

Islands of Innovation: Defense Microelectronics Contracting and the Limits
of Geographic Spillovers

Daniel Ju

PSCI 3192W

Professor. White

1. Introduction

Semiconductors occupy a central position in contemporary struggles over economic growth, national security, and technological leadership. From data centers that train large language models to missile guidance systems and other elements of critical infrastructure, advanced microelectronics function as indispensable inputs across both civilian and military domains. Recent legislative efforts, most notably the CHIPS and Science Act, reflect an increasingly common view in Washington that semiconductor production and innovation now resemble strategic infrastructure rather than just a conventional commercial industry (136 Stat. 1637).

The historical roots of this perception run deep. The Department of Defense (DoD) played a formative role in shaping the early trajectory of the U.S. semiconductor sector, with programs such as Apollo and Minuteman generating reliable early demand for integrated circuits. These initiatives helped bind defense laboratories, research universities, and private contractors into what scholars have described as a “triangular alliance,” widely regarded as a central engine of postwar technological progress. The underlying logic of this arrangement suggests that mission-driven public research and development (R&D) can cultivate private innovation by producing new technical knowledge, expanding the scientific labor force, and reinforcing localized innovation ecosystems.

However, the geographic reach of these benefits remains unclear. Much of the existing evidence on public R&D focuses on national aggregates or broad regional patterns, leaving unresolved whether innovation gains diffuse beyond the specific organizations or cities that receive federal investment. In semiconductors, this uncertainty is particularly strong, as the

industry is characterized by massive capital requirements (Platzter and Sargent Jr. 9) and a reliance on highly specialized labor pools, structural features that historically favor established hubs over new locations. Consequently, policymakers face a risk that public subsidies may merely reinforce these existing centers of dominance rather than broadening the spatial distribution of U.S. innovation - a tension central to the implementation of the CHIPS and Science Act.

This tension drives the core inquiry examined here: whether defense oriented microelectronics investments generate observable innovation spillovers for nearby private inventors that do not themselves receive DoD funding. If public R&D produces locally embedded knowledge that diffuses through collaboration, labor mobility, or firm networks, increases in nearby private patenting should be detectable. If, instead, the sector's structural concentration dominates, these effects may remain confined to existing hubs.

To evaluate this possibility, the paper proceeds as follows. Section 2 formalizes the research inquiry by defining the puzzle of localized spillovers in a structurally concentrated industry. Section 3 reviews the existing scholarship on defense innovation and regional clustering, contrasting optimistic theories of public R&D spillovers with evidence of persistent geographic concentration in semiconductors. Building on this literature, Section 4 develops a theoretical framework centered on project-level knowledge flows and regional labor markets to generate testable hypotheses. The subsequent sections outline the empirical strategy and present results from a newly constructed county-year dataset assessing the impact of DoD

microelectronics R&D on private, non-recipient semiconductor patenting, followed by a discussion of the findings.

2. Research Question

Building on the preceding discussion, this section specifies the relationship under examination. Public R&D is often associated with downstream innovation, yet the semiconductor industry's extreme geographic concentration creates uncertainty about whether such gains extend beyond the counties that directly receive federal support. The guiding question of this paper is therefore:

Does increased DoD microelectronics R&D funding generate localized spillover effects that raise semiconductor patenting by private, non-recipient inventors in nearby counties?

Conceptually, the outcome of interest is county level semiconductor patenting by private inventors, while the exposure of interest is DoD microelectronics R&D obligations over the subsequent one to three years. Throughout this study, "microelectronics R&D" refers to federal funding obligations targeting integrated circuits and related hardware, while "semiconductor patenting" measures the resulting commercial invention by private industry. The question operates along two related dimensions. First, it asks whether counties receiving defense microelectronics funding experience subsequent increases in private inventive activity. Second, it considers whether nearby counties that do not receive such funding exhibit comparable changes, which would signal the presence of spatial spillovers.

These alternatives reflect competing intuitions about how public R&D operates in highly concentrated industries. Spillovers may emerge if public investment expands local knowledge stocks,

scientific labor markets, or collaborative networks. In contrast, where clustering is entrenched and capital barriers are high, innovative activity may remain confined to established hubs. To contextualize these competing expectations, the following section reviews the existing scholarship on defense innovation, public R&D spillovers, and regional clustering. By synthesizing what is already known about how publicly funded research shapes private inventive activity, the literature review provides the foundation for developing the theoretical mechanisms and hypotheses evaluated in the empirical analysis.

3. Literature Review

3.1 Public R&D, Defense Innovation, and Early Microelectronics

Public investment in R&D has been central to technological change in the United States, particularly in early microelectronics, where mission-driven defense R&D played a decisive role in expanding local stocks of technical knowledge (Mowery, “Military R&D” 1236-1237). During the Cold War, military and space programs generated an unprecedented surge in demand for reliable integrated circuits (IC), accelerating the industry’s early growth. The Apollo space program and the Minuteman intercontinental ballistic missile program illustrate this dynamic - in the early 1960s, these initiatives “bought almost the entire supply of chips for several years,” effectively jumpstarting the initial scaling of the microelectronics industry (Tomayko 131). This demand linked federal agencies, university laboratories, and private firms into a tightly coordinated innovation system in which mission requirements shaped both research agendas and production capabilities. The resulting structure closely resembles the “triangular alliance” among government, academia, and industry

identified by Goodman and Gerstel, in which public procurement and research contracts jointly advanced scientific knowledge and industrial capacity (5).

Procurement policy further reinforced this process - the introduction of military “second source” requirements during the Cold War compelled competing firms to exchange design and process knowledge, pushing semiconductor technologies beyond early experimentation and into commercial viability (Mowery, “Military R&D” 1248). Defense contracts funded advances in diffused transistors and ICs, while associated technical assistance and inter firm collaboration standardized production practices and proliferated process knowledge. Between 1958 and 1974, the federal government used “production preparedness” contracts to subsidize manufacturing capacity, channeling approximately \$930 million into semiconductor R&D (Linville and Hogan 1011). Collectively, this literature establishes a foundational expectation for later works: that defense oriented public R&D tends to crowd in private innovation rather than displace it.

3.2 Geographic Spillovers and Regional Clustering

While the functional relationship between public and private R&D suggests positive spillovers, the spatial distribution of these effects remains contested. Though public R&D can stimulate innovation, its effects are often geographically uneven. A 2019 OECD analysis of over 21,000 research institutions across 36 countries found that the proximity of industry to these institutions is strongly associated with higher rates of corporate patenting, where nearly 69 percent of all inventive activity by industry in the sample occurred within just 20 kilometers of a research institution (Paunov et al. 22). This evidence aligns with classic theories of agglomeration and regional

innovation clusters, which argue that the spatial concentration of skilled labor, specialized suppliers, and research institutions generates self-reinforcing local ecosystems (Paunov et al. 11). Clustering also characterizes contemporary defense-funded innovation in semiconductor adjacent fields. By examining the DoD's Small Business Innovation Research (SBIR) awards in artificial intelligence (AI), Chowdhury et al. find that funded projects are disproportionately concentrated in states hosting top tier research universities (772). Although focused on AI rather than microelectronics, this pattern suggests that defense R&D programs often reinforce existing research ecosystems rather than disperse innovative activity geographically.

