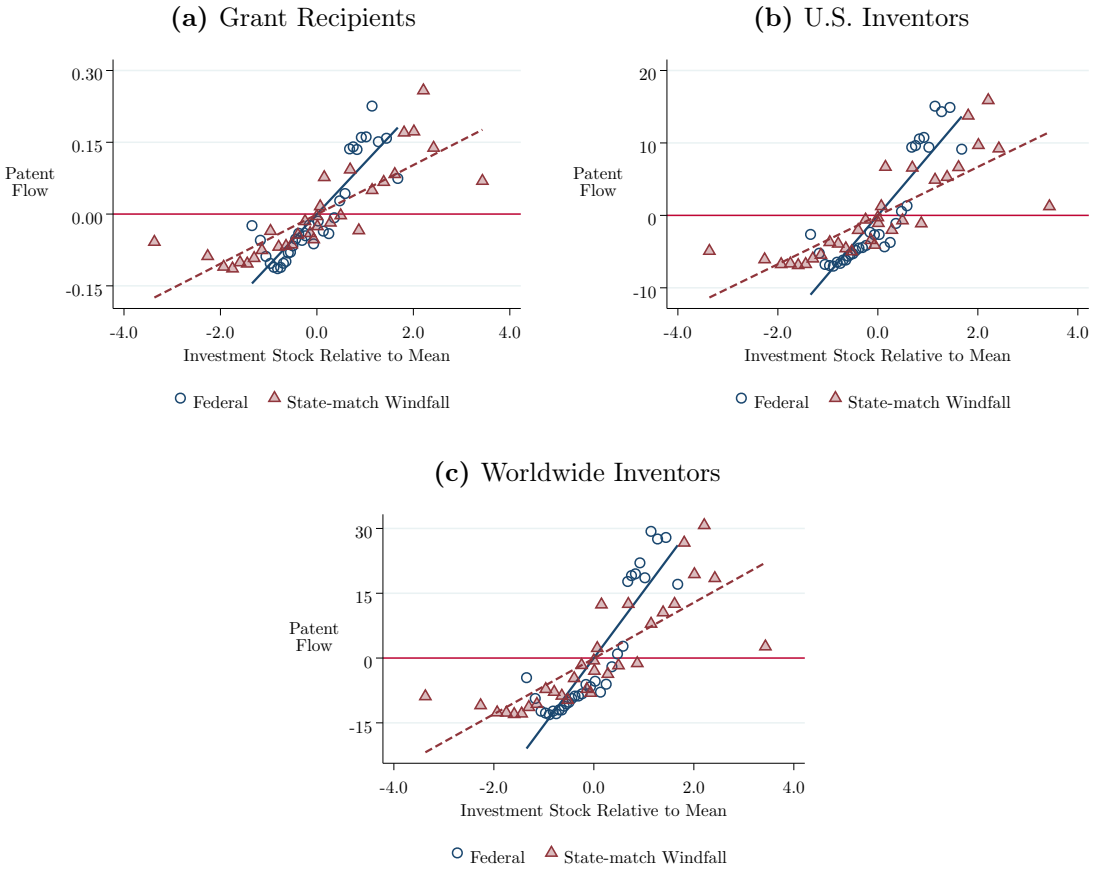

Notes: This figure shows equally binned scatterplot of investment stocks and patent flows; stocks are demeaned and both variables are residuals from year fixed effects. Linear fit lines are also plotted (solid and dashed lines correspond to federal and windfall stocks, respectively).

Panels B and C of Table 1 examines all USPTO patents assigned to U.S.-based inventors and then inventors worldwide. In both cases, we again find that the elasticity identified using the state match windfall is significantly smaller than the straightforward approach of using federal or total investments. We estimate that the average marginal patent requires about $220,000 of additional investment from the DoE if all U.S. inventors are counted and only about $170,000 if the location of the inventor is irrelevant. Comparing this to the $790,000 required to spur a patent by the grant recipients themselves suggests that geographic (across firm) spillovers could be upwards of three- to fivefold. In other words, these estimates suggest that grant recipients are responsible for only about 20% of the net patent output spurred by investments via

the DoE's SBIR program. And U.S. inventors appear responsible for roughly 75% of the net patent output. Here, unlike the narrow case of grant recipients, there does not seem to be a premium for citations, as the cost of citations and patents are very similar.

Table 1: Output Elasticity and Implied Cost Estimates for Boundary Cases

	Patent Count			Cites
	(1)	(2)	(3)	(4)
		Panel A: Grant Recipients		
Elasticity (θ)	1.146	0.966	0.485	0.351
	(0.193)	(0.228)	(0.115)	(0.080)
\$M per +1	0.333	0.395	0.787	1.086
Patent or Cite	[0.25–0.50]	[0.27–0.74]	[0.54–1.47]	[0.75–1.96]
		Panel B: U.S. Inventors		
Elasticity (θ)	0.915	0.674	0.382	0.401
	(0.102)	(0.132)	(0.077)	(0.081)
\$M per +1	0.094	0.128	0.226	0.215
Patent or Cite	[0.08–0.12]	[0.09–0.21]	[0.16–0.37]	[0.15–0.35]
		Panel C: Worldwide Inventors		
Elasticity (θ)	0.932	0.685	0.384	0.404
	(0.090)	(0.128)	(0.083)	(0.084)
\$M per +1	0.070	0.095	0.170	0.162
Patent or Cite	[0.06–0.09]	[0.07–0.15]	[0.12–0.30]	[0.11–0.27]
N	114,050	114,050	114,050	114,050
Investment Stock	Federal	Federal+State	Windfall	Windfall

Notes: Bootstrapped standard errors, clustered at the aggregated CPC level and years, are reported in parentheses. "\$M per +1 Patent or Cite" presents the average marginal costs implied by the elasticity estimates, in millions, with 95% confidence interval bounds (based on the standard errors of the elasticity) in brackets.

To explore the sensitivity of our results to our data construction and modeling decisions, Appendix G recreates multiple versions of our main instrumental variable results varying multiple aspects of our data construction and estimation procedures. We obtain very similar elasticities across the alternative models. Introducing nonzero discount rates for investments yields smaller average marginal costs, largely by construction, though we prefer our more conservative estimates because they better capture the fact that the policymaker must trade off current investment flows for future patent flows. Across alternative ways of handing the similarity scores and setting

the threshold of technological spillovers, we obtain results that are all qualitatively similar to what we report here. Altogether, this gives us confidence that these results are not driven by any of the choices we make in the data construction or estimation.

Appendix G also reports the results of our approximate permutation tests. For the boundary cases of grant recipients and worldwide inventors, we estimate p-values of 0.027 and 0.036, respectively (i.e., only 2.7% of the estimates from permutations of observed match policies were larger than our estimates based on the true observed match policies). These results reject the null at conventional levels and give us confidence that the effect of the windfall estimates is unlikely to be due to chance assignment of these policies to particular states in particular years.

We also conduct an exercise to further investigate the validity of our key identification assumption that the state match windfall investments are uncorrelated with any (unobservable) shocks to the patent flows. We relate our approach to the control function methodology for estimating production functions (Olley and Pakes 1996) and note that, under some stricter assumptions, it may be reasonable to consider federal investment stocks as a proxy for these unobservable shocks. In other words, if we assume that the DoE and firms observe, forecast, and respond to ω_{jt}, then we could proxy for these shocks using a flexible function of federal investment stocks, K_{jt}^{federal}. The results of these regressions where flexible functions of K_{jt}^{federal} are included as covariates are shown in Appendix G. They yield smaller, but still relatively precise, estimates of the coefficient on windfall investments, θ. Direct comparisons of these coefficients is difficult given changes to the functional form and underlying assumptions of the model. Rather, we interpret the fact that windfall investments are relevant even when conditioning on federal investments as further evidence that the variation in these windfalls is plausibly uncorrelated with any unobservable determinants of patenting rates.

6.3 Spillovers, Two Ways

The prior results suggest that the cost to the DoE of spurring "one patent's worth" of any technology is much lower when considering how their investments lead to spillovers across firms and geographic space. But are those net geographic spillovers concentrated among firms and inventors in nearby locations? And how closely do these new patents align with the DoE's objectives as expressed in their FOAs? Most

generally, how does the cost of spurring a patent change depending on one's perspective as to which firms' patents and patents of which technology classes count?

Figure 6 illustrates the distribution of costs depending different thresholds for counting patents generated by different sets of firms or inventors (y-axis) and patents that are more or less similar to the DoE's objectives (x-axis). We report cumulative cost metrics, in the sense that the cost at any point (y, x) indicates how much investment via an FOA must be increased to spur a patent by a firm that is at most y-distance away from grant recipients, and that patent is in a technology class that is at most x-similar to the FOA's objectives. In other words, the cost estimate displayed in the bottom-left corner of the plot – roughly \$5.4 million – is the cost to be expected if the DoE wanted their grant recipients to develop a technology that is extremely similar to the FOA's objectives. Conversely, the cost estimate in the top-right corner of the plot – roughly \$175,000 – is the cost to be expected if the DoE wanted any firm or inventor in the world to develop a technology that is in any way related to the objectives of the FOA (this also approximates the estimate from Table 1, Column (3)).[29]

Overall, the range of cost estimates suggest that technological spillovers can be on the order of four- to sevenfold, and geographic spillovers can be on the order of four- to fivefold. This implies that (1) the costs of spurring patents that are very much in line with the technological objectives of an FOA can be four to seven times larger than the costs of spurring any patent, and (2) the costs of spurring a patent from a SBIR grant recipient can be four to five times larger than the costs of spurring a patent from any inventor in the world.

Figure 7 illustrates the same underlying results but uses an alternative transformation of the marginal costs – the cumulative share of net patent output that arises along the two dimensions. The base of the net output is defined as the output if technology and geographical spillovers are fully appreciated – the upper right-hand corner of the figure where we count all patents that are within the technology spillover boundary.

At the top and bottom of Figures 6–7 we focus on the boundaries of grant recipients and worldwide inventors. Recall that each of these horizontal slices are based on a single regression and report cumulative costs. Interestingly, compared to the case

[29]This cumulative perspective reflects the fact that, in practice, these investments are allocated via FOAs and awards that always affect multiple classes (per spillovers).

where we constrain investments to have the same average effect across the similarity percentile bins (as shown in Table 1), here, where we allow these output elasticities to vary, we identify larger net costs per patent. At the border of technological spillovers, this more flexible model implies a marginal cost of roughly $940,000 per patent (the bottom-right corner of the plot), compared to roughly $790,000 in the case where only a single average effect is estimated.

The middle portion of the figure is based on a series of 50 regressions (each with 9 coefficients) where we concentrically include inventors and firms located in U.S. counties that are within a given threshold of average travel costs to the SBIR recipient firms. Thus, the top horizontal slice of the middle portion of the figure represents the cost and share metrics if all U.S.-based inventors or firms are included in the output counting.

Figure 6: Cumulative Marginal Patent Cost per Geographic and Technological Spillovers

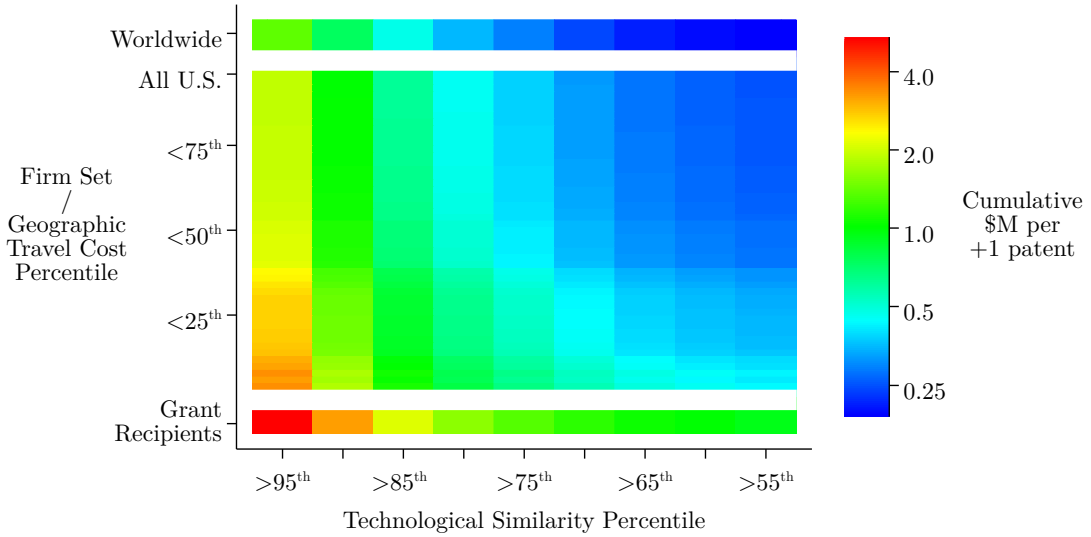

Notes: This figure plots the implied cumulative cost, in millions, of spurring one marginal patent as a function of (1, y-axis) the travel costs between the location (of the firm) where funds are invested and where any patent is generated, or (2, x-axis) the technological similarity between the objectives of the investments and the content of the patent.

Because Figures 6 and 7 do not convey the uncertainty in our estimates (i.e., the standard errors), Appendix G includes a version of these plots where, instead of plotting costs or output shares, we plot the z-statistics from the negative binomial regressions underling these estimates. An interesting pattern emerges, where we ob-

tain smaller test statistics for both the very high and very low technological similarity investment bins. The return on investments appears more uncertain when focusing on technologies that are very closely aligned with an FOA or are very distant from the FOA's objectives.

Figure 7: Cumulative Share of Net Patent Output per Geographic and Technological Spillovers

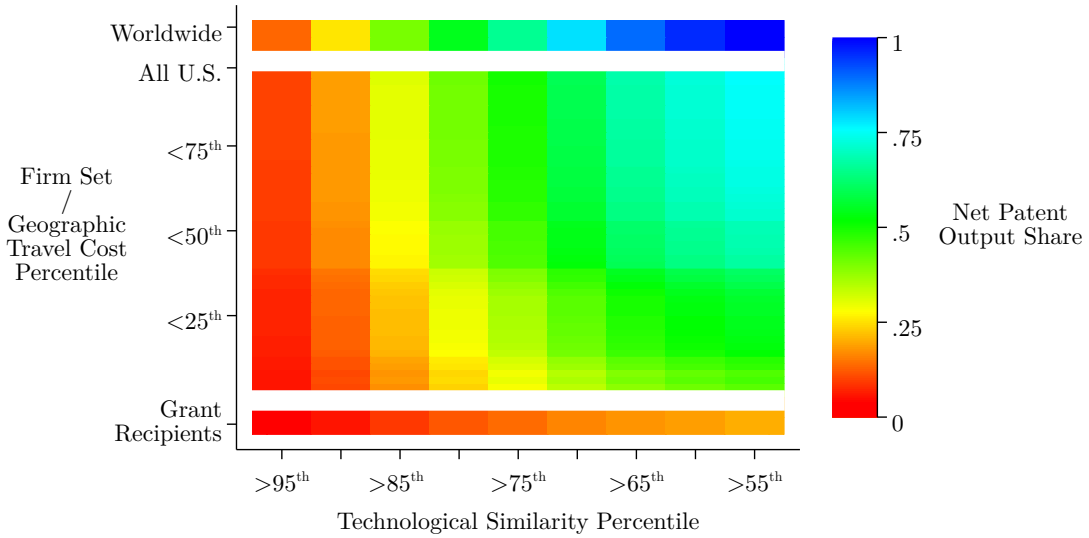

Notes: This figure plots the cumulative share of patent output as a function of (1, y-axis) the travel costs between the location (of the firm) where funds are invested and where any patent is generated, or (2, x-axis) the technological similarity between the objectives of the investments and the content of the patent. By construction, the share of net output incorporating all geographic and technological spillovers – the upper-most right-hand corner of the plot – is set to 1.