At the same time, this clustering dynamic raises concerns about the spatial reach of innovation gains. Geographic proximity may amplify the returns to public R&D within existing hubs while doing little to extend innovation outward. In the semiconductor industry, this pattern is especially pronounced - National Science Foundation data shows Santa Clara County alone accounted for roughly one-fifth of all U.S. semiconductor-related patents granted between 2000 and 2020 (Robbins and Cosby 1). Institutional evaluations of federal innovation programs echo this uneven spatial diffusion. A comprehensive National Research Council assessment of the SBIR program found that, as of 2008, 54 percent of SBIR grants were captured by just 6 states, while the lowest winning eight states accounted for just 1 percent (99). Despite SBIR's explicit goal of expanding innovation beyond incumbent contractors (Council 159), its geographic footprint remains shaped by preexisting regional capacity, raising questions about whether spillovers extend meaningfully to nearby but less technologically embedded counties while amplifying pre-existing ecosystems.

3.3 Institutional Mechanisms of Defense-Funded Spillovers

Beyond their historical role in early microelectronics, specific procurement rules reflect a broader strategy to shape market structure. Second-source requirements illustrate this clearly: by requiring firms to externalize process knowledge so that alternative suppliers can replicate production, the government embedded diffusion directly into procurement contracts. Historically, this strategy drove rapid industry-wide learning, with cumulative output doubling associated with a 20 to 30 percent decline in production costs (Mowery, “Federal Policy” 166). Modern reforms to solicitation rules operate on a similar logic. Howell et al. find that shifting to more open solicitation designs in Air Force SBIR programs lowered entry barriers for nontraditional firms, raising the probability of subsequent venture capital investment by 12 percentage points and patenting by 8.9 percentage points (3608). These procurement rules, however, operate within broader institutional environments that actually condition how effectively such diffusion occurs.

Organizational capacity within defense agencies further conditions these outcomes. Studies emphasize that government officials are not passive administrators but active shapers of innovation. Decarolis et al. show that the loss of buyer expertise - measured through the unexpected death of a program manager - reduces patenting per contract by 32.3 percent (697). This importance of management extends to the structure of research communities themselves. Fuchs’ analysis of DARPA shows that program managers deliberately “re-architect social networks” by mandating early-stage knowledge sharing across firms and universities, accelerating diffusion beyond initial project boundaries (1144–45). Unlike second-sourcing, which mandates the transfer of finished production

specs, this approach forces upstream collaboration among otherwise disconnected researchers. Beyond formal procurement and program design, diffusion can also occur through less centralized channels that bypass contracting structures altogether.

Finally, labor mobility and professional networks constitute distinct channels of diffusion that operate independently of formal contracts. Research on inter-firm worker flows indicates that employees transmit tacit knowledge across organizational boundaries, with firms hiring workers with prior R&D experiencing measurable productivity gains even when those workers are not assigned to research roles (Maliranta et al. 1161). Similarly, institutional evaluations of the SBIR program reveal that over one-third of awardees collaborate with university faculty, and more than half publish peer-reviewed research (Council 163–65). These findings imply that knowledge diffuses primarily through professional networks and labor markets rather than simple geographic proximity - a distinction that motivates the empirical focus of this study.

Collectively, the literature shows that defense R&D generates both codified and tacit knowledge through multiple institutional channels. However, knowledge may diffuse along professional, contractual, or supply chain networks that span distant locations while bypassing geographically proximate areas without relevant industrial capacity, raising the possibility that even substantial defense R&D may fail to generate observable spillovers at the county level.

3.4 Constraints on Spillover Diffusion in Semiconductors

Despite strong theoretical and empirical support for crowding-in effects, a growing body of work suggests that industries like semiconductors may resist broad geographic diffusion even under

sustained public support. As the microelectronics industry matured, defense to civilian technology spillovers declined in significance and, in some cases, reversed as technical requirements diverged, isolating defense innovation from commercial markets (Mowery, “Federal Policy” 169). In other cases, the specific dynamics of regional economic development may exacerbate unevenness. Reinsch et al. caution that poorly managed public R&D can generate corporate favoritism and capture, distorting markets and suppressing innovation by smaller firms or competitors outside favored networks (2). In semiconductors, this pattern is reinforced by the industry’s fragmented production structure. Design, fabrication, and packaging are frequently spatially separated, reflecting extreme capital intensity and specialization, with a single advanced fabrication facility costing up to \$10 billion (Platzer and Sargent Jr. 9).

Innovation, in complex, capital-intensive sectors, tends to concentrate in established hubs where firms can access “a diverse set of skills and complementary services” and benefit from high-coordination agglomeration economies (Pallante et al. 4). Historical analyses of defense linked microelectronics demonstrate this tendency. Acosta et al. show that firms engaged in dual use technologies are statistically distinct, exhibiting substantially higher sales, employment, and patenting activity than their peers, consistent with innovation remaining concentrated among large incumbents (22). Rather than diffusing outward, public investment may reinforce firms that already possess the complementary assets required to absorb and exploit new knowledge.

Because these stages are geographically dispersed, innovation networks rarely align with county boundaries. Knowledge generated in one location may benefit distant partners embedded in the same supply chain while bypassing nearby regions lacking relevant industrial capacity. Under such

conditions, increases in defense related R&D may deepen existing centers of capability without producing detectable gains in adjacent or less developed counties, particularly at fine geographic scales.

3.5 Defense R&D at Larger Scales and the Gap in Local Evidence

A substantial empirical literature evaluates the effects of defense-related R&D at national or state scales, generally finding evidence consistent with crowding-in rather than displacement of private innovation. Using cross-country panel data, Moretti et al. estimate that a 10 percent increase in government-funded R&D raises private R&D investment by approximately 4.3 percent (4). At a subnational level, Pallante et al. show that U.S. states with greater exposure to military R&D experience significant crowding-in effects over a four-to-five year horizon, estimating that each dollar of government spending generates about 0.57-0.72 dollars of additional private R&D (7). Together, these studies suggest that defense R&D generates positive aggregate spillovers and contributes meaningfully to private innovative effort.

Yet this evidence is largely silent on how defense-funded innovation diffuses across space at finer geographic scales. Even where aggregate crowding-in effects are well documented, existing studies do not examine whether such gains extend to nearby private actors at the county-year level, where proximity, industrial structure, and local absorptive capacity may sharply condition spillover dynamics.

The absence of local evidence is consequential. Many of the mechanisms emphasized in prior research, including labor mobility, supplier relationships, and inter firm knowledge transfer, are spatially mediated and unlikely to operate uniformly within states. Aggregated estimates may therefore

mask substantial heterogeneity in where spillovers materialize, particularly in industries such as semiconductors that exhibit extreme geographic concentration and fragmented production networks. Whether defense-funded R&D generates measurable patent-based innovation effects among nearby private actors at the county-year level thus remains an open empirical question.

3.6 Synthesis: Theoretical Expectations and the Puzzle of Local Spillovers

Taken together, the literature offers a coherent but unresolved picture. Public R&D generally stimulates private innovation. Proximity frequently amplifies innovative gains. Defense-funded projects generate codified and tacit knowledge capable of supporting downstream invention. Yet semiconductor production is highly clustered, capital-intensive, and regionally locked in - features that may inhibit the spatial diffusion mechanisms predicted by canonical spillover theory. The unresolved tension between these expectations provides the foundation for this paper's contribution. This study tests whether the positive patterns documented at national or regional levels reappear at the granular county-year level in a sector whose structure plausibly restricts local spillovers. In doing so, this study directly addresses the empirical gap noted across multiple literatures: whether defense-funded microelectronics R&D benefits nearby private innovators or remains confined within established centers of semiconductor activity.

4. Theory

DoD microelectronics R&D contract obligations (IV) are causally related to the rate of private semiconductor patenting by non-recipient inventors (DV) because public research activity generates locally embedded knowledge and technical capacity circulating beyond contract recipients that private

firms can draw upon. As the literature supports, where such knowledge diffuses through labor markets or professional networks, private inventive activity increases. However, this relationship is conditioned by structural constraints; if innovation remains concentrated within established hubs or defense procurement geographies do not align with civilian production networks, these spillovers will fail to materialize.

4.1 Hypotheses

H1. Temporal Causal Hypothesis: Increases in Department of Defense microelectronics R&D contract obligations are followed by increases in local semiconductor patenting by private, non-recipient inventors after one to three years.

H2. Geographic Spillover Hypothesis: Non-recipient inventors located closer to organizations receiving Department of Defense microelectronics R&D contracts increase semiconductor patenting more than comparable non-recipient inventors located farther away.

H0. Null Hypothesis: Department of Defense microelectronics R&D contract obligations have no measurable temporal or spatial effect on semiconductor patenting by private, non-recipient inventors.

5. Methodology

5.1 Overview of Research Design

This study utilized a county-year panel data design to evaluate whether DoD microelectronics R&D contract obligations generate localized spillovers, measured by increases in private, non-recipient semiconductor patenting in the same county or nearby counties. A localized spillover refers to an

indirect effect of DoD microelectronics R&D activity on private innovation that accrues to nearby, non-recipient actors. While the literature identifies multiple mechanisms through which defense-funded research may diffuse, including labor mobility, professional networks, and supplier relationships, these channels are not observed directly in this paper. Instead, spillovers are operationalized as increases in private, non-recipient semiconductor patenting in county-years that are spatially proximate to DoD-funded microelectronics activity, relative to less exposed county-years, conditional on county and year fixed effects. This definition treats spillovers as observable outcomes of localized exposure to defense R&D rather than as direct measurements of specific diffusion mechanisms.

The county-year is the unit of analysis because it represents the finest geographic scale at which federal procurement data, patenting outcomes, and key economic covariates can be consistently harmonized across sources while preserving nationwide coverage and temporal continuity. Counties also provide a natural unit for measuring localized exposure, as they are small enough to capture spatial heterogeneity in innovation activity but large enough to support stable annual measurement.

The empirical analysis proceeds sequentially, moving from a baseline bivariate regression to assess unconditional association, to multivariate panel regressions that address observable confounding and unobserved heterogeneity. To isolate the causal effect of defense obligations on local innovation, the design explicitly addresses four distinct inferential challenges:

Hurdle 1. Credible causal mechanism: The research design is grounded in a well-established theoretical mechanism linking defense-funded research and private innovation. The literature identifies labor mobility, professional networks, supplier relationships, and shared infrastructure as

channels through which publicly funded research can diffuse beyond direct recipients. The literature also demonstrates DoD microelectronics R&D generates codified and tacit technical knowledge, expands specialized labor pools, and reinforces localized innovation ecosystems. These mechanisms motivate the expectation that nearby private, non-recipient inventors may respond to DoD-funded activity with increased inventive output, establishing a credible causal pathway from funding to patenting.


Hurdle 2. Reverse causality: To ensure that observed associations do not reflect innovation-driven targeting of DoD contracts, the design enforces temporal ordering between funding and outcomes. To establish temporal precedence and mitigate reverse causality, the model adopts the logic of lagging used by Moretti et al., which regresses private R&D outcomes on government-funded R&D from the previous year (6–7). DoD microelectronics obligations are lagged by one, two, and three years, reflecting the cumulative and time-intensive nature of semiconductor innovation. Patents are assigned to grant years, while funding is assigned using award action dates, ensuring that DoD spending precedes measured patent outcomes and reducing the plausibility that contemporaneous patenting influences funding decisions.

Hurdle 3. Covariation between X and Y: Before introducing controls or fixed effects, the empirical strategy evaluates whether a basic association exists between DoD microelectronics R&D exposure and subsequent private semiconductor patenting. A baseline bivariate regression assesses unconditional covariation between lagged funding and patent outcomes, before proceeding to more demanding specifications.

Hurdle 4. Confounding from observable and unobservable factors: The primary threat to causal inference arises from confounding variables that influence both the placement of DoD contracts and local innovative capacity. To address observable confounding, the models incorporate a defined set of economic and innovation controls capturing local economic conditions and baseline inventive capacity. Following the econometric approach detailed by Moretti et al., the analysis employs fixed effects to control for heterogeneity; specifically, country-industry fixed effects absorb differences in investment propensity, while time dummies capture temporal variations such as technology shocks and product demand shocks (10). Together, these design elements limit identification to within-county changes in exposure, net of both observed and unobserved confounders.

Data preprocessing and panel construction were conducted in Python, while linear regression analysis was carried out in Stata. Together, this research design produced a spatially referenced county-year panel in which each observation contained (1) a population-normalized non-recipient semiconductor patent outcome, (2) DoD microelectronics obligations and their lags, (3) a defined set of county-year covariates capturing economic and innovation conditions, and (4) annual distance-to-treated measures used to test geographic attenuation consistent with spillover theory.

5.2 Dataset Selection and Sources

Fig. 1. Data Sources and County-Year Workflow

This study constructed a county-year panel by integrating multiple nationally consistent administrative datasets. All datasets were selected based on three criteria: nationwide county-level coverage, annual temporal resolution, and suitability for spatial harmonization using standardized FIPS codes. A full technical description of data origins, filtering decisions, and join procedures is provided in Appendix A.

5.2.1 Patent Data

Patent data used to construct both the dependent variable and patent stock control were obtained from the PatentsView bulk data filed by the U.S Patent and Trademark Office (USPTO). PatentsView is selected because it uniquely links granted patents to disambiguated inventor locations at the county level, enabling geographic assignment across the study period.

5.2.2 DoD R&D Contracts

Data on DoD microelectronics R&D obligations were drawn from USAspending.gov Prime Award Summaries. This source was selected because it provides authoritative contract-level data on federal obligations, including detailed product and service classifications and precise place-of-performance geography. These features allow for accurate identification and geographic assignment of defense-funded microelectronics R&D activity at the county level.