As is often the case, we find that when it comes to the costs of spurring invention via this program, "it depends." It depends on the policymaker's preferences over how closely these inventions must align with their objectives and on whether or not the policymaker has preferences over which inventors are responsible for these inventions.

Taken altogether, these results make clear that although the SBIR grant recipients can convert their new funds into (privately) valuable patents, the work that these funds enable also makes a tremendous amount of new invention possible. Across the distribution of technology spillovers, it appears that SBIR grant recipients are only *directly* responsible for about one fifth all USPTO patents spurred by these grants.

A lax interpretation of this result would conclude that the private return (to recipients) on these R&D investments are somewhere around four times lower than

the social returns (to all other firms and inventors). As a first note on the plausibility of this magnitude, the best evidence to date has arrived at estimates very much in this range. Jones and Williams (2000), Bloom et al. (2013), and Arque-Castells and Spulber (2019) all estimate the social returns to R&D to be roughly two to four times as large as the private returns. Our estimates are likely to be overestimates for two reasons.

First, we cannot capture any sense of the magnitude of product market spillovers that occur downstream from the inventions we study. We cannot determine what share of downstream revenues these patents are "creating" new value versus "stealing" from it from existing products. Second, we cannot observe any of the market for technology, which Arque-Castells and Spulber (2019) show to be an important source of surplus transfer within innovating firms. In their analyses, they find that not accounting for the licensing of technologies across firms can overstate the private-social wedge by roughly twofold. Still, our results provide further evidence that the wedge between private and social returns to R&D efforts is likely manyfold.

6.4 Comparison to Citation Linkages

Our research design also allow us to investigate the usefulness of traditional front-page citation metrics in terms of capturing spillovers.[30] Though to be clear, we cannot investigate what share of citations are likely reflective of "real" spillovers. Instead, we can investigate what share of these real spillovers (at least, as identified by our models) are captured in these citations.

Table 2 redisplays our main results for two boundary cases and also shows estimates from using citation trails to build the set of relevant patents for our dependent variable. In Column (3), we use the most common approach of counting one degree citations as evidence of a spillover, and we count all patents granted to or citing SBIR grant recipients (3% of all patents in the sample). In Column (4), we extend this approach to include all degrees of citations such that our dependent variable includes any patent that can be connected by any citation trail to a patent from SBIR grant recipients (20% of all patents in the sample). These approaches, in particular the "all-degree" method, provide cost estimates that are well within the bounds of

[30] Although survey evidence suggests citations can reflect meaningful knowledge flows (Jaffe et al. 2000), econometric analyses suggest a significant portion of citations are strategic and/or are in response to unobservable shocks (Arora et al. 2018).

Table 2: Comparison to Formal Citation Connections

	Grant Recipients (1)	Worldwide Inventors (2)	1° Citations (3)	All^o Citations (4)
Elasticity	0.485	0.384	0.452	0.500
	(0.115)	(0.083)	(0.108)	(0.122)
$M per	0.787	0.170	0.238	0.168
+1 Patent	[0.54–1.47]	[0.12–0.30]	[0.16–0.45]	[0.11–0.32]
N	114,050	114,050	114,050	114,050
% of Patents	0.6	100	3.3	20.1
Investment	Windfall	Windfall	Windfall	Windfall

Notes: Bootstrapped standard errors, clustered at the aggregated CPC level and years, are reported in parentheses. "$M per +1 Patent or Cite" presents the average marginal costs implied by the elasticity estimates, with 95% confidence interval bounds in brackets. 1° citations include Grant Recipients' patents and any patents that cite those particular patents in the dependent variable. All^o citations includes grant recipients' patents and any patent that can be traced through any number of citation linkages to a patent by a grant recipient. % of Patents reports the percentage point of patents attributable to each set, where "Worldwide Inventors" is 100 by definition.

our most unconstrained approach implemented in Column (2). In other words, when focusing on these much smaller subsets of the patent record, we obtain very similar results. This finding suggests that the spillovers we identify are largely reflected in citation trails.

6.5 Return on Investment

Ideally, we could use our estimates to estimate a return on investment or even the net value of the output that is generated. Unfortunately, precise valuations of patents spanning the full patent record are difficult to generate. Kogan et al. (2017) use stock market reactions to patent announcements and estimate median valuations upwards of $10 million (in 2017 dollars); the tenth percentile is nearly $300,000, and one additional citation is associated with an increase in anywhere from $40,000 to about $1.3 million. Their sample includes only large, public firms, so it is unclear how to extrapolate their results to this setting. Still, appreciating spillovers in our setting suggests that spurring (any) patent requires an additional $200,000 of investment via the SBIR program. This suggests a high likelihood that value is being created.

7 The Rate and Direction of Spillovers across U.S. Counties

The main results indicate there are substantial spillovers across geographic space. In the following analyses, we zoom in on the U.S. and examine which counties appear to contribute most to these spillovers and it what way. To unpack which counties contribute most to the net spillovers, we perform about 3,100 regressions where the dependent variable in each regression is the sum of patent output from all firms receiving DoE SBIR awards plus all patent output from one of the roughly 3,100 U.S. counties. From these regressions, we can estimate the implied cost to spur an additional patent from either grant recipients *or* any other inventor in a specific county. Thus, counties where this cost is lower are implicitly responsible for a larger share of the within-U.S. spillovers documented in the prior section. And, just as before, we can obtain cumulative cost estimates where only high-match technologies are counted (i.e., technological similarity bin $b \geq 95^{th}$) or where we only impose the minimum observed similarity (i.e., $b \geq 55^{th}$).

Figure 8a plots the distribution of these county-level costs focusing on the case where we allow for all technological spillovers ($b \geq 55^{th}$). Clearly, the vast majority of counties in the U.S. do not contribute to spillovers (note the log scale of the y-axis).[31] Recalling that the implied cost per patent is roughly \$940,000, only about 100 counties have output that can lower this cost below \$900,000 and only for 36 does this estimate drop below \$800,000. Despite the wide range of technological spillovers we observe, this indicates that the vast majority of these spillovers are concentrated in a very small subset of U.S. counties.

We are also interested in the direction of these spillovers – that is, when inventors in a county generate these cumulative inventions, are these inventions more or less technologically similar to the DoE's initial objectives as stated in the FOAs? In other words, we are interested in how "explorative" or "exploitative" a county is. If policymakers are interested in new technologies writ large, then funneling funds towards explorative counties may be worthwhile. On the other hand, if the policymaker has strong preferences over the nature of a technology, then focusing on exploitative counties could help ensure cumulative inventions are closer to the technological objectives.

[31] This largely stems from the high concentration of inventive activity among a few counties. For more, see the Lorenz curves in Forman et al. (2014).

We quantify where each county exists on this explore-exploit spectrum by comparing the implied costs of spurring any technology versus only high-match technologies. We define this index as follows:

$$\begin{aligned}\text{Explore-Exploit Index} &= std\left(\frac{\text{Patent Cost, } b \geq 55^{th}}{\text{Patent Cost, } b \geq 95^{th}}\right) \\ &\approx \frac{\text{Cost of spurring "anything"}}{\text{Cost of spurring "specific things"}},\end{aligned} \tag{5}$$

where *std* indicates standardization. The resulting measure is negative for counties where it is relatively cheaper to produce "any" patents compared to producing "specific" patents, and vice versa.

Figure 8 plots these values and shows how they covary with the aforementioned county-specific patenting costs. First, we note the mass of counties with an explore-exploit index of zero is due to the fact that nearly three quarters of counties have very low to no patent output. Second, we find a clear relationship whereby more explorative counties tend to be lower cost counties.

Figure 8: Across County Differences

(a) Implied Marginal Patent Cost, all Tech. **(b)** County-level Explore-Exploit Index

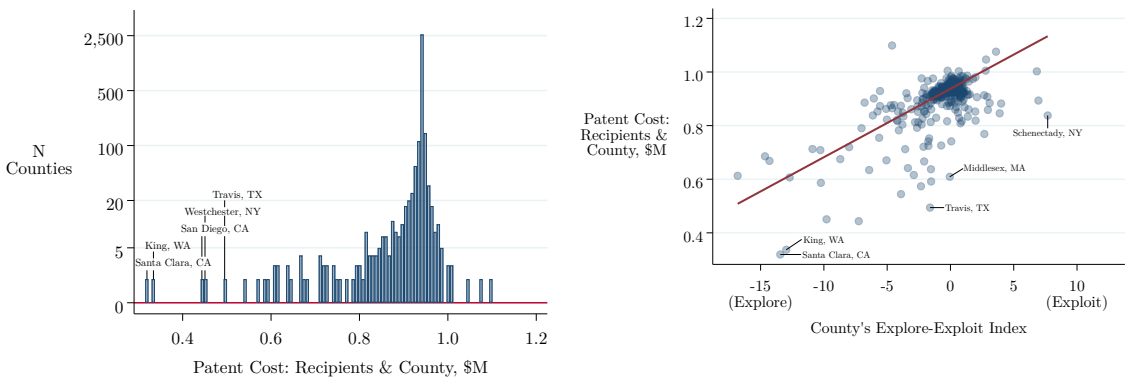

Notes: Panel 8a plots the distribution of county-level costs of spurring a patent from either grant recipients, or any inventor in the respective county. The least stringent technological similarity threshold ($b \geq 55^{th}$) is used. Panel 8 plots this same measure (y-axis) against the same county's explore-exploit index score, which is expressed in standard deviations.

Plainly put, it is easier to spur patents in counties where inventors seem more willing to pursue a wider range of cumulative inventions. It is also interesting to identify counties with large negative residuals from this trend, which indicates inventors in that area are more productive than their choice of strategies (explore versus

exploit) would typically suggest. The counties on the explore end of the spectrum are the usual suspects for centers of highly-productivity innovative output: Santa Clara County, CA (incl. Palo Alto) and King County, WA (incl. Seattle). Some of the high-productivity "balanced" counties in the middle of the spectrum are also rather unsurprising: Travis County, TX (incl. Austin) and Middlesex County, MA (incl. Boston). Most interestingly is the presence of a few highly productive counties that take a predominately exploitative stance. For instance, Schenectady County, NY appears much more productive than its degree of exploitation would suggest. It is likely not a coincidence that Schenectady has significant public-private partnerships focused on the energy sector, and solar power in particular.[32]

Next, we attempt to identify what specific features of these counties might be most predictive of the spillover response. To explore this, we use 52 variables from the Cluster Mapping project (Delgado et al. 2016) that describe a wide range of performance, business environment, and socio-demographic indicators over periods generally ranging from 1998 to 2016 (i.e., labor force productivity, establishment growth rates, venture capital, manufacturing intensity, population, etc.).[33] We supplement these covariates with two important characteristics: (1) the average cost of traveling from the focal county to each of the DoE SBIR grant recipients (the same variable used to draw our concentric geographic bands); and (2) the dollar amount of DoE SBIR grants awarded to the focal county.

The focal data set for this exercise consists of nine cumulative cost estimates (from each of the nine technological similarity bins) for each of the approximately 3,100 counties and these 54 regional covariates. Using a Lasso regression, we identify which of these covariates are the best predictors of these costs.[34] The Lasso algorithm selects only two covariates, both being predictors of lower costs (larger spillovers): (1) the pre-sample annual patenting rate in the county; and (2) the rate of venture capital investment in the county.[35] This result is completely in line with one of the earliest studies of the SBIR program by Lerner (2000), who finds that (when only focusing on

[32]For example, see https://invent.ge/2PzOIDn, or http://bit.ly/2VtRxcH.

[33]Some metrics in this set are based on patent outputs, so for only these particular metrics we restrict the measurement period to be prior to our main estimation sample.

[34]We use log transformations of the cost estimates as the dependent variable, either log or inverse-hyperbolic-sine transformations of the independent variables (the latter if the variable contains zero values), and include intercepts for each of the nine technological thresholds.

[35]If any covariates based on patenting rates are excluded, Lasso only selects venture capital flows. The results of these Lasso regressions are available upon request.

grant recipients and matched controls), success was most often confined to recipients in areas with strong private venture capital support.

We find no evidence that the travel costs of traversing geographic distances constrains spillovers (at least across counties). This is consistent with work such as Forman and Zeebroeck (2012), Forman et al. (2014), and Kantor and Whalley (2019) which find that the rise of the internet (prior to our sample) likely reduced the costs of knowledge transmission and collaboration across geographies. Certainly, it is still possible that on much smaller scales the costs associated with micro-geographies is still relevant (Roche 2019). Returning to our map of cost estimates, Figure 6, it appears that the apparent diffusion across geographic space is less due to any geographic-based constraint (i.e., travel costs) but simply reflects the spatial correlation of SBIR investments and the features that may amplify these investments (i.e., private venture capital).

We perform a similar Lasso-based approach to identify what covariates predict a county's explore-exploit index. Again we find the pre-sample patenting rate and venture capital to be important predictors, with the Lasso also selecting the employment rate in the county. For all three of these covariates, larger values predict a more explorative orientation. Focusing on the venture capital component, our finding is in line with prior evidence that suggests a complementarity between public and private investments in this setting (Lerner 2000; Gans and Stern 2003; Howell 2017). Furthermore, when combined with the prior result, this indicates that while public investments in areas saturated with venture capital may be more productive in the sense of producing any patents, it is also more likely that any spillovers generated by these investments are unrelated to the technological objectives of these investments. Depending on the policymaker's preferences, this could present a tradeoff.