5.2.3 Economic and Institutional Controls

To address potential confounding factors that may influence the geographic allocation of DoD R&D activity and county semiconductor patent output, the analysis incorporates county-level economic and institutional data from three nationally standardized administrative sources. Real gross domestic product (GDP) data is drawn from the Bureau of Economic Analysis (BEA) and provides a consistent measure of regional economic capacity. County-level unemployment rates are obtained from the Bureau of Labor Statistics (BLS) Local Area Unemployment Statistics (LAUS) program, the authoritative source for local labor market conditions. Measures of university R&D expenditure and location are drawn from the National Science Foundation's Higher Education Research and Development (HERD) Survey. These sources were selected for their annual temporal resolution, nationwide county-level coverage, and standardized methodologies, ensuring consistency across time and space.

5.2.4 Spatial and Demographic Data

To enable spatial exposure measurement and population normalization, the dataset incorporates auxiliary geographic and demographic data from the U.S. Census Bureau. County

boundary geometries were obtained from the 2024 Census TIGER/Line cartographic boundary files, which provide the federal standard for county definitions and support consistent spatial harmonization across datasets. These shapefiles supply the geometric foundation of the analysis by allowing the computation of county centroids and the calculation of geographic distances between treated and untreated counties.

Annual county-level population estimates were compiled from Census Bureau intercensal and postcensal population series spanning 2000-2024. Population data was used to normalize patent counts, allowing innovative output to be expressed per 100,000 residents. This normalization ensures that observed differences in patenting intensity are not mechanically driven by variation in county population size.

5.3 Measurement of Variables

5.3.1 Dependent Variable: Unfunded Private Semiconductor Patenting (Y_{it})

The dependent variable Y_{it} measures county private-sector innovation in semiconductor technologies, operationalized as the annual count of granted patents constructed at the county-year level not directly funded federally. The objective is to capture spillover patenting activity by private inventors, excluding patents owned by federally funded or federally owned organizations.

Semiconductor patents are identified using CPC subclasses H01L (semiconductor devices), H10 (semiconductor technology), and G11C (semiconductor memory). Patents may carry multiple CPC codes; classification requires only one qualifying CPC subclass. Because CPC classifications are stored as one row per patent-classification pair, patents with multiple CPC assignments may appear multiple

times in the classification data. Each patent is therefore deduplicated after filtering so that it enters the dependent variable only once, regardless of the number of qualifying CPC subclasses it carries.

Identification of Non-Recipient Patents

To isolate private inventive activity, all patents assigned to the DoD or other federal entities and contractors are excluded through the government interest tag in the PatentView dataset. Any patent with at least one assignee identified as federally owned is removed in full. This exclusion ensures that the dependent variable does not mechanically capture the direct inventive output of government laboratories or defense agencies. Federal ownership is identified using a precision-oriented screening procedure applied to assignee records. This includes legally standardized assignee formulations indicating U.S. government ownership, as well as institutional names corresponding to DoD agencies and laboratories. Patents with any such assignee are excluded, ensuring that the dependent variable captures only private inventive activity not directly funded by the federal government.

Geographic Assignment

Patents are assigned to counties based on inventor location information. Inventor-patent records are merged with inventor location data to obtain county-level FIPS codes. A patent is attributed to a county if at least one inventor associated with the patent is located in that county. When multiple inventors on the same patent reside in the same county, the patent is counted once for that county. When inventors reside in different counties, the patent is attributed once to each inventor's county of residence.

This assignment rule reflects the geographically distributed nature of inventive activity by recognizing that knowledge production and technical labor contribute to innovation in each

inventors' location. Counting patents across inventor counties avoids imposing an arbitrary single-location attribution and is consistent with treating innovation as a spatially embedded process.

Duplicate patents-county pairs are removed prior to aggregation. Patents lacking any inventor with a valid U.S. county identifier are excluded. These cases primarily reflect foreign inventors or addresses that cannot be geocoded to the county level and therefore fall outside the geographic scope of the analysis.

Aggregation and Normalization

Let counties be indexed by i and years by t . Define $P_{i,t}$ as the raw count of unique, unfunded semiconductor patents attributed to county i in year t . This count is constructed by aggregating all qualifying patents granted in year t that are geographically attributed to county i . No additional weighting is applied at this stage; $P_{i,t}$ is a non-negative integer count. To account for differences in county population size, the dependent variable is normalized on a per-capita basis. Let $\text{Pop}_{i,t}$ denote the population of county i in year t . The final dependent variable is defined as:

$$Y_{i,t} = \frac{P_{i,t}}{\text{Pop}_{i,t}} \times 100,000$$

This transformation yields the number of unfunded semiconductor patents per 100,000 residents, allowing meaningful comparisons across counties of different sizes.

Temporal Coverage

Because inventor geocoding and CPC subclass coverage in PatentsView are stable beginning in 2013, and because the empirical models incorporate up to three years of lagged defense R&D

exposure, the dependent variable is constructed for the period 2013-2024. This window defines the effective study period for all analyses using the dependent variable.

5.3.2 Independent Variable: DoD Microelectronics R&D (X_{it})

The independent variable X_{it} captures local exposure to a public research stimulus in the form of DoD microelectronics R&D funding. It is intended to measure a time-varying county-level shock to publicly funded R&D activity that may generate downstream private-sector innovation through localized spillovers. The independent variable is constructed as the total inflation-adjusted obligated dollar amount for DoD microelectronics R&D contracts, aggregated to the county-year level. Each obligation reflects funding committed under a DoD prime contract for microelectronics-related R&D activity performed in a specific county during a given year.

Filtering and Scope Restrictions

The independent variable pre-filters aggregate obligations for Product Service Codes (PSC) corresponding to R&D stages (AC61-AC64) and semiconductor manufacturing and hardware (5961, 5962, 3670). Only definitive contracts are retained; grants, including SBIR awards, are excluded because DoD microelectronics R&D is overwhelmingly executed through procurement contracts rather than grant mechanisms (Nicastro), and grant records lack reliable place-of-performance geography - further filtering details are specified in Appendix A2.

Geographic Assignment

Each contract is assigned to a county using the place-of-performance county FIPS code reported in the procurement record. Only awards with valid county identifiers are retained. This

approach assigns funding to the location where R&D activity is performed rather than to contractor headquarters.

Aggregation

Let counties be indexed by i and years by t , where $X_{i,t}$ represents the total value of DoD microelectronics R&D obligations performed in county i during year t . This measure is constructed by summing obligated amounts across all qualifying contracts assigned to that county-year:

$$X_{i,t} = \sum_{a \in \mathcal{A}(i,t)} \text{Obligation}_a$$

where $\mathcal{A}(i, t)$ denotes the set of included DoD microelectronics contract awards with place of performance in county i and year t . The variable is measured in U.S. dollars. To test the sensitivity of results to functional form, the model utilizes a population-normalized measure of exposure, defined as inflation-adjusted DoD microelectronics obligations per 100,000 residents (X_{it}^{norm}). This allows for symmetrical specifications where both treatment and outcome are expressed as intensity rates.

Lag Construction

Temporal ordering is central to the research design. DoD obligations in year t cannot plausibly influence patents granted in the same year, because invention, drafting, application, and examination require time. Theoretical expectations and empirical studies of patenting cycles suggest that effects, if present, emerge within several years (Howell et al. 3617–26). Therefore, the analysis constructs one-, two-, and three-year lags:

$$X_{i,t-k} \quad \text{for } k \in \{1, 2, 3\}$$

These lagged values represent the treatment variables used in the empirical models.

5.3.3 Control Variables ($Z_{i,t}$)

Four control variables, noted collectively as $Z_{i,t}$, are operationalized to isolate the impact of defense funding from confounding economic and structural factors. These measures address distinct pathways through which DoD funding and private patenting may be jointly influenced: local economic scale, labor market health, academic research capacity, and accumulated innovative advantage. Finally, county population size (Pop_{it}) is included as a control variable. Although the dependent variable is expressed as a per-capita rate, population is included to account for agglomeration effects whereby larger regions may exhibit higher innovation intensity due to scale, network density, or specialization that are not fully captured by per-capita normalization alone.

Local Economic Conditions

Local economic conditions are represented using BEA Real County Gross Domestic Product (GDP). Real County GDP provides a broad measure of economic scale and productive activity, capturing variation in regional industrial composition and economic capacity that is systematically associated with innovative output and with the geographic distribution of federal R&D activity. Prior research defines innovation outputs to include knowledge creation such as patenting and documents a strong association between socioeconomic development and innovative activity (Suzuki and Demircioglu 756). As a result, county-level real GDP in chained 2017 dollars is incorporated as a standardized measure of local economic context, where using inflation-adjusted GDP by the source agency ensures comparability across years.

Labor Market Conditions

Labor-market conditions are measured using BLS county-year unemployment rate. The unemployment rate reflects short-run labor-market conditions and local economic slack, which function as determinants of regional innovative capacity. Corrocher et al. establish that labor-market conditions directly influence the nature of technological change, shaping a region's ability to generate innovative new products by determining the stability and accumulation of tacit knowledge required for complex technical work (16). This structural influence is empirically quantified by Bendig et al., whose multilevel analysis demonstrates that higher unemployment acts as a macro-level constraint that significantly suppresses "ambitious entrepreneurship" - the specific subset of opportunity-driven, high-growth activity that drives patenting (6, 9). Including the unemployment rate therefore controls for short-run economic conditions that may independently affect both patent output and the placement of defense R&D funding.

Academic Research Capacity

Academic research capacity is captured using county-level university R&D expenditures from HERD in thousands of nominal dollars. Ponds et al. provide direct empirical evidence that the presence of universities has a large and statistically significant effect on regional innovation outcomes through localized spillovers from academic research independently of defense funding (231). This ensures that the model accounts for baseline academic research intensity that may attract both federal contracts and private-sector innovation.

Because university data is often reported by institution ZIP code, a spatial crosswalk is applied. Let $U_{i,t}$ represent the set of universities mapped to county i in year t using the U.S Department of

Housing and Urban Development (HUD) ZIP-to-county crosswalk. The county-level variable is calculated as:

$$\text{UnivRD}_{i,t} = \sum_{u \in U_{i,t}} \text{HERD_Total}_{u,t}$$

Because regions with strong academic research bases may both generate more private patents and attract federal research investments, controlling for university R&D ensures that the estimated effect of DoD funding is not confounded by pre-existing or evolving academic research capacity.

Accumulated Innovative Capacity

Accumulated innovative capacity is controlled for using the county's prior patent stock derived from PatentsView data. Innovation is path-dependent; regions with historical advantages such as established engineering talent, firm concentration, and production routines tend to sustain higher levels of innovative activity through self-reinforcing dynamics, even in the absence of new external stimuli (Kenney and Burg 94). The inclusion of a patent stock controls for this accumulated advantage and isolates the marginal effects of defense R&D from baseline innovative capacity, preventing the model from attributing historical advantages to contemporaneous funding. Hausman et al. provide the econometric precedent for this approach, interpreting unit-specific patent histories as proxies for the accumulated stock of knowledge that shapes future inventive capacity (920).

Moretti et al. formalize this relationship by modeling productivity as dependent on lagged R&D knowledge stocks (10). Following this framework, the model builds the stock using the Perpetual Inventory Method (Moretti et al. 31). Adopting the standard Hall-Jaffe-Trajtenberg

specification (Hall et al. 25), $S_{i,t}$ represents the patent stock in county i at time t , and $P_{i,t}$ as the flow of new patents granted in that year. The stock evolves recursively:

$$S_{i,t} = (1 - \delta)S_{i,t-1} + P_{i,t}$$

where δ represents the standard depreciation rate of knowledge, set to 0.15 (Hall et al. 21). For the first observed year of any county, the stock is initialized at 0, with the recursive calculation building the stock over the pre-sample period to minimize initialization bias.

5.4 Spatial Exposure and Distance Measures

Because the hypotheses concern localized spillovers, spatial measurement is a central component of the design. The spatial framework proceeds in three steps: centroid construction, great-circle distance calculation, and distance-band assignment.

5.4.1 County Centroids

To compute meaningful centroid coordinates, counties are projected into an equal-area coordinate system (EPSG:5070), centroids are calculated, and the results are transformed back into latitude and longitude in WGS84 (EPSG:4326). These coordinates serve as anchor points for distance calculations across all years.

5.4.2 Distance to Treated Counties

A county-year is defined as treated if the county receives any positive amount of DoD microelectronics R&D obligations in that year:

$$\text{Treated}_{i,t} = \mathbf{1}\{X_{i,t} > 0\}$$

where $X_{i,t}$ is the total DoD microelectronics R&D obligation defined in Section 5.3.2. This binary indicator captures the presence of active DoD research funding without imposing arbitrary thresholds on award size. For each county-year, the Haversine formula is used to calculate the shortest distance between two points on a sphere using their latitude and longitude:

$$d = 2r \arcsin \left(\sqrt{\sin^2 \left(\frac{\Delta\phi}{2} \right) + \cos(\phi_1) \cos(\phi_2) \sin^2 \left(\frac{\Delta\lambda}{2} \right)} \right)$$

where ϕ denotes latitude, λ denotes longitude, and r is Earth's radius. Great-circle distance is necessary because Euclidean distance would introduce distortion across large geographic areas. The variable of interest, $\text{Dist}_{i,t}$, is defined as the minimum distance to the nearest treated county:

$$\text{Dist}_{i,t} = \min\{d_{ij} \mid \text{Treated}_{j,t} = 1\}$$

For treated counties themselves ($i = j$), $\text{Dist}_{i,t} = 0$.

5.4.3 Distance Bands

To allow for flexible and interpretable estimation of spatial decay, distances are discretized into ordered geographic bands:

$$\text{Treated: } \text{Dist}_{i,t} = 0\text{km}$$

$$0\text{-}25 \text{ km: } 0\text{km} < \text{Dist}_{i,t} \leq 25\text{km}$$

$$25\text{-}50 \text{ km: } 25\text{km} < \text{Dist}_{i,t} \leq 50\text{km}$$

$$50\text{-}100 \text{ km: } 50\text{km} < \text{Dist}_{i,t} \leq 100\text{km}$$

$$100 + km : \text{Dist}_{i,t} \geq 100\text{km}$$

These bands allow the spillover effect to vary flexibly across distances, testing the premise that knowledge diffusion varies with distance.