To summarize our findings and illustrate the geographic patterns we uncover, Figure 9 plots the distributions of SBIR investments and venture capital penetration, along with our estimates of the rate of spillovers (per the change in patenting costs) and their direction (per the explore-exploit index). The geographic sparsity and correlation of these variables is clear. One interesting pattern we observe is the co-location of explore- and exploit-oriented counties. Such a pattern might reflect the fact that explore-oriented counties provide a wellspring of ideas for exploit-oriented counties to bend towards their direction of focus. But perhaps just as likely, this could reflect explore-oriented counties building in many new directions off of the

Figure 9: Within-U.S. Geographic Distributions

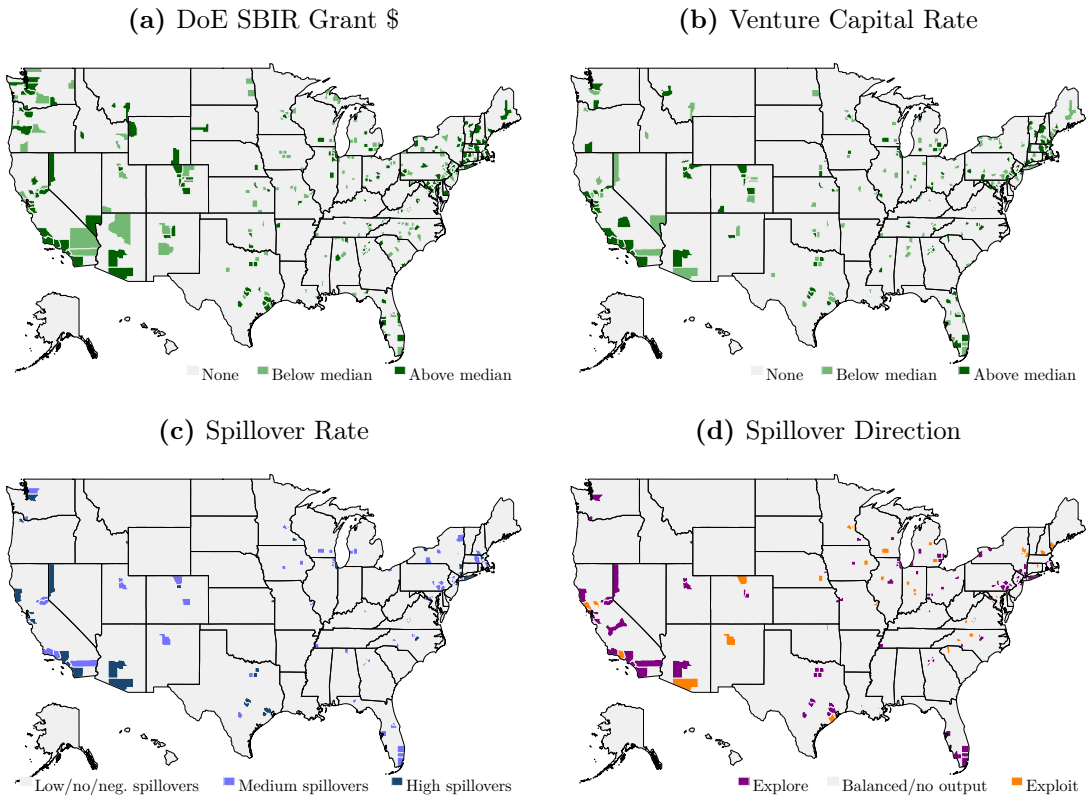

Notes: "DoE SBIR Grant $" represents total DoE SBIR investments and the median is conditioned on non-zero investments; "Venture Capital Rate" represents total venture capital (across all sectors) per GDP and the median is conditioned on non-zero venture capital; "Spillover Rate" refers to counties where the implied cost-per-patent falls below $900,000 (medium spillover) or $800,000 (high spillover); "Spillover Direction" is explorative exploitative if the explore-exploit index for the county is below or above 1 s.d. from the mean.

specific ideas generated by nearby exploit-oriented counties. Further work on these interactions across geographic and technological space could prove relevant to how macroeconomic policies manage interactions of innovators across different regions.

8 Discussion

In this paper, we embrace the targeted nature of SBIR research subsidies and estimate how they influence not just the firms that receive them, but also the state of science they were intent on pushing forward. By leveraging the plausibly random allocation of state matching policies, we circumvent the likely endogeneity of federal investments.

And by evaluating outcomes at the level of technological areas, we identify how these grants lead to new patents that are, or are not, in line with the objectives of the government.

Due to technological spillovers, the specific technologies ultimately invented because of this program can vary widely in terms of how closely they align with the stated objectives. In particular, spurring patents that are very similar to an FOA's objectives costs upwards of four to seven times as much as the costs to spur any patent at all. This means that if policymakers have strong preferences over the technological content of the program's outputs (e.g., in order to support the specific needs of a national laboratory), then they will have to pay a large premium for those particular patents.

Due to geographic spillovers, we find that most of the intellectual property spurred by the DoE's SBIR program is ultimately awarded to inventors that do not directly participate in the program. Of all of the patents spurred by this program, only about 20% are assigned to grant recipients. This estimate suggests two conclusions. First, traditional firm-level analyses of research subsidies that only examine program participants may significantly understate the net impact of the program. Second, this supports estimates that suggest the wedge between the private and social returns to research could be upwards of one- to threefold (Jones and Williams 2000; Bloom et al. 2013; Arque-Castells and Spulber 2019).[36]

We also find that 75% of the patents spurred by this program are assigned to U.S. inventors, with the remainder 25% assigned to foreign inventors. To our knowledge, this is one of the first microeconomic estimates of the domestic share of returns to R&D subsidies. Macroeconomists have long used equilibrium models and aggregate data to estimate how much foreign countries benefit from domestic technological advances (e.g., Coe and Helpman 1995; Eaton and Kortum 1999, 2002; Coe et al. 2009). We cannot directly map our estimates to these models. But to the extent that these patents are indicative of a flow of value from the U.S. to foreign countries, this 75/25 split of domestic-foreign returns seems to be on the order of magnitude that one would expect given the macroeconomic evidence.[37]

[36]But as noted in Section **??**, since we cannot observe private investments and downstream product market outcomes – which are relevant for social welfare – we caution against interpreting these results as definitive, but instead present them as plausible upper bounds.

[37]For example, Eaton and Kortum (2002) estimate that when the state of U.S.-based technology increases the relative magnitude of welfare improvements for foreign countries is typically between

Our decomposition of the county-level responsiveness to these investments suggests that policymakers may face a tradeoff with respect to their ability to affect the rate or direction of innovation. The counties that seem to amplify the government's investments the most – those counties with high venture capital penetration – seem to do so by pursuing many different directions for cumulative inventions. On the other hand, there appear to be counties that are less productive in the sense of producing *any* patent, but the patents they do produce tend to be more closely aligned with the technological objectives of the investments. Thus, the optimal distribution of research subsidies may depend on whether the government is intent on spurring either the rate or direction of innovation.

Overall, we have shown that the cost effectiveness of public research grants can vary widely on the extent to which spillovers across technological and geographic space are appreciated. We hope future empirical studies that evaluate the efficiency of R&D subsidies adopt the technology-level unit of analysis to more closely align with the theoretical motivation for these subsidies.

10% to 20% of the U.S. welfare improvements. Similarly, Coe and Helpman (1995) predict that, for G7 countries, roughly 25% of the total benefits of R&D investment (in terms of GDP) are accrued to the country's trade partners.

References

Acemoglu, D., Akcigit, U., Hanley, D., and Kerr, W. (2016). Transition to clean technology. *Journal of Political Economy*, 124(1):52–104.

Aghion, P., Dechezleprêtre, A., Hemous, D., Martin, R., and Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1):1–51.

Aghion, P. and Jaravel, X. (2015). Knowledge spillovers, innovation and growth. *The Economic Journal*, 125(583):533–573.

Agrawal, A., Galasso, A., and Oettl, A. (2017). Roads and innovation. *Review of Economics and Statistics*, 99(3):417–434.

Arora, A., Belenzon, S., and Lee, H. (2018). Reversed citations and the localization of knowledge spillovers. *Journal of Economic Geography*, 18(3):495–521.

Arque-Castells, P. and Spulber, D. F. (2019). Measuring the private and social returns to R&D: Unintended spillovers versus technology markets. *Mimeo*.

Arrow, K. (1962). *The Rate and Direction of Inventive Activity: Economic and Social Factors*, chapter Economic Welfare and the Allocation of Resources for Invention, pages 609–626. Princeton University Press.

Audretsch, D. B. and Feldman, M. P. (1996). R&D spillovers and the geography of innovation and production. *American Economic Review*, 86(3):630–640.

Azoulay, P., Graff Zivin, J. S., Li, D., and Sampat, B. N. (2018). Public R&D investments and private-sector patenting: Evidence from NIH funding rules. *Review of Economic Studies*, 86(1):117–152.

Bell, A., Chetty, R., Jaravel, X., Petkova, N., and Van Reenen, J. (2019). Who becomes an inventor in america? the importance of exposure to innovation. *The Quarterly Journal of Economics*, 134(2):647–713.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics*, 119(1):249–275.

Bloom, N., Schankerman, M., and Van Reenen, J. (2013). Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393.

Bryan, K. A. and Lemus, J. (2017). The direction of innovation. *Journal of Economic Theory*, 172:247–272.

Budish, E., Roin, B. N., and Williams, H. (2015). Do firms underinvest in long-term research? Evidence from cancer clinical trials. *American Economic Review*, 105(7):2044–85.

Coe, D. T. and Helpman, E. (1995). International R&D spillovers. *European Economic Review*, 39(5):859–887.

Coe, D. T., Helpman, E., and Hoffmaister, A. W. (2009). International R&D spillovers and institutions. *European Economic Review*, 53(7):723–741.

Delgado, M., Porter, M. E., and Stern, S. (2016). Defining clusters of related industries. *Journal of Economic Geography*, 16(1):1–38.

Eaton, J. and Kortum, S. (1999). International technology diffusion: Theory and measurement. *International Economic Review*, 40(3):537–570.

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5):1741–1779.

Ellison, G., Glaeser, E. L., and Kerr, W. R. (2010). What causes industry agglomeration? Evidence from coagglomeration patterns. *American Economic Review*, 100(3):1195–1213.

Feyrer, J., Mansur, E. T., and Sacerdote, B. (2017). Geographic dispersion of economic shocks: Evidence from the fracking revolution. *American Economic Review*, 107(4):1313–34.

Forman, C., Goldfarb, A., and Greenstein, S. (2014). Information technology and the distribution of inventive activity. In *The Changing Frontier: Rethinking Science and Innovation Policy*, pages 169–196. University of Chicago Press.

Forman, C. and Zeebroeck, N. v. (2012). From wires to partners: How the internet has fostered r&d collaborations within firms. *Management Science*, 58(8):1549–1568.

Gans, J. and Stern, S. (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.

Griliches, Z. (1981). Market value, R&D, and patents. *Economics Letters*, 7(2):183–187.

Hoberg, G. and Phillips, G. (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5):1423–1465.

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

Igami, M. and Subrahmanyam, J. (2019). Patent statistics as an innovation indicator? evidence from the hard disk drive industry. *Japanese Economic Review*, 70(3):308–330.

Jaffe, A. (1986). Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits, and market value. *American Economic Review*, 76(5):984–1001.

Jaffe, A. B., Trajtenberg, M., and Fogarty, M. S. (2000). Knowledge spillovers and patent citations: Evidence from a survey of inventors. *American Economic Review*, 90(2):215–218.

Jaffe, A. B., Trajtenberg, M., and Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics*, 108(3):577–598.

James, A., Smith, B., et al. (2019). The geographic dispersion of economic shocks:

Evidence from the fracking revolution: Comment. *American Economic Review*, (Forthcoming).

Johnstone, N., Haščič, I., and Popp, D. (2010). Renewable energy policies and technological innovation: Evidence based on patent counts. *Environmental and Resource Economics*, 45(1):133–155.

Jones, C. I. (2005a). Growth and ideas. In *Handbook of economic growth*, volume 1, pages 1063–1111. Elsevier.

Jones, C. I. (2005b). The shape of production functions and the direction of technical change. *Quarterly Journal of Economics*, 120(2):517–549.

Jones, C. I. and Williams, J. C. (1998). Measuring the social return to R&D. *Quarterly Journal of Economics*, 113(4):1119–1135.

Jones, C. I. and Williams, J. C. (2000). Too much of a good thing? The economics of investment in R&D. *Journal of Economic Growth*, 5(1):65–85.

Kantor, S. and Whalley, A. (2019). Research proximity and productivity: long-term evidence from agriculture. *Journal of Political Economy*, 127(2):819–854.

Kline, P. and Santos, A. (2012). A score based approach to wild bootstrap inference. *Journal of Econometric Methods*, 1(1):23–41.

Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017). Technological innovation, resource allocation, and growth. *Quarterly Journal of Economics*, 132(2):665–712.

Lanahan, L. and Feldman, M. P. (2015). Multilevel innovation policy mix: A closer look at state policies that augment the federal SBIR program. *Research Policy*, 44(7):1387–1402.

Lanahan, L. and Feldman, M. P. (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, J. (2000). The government as venture capitalist: The long-run impact of the SBIR program. *Journal of Private Equity*, 3(2):55–78.

Link, A. N. and Scott, J. T. (2010). Government as entrepreneur: Evaluating the commercialization success of SBIR projects. *Research Policy*, 39(5):589–601.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *Review of Economic Studies*, 60(3):531–542.

Myers, K. (2020). The elasticity of science. *American Economic Journal: Applied Economics*, (Forthcoming).

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

Nelson, R. R. (1959). The simple economics of basic scientific research. *Journal of Political Economy*, 67(3):297–306.

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263.

Popp, D. (2002). Induced innovation and energy prices. *American Economic Review*, 92(1):160–180.

Roche, M. P. (2019). Taking innovation to the streets: Microgeography, physical structure and innovation. *Review of Economics and Statistics*, pages 1–47.

Roodman, D., Nielsen, M. Ø., MacKinnon, J. G., and Webb, M. D. (2019). Fast and wild: Bootstrap inference in stata using boottest. *Stata Journal*, 19(1):4–60.

Thompson, P. and Fox-Kean, M. (2005). Patent citations and the geography of knowledge spillovers: A reassessment. *American Economic Review*, 95(1):450–460.

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

Toulis, P. and Kao, E. (2013). Estimation of causal peer influence effects. In *International conference on machine learning*, pages 1489–1497.

Wallsten, S. J. (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.

Williams, H. L. (2013). Intellectual property rights and innovation: Evidence from the human genome. *Journal of Political Economy*, 121(1):1–27.

For Online Publication

Appendices

A Funding Examples and Summary Statistics

A.1 FOA Topic Examples

Figure A.1: FOA Example #1–Solar Energy

(a) FOA Text

2. ADVANCED SOLAR TECHNOLOGIES

Solar energy is our largest energy resource and can provide clean, sustainable energy supplies, including electricity, fuels, and thermal energy. The President's economic recovery package emphasized solar energy, among others, as a key element in combating global climate change. However, the cost-effective capture of the enormous solar resource is problematic. This topic seeks to develop novel, commercially feasible, solar systems and production techniques.

Grant applications submitted in response to this topic should: (1) include a review of the state-of-the-art of the technology and application being targeted; (2) provide a detailed evaluation of the proposed technology and place it in the context of the current state-of-the-art in terms of lifecycle cost, reliability, and other key performance measures; (3) analyze the proposed technology development process, the pathway to commercialization, the large potential markets it will serve, and the attendant potential public benefits that would accrue; and (4) address the ease of implementation of the new technology.