5.5 Quantitative Analysis Method

To mitigate the effect of omitted variable bias and account for unobserved heterogeneity, the analysis relies on a multivariate Two-Way Fixed Effects (TWFE) estimation strategy. As Imai and Kim note, this approach has become the “default methodology for estimating causal effects from panel data” as it allows for the adjustment of “unobserved unit-specific and time-specific confounders at the same time” (1). The general form of the TWFE model is defined as:

$$Y_{it} = \alpha_i + \gamma_t + \beta X_{it} + \epsilon_{it}$$

where Y_{it} represents the dependent variable for unit i at time t , and X_{it} represents the independent variable. All models crucially include unit fixed effects (α_i) to control for time-invariant characteristics, in this case, specific county fixed effects like geography, historic industrial clusters, and county characteristics. All models also include time fixed effects (γ_t) to absorb national shocks affecting all counties simultaneously, in this case, year fixed effects such as macroeconomic cycles and federal semiconductor policy. Standard errors are clustered at the county level, allowing the error term ϵ_{it} to be correlated within counties across years and ensure correct inference under serial correlation.

5.5.1 Model 1: Spatial Exposure Baseline (Twoway FE)

This paper adapts the general framework to evaluate the spatial spillover hypothesis (H2) by testing whether geographic proximity to a treated county is sufficient to generate innovation, independent of funding magnitude. The model regresses semiconductor patenting rates on the discretized distance bands defined in 5.4.3,

$$Y_{it} = \alpha_i + \gamma_t + \sum_b \beta_b \mathbb{I}(\text{Dist}_{it} \in b) + \theta \mathbf{Z}_{it-1} + \epsilon_{it}$$

where b indexes the distance bands relative to the reference category of counties located more than 100 km from a recipient. A positive and significant β_b would indicate that proximity alone produces a spillover advantage, representing the marginal effect of spatial proximity on patenting rates. $\mathbf{Z}_{it-1}$ represents a vector of time-varying control variables lagged by one year, while θ defines the vector of coefficients for the control variables. This model serves as a baseline to determine if the mechanism of innovation is purely geographic.

5.5.2 Model 2: Funding Intensity (Twoway FE)

To test the funding intensity hypothesis (H1), the second model replaces the passive distance bands with the continuous measure of active DoD investment. This specification estimates the marginal effect of funding dollars on innovation output:

$$Y_{it} = \beta_1 X_{i,t-1}^{norm} + \theta \mathbf{Z}_{it-1} + \alpha_i + \gamma_t + \epsilon_{it}$$

where $X_{i,t-1}^{norm}$ represents the inflation-adjusted DoD microelectronics obligations per 100,000 residents, lagged by one year. This model distinguishes whether innovation is driven by the intensity of the fiscal stimulus rather than simple geographic proximity.

5.5.3 Model 3: Dynamic Distributed Lag (Twoway FE)

The primary causal specification extends Model 2 into a distributed lag structure to characterize the temporal dynamics of the effect and guard against reverse causality. The model includes three lags of the treatment variable:

$$Y_{it} = \sum_{k=1}^3 \beta_k X_{i,t-k}^{norm} + \theta \mathbf{Z}_{it-1} + \alpha_i + \gamma_t + \epsilon_{it}$$

This structure allows us to test whether the effect of funding is immediate ($t - 1$), persistent

($t - 2, t - 3$), or only reflective of pre-existing trends. Finding a significant effect at $t - 1$

with null effects at further lags would support a "pulse" stimulus interpretation, whereas persistent lags might suggest endogenous trends or reverse causality.

5.5.4 Model 4: Intensity with Mega-County Exclusion

To determine whether the estimated funding effects may be driven by a small number of industrially dense and/or extreme innovation hubs (Silicon Valley or major metropolitan areas), a final specification re-estimates the funding intensity model (Model 2) after excluding the top 1 percent of counties by population, which includes major tech hubs the paper is concerned with (Santa Clara, San Francisco, Alameda, etc). The specification is otherwise identical to Model 2 and is estimated using the same equation on the restricted sample.

With these specifications defined to isolate funding intensity from geographic proximity, the following section presents the empirical results.

6. Results and Evaluation


The analysis proceeds in two stages. First, the study evaluates whether geographic proximity to defense-funded counties increases private semiconductor patenting per H2. Second, the analysis examines whether the intensity of defense R&D funding produces a temporally ordered response in patenting activity, using the lag specifications in 5.3.2 to test H1. All models are estimated using county and year fixed effects, with standard errors clustered at the county level. The dependent


variable is semiconductor patents per 100,000 residents, denoted Patents_{it} . Full-size versions of all figures are provided in Appendix B.

6.1 Spatial Distribution and Clustering of Innovation

Descriptive spatial patterns reveal a significant concentration of semiconductor patenting activity. County-level maps show that most U.S. counties generate zero semiconductor patents in a given year, while a small subset of counties exhibit persistently high patenting rates (Fig. 2). These high-output counties are overwhelmingly located in established technology corridors, consistent with strong path dependence and limited capacity. Maps of distance-to-treated counties further show that DoD microelectronics R&D contracts are geographically widespread but unevenly distributed across

Fig. 2. Intensity of Private Semiconductor Patenting, 2017-2020.







Note: Maps display the number of unfunded private semiconductor patents granted per 100,000 residents. Darker blue regions indicate high-intensity innovation clusters. The persistence of high-output regions (e.g., California, Texas, Massachusetts) aligns visually with the "Treated" regions in Fig. 2, suggesting spatial agglomeration.

time (Fig. 3). Visually, treated counties often overlap with or border regions of high private patenting activity, suggesting a correlation between defense R&D exposure and innovation. However, this visual co-location does not imply a causal spillover effect.

Fig. 3. Geographic Distribution of DoD Microelectronics R&D Obligations, 2012-2024.





Maps illustrate the spatial stability of defense funding over a 12-year period. "Treated" counties (red) represent locations receiving direct DoD microelectronics contracts. The surrounding bands (orange to blue) indicate 25 km, 50 km, and 100 km proximity zones used in the distance-band regression analysis.

Table 1, Model (1) formally tests the geographic spillover hypothesis (H2) using distance-band indicators. After controlling for county fixed effects, year fixed effects, accumulated patent stock, and local economic conditions, proximity alone does not predict higher private semiconductor patenting. The coefficients for counties located within 0-25 km, 25-50 km, and 50-100 km of a treated county are all statistically indistinguishable from zero, with wide confidence intervals overlapping zero. This finding is substantively informative because it supports the proximity-challenging literature in 3.3 and 3.4, while rejecting the strongest forms of geographic spillover theory in 3.2. Geographic proximity to defense-funded activity, absent direct exposure to funding, does not independently increase innovation output.