Phase I should include (1) a preliminary design; (2) a characterization of laboratory-scale devices using the best measurements available, including a description of the measurement methods; and (3) a road map with major milestones, leading to a production model of a system that would be built in Phase II. In Phase II, devices suitable for near-commercial applications must be built and tested, and issues associated with manufacturing the units in large volumes at a competitive price must be addressed.

Grant applications are sought in the following subtopics:

a. Manufacturing Tools for Reliability Testing—Grant applications are sought for the development of tools that can be used to conduct reliability testing in PV module manufacturing environments. For example, tools such as light soaking equipment are used to prepare modules or components for accelerated lifetime testing, which is frequently conducted in-house at the module manufacturing facility or by service companies before sending for official third party certification. New tools are needed for the testing of components (e.g., modules, inverters) or subcomponents (e.g., cells, microinverters, individual layers of a module), and should combine high performance, low cost, and a small floor footprint.

Questions – contact: Alec Bulawka (Alec.Bulawka@ee.doe.gov)
James Kern (James.Kern@ee.doe.gov

b. Module and System Manufacturing Metrology and Process Control—The rapid scale-up of the manufacturing of photovoltaics, particularly for new thin-film technologies, is challenging the possibility of using conventional technologies to make real-time non-destructive measurements of material characteristics in high-volume, high-production-rate environments and then using this information to implement real-time process control of the manufacturing process. Therefore, grant applications are sought for the development of novel, advanced, real-time non-destructive materials characterization tools for use in high-volume manufacturing lines for photovoltaic systems.

Questions – contact: Alec Bulawka (Alec.Bulawka@ee.doe.gov)
James Kern (James.Kern@ee.doe.gov

c. Photovoltaics (PV) System Diagnostic Tools—The current rapid growth of the PV industry has led to diverse and innovative product designs, which frequently require non-traditional tests for reliability and performance. Examples of these non-traditional tests include performance testing and tracking requirements for concentrating PV modules, and software-based system diagnostic tools. Grant applications are sought for innovative methods to monitor PV system and component performance, in order to identify failures and loss mechanisms and to minimize system down time. Approaches of interest include the development of diagnostic tools that are process-oriented and internal to the system components, or those that can be integrated – i.e, "piggy-backed" – through ancillary application.

Questions – contact: Alec Bulawka (Alec.Bulawka@ee.doe.gov)
James Kern (James.Kern@ee.doe.gov

(b) Titles of Relevant CPC Classes

Similarity Percentile	Example CPC Titles
100-96	Generation of electric power by conversion of infra-red radiation, visible light or ultraviolet light, e.g. using photovoltaic modules Modulation [of electromagnetic waves]
95-91	Circuit arrangements or systems for supplying or distributing electric power Apparatus for [electric power] conversion ... conversion of DC or AC input power into surge output power...
90-86	Climate change mitigation technologies related to buildings, e.g. housing, house appliances or related end-user applications Control or regulation of electric motors, electric generators or dynamo-electric converters; controlling transformers, reactors or choke coils Electric equipment or propulsion of electrically-propelled vehicles; magnetic suspension or levitation for vehicles; electrodynamic brake systems
85-81	Reduction of greenhouse gas emissions Guiding railway traffic Control systems specially adapted for hybrid vehicles

Notes: Topic #2 from the FY2010 Release 1 Funding Opportunity Announcement. Titles are based on Level 4 of the CPC Scheme ("Groups"), averaging the similarity scores of Level 5 CPC codes.

Figure A.2: FOA Example #2–Geothermal Energy

(a) FOA Text

4. GEOTHERMAL ENERGY TECHNOLOGY DEVELOPMENT

This topic is focused on the development and innovation required to achieve technical and commercial feasibility of EGS. Because of the complexity of these systems, grant applications are expected to focus on a component or supporting technology of EGS development that would enable improvements to the overall system. The unique function and innovation of the targeted subsystem or supporting technology must be clearly described and its function in relationship to the greater EGS system must be expressed clearly. Approaches can be targeted at any of the multi-step project stages for technology development: from design concept, through scale model development (if applicable), to laboratory testing, field testing, and commercial scale demonstrations.

Grant applications are sought in the following subtopics:

a. High Temperature Downhole Logging and Monitoring Tools—Challenging subsurface conditions are one of the barriers to an accelerated ramp-up of geothermal energy generation. To address this challenge, grant applications are sought to develop logging and monitoring tools that are capable of tolerating extreme environments of high temperatures and pressures. The instruments of interest include, but are not limited to, temperature and pressure sensors, flow meters, fluid samplers, inclination and direction sensors, acoustic instruments (high and low frequency), resistivity probes, natural gamma ray detectors, epithermal neutron scattering gauges, rock density gauges (gamma and sonic), casing monitoring devices (e.g. cement bond logs and casing collar locators), fluid conductivity, pH indicators and well dimension probes (caliper). The target temperatures and pressures for these logging and monitoring tools should be supercritical conditions (374° C and 220 bar for pure water), and the tools may be used at depths of up to 10,000 meters.

Questions – Contact Raymond Fortuna, 202-586-1711, raymond.fortuna@ee.doe.gov.

b. Cements for EGS Applications—While conventional geothermal wells experience large temperature rises during production, EGS wells experience large temperature drops at the bottom of the well during the stimulation process, due to the cooling effect of the injected water. This temperature drop may be in the neighborhood of 350°F. This unique situation causes significant stress and potential failure of the cement sheath if conventional cement systems are utilized. To address this issue, grant applications are sought for the research, design, development, testing, and demonstration of a cement system for the high temperature and stress conditions of an EGS wellbore. Proposed approaches may define cement formulations that would be used by the geothermal industry to place the cement within a long string of casings; such approaches should focus on preventing a premature set and maintaining a strong seal at the shoe (so that stimulations may be performed through the casing).

Questions – Contact Raymond Fortuna, 202-586-1711, raymond.fortuna@ee.doe.gov.

c. Drilling Systems—High upfront costs, largely due to high drilling costs, are a major barrier to expanded geothermal energy production in the United States. Therefore, grant applications are sought to reduce drilling costs by developing a drilling technology (horizontal and/or directional) that is capable of drilling three times faster than conventional rotary drilling. Approaches of interest include, but are not limited to the design and development of improved drilling fluids (to reduce frictional viscosity and remove cuttings), high-performance bottom-hole assemblies (e.g., collars, bent subs, drill bits), and downhole motors (to control wellbore orientation). Proposed approaches must demonstrate reliable operation and equipment durability that exceeds the performance of conventional equipment at depths up to 10,000 meters and temperatures up to 300° C.

Questions – Contact Raymond Fortuna, 202-586-1711, raymond.fortuna@ee.doe.gov.

d. Fracture Characterization Technologies—Subsurface imaging is an important part of creating a productive EGS reservoir, which requires visualization before, during, and after creation. In order to advance technology and reduce the upfront risk to geothermal projects, more robust subsurface imaging technologies must be developed. Grant applications are sought to develop improved downhole and remote imaging methods to characterize fractures. Fracture characterization includes prediction of fracture and stress orientation prior to drilling (needed to properly orient horizontal wells); determination of fracture location, spacing, and orientation (while drilling); and determination of the location of open fractures (after stimulation), in order to identify the location of fluid flow pathways within the enhanced geothermal reservoir. Proposed approaches should address robust methods for interpreting and imaging the subsurface, including but not limited to, the development of active or passive seismic, processing software, and joint inversion of geophysical techniques.

Questions – Contact Raymond Fortuna, 202-586-1711, raymond.fortuna@ee.doe.gov.

e. Working Fluids for Binary Power Plants—Binary power plants are rapidly becoming a major part of the geothermal industry, due to increased development of lower temperature geothermal resources. To address cost barriers associated with the working fluids in these binary power plants, grant applications are sought to (1) identify non-azeotropic mixtures of working fluids for improved utilization of available energy in subcritical cycles; (2) characterize the composition and thermophysical and transport properties of those mixtures; (3) identify working fluids for supercritical cycles and trilateral cycles; and (4) characterize the composition, thermophysical, and transport properties of those working fluids. Proposed approaches may address working fluids or mixtures of working fluids with the potential for greater energy conversion efficiency than conventional working fluids, such as isobutane or refrigerants.

Questions – Contact Raymond Fortuna, 202-586-1711, raymond.fortuna@ee.doe.gov.

f. GHP Component R&D—High initial costs have been identified as a key barrier to widespread GHP deployment. To address this barrier, applications are sought to improve GHP components to increase efficiency as well as energy savings as compared to conventional systems. Applications may address but are not limited to: variable-speed (VS) components, advanced sensors and controls (including water flow sensing), electronic expansion valves, heat exchange (HX) design and fluids, system optimization, unit control algorithms, and load management tools.

Questions – Contact Raymond Fortuna, 202-586-1711, raymond.fortuna@ee.doe.gov

g. Innovative System/Loop Designs—One of the main barriers in GHP technology is the high cost of drilling and loop installation. Applications are sought for innovative system/loop designs that reduce the costs of system and/or loop installation, through new design layouts, system components, materials, and/or methods.

Questions – Contact Raymond Fortuna, 202-586-1711, raymond.fortuna@ee.doe.gov

(b) Titles of Relevant CPC Classes

Similarity Percentile	Example CPC Titles
100-96	Geothermal collectors; geothermal systems Earth drilling, e.g. deep drilling; obtaining oil, gas, water, soluble or meltable materials or a slurry of minerals from wells
95-91	Heat-exchange apparatus Pipe-line systems; pipe-lines Preheating, or accumulating preheated, feed-water; controlling water level; auxiliary devices for promoting water circulation within boilers
90-86	Reclamation of contaminated soil Air-conditioning; air-humidification; ventilation; use of air currents for screening Turning; boring [drilling machines and equipment]
85-81	X-ray techniques Removing bark or vestiges of branches; splitting wood Details, components, or accessories for machine tools

Notes: Topic #4 from the FY2010 Release 1 Funding Opportunity Announcement. Titles are based on Level 4 of the CPC Scheme ("Groups"), averaging the similarity scores of Level 5 CPC codes.

Figure A.3: FOA Example #3–Data Management

(a) FOA Text

38. DATA MANAGEMENT AND STORAGE

a. Green Storage for HPC with Solid State Disk Technologies: From Caching to Metadata Servers—Most solid-state storage devices (SSDs) use non-volatile flash memory, which is made from silicon chips, instead of using spinning metal platters (as in hard disk drives) or streaming tape. By providing random access directly to data, the delays inherent in electro-mechanical drives are eliminated. The common consumer versions, known as flash drives, are compact and fairly rugged. Advantages attributed to SSDs include higher data transfer rates, smaller storage footprint, lower power and cooling requirements, faster I/O response times (up to 1000 times faster than mechanical drives), improved I/O operations per second (IOPS), and less wasted capacity.

Furthermore, upcoming processor chip designs from Intel and AMD will include SSD/FLASH controllers built on-board the CPU chip, in order to improve integration for laptop and embedded applications. Such technology is likely to enable a localized checkpoint-restart capability to mitigate increased transient failure rates on future ultra-scale computing systems. This increased level of hardware integration makes it clear that x86 server nodes, which incorporate SSD directly onto the node, are on the horizon.

In view of these developments, the DOE seeks to improve its understanding of the implications of SSDs for large-scale, tightly-coupled systems in High Performance Computing (HPC) environments. Therefore, grant applications are sought to further develop SSD technology as a cost-effective and productive storage solution for future HPC systems, including, but not limited to:

1) **Categorization of SSD failure modes** - The rate of deployment of SSDs in HPC environments will be artificially slowed until a better understanding of the failure modes of this new class of storage is achieved. Proposed approaches should categorize the type of failure (wire bond, cell wear-out, or other failure) and determine how the failures would be detected and/or repaired in a composite device fielded in an HPC environment.

2) **Use of SSD for node-local storage, for faster (localized) checkpoint/restart (CPR)** - If transient failures cause nodes to die, then SSD could be a viable approach for fault-resilience. However, for nodes subjected to hard-failures, the use of SSD could produce an even higher node failure rate, due to the inherent failure characteristics of the SSD; in this case, the SSD approach would not be viable for CPR. Approaches of interest should collect and analyze data on the known failure modes of existing SSD components vis-a-vis node failure modes, in order to determine if SSD presents an effective alternative to the checkpoint/restart of a shared file system.

3) **Use of SSD for scalable out-of-core applications** - Although node-local disk systems have been used to support some applications that use out-of-core algorithms (such as some components of NWChem), the failure rates of spinning disks have rendered this practice unfeasible. Rather, central file systems are used to support these out-of-core applications, greatly affecting their scalability. Approaches are sought to determine whether local SSD might be reliable enough to enable a scalable approach to out-of-core processing.

4) **Use of SSD for metadata servers -** Metadata servers subject disk subsystems to many very small transactions, a feature that is very difficult to support with existing mechanical/spinning-disk based systems. SSDs might respond better to the random-access patterns required for metadata servers, but may not perform as well for write functions. Approaches of interest should analyze the data access patterns of a typical HPC Lustre metadata server and, using an SSD performance model, determine how well an SSD-based system would respond to a metadata server load.

5) **Use of SSD for accelerated caching for the front-end of large-scale disk arrays** – The use of SSDs in caching for large-scale disk arrays is an emerging technology that is not well understood. Approaches are sought to determine of both its performance potential when subjected to real workloads and its fault resilience.

b. Data Management Tools for Automatically Generating I/O Libraries—Database-like, self-describing, portable binary file formats, such as Network Command Data Form (NetCDF) and Hierarchical Data Format (HDF), greatly enhance scientific I/O systems by raising the level of abstraction for data storage to very high-level semantics (of data schemas and relationships between data objects stored) rather than low-level details of the location of each byte of the data stored in the file. However, both NetCDF and HDF5 still rely on very complex APIs to describe the data schema, and many performance pitfalls can arise if the APIs are not used in an optimal manner. Consequently, application developers must invest considerable effort in creating their own "shim" I/O APIs that are specific to their applications, in order to hide the complexity of the general-purpose APIs of NetCDF and HDF5.

Grant applications are sought to develop software tools that not only would enable rapid prototyping of high-level data schemas but also would automatically generate a high-level API for presentation to application developers, thereby hiding the complexity of the low-level NetCDF and HDF5 APIs for managing the file format. Such tools also might use auto-tuning techniques to find the best performing implementation of an I/O method.

c. Integration of Scientific File Representations with Object Database Management Systems—Scientific file formats like Network Command Data Form (NetCDF) and Hierarchical Data Format (HDF5) have capabilities that closely match those of commercial Object Database Management Systems (ODBMS); yet, commercial ODBMSs provide much more sophisticated data management tools than are available to users of NetCDF and HDF5. Unfortunately, ODBMSs are not designed to accommodate parallel writes to the same data entry from multiple parallel writers. Furthermore, database storage formats are opaque and non-portable, and no file standard exists to facilitate the movement of data from one database system to another. By contrast, NetCDF and HDF5 both offer open, standardized formats and portable, self-describing binary formats for storing data represented as Object Databases.

(b) Titles of Relevant CPC Classes

Similarity Percentile	Example CPC Titles
100-96	Electric digital data processing
	Climate change mitigation technologies in information and communication technologies, i.e. information and communication technologies aiming at the reduction of their own energy use
	Wireless communications networks
95-91	Pictorial communication
	Recognition of data; presentation of data; record carriers; handling record carriers
90-86	Speech analysis or synthesis; speech recognition; speech or voice processing ...
	Holographic processes or apparatus
	Printed circuits; casings or constructional details of electric apparatus
85-81	Magnets; inductances; transformers
	Safes or strong-rooms fo valuables; safety transaction partitions

Notes: Topic #38 from the FY2010 Release 1 Funding Opportunity Announcement. Titles are based on Level 4 of the CPC Scheme ("Groups"), averaging the similarity scores of Level 5 CPC codes.

A.2 Summary Statistics

Table A.1: Main Sample Summary Statistics

Panel A: Investment Stocks

	p10	p25	p50	p75	p90	mean	s.d.
Federal Stock, 100-96th	0	0	8,639	39,308	123,399	39,603	76,260
..., 95-91th	0	1,055	13,314	50,143	112,507	36,687	52,857
..., 90-86th	0	2,618	16,229	52,636	99,598	34,907	44,062
..., 85-80th	0	3,944	19,433	52,205	90,233	34,000	38,915
..., 80-76th	68	5,496	21,822	51,334	83,624	33,515	35,486
..., 75-70th	794	6,853	23,590	49,777	79,474	33,255	33,291
..., 70-66th	1,542	8,099	23,671	48,088	75,413	32,346	30,815
...,65-60th	1,831	8,304	23,530	46,046	71,260	31,255	29,096
..., 60-55th	1,784	8,662	23,259	45,098	69,625	30,561	27,956
Federal Stock, 100-55th	21,626	74,174	242,322	490,831	691,585	306,129	258,315
Windfall Stock,100-95th	-101	-11	0	21	164	16	206
..., 95-91th	-131	-24	0	25	159	2	191
..., 90-86th	-132	-27	0	31	173	7	186
..., 85-81th	-134	-29	0	39	193	13	183
..., 80-76th	-128	-29	0	54	219	24	186
..., 75-71th	-133	-30	0	62	227	26	188
..., 70-66th	-129	-32	0	65	227	28	185
..., 65-61th	-129	-31	0	69	233	30	188
..., 60-55th	-128	-31	0	71	232	31	188
Windfall Stock, 100-55th	-315	-25	95	465	807	176	436

Panel B: Patent Flows

	p10	p25	p50	p75	p90	mean	s.d.
Grant Recipients	0.000	0.000	0.000	0.000	0.091	0.121	0.91
U.S. Firms/Inventors	0.000	0.125	0.942	4.242	14.824	9.317	56.00
All Firms/Inventors	0.000	0.325	1.783	7.990	27.490	17.483	95.67
Grant Recip. + 1° Cite	0.000	0.000	0.000	1.000	19.105	36.780	502.25
Grant Recip. + All^o Cite	0.000	0.000	0.250	10.190	73.588	100.818	1,038.21

Panel C: Citation-weighted Patent Flows

	p10	p25	p50	p75	p90	mean	s.d.
Grant Recipients	0.000	0.000	0.000	0.000	0.286	0.787	6.73
U.S. Firms/Inventors	0.000	0.442	5.757	33.900	142.684	125.009	1,077.90
All Firms/Inventors	0.000	1.158	9.255	50.079	203.803	168.675	1,288.26
Grant Recip. + 1° Cite	0.000	0.000	0.000	1.000	19.105	36.780	502.25
Grant Recip. + All^o Cite	0.000	0.000	0.250	10.190	73.588	100.818	1,038.21

Notes: Based on the main estimation sample, which consists of 114,050 observations at the class j year t level, with 8,885 classes and years from 2006 to 2017. This sample is based on the criterion that the class-year observation must have a non-zero federal investment stock (that is, *Federal Stock*, 100-55th must be > 0).

A.3 Major Investment Classes

Table A.2: Top CPCs Invested In, Aggregated

Rank Federal [Windfall]	CPC Title	Rank Windfall [Federal]	CPC Title
1 [1]	G01: measuring; testing	1 [1]	G01: measuring; testing
2 [84]	H01: basic electric elements	2 [7]	C10: petroleum, gas or coke industries; technical gases containing carbon monoxide...
3 [5]	H02: generation; conversion or distribution of electric power	3 [4]	H03: basic electronic circuitry
4 [3]	H03: basic electronic circuitry	4 [21]	F25: refrigeration or cooling; combined heating and refrigeration systems...
5 [69]	H04: electric communication technique	5 [3]	H02: generation; conversion or distribution of electric power
6 [53]	G06: computing; calculating; counting	6 [26]	H05: electric techniques not otherwise provided for
7 [2]	C10: petroleum, gas or coke industries; technical gases containing carbon monoxide...	7 [20]	E21: earth drilling; mining
8 [46]	F16: engineering elements and units...	8 [25]	B65: conveying; packing; storing; handling thin or filamentary material
9 [93]	C12: biochemistry; beer; spirits; wine; vinegar; microbiology; enzymology...	9 [18]	F01: machines or engines in general; engine plants in general; steam engines
10 [107]	B60: vehicles in general	10 [38]	C25: electrolytic or electrophoretic processes; apparatus therefor
11 [106]	F02: combustion engines; hot-gas or combustion-product engine plants	11 [40]	F22: steam generation
12 [13]	B01: physical or chemical processes or apparatus in general	12 [45]	G02: optics

Notes: Displays the top CPC classes (based on 3-digits) based on the amount of direct federal investments from the DoE's SBIR grants (Federal Rank) and the state match windfalls (State Rank) over the full sample period. For comparison, each rank number has the ranking of that same CPC class for the opposite investment in brackets alongside (e.g., a Federal Rank displayed as "1 [10]" indicates that CPC class is the most invested in by the DoE directly, and receives the 10th largest amount of state match windfall).

B Dataframe Construction

This section describes how we convert the patent and SBIR data into an input-output data set. For reference, the level of observations and key variables in the raw data sets are as follows:

1. Patent Record
   - Observation level: Patent–Inventor or Assignee–CPC classes
   - Key variables: year application submitted and granted; inventor and assignee (firm) location; backward and forward citations
2. SBIR Award Data
   - Observation level: Year–Funding Opportunity Announcement (FOA) Topic–Firm
   - Key variables: dollar amount of grant
3. SBIR FOA Data
   - Observation level: Year–Topic
   - Key variables: text of Topic description

The overall flow of the data construction is as follows:

1. Standardize CPC classes
2. Compute similarity of patent abstract text to FOA Topic text
3. Collapse patent-Topic similarity scores to CPC-Topic similarity scores
4. Allocate funds awarded via each Topic (to firms) into each CPC class as a function of CPC-Topic similarities
5. Collapse patent flows to CPC classes

B.1 Choosing a CPC Class Level

The first major decision is choosing a level of aggregation of the CPC. Each CPC code has what is referred to as a "main trunk", which consists of five units of the form:

[1 letter][2 numbers][1 letter][1-3 numbers](/)[2-6 numbers],

i.e., A01B33/08. We could, in theory, use the full code or splice the main trunk at any of the four breaks to generate different levels of aggregation of the hierarchy. To get a sense of the range of aggregation possible, there are nine different 1-digit codes, 128 different combinations for the three digit codes, 662 combinations for the four digit codes, about 10,000 codes if spliced at the "/", and over 220,000 codes if all digits are used. For simplicity, we refer to these as Level 1-5 codes, respectively.

For example, we could group all patents together if they are labeled with the Level 2 code for "Basic Electrical Elements". Or, we could separate these patents out into the fourteen Level 3 codes within, relating to "Cables", "Resistors", "Magnets", etc. Or, we could further separate out, for example, "Magnets" into another seventeen Level 4 codes, each of which has yet another dozen or so Level 5 codes within.

If we aggregate less, then we will have a larger sample size, we will rely less on the idiosyncrasies of the CPC hierarchy. Conversely, the traditional advantage to aggregation is related to the Stable Unit Treatment Value Assumption (SUTVA) necessary to make causal statements from statistical models. The idea is that if the researcher aggregates units together which, for instance, are most likely to experience spillovers or substitution from treatment, then the SUTVA should hold for the newly aggregated set of units.[1]

But because we are intent on identifying the magnitude of across-technology class spillovers, we lean towards less aggregation to preserve sample size. We choose to work with the Level 4 codes of the CPC. We think suitably balances the need to avoid reliance on the CPC hierarchy (Thompson and Fox-Kean 2005), without dividing the data into units so small that patent counts are too rare to prove useful in our analyses. Explorations using Level 3 codes as the main unit of analysis result in qualitatively similar patterns and conclusions.

B.2 Mapping Grants to CPC Classes

Thankfully, all patents are automatically labeled with CPC codes by the UPSTO. Our challenge then is to determine how each SBIR grant maps to a particular CPC code – we need to identify what technologies the government invests in. We leverage the fact that we can connect each SBIR grant to the corresponding FOA that the grant application responded to. The text of these FOAs is the key source of data we

[1] Although, this approach does not permit the researcher to tease out the extent to which treatment effects are "direct" or driven by spillovers across units within these aggregations.

use to, in effect, answer the following: if a grant recipient patented an invention that was in line with the stated goals of the FOA they responded to, what CPC classes would that patent be assigned? If we know the answer to this question, then we will claim that the dollar value of all grants awarded through each FOA has been "invested" in these corresponding CPC classes.

We tackle this prediction problem in a series of steps that include textual similarity analysis and some simple summations and averaging. What follows are the steps of how we assign the grant dollars awarded via each FOA indexed by o, into the relevant CPC classes indexed by j. Following the outline, we motivate and describe the steps in further detail.

1. Estimate the text similarity S between the description of each FOA Topic o and the abstract of each patent p: $S_{op} \in [0, 1]$

2. Calculate the mean similarity $\overline{S}$ for each FOA Topic to each class j based on the classes assigned to each patent, where $\mathcal{P}_j$ is the set of patents assigned class j and $N_{\mathcal{P}_j}$ is the number of patents in that set:

$$\overline{S}_{oj} = \frac{\sum_{p \in \mathcal{P}_j} S_{op}}{N_{\mathcal{P}_j}} \tag{1}$$

3. Adjust $\overline{S}_{oj}$ using Bayesian shrinkage

4. Calculate percentile bins b of $\overline{S}_{oj}$, summing to create $\overline{S}^b_{oj}$, using either:

   - $\overline{S}_{oj}$ values de-meaned at the FOA level, or
   - FOA-specific ranks of $\overline{S}_{oj}$

5. Calculate the twenty ventiles b (five percent groupings) of the $\overline{S}^b_{oj}$ distribution, and assume that below some percentile bin $\overline{b}$ threshold, spillovers do not occur across technology classes.

6. Evenly divide the total amount of SBIR awards given out via each FOA Topic, I_o, to all ventiles b above the percentile threshold $\overline{b}$, to give the topic-class level investment I^b_{oj}

7. Sum I^b_{oj} investments to the b and j level to obtain I^b_j – the total amount invested into class j from awards given via FOA topics that are in percentile b of similarity.

(Step 1) Text Similarity between FOA Topics & Patents: Our key assumption for this exercise is that if a patent abstract and an FOA use the same terminology, and especially if few other documents use that terminology, they are likely referring to the same technologies. This approach of exploiting the similarities between texts to link units of data has become increasingly common in economics. A number of studies leverage this approach to map "scientific space" by comparing the similarity of words used in publication abstracts (Azoulay et al. 2018, Myers 2020) and "product space" by comparing the similarity of words used in product descriptions (Hoberg and Phillips 2016). We follow the norms of modern natural language processing approaches. This includes removing "stop words" (i.e., a, the, and, etc.) and "stemming" words to remove common prefixes and suffixes. We use the commonly employed cosine function to calculate the similarity between text pairs. We also follow norms in using n-grams to identify terms, and weight these terms using the term-frequency-inverse-document-frequency (tfidf) method. We use 1-, 2-, and 3-grams in all of our specifications, which creates terms from all unique 1- 2- and 3-word combinations. Beyond 3-grams, we approached computational challenges given the size of the matrices created.[2]

(Steps 2–3) Averaging and Shrinking Similarity Scores: Ideally, Step 1 would have been to relate FOA text directly to some description of each CPC class. However, the definitions in the CPC scheme tend to be very short pieces of text not well suited to this sort of similarity analysis. Hence, we rely on the patent abstracts. To account for the fact that these CPC-level average scores are generated from a wide range of patent documents (i.e., some CPC classes are assigned to one patent, some to thousands), we employ a standard Bayesian shrinkage estimator that compresses CPC-level means with high variances towards the overall mean.

(Step 4) Accounting for Spurious Text Correlations: Interpreting the cardinality of these scores leans rather heavily on assumptions about linguistic choices across FOAs. To avoid making inferences based on any spurious use of texts across FOAs, we undertake two alternative strategies. One is to demean the similarity scores

[2]To avoid the endogenous use of terminology by patenters or the DoE, we use patents from before our sample, 2001-2004, to estimate the similarity scores.

($\overline{S}_{oj}$) at the FOA level (o) and use the residuals to form the similarity percentiles. The other is to use the FOA-specific ordinal score rankings to determine the similarity percentiles. These approaches remove any variation in similarity connections that might arise purely based on how certain DoE program managers or offices write the text of the Topic descriptions. The demeaning process is our preferred approach as it does not eliminate all of the variation in scores. Though results shown in Appendix G show that using the rank approach yields very similar findings.

(Step 5) Setting the Spillover Threshold: As discussed in Section 5, this sort of assumption is necessary to identify the peer effects as shown by Manski (1993) and Toulis and Kao (2013).

(Steps 6–7) Aggregating to Similarity Bins at the Class Level: Clearly, we need to aggregate investments to the CPC class level because this is the level of our outcome (patent flows). We use the similarity bins to identify heterogeneity in the degree to which across technology spillovers occur. We have good reason to think that, even within the assumed threshold of spillovers, the magnitude of these spillovers is likely to be an increasing function of the similarity between two classes. With these separate bins of investments, we can include multiple stocks in the production function and recover similarity-bin-specific estimates of the returns to investment, which can then be used to quantify the magnitude and shape of spillovers.

As a simple example, consider the following scenario:

- there is one FOA topic (o=1),
- there are one hundred unique CPC classes ($j = \{1, 2, ..., 100\}$),
- one grant of \$50,000 is awarded via the topic (I_o=\$50,000),
- we use two bins of size 50 ($b = \{100 - 51, 50 - 1\}$),
- we set the spillover percentile threshold to $\underline{b} = 51$, i.e., only the classes in $b = \{100 - 51\}$, the most similar 50% of classes, are allowed the possibility of spillovers.