6.2 Funding Intensity and Dynamic Treatment Effects

When the analysis shifts from geographic proximity to funding intensity, the estimated effects change meaningfully. Table 1, Models (2)-(4) report estimates of the association between DoD

microelectronics R&D obligations, measured as dollars per 100,000 residents and lagged over time, and subsequent private semiconductor patenting.

Model (2) indicates that lagged funding intensity at $t - 1$ is positively associated with private semiconductor patenting. An additional \$1 million in DoD microelectronics R&D per 100,000 residents in the previous year is associated with an increase of approximately 0.106 patents per 100,000 residents ($p < 0.10$), indicating a marginal but detectable short-run association.

Table 1. Effect of DoD Microelectronics R&D on Private Semiconductor Patenting

	Model (1)	Model (2)	Model (3)	Model (4)
	Distance Bands (<i>Spatial Exposure</i>)	Funding Intensity (<i>t - 1, Two-Way FE</i>)	Distributed Lag (<i>Two-Way FE</i>)	Funding Intensity (<i>Excl. Top 1%</i>)
Funding Intensity ($t - 1$) \$1M per 100k residents		0.106* [-0.005,0.217]	0.043* [-0.005,0.091]	0.158** [0.021,0.295]
Funding Intensity ($t - 2$) \$1M per 100k residents			-0.007 [-0.038,0.025]	
Funding Intensity ($t - 3$) \$1M per 100k residents			0.013 [-0.018,0.043]	
Distance (0-25 km)	0.020 [-2.009,2.050]			
Distance (25-50 km)	2.190 [-2.525,6.906]			
Distance (50-100 km)	-1.608 [-6.119,2.904]			
Distance (100+ km)	-1.187 [-4.393,2.019]			
Patent Stock	0.019*** [0.010,0.029]	0.023*** [0.009,0.036]	0.014 [-0.006,0.033]	0.022*** [0.007,0.036]
Unemployment Rate	0.740* [-0.037,1.518]	0.857 [-0.629,2.342]	0.634 [-0.765,2.033]	0.896 [-1.005,2.796]
GDP & Univ R&D	Yes	Yes	Yes	Yes

Table 1. Effect of DoD Microelectronics R&D on Private Semiconductor Patenting

	Model (1)	Model (2)	Model (3)	Model (4)
	Distance Bands (<i>Spatial Exposure</i>)	Funding Intensity (<i>t - 1, Two-Way FE</i>)	Distributed Lag (<i>Two-Way FE</i>)	Funding Intensity (<i>Excl. Top 1%</i>)
County Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
Observations	1974	1126	883	927
Counties	295	202	138	181

Note: Model (1) tests the spatial spillover hypothesis using distance bands; the reference category is counties located >100 km from the nearest treated county. Models (2)-(4) test the funding intensity hypothesis using DoD microelectronics R&D obligations, scaled in increments of \$1 million per 100,000 residents. Model (4) is a static specification ($t - 1$) excluding the top 1% of counties by population to control for agglomeration outliers. Standard errors clustered at the county level are in parentheses; 95% confidence intervals are reported in brackets. All models include county fixed effects, year fixed effects, and controls for patent stock, GDP, and university R&D (full output reported in Appendix B).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$


Model (3) extends this specification to a distributed lag framework by including funding intensity at $t - 1$, $t - 2$, and $t - 3$. The distributed lag estimates are summarized in Fig. 4, which shows a positive point estimate at the one-year lag and near-zero estimates at longer horizons, with confidence intervals that widen and overlap zero beyond $t - 1$. As Fig. 4 makes clear, the estimated response is temporally concentrated rather than cumulative, consistent with a short-run adjustment rather than a persistent innovation trend. The coefficient on funding intensity at $t - 1$ remains positive and statistically significant at the 10 percent level, while the coefficients at $t - 2$ and $t - 3$ are small and statistically indistinguishable from zero. The estimated magnitude implies that an additional \$1 million in funding per 100,000 residents generates approximately 0.04 additional patents per 100,000 residents in the following year.

This temporal pattern provides suggestive but limited evidence of a short run association between DoD microelectronics R&D funding and subsequent semiconductor patenting, offering only weak support for the Temporal Causal Hypothesis (H1). Point estimates are consistently positive and concentrated at a one-year lag, but they are statistically imprecise and sensitive to specification, warranting caution in causal interpretation. In Table 1, Model (2), the one-year lag coefficient is positive and marginally significant ($p < 0.10$), with a 95 percent confidence interval that narrowly includes zero, while Model (3) yields a smaller but similarly timed estimate. The absence of detectable effects at longer lags further weakens reverse causality explanations - if DoD funding merely followed pre-existing innovation trends, one would expect persistent or anticipatory effects across multiple lag periods. As shown in Figure 4, longer lags are small and statistically indistinguishable from zero. This timing pattern is consistent with a short-run response and makes a purely trend-following explanation less compelling, though the wide confidence intervals warrant caution in causal interpretation.

Table 1, Model (4) further tests the robustness of this relationship by excluding the top 1 percent of counties by population. This restriction addresses the concern that observed effects might be driven purely by mega-county agglomeration economies or urban scaling dynamics inherent to the largest metropolitan areas. Fig 2. and Fig. 3 demonstrate that a handful of dominant hubs account for a large share of national inventive activity, and are also more likely to receive large defense R&D allocations. In pre-selection models, aggregate estimates can be dominated by a few counties, making it important to check whether results persist once the extreme right tail is trimmed. Consistent with this concern, Model (4) is the only specification in Table 1 where the one-year lag of funding intensity reaches conventional statistical significance at the 5 percent level. After excluding the top 1 percent of

counties, the estimated $t - 1$ coefficient rises to 0.158 ($p < 0.05$) with a 95 percent confidence interval of 0.021 to 0.295. The sensitivity of the estimated coefficient to the exclusion of extreme counties indicates that the observed temporal association is not uniform across the county distribution.

Fig. 4. Distributed Lag Coefficient Plot



Note: Bars represent 95% confidence intervals. Standard errors clustered by county.

Rather than reflecting broad, geographically diffuse spillovers, the shift in magnitude and precision suggests that responses to defense-linked R&D differ inside dominant innovation clusters relative to typical counties. This pattern reinforces the paper's broader conclusion that semiconductor innovation is structurally concentrated, and that any measurable response to defense R&D operates through cluster-specific dynamics, institutional access, or contracting ecosystems rather than through passive spatial diffusion.

Taken together, Models (2)-(4) provide consistent, though specification-sensitive, and limited evidence in support of the Temporal Causal Hypothesis (H1). Across specifications, lagged DoD microelectronics R&D funding is associated with a short-run increase in private, non-recipient semiconductor patenting, with the strongest evidence emerging in specifications that exclude extreme population outliers.

6.3 Controls, Fixed Effects, and National Innovation Waves

Across all specifications, accumulated patent stock enters positively and is frequently statistically significant, confirming the path-dependent nature of regional innovation emphasized in the methodology. The year fixed effects in Fig. 5 are large and highly significant for the period 2014-2019. This pattern aligns closely with the national semiconductor patenting time series (Fig 5.), which shows a pronounced nationwide surge beginning around 2014-2015, followed by stabilization and decline. These year effects absorb common shocks associated with industry-wide developments, including the deep learning boom and broader technological consolidation.