Here, I_j^b =\$1,000 for $b = \{100 - 51\}$ for the 50 classes most similar to the topic and \$0 otherwise, and I_j^b for $b = \{50 - 1\}$ is \$0 for all classes.

C Other Estimation & Data Notes

R&D Stock Construction: We use standard perpetual inventory methods to construct the stock of R&D investments: $K_{jt} = I_{jt} + (1 - \rho)K_{j(t-1)}$ where I_{jt} is the investment flow and ρ is the discount rate.[3] In our preferred specifications we use no discounting, as it gives us the most conservative estimates and also acknowledges the fact that technologies in the future are worth a discounted value at present. All dollar values are scaled used the 2017 Consumer Price Index.

Patent Location Attribution: Most patents include some combination of multiple inventors and/or multiple firm assignees, each with a corresponding location. As a simplification, we divide the patent evenly between the locations of all parties involved. This implies each party's contribution to the value of the invention is equal, each party is a perfect substitute.

Dispersion & Error Terms: In all of our regressions where we estimate the dispersion parameter α instead of assuming it is zero (as in a Poisson regression model), we can very safely reject the null. Thus, we prefer the flexibility of allowing for this non-zero dispersion as a function of the conditional mean. As a sensitivity test, in Appendix G, we report estimates of our main results using models where we assume constant dispersion in the negative binomial model, as well as two versions of Poisson regressions where the error term is assumed to be either additive or multiplicative. All of these alternative models yield very similar estimates.

Predicting Implied Average Marginal Costs: Throughout the analyses we make inferences based on differences in the average marginal cost of producing one additional patent implied by the elasticities we estimate. This metric is inherently policy-relevant (i.e., for comparing the effectiveness of multiple R&D subsidy programs), but also serves the basis of the spillover metrics we create. The two minor challenges when imputing these costs are handling (1) the presence of zero patent flows, and (2) converting the cost of increasing the annual patent flow by one to a more interpretable cost of producing a single patent. Formally, given an estimated elasticity ε, the average marginal cost of increasing the annual patent rate by one is given by the sample average of $(1/Y_{jt}) \times (1/\varepsilon) \times K_{jt}$ (which is undefined if $Y_{jt} = 0$).

[3] Initial conditions are also solved used the traditional approach of assuming steady-state and dividing the average of the earliest investments we observe (in first five years) by the sum of the discount rate d and the pre-sample growth rate, which ranges from 8-10% depending on the types of similarity metrics and thresholds used.

We handle issue (1) by assuming that this left-censoring process occurs at random (because we do not model this process) and substitute the sample mean of Y_{jt} for any zero. We handle issue (2) by scaling this estimate by an additional term, T_t, that is the number of years between 2017 (the final year of our sample) and year t plus one, because we are able to observe every investment flow from the year of investment until 2017.

Estimating Approximate Elasticities with Negative Numbers: The following demonstrates the usefulness of the "demeaning" transformation to handle the fact that our instrumental variable, the stock of state match windfall investments, can take on negative values and we are interested in estimating elasticities directly in the negative binomial regression. For the purpose of clarity, consider the following analogous "log-log" linear regression models Eqs. 2–3.

$$\begin{aligned} \log(Y_j) &= \alpha + \log(X_j)\beta + \epsilon_j, \\ \beta &= \frac{\partial \log(Y_j)}{\partial \log(X_j)} = \frac{\partial Y_j}{\partial X_j}\frac{X_j}{Y_j}, \end{aligned} \tag{2}$$

$$\begin{aligned} \log(Y_j) &= \tilde{\alpha} + \frac{X_j}{\overline{X}}\theta + \tilde{\epsilon}_j, \\ \theta &= \frac{\partial \log(Y_j)}{\partial \frac{X_j}{\overline{X}}} = \frac{\partial Y_j}{\partial X_j}\frac{\overline{X}}{Y_j}, \end{aligned} \tag{3}$$

Below the regression models, which relate the dependent variable Y to the independent variable X, we also define the coefficients of interest: β and θ. Eq. 2 shows the useful result that average elasticities are estimated directly when using the log-log transformation. In the case of Eq. 3, a similar coefficient is estimated, although now instead of estimating a "mean elasticity" given by β, the θ coefficient describes the elasticity across all values of Y_j, but at the sample mean of X_j, denoted here by $\overline{X}$.

In practice, we assume that any difference between these two parameters is negligible. To motivate this assumption, the following shows that θ approximates β. Substituting a first order Taylor series approximation of $\log(X_j)$ around $\overline{X}$, $\log(\overline{X}) + \frac{X_j - \overline{X}}{\overline{X}}$ into Eq. 2 yields:

$$\begin{aligned} \log(Y_j) &\approx \alpha + \Big(\log(\overline{X}) + \frac{X_j - \overline{X}}{\overline{X}}\Big)\beta + \epsilon_j, \\ \log(Y_j) &\approx \big[\alpha + \log(\overline{X})\beta - \beta\big] + \frac{X_j}{\overline{X}}\beta + \epsilon_j, \end{aligned} \tag{4}$$

where $\left[\alpha + \log(\overline{X})\beta - \beta\right] \approx \tilde{\alpha}$, mapping to Eq. 3.

Approximating & Imputing Travel Costs: Although geographic distance separates inventors, the implication of much empirical work on the geographic distribution of invention (e.g., Agrawal et al. 2017) is that the costs of human travel, not geographic distance per se, constrains the flow of ideas. It is beyond the scope of this paper to estimate highly accurate travel costs given the high dimensionality of the data and numerous modes of transportation possible. However, we make new strides in this direction by: (1) focusing on U.S. county-to-county pairs as semi-dense yet computationally tractable set of regions to focus on; (2) using U.S. Internal Revenue Service (IRS) driving mileage rates to approximate driving costs between all counties[4]; (3) using the Department of Transportation (DoT) Airline Origin and Destination Survey (DBIB) to obtain negotiated airfare rates between all U.S. airports; (4) NBER Place Distance Database; and (5) solving for the minimum cost of traveling between each county pair in the U.S. using the minimum of either these approximate costs of driving directly, or driving to the nearest airports and flying.[5] Then, by taking the set of counties where DoE SBIR grants (that focus on a particular technology class) are awarded as the focal set of counties, we can calculate the average cost of making a round trip to these focal counties for individuals in all other counties.

We define the mode of transit based on the proximity of the origin and destination county to an airport. This defines four possible paths: (i) origin county with an airport to destination county with an airport; (ii) origin county without an airport to destination county with an airport; (iii) origin county with an airport to destination county without an airport; and (iv) origin county without an airport to destination county without an airport. For (i), we simply use the average annual market fares reported in DB1B to compute the travel cost. For the paths that include a county without an airport, we add the cost of driving from the center of a county without an airport to the center of the closest county with an airport. We rely on the mileage – as reported by NBER – and the IRS standard mileage reimbursement rate to compute this driving cost. For the total cost, we add the ground and air transportation accordingly. For all cost measures, we adjust for inflation using the 2017 CPI adjusted index.

[4]As well as an estimate of the conversion between geographic distance and driving distance.

[5]DoT data is available at: http://bit.ly/2RSwkG1. NBER data is available at: http://bit.ly/2U6TtHk. IRS data is available at: http://bit.ly/37wuVvl.

Because fares are not observable for all possible airport-pairs, we impute fares using observed data as follows: for airports a and b, we estimate the cost of a round trip flight using the following regression:

$$\text{fare}_{ab} = \alpha_a + \beta_b + \text{geographic distance}_{c(a),c(b)}\delta + \epsilon_{ab}, \tag{5}$$

where α and β are airport fixed effects, $c(\cdot)$ defines the county that an airport is located in, and the parameter δ describes how fares grow with the distance traveled. Then we use the estimated values of α, β, and δ to predict fares both in-sample (where we observe fare_{ab}) and out-of-sample (where we do not observe fare_{ab}, but we do observe at least two fares for a and/or b). These imputed fares are then used in our main analyses.

Figure C.1 plots the joint distribution of these county-level average travel costs and populations. We see that SBIR awards disproportionately are directed to high-population and low-travel cost ("easy to get to") counties relative to the full U.S. distribution. Although, as shown by the similarity in the fitted lines, the travel cost population relationship amongst SBIR-receiving counties is approximately the same as the full distribution, suggesting that awards do not go to counties with low travel costs conditional on their population.

Figure C.1: Joint Distribution of County-level Travel Costs and Population

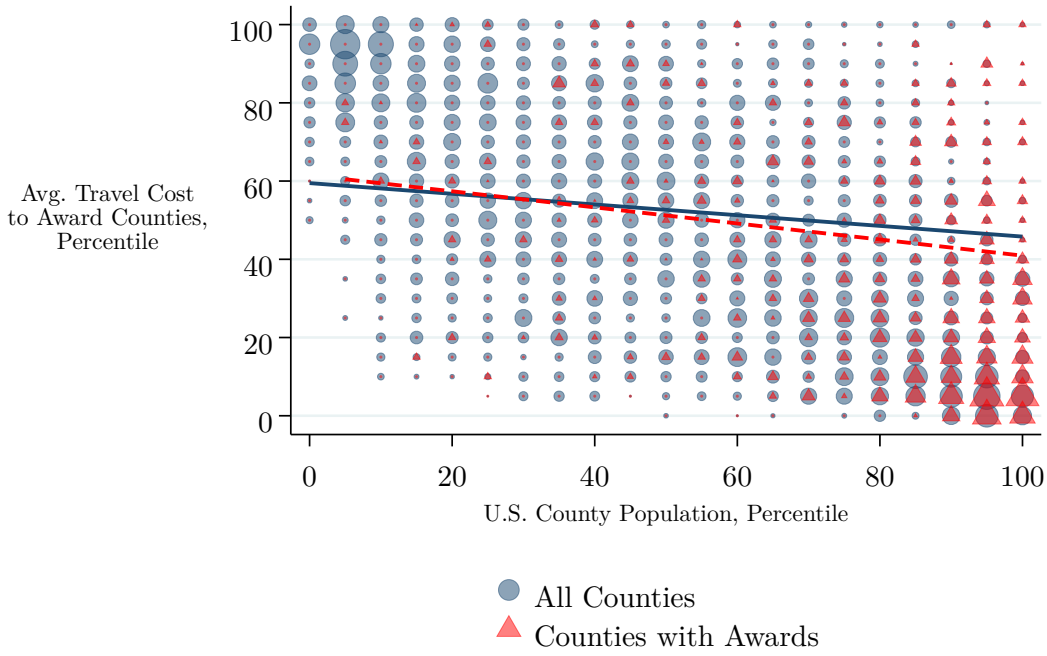

Notes: Plots the joint distribution of county-level (1) average travel cost to all other counties, and (2) population. Marker sizes are weighted by the number of counties; linear fits are plotted (solid blue for all counties, dashed red for counties ever receiving SBIR awards in our estimation sample).

D Comparison of States with and without SBIR Match Policies

To test whether states enacting SBIR grant matching policies tend to be receiving more (or less) federal SBIR investments, we estimate the following regression

$$Y_{st} = \sigma_t + X_{st}\beta + \epsilon_{st}, \tag{6}$$

where st indexes the state-year observations, Y is alternative transformations of federal SBIR investments, and X is either a dummy variable indicating that state s has a matching program in place in year t or is a vector of two dummy variables indicating that there is either a high- or low-match-rate program in place (per splitting the sample of match rates at the median). β describes the extent to which mean investment flows vary as a function of the state's match program.

Table D.1 reports the results of these regressions, exploring federal investments just from the DoE, from the other major federal agencies, and the DoE's share of all federal investment. Across the alternative models, we find no evidence that more funds per capita are flowing into state-years with matching programs.

Table D.1: State Comparison

Dep Var: federal SBIR $\$_{st}$

	log(DoE $ per capita)		log(Other Agency $ per capita)		DoE Share	
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbf{1}\{\text{Any Match}_{st}\}$	-0.459		0.0215		-0.00945	
	(0.279)		(0.219)		(0.0111)	
$\mathbf{1}\{\text{Low Match}_{st}\}$		-0.452		0.226		-0.0219
		(0.342)		(0.225)		(0.0124)
$\mathbf{1}\{\text{High Match}_{st}\}$		-0.439		-0.161		-0.0103
		(0.420)		(0.266)		(0.0182)
N State(s)-Year(t) Obs.	474	474	612	612	612	612
R^2	0.0698	0.0686	0.0358	0.0429	0.102	0.107
mean D.V.	.786	.786	7.392	7.392	.073	.073
Year F.E.	Y	Y	Y	Y	Y	Y

Notes: Standard errors clustered at the state level in parentheses. Any Match is a dummy variable indicating whether a state has a matching policy, with High and Low matches indicating the match policies are above or below the median, respectively. The dependent variable in columns (1-4) is the log-transformed amount of SBIR $ per capita going to firms in each state each year. The "DoE Share" examined in columns (5-6) is the share of all federal SBIR $ going to firms in the state coming from the DoE.

To explore whether the enactment of state matching programs can influence firms'

location decision, we estimate regressions of the form:

$$Y_{ist} = \sigma_t + X_{st}\beta^{f(i)} + [\alpha_i + \delta_s +]\epsilon_{st}, \tag{7}$$

where firm-state-year observations are indexed by ist, Y is a dummy variable indicating that firm i resides in state s in year t, X is again a dummy variable indicating whether state s has a matching program in place in year t, and σ_t, α_i, δ_s are year, firm and state fixed effects sometimes included in the model.

All firms in the sample win an SBIR award at some time from the DoE or any of the other major federal agencies. We can observe firms' locations from the time they first appear in the National Establishment Time Series data onwards, and can match about 80% of ever-SBIR-winners. We allow firms response to vary as a function of whether they ever receive a DoE award or not, as indicated by the $\beta^{f(i)}$ parameter, where $f(i)$ indicates from which agencies the firm i won a SBIR award.

Table D.2 reports the results of these regressions. Regardless of how much we saturate the model (or not) with fixed effects, we find no meaningful association between the movement of these award winners and the state matching policies. This suggests that the firms winning awards in states with matching policies likely did not relocate their firm to the state because of the matching policies or any of the underlying economic or political forces that motivated the enactment of those policies.

Table D.2: Firm Location Analysis

Dep Var: $\mathbf{1}\{$firm i resides in state s in year $t\}$

	(1)	(2)	(3)
Ever non-DoE SBIR$_i$	-0.00908	-0.000353	-0.000332
$\times$ Any Match$_{st}$	(0.00586)	(0.000389)	(0.000404)
Ever DoE SBIR$_i$	-0.00991	-0.00118	-0.00136
$\times$ Any Match$_{st}$	(0.00572)	(0.00131)	(0.00139)
R^2	0.000654	0.0458	0.0458
Year F.E.	Y	Y	Y
State F.E.		Y	Y
Firm F.E.			Y

Notes: Standard errors clustered at the state level in parentheses. The mean of the D.V. is $\frac{1}{50}$ for all models because firms can only reside in one state each year. All models have 5,505,250 firm-state-year obs., based on 1,304 firms that ever receive a DoE SBIR grant (67 of which move) and 11,820 other small businesses that receive an SBIR award from the other four major agencies (662 of which move), only using firm-year obs. including and after the year of founding.

E Instrumental Variable Construction

The total amount invested in each class each year I_{jt} is the sum of federal awards I^f_{jt} and the matching funds awarded by certain states I^m_{jt}. The federal investments may be endogenous, in that they may be correlated with unobservable productivity shocks. And because the state match is a function of federal investments, we cannot simply use the total amount of state matches as an instrumental variable without the same endogeneity concern. Instead, we isolate the residual amount of state match investment that arises due the distribution of SBIR grant winners across states with and without matching policies. To do this, we first note that we can write I^m_{jt} as a function of federal awards:

$$I^m_{jt} = \alpha + I^f_{jt}\gamma_t + W_{jt} \text{ if } I^f_{jt} > 0, \tag{8}$$

where α is a constant and the parameter γ_t captures the fact that all matching programs are a linear transformation of federal awards (i.e., if all have states a 50% matching rate and received the same level of federal investments, then $\gamma = 0.5$).[6] Should we estimate Eq. 8 via OLS, our estimate of γ_t will be the average of match rates across the country in year t, weighted by the amount of I^f_{jt} invested in each state. The residual W_{jt} – which we term the state match windfall – arises because firms working on different technologies are differentially concentrated in states with certain matching programs. Therefore, our key identification assumption is that this distribution of firms and state policies is not related to any unobservable productivity shock. In other words, it cannot be the case that firms working on technologies with larger (smaller) unobservable productivity shocks are concentrated in states with larger (smaller) matching rates.

In general, these windfall residuals are 1-10% of the federal investment amount. Figure E.1 plots the joint distribution of total federal DoE investments and these windfall residuals.

[6] In practice, most states programs operate as lump sum matches. But because the size of Phase I and Phase II awards is largely standardized, this is effectively equivalent to the states setting a match rate.

Figure E.1: Total Federal and State-match Windfall Investments

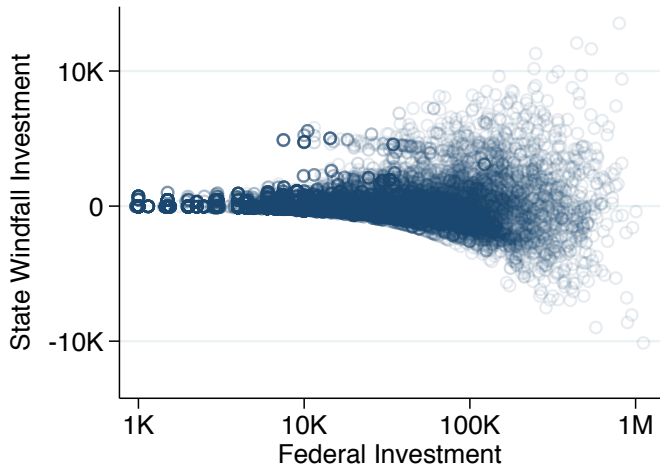

Notes: Scatterplot of federal investments and the resulting state match windfall investments for each class-year jt observation (K = thousands and M = millions of \$).

Other Notes on the Instrumental Variable Approach

We have assumed fixed returns to scale – β does not vary – which may not accurately reflect the impact of marginal investments. If, for reasons perhaps related to fixed costs, the returns to scale are dependent on the level of investment, then the support of our instrument would likely not suffice to identify a richer model of output. Still, we think our approach is a useful step forwards towards understanding the technology-level costs and benefits of these targeted subsidies.

We also note that leveraging these marginal investments from states identifies the marginal effect of increasing the amount of funds granted to SBIR awardees. We do not identify the effect of increasing the number of SBIR awards, but clearly, understanding the effect of changing the sizes of these awards is policy relevant.

It is possible that part of the effect we identify is driven by changes in firms' choices to apply for SBIR grants. If firms are aware of these matching policies, then the expected value of applying is larger, and this may incentivize new applicants and/or different applications from old applicants. This is the same rationale for why we conducted the prior firm movement analysis. We think it is important to show that across-state movers are not responsible for the effects we observe, because that could imply that some portion of the spillovers we observe are due to this flow of

firms and the not flow of new ideas.

However, we are less concerned with entirely excluding this non-mover selection effect that the additional state dollars might induce. If other states were to increase the size of SBIR awards then this policy change would result in similar selection effects where firms previously on the margin may start to apply. Certainly, if the federal program made a large, publicized increase in the size of federal awards, then other general equilibrium forces may become relevant and more significantly change the pool of applicants compared to the null response we observe due to the state programs. So although we cannot rule out the possibility that a portion of the productivity effects we identify are based on changes in the composition of firms applying for grants, any policy that increases the size of these grants would suffer or benefit from this same compositional change.

It is also important to point out that we cannot address the selection effect that constrains us to only examine the productivity of investments conditional on the DoE investing non-zero amounts into a CPC class. We lack an instrument that influences which CPC classes receive any federal investments, but not the amount of investment. Still, our results are useful for understanding the effects of increasing the size of SBIR grants, holding fixed the processes by which the DoE chooses topics to pursue.

Lastly, because we do not have actual data on the dollar amount of state matching funds that were acquired by firms in practice, our instrumental variable approach should be viewed through the lens of an intent-to-treat analysis. Still, we do not have good reason to think that using the policy rules to estimate the size of these windfalls would severely overestimate the flow of state funds because our discussions with state managers largely suggests the transaction costs of obtaining match funds tend to be very low relative to their value.

F Spillover Threshold Search

As described in the main text, we must make an assumption about the similarity threshold between technology classes and an FOA's objectives beyond which investments via the FOA will not influence patent flows in the class. Formally, we must constrain the set of similarity investment bins $\mathcal{B}$ to not include all possible bins (i.e., investment from each FOA cannot be allowed to influence all CPC classes).

Recall the main estimating equation

$$\mathbb{E}[Y_{jt}^d \,|\, \mathbf{X}_{jt}] = \exp\Big(\sum_{b\in\mathcal{B}} \frac{K_{jtb}}{\overline{K}}\theta_b^d + \sigma_t^d\Big), \tag{9}$$

where $\mathcal{B}$ is the set of similarity bins within the assumed boundary of technological spillovers, and $b = \{100-96, 95-91, ..., 5-1\}$ index these bins, such that $K_{jt,b=100-96}$ describes the stock of (state match windfall) investments in class j at time t originating from FOA topics where class j is in the top 5% of similarity scores.

To search for the point at which spillovers stop (θ_b^d can safely be assumed to be 0), we estimate a series of models where we successively increase the size of $\mathcal{B}$ to include CPC classes less and less similar to each FOA. That is, we first estimate a model where only $b = \{100-96\}$, yielding a single θ_b^d parameter. Then we expand $\mathcal{B}$ so that $b = \{100-96, 95-91\}$, yielding two θ_b^d parameters, etc.

Figure F.1 plots the vector of θ_b^d based on regressions up until the lowest similarity bin we explore is $b = \{45-40\}$ (the markers for coefficients from each model are shaded the same). First, we see that the effect of high-similarity investments becomes much larger as less-similarity investments are included in the regression. In other words, not conditioning on investments into "nearby" classes drastically underestimates the effect of investments. This pattern of negative (technological) spatial correlation is consistent with the DoE *not* investing in areas where other investments are likely to have spillovers. This stands in contrast to typical network behavior where agents (like private profit-maximizing firms) typically have incentives to generate positive spatial correlations in their investments in order to capitalize on output.

For the purposes of making our assumption about the boundary of spillovers, we note that starting near the similarity bins $b = \{55-50\}$, we cannot consistently reject a null of no effect. This boundary appears to be the same whether counting patent flows from only grant recipients, or all firms and inventors. Furthermore, the

change in the coefficients on the high-similarity bins diminishes significantly at this point. These patterns motivate us to use $b = \{55 - 50\}$ as the boundary of $\mathcal{B}$ in all of our preferred specifications. In Appendix G we report results using alternative boundaries near $b = \{55 - 50\}$ and find qualitatively similar results.

Figure F.1: Technological Spillover Search

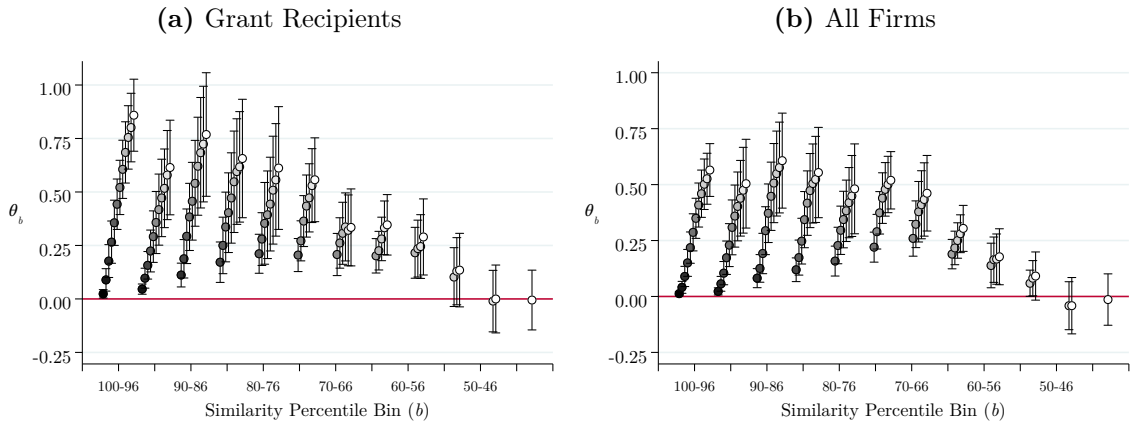

Notes: Plots the θ_b^d coefficients for d =“Grant Recipients” and d =“All Firms” using successively larger sets of investment bins. Coefficients from the same regression model are shaded the same color.

G Alternative Specifications, Robustness Tests, & Additional Results

G.1 Permutation Placebo Tests

We follow Bertrand et al.'s 2004 approach to generating a null distribution of effects (elasticities) to use in the calculation of p-values. As Bertrand et al. (2004) note, in non-experimental settings where the (quasi-)randomization of policies – here, the assignment of match policies to state-year observations – is unknown, the researcher must make assumptions about the policy generating process. In the spirit of Fisher's exact test we generate permutations by randomly reassigning the observed distribution of policies.[7] Formally, we approximate the null distribution as follows:

1. Index the set of policies observed across our sample of state-year observations, with each policy characterized by the Phase I and Phase II match amount awarded to any SBIR grant recipients in that state-year

2. Randomly re-assign policies across the sample of state-years

3. Construct placebo firm-level grant amounts per the sum of the true federal grant amounts and the placebo state match amounts

   Note: For example, if the randomization process replaced Texas's 2010 policy of no matching program with North Carolina's 2007 policy of a $50,000 Phase I match and a $100,000 Phase II match, then any grant recipient located in Texas in 2010 would receive a placebo match of either $50,000 or $100,000 depending on its Phase.

4. Construct placebo CPC-year level investment flows from the placebo grant amounts

5. Estimate placebo state match windfall amounts

6. Construct CPC-year level stocks of placebo windfalls

[7] The implicit assumption here is that we are certain that the distribution of observed match policies would occur in this sample, but some unobservable random process decided which state would adopt which policy in a particular year.

7. Estimate the main production function using the placebo windfalls and record the estimated coefficient

   Note: Due to the complexity of re-constructing our entire data set for each permutation, we focus on the simpler cases of estimating the model at the two geographic boundary cases which do not decompose technological spillovers. This corresponds to the elasticity estimates reported in Table 1 for Grant Recipients (θ=0.485) and Worldwide Inventors (θ=0.384).

8. Repeat (2–7) 2,500 times to approximate the null distribution

Then we define our p-values as the share of the null distribution that is more extreme than our main estimate obtained from the observed distribution of match policies. Unlike many traditional permutation tests, our empirical null distribution is not symmetric around zero because the permutations manipulate investment flows and the data construction process creates stocks of these permutated flows. Thus, each stock has a non-zero probability of containing an investment flow from a state-year that was randomly assigned its true policy. Figure G.1 displays the results of the

Figure G.1: Placebo Tests at Geographic Boundaries

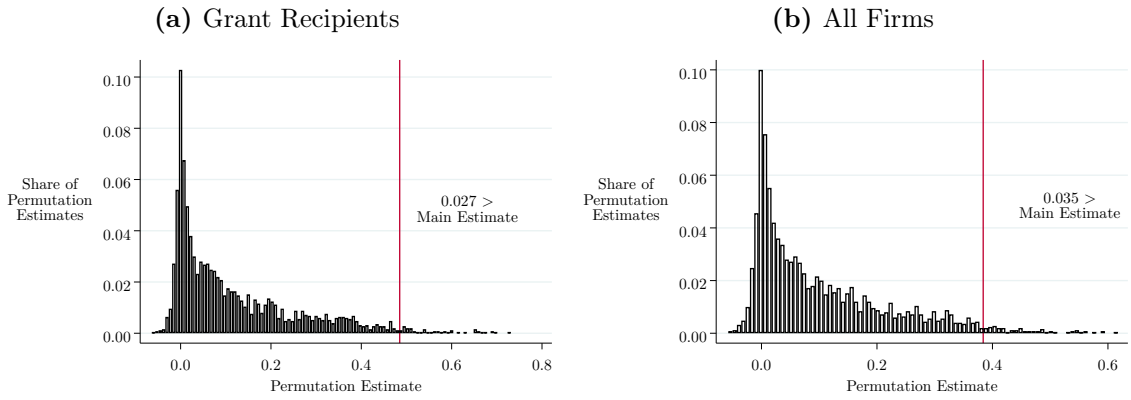

Notes: Plots the approximate null distributions based on 2,500 permutations of the observed match policies. Vertical red lines indicate the main estimates per the true observed match policies.

placebo tests and shows that in the case focusing only on grant recipients, only 2.7% of the estimates from the 2,500 placebo tests were larger than our main estimate. Similarly, only 3.5% of the placebo estimates were larger than our main estimate when counting patents from inventors worldwide. These results reject the null at

conventional levels, and give us confidence that the effect of the windfall estimates is unlikely to be due to chance assignment of these policies to particular states in particular years.

G.2 Alternative Specifications

The following tables replicate versions of the main instrumental variable results using alternative estimation methods or different assumptions in the data construction process (i.e., discount rates, similarity score adjustments). The table notes provide further detail.

Table G.1: Output Elasticities and Implied Costs using Alternative Count Models for Patent Counts at Boundary Cases

	Grant Recipients			Worldwide Inventors		
	(1)	(2)	(3)	(4)	(5)	(6)
θ	0.485	0.479	0.428	0.384	0.398	0.382
	(0.115)	(0.203)	(0.065)	(0.083)	(0.134)	(0.053)
$\log(\alpha)$	1.699			1.239		
	(0.247)			(0.0675)		
N	114,050	114,050	114,050	114,050	114,050	114,050
Spec.	Neg. Bin.	Poisson	Poisson	Neg. Bin.	Poisson	Poisson
Dispersion	Mean	0, Add.	0, Mult.	Mean	0, Add.	0, Mult.
Method	MLE	MLE	GMM	MLE	MLE	GMM
I.V.	Y	Y	Y	Y	Y	Y

Notes: Replicates the results from the main table of results for the geographic boundary cases, including alternative count data regression models. Cols. (1) and (4) replicate Column (3) from Panels A and C of Table 1, respectively, here also reporting the dispersion parameter α. "Poisson" with "0, Add." indicates a Poisson regression model with an additive error term of the generic form: $Y = \exp(X\beta) + \epsilon$. "Poisson" with "0, Mult." indicates a Poisson regression model with a multiplicative error term of the generic form: $Y = \exp(X\beta)\epsilon$. "MLE" models are estimated via maximum likelihood, "GMM" models are estimated via generalized method of moments. Standard errors clustered at the aggregated CPC level reported in parentheses.

Table G.2: Output Elasticities and Implied Costs using Alternative Discount Rates for Patent Counts at Boundary Cases

	(1)	(2)	(3)	(4)
		Panel A: Grant Recipients		
Elasticity (θ)	0.485	0.954	0.800	0.464
	(0.115)	(0.144)	(0.132)	(0.100)
\$M per +1 Patent	0.787	0.689	0.462	0.455
	[0.54 - 1.47]	[0.53 - 0.98]	[0.35 - 0.68]	[0.32 - 0.79]
		Panel B: Worldwide Inventors		
Elasticity (θ)	0.384	0.832	0.679	0.347
	(0.083)	(0.088)	(0.108)	(0.109)
\$M per +1 Patent	0.170	0.127	0.088	0.099
	[0.12 - 0.30]	[0.11 - 0.16]	[0.07 - 0.13]	[0.06 - 0.26]
N	114,050	114,050	114,050	114,050
Discount Rate	0.00	0.05	0.10	0.20
Investment	Windfall	Windfall	Windfall	Windfall

Notes: Bootstrapped standard errors clustered at the aggregated CPC level and years reported in parentheses. "\$M per +1 Patent" presents the average marginal costs implied by the elasticity estimates, in millions, with 95% confidence interval bounds (based on the standard errors of the elasticity) in brackets.

Figure G.2: Implied Marginal Costs at Boundary Cases using Alternative Similarity Methods

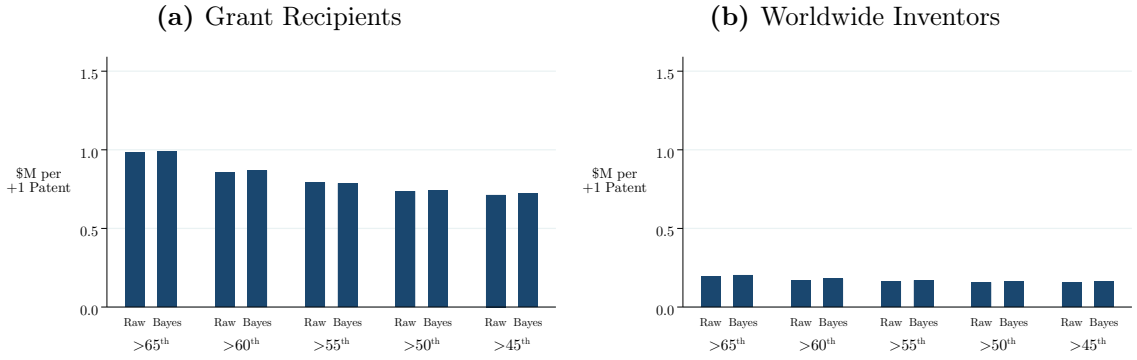

Notes: Plots the implied average marginal costs of producing a patent at the geographic boundary cases of counting only patents produced by grant recipients or counting all patents. Each bar (within each panel) corresponds to an alternative approach used to generate the final data set. "Raw" and "Bayes" refer to whether or not a Bayesian shrinkage approach was used when collapsing the patent-FOA similarity scores to the class-FOA level. The cumulative percentile numbers indicate the class-FOA similarity threshold after which spillovers are assumed to be zero. The preferred approach used throughout the main text is "Bayes" and "$> 55^{\text{th}}$".

G.3 Three Dimensional Plot

Figure G.3 is a 3D version of Figure 6.

Figure G.3: Cumulative Marginal Patent Cost per Geographic and Technological Spillovers

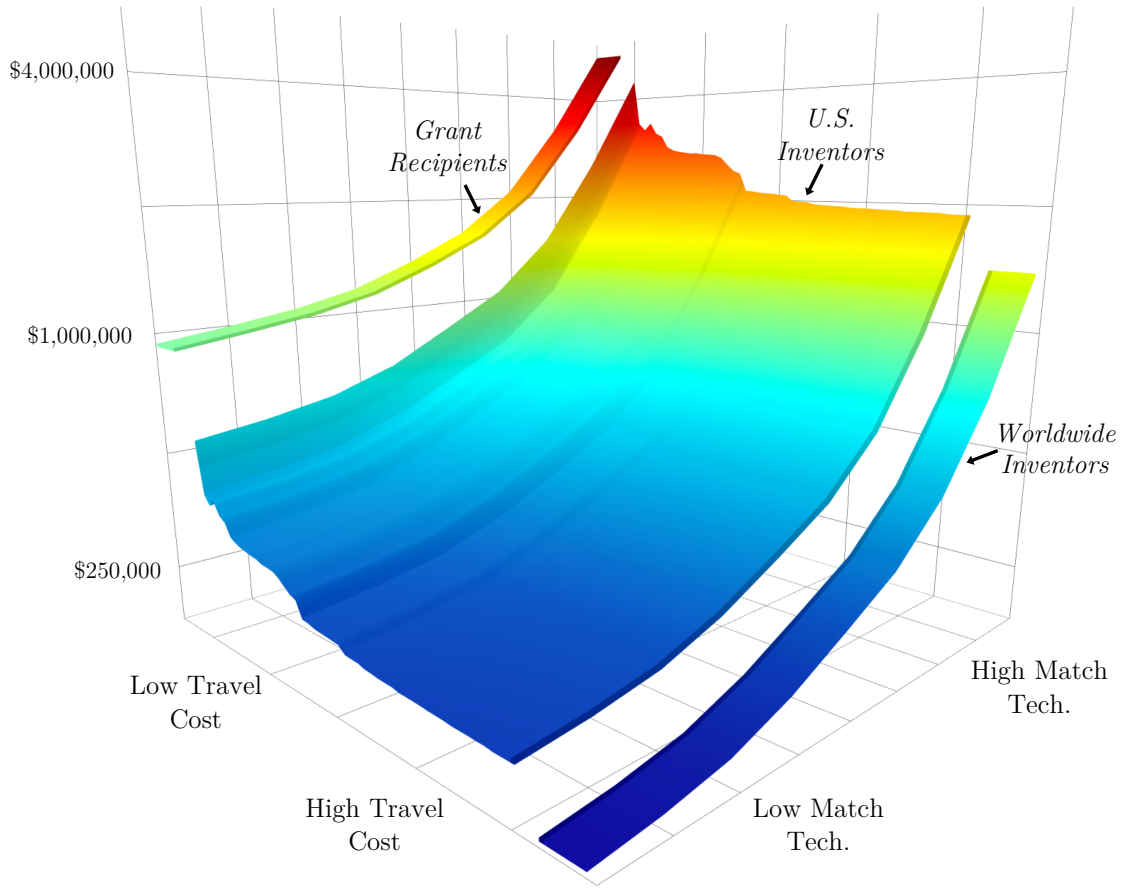

Notes: High and Low Match Tech. refers to higher and lower values of technological similarities.

G.4 Psuedo-Control Function Approach

Ignoring spillovers momentarily, recall that the main estimating equations are of the form:

$$\mathbb{E}[Y_{jt} \,|\, \mathbf{X}_{jt}] = \exp\Big(\frac{K_{jt}}{\overline{K}}\theta + \sigma_t + \omega_{jt}\Big), \tag{10}$$

where we assume that when state match windfall investment is used for K, then there is no correlation between K and any unobserved state variable(s) ω that are also correlated with patent flows Y. Our approach to recovering the windfall investments does not take a strict stance on the structure of ω (e.g., whether it is a single scalar variable or a vector, etc.).

Motivated by the control function approach to estimating production functions,[8] we investigate the validity of this assumption. Control function approaches leverage assumptions about the structure and timing of unobservable shocks (here, ω) to motivate the use of certain observable variables (e.g., intermediate inputs) as proxies for these correlated shocks.

We do not write down an explicit model of these information shocks, but consider the following: let us assume that ω_{jt} is in fact a single scalar variable. Furthermore, let us assume that, in any given year, firms and DoE policymakers make expectations about all future values of this shock. In this model, firms' and policymakers' decisions (per their expectations) in a given year determine the DoE's investment flows, and therefore the stock of federal investments. Thus, some portion of the investment flow in class-year jt is determined by the expected realization of $\omega_{j(\tau>t)}$. This line of reasoning motivates the use of (a flexible transformation of) the stock of federal investments as a proxy for the realization of ω_{jt}.

If (i) our assumptions about the windfall investments being uncorrelated with any unobservable variables that co-determine patent flows, and (ii) these unobservable variables can be sufficiently captured by flexible transformations of federal investment levels, as motivated by the logic above, then including these transformations in the main regression should yield similar estimates of the focal elasticity parameter θ. Substituting the stock federal investments for ω_{jt} in Eq. 10, we estimate models of the form

$$\mathbb{E}[Y_{jt} \,|\, \mathbf{X}_{jt}] = \exp\Big(\frac{K_{jt}^{\text{windfall}}}{\overline{K}^{\text{windfall}}}\theta + \sigma_t + f(K_{jt}^{\text{federal}})\Big), \tag{11}$$

[8] As initially developed by Olley and Pakes (1996).

where f is a flexible function such as a polynomial or restricted cubic spline.

Because we are not formally modeling the relationship between K_{jt}^{federal} and ω_{jt} and the functional form of the model has changed, we estimate versions of Eq. 11 not to see if we can recover the exact same θ parameter with or without including $f(K_{jt}^{\text{federal}})$ in the model, but rather to ensure that upon including this term, there is still identifying variation left in K_{jt}^{windfall} that yields reasonable parameters with reasonable standard errors. This is why we refer to this exercise as a psuedo-control function approach.

Table G.3 presents the results of estimating 11 using a variety of functions in place of f. Compared to our main specifications where no pseudo-control variables are included (Col. 1), the inclusion of the pseudo-control functions yield θ estimates that are smaller, but in all cases are statistically significant using traditional conventions. This pattern holds for both of our geographic boundary cases. We take this as evidence supporting our key assumption that the variation in state match windfall investments is uncorrelated with any unobservable (to us) shocks that co-determine patent flows.

Table G.3: Output Elasticities and Implied Costs via Psuedo-Control Function at Boundary Cases

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				A: Grant Recipients			
θ	0.485	0.159	0.178	0.188	0.183	0.174	0.162
	(0.0613)	(0.0635)	(0.0368)	(0.0693)	(0.0681)	(0.0670)	(0.0683)
Federal Stock Control(s)		3.227	0.705	n.r.	n.r.	n.r.	n.r.
		(0.554)	(0.0988)				
				B: Worldwide Inventors			
θ	0.384	0.119	0.186	0.132	0.132	0.120	0.112
	(0.0531)	(0.0421)	(0.0371)	(0.0502)	(0.0494)	(0.0473)	(0.0502)
Federal Stock Control(s)		2.784	0.387	n.r.	n.r.	n.r.	n.r.
		(0.371)	(0.0498)				
N	114,050	114,050	114,050	114,050	114,050	114,050	114,050
Control Function		Linear	Log	Poly., 3	Poly., 4	Spline, 3	Spline, 4
I.V.	Y	Y	Y	Y	Y	Y	Y

Notes: n.r.–not reported. Control Function transformations of the federal investment stock are as follows: "Linear" is the stock level (in $ millions); "Log" is the logarithm; "Poly., 3/4" is a polynomial raised to the third or fourth degree, respectively; "Spline, 3/4" is a restricted cubic spline with either three or four dimensions (that is, four or five knots), respectively. Standard errors clustered at the aggregated level-3 CPC class reported in parentheses.

G.5 Test Statistics for Two-way Spillovers

Figure G.4 plots the score bootstrap z statistics from the negative binomial regressions underlying the implied cost estimates plotted in Figures 6 and 7. These are generated using the wild score bootstrap method (Roodman et al. 2019), clustering at the aggregated level-3 CPC class and year level. Recall, the variation across values of the x-axis in these figures are generated by the different coefficients recovered on the similarity-bin investment stocks–test scores in the same "row" come from the same regression but different independent variables. Variation across values of the y-axis in these figures are generated by the different coefficients recovered on the the similarity-bin investment stock using different sets of patents in the dependent variable–test scores in the same "column" come from different regressions but the same independent variable. Unlike in Figures 6 and 7, which plot cumulative costs and shares of net output, there is no cumulative nature to the z statistics.

The pattern of z statistics suggests that there is more uncertainty about the impact of investments on technologies very close and very far from a given FOA's objectives, relative to that for technologies in the middle of the similarity score distribution.

Figure G.4: z-statistics for Two-way Spillovers Results

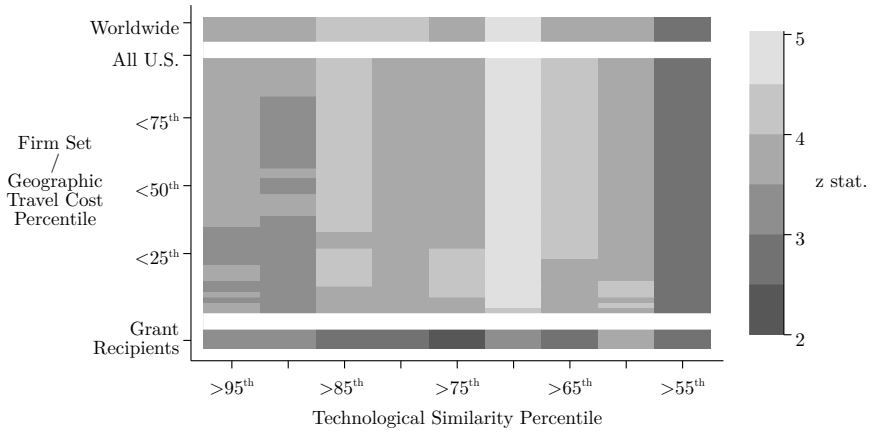

Notes: Plots the z statistics for the coefficients underlying Figures 6 and 7.