Two control variables - population and university R&D expenditures - are not precisely estimated once county fixed effects are included. This reflects the highly persistent and structural nature of both population size and academic research activity, rather than a lack of substantive importance. Population primarily captures long-run scale differences across counties, while university R&D reflects slow-moving academic research capacity. Both are largely time-invariant within counties and are therefore absorbed by the fixed-effects structure and the patent stock control. Their


imprecision is thus a consequence of the identification strategy rather than a shortcoming of the model.

6.4 Evaluation of Hypotheses and Interpretation of Findings

The geographic spillover hypothesis (H2) is not supported: proximity to defense-funded counties does not independently increase private semiconductor patenting. The null hypothesis (H0) is therefore rejected only in part. In contrast, the temporal causal hypothesis (H1) receives limited and specification-sensitive support. Funding intensity matters, and its effect is temporally ordered, short-lived, and demonstrable across changes in DV modeling and sample restrictions. The evidence indicates that defense R&D obligations generate a localized innovation response through direct exposure to funding rather than passive spatial diffusion.

Substantively, this implies that money matters more than geography. The results suggest that defense R&D stimulates private semiconductor innovation not through passive geographic spillovers, but through mechanisms that meaningfully reduce the cost of discovery and accelerate experimentation among proximate private actors. While these mechanisms are not directly observed in the regression framework, they are well-documented in the institutional and historical evidence reviewed earlier in the 3.1 and 3.3, and consistent with the temporal pattern identified in Section 6.

Fig. 5. National Time Series Trend



Note: The solid blue line represents the total count of U.S. semiconductor utility patents granted annually. The dashed teal line represents the patenting rate per 100,000 residents. Annotations indicate major technological periods corresponding to the major fixed year effects observed in the regression analysis.

7. Limitations and Future Research

This study faces several limitations that shape how the results should be interpreted. These constraints do not undermine the consistency of the empirical design, but they constrain the scope of the conclusions that can be drawn about the local spillover effects of DoD microelectronics R&D.

7.1 Measurement of Innovation Outcomes

A central limitation of the analysis is its reliance on patent counts as a proxy for semiconductor innovation. Patents are a widely used and observable measure of inventive activity, but they do not distinguish between incremental and transformative inventions, nor do they weight innovations by economic value or technological significance. Prior research shows that patent-based measures may obscure substantial heterogeneity in novelty and impact, especially in high-variance sectors such as semiconductors (Howell et al. 3618). As a result, the dependent variable captures the quantity of patented output rather than the quality, originality, or importance of innovations.

This limitation is especially important with semiconductors, where significant advances may occur through process improvements, design optimization, or manufacturing knowledge that is never actually patented or is retained as trade secrets. Consequently, the analysis may understate innovative activity that occurs outside formal intellectual property filings, or manifest primarily through changes in patent quality rather than quantity. The null spillover results should therefore be interpreted as evidence about the absence of geographically diffused effects on patent counts, not as definitive evidence that DoD microelectronics R&D totally fails to influence innovation broadly.

7.2 Spatial Aggregation

A further limitation concerns spatial aggregation. Counties are an imperfect unit for analyzing innovation spillovers, which may operate through metropolitan labor markets, firm networks, and institutional linkages rather than administrative boundaries (Corrocher et al. 16; Fuchs 1445).

Distance bands constructed from county centroids and great-circle distance provide a proxy for geographic proximity, but they can only capture how far apart two places are, not the organizational and relational structures through which knowledge diffusion typically occurs in high-technology industries. For example, two counties could be 10 km apart and have zero interaction, while two firms 2,000 km apart may exchange designs weekly under the same defense program - centroid distance cannot see that. This limitation of distance bands is especially relevant in analyzing semiconductor innovation, which is distributed across geographically separated stages of design, fabrication, packaging, and equipment production (Platzer and Sargent Jr. 9). As a result, spillovers may operate through sector-wide labor markets, shared platforms, or interfirm collaboration rather than through localized spatial adjacency. Distance-based exposure measures will therefore miss diffusion channels that cut across counties or operate at broader spatial scales than those required for identification in a county-year panel.

In addition, the reliance of county centroids for calculating spatial exposure introduces significant measurement error for the "Geographic Spillover" analysis. As noted in Appendix A7, using geometric county centroids to measure inter-county proximity introduces measurement error in the spatial exposure variables, particularly for large, irregular counties where centroids can lie far from population and employment centers. This can misclassify counties across distance bands (e.g. 0-25 km

vs 25-50 km), which would attenuate estimated proximity effects toward zero and reduce statistical power. Accordingly, the imprecision of the 0-25 km estimate in Table 1, Model (1) may reflect either a true absence of geographic spillovers or dilution arising from distance mismeasurement. Robustness checks using population-weighted centroids or boundary-to-boundary distances would help distinguish between these interpretations.

7.3 Causal Interpretation and Policy Implications

Although the model incorporates county and year fixed effects, lagged treatment variables, and a rich set of controls, the design relies on observed variation in DoD R&D allocations rather than on random assignment. Funding decisions are not random and may respond to evolving strategic priorities and preexisting regional capabilities that are only partially observed. In addition, time-varying shocks to local innovation ecosystems, such as changes in private investment, firm organization, or industrial structure, may not be fully captured by the included covariates. Together, these features of the data complicate causal interpretation and warrant caution in treating null spillover estimates as definitive evidence that no diffusion occurs.

From a policy perspective, however, the results speak directly to the distributional concern raised in the introduction. If DoD microelectronics R&D does not reliably increase semiconductor patenting among nearby non-recipient counties, then defense R&D alone is unlikely to generate new innovation hubs at the county level. Instead, the evidence suggests that federal microelectronics funding primarily reinforces innovative activity within a small number of entrenched clusters that already host defense contractors and civilian semiconductor firms. This interpretation is consistent with Fig. 2 and Fig. 3, which show that the geography of semiconductor patenting remains highly

unequal even as Department of Defense microelectronics contracts are distributed across a wide range of locations.

7.4 Unobserved Confounders and Omitted Variable Concerns

Despite the inclusion of county and year fixed effects, the TWFE model cannot account for all potentially relevant sources of time-varying confounding. In particular, due to data availability and scope constraints, the models do not directly control for county-level venture capital (VC) inflows or for changes in local industrial composition, such as the share of manufacturing or high-technology activity.

VC represents a key channel through which non-recipient private-sector innovation is financed and commercialized, particularly in technology-intensive sectors. Estimates suggest that a doubling in the supply of venture capital within a region generates roughly four additional patents per year on average, but this figure rises to more than ten additional patents in regions with high levels of public research funding, emphasizing the catalyst effect of VC (Samila and Sorenson 1353). This may drive DoD R&D funding to disproportionately flow to regions that already host entrepreneurial ecosystems and promising technological capabilities, creating a confounding pathway if not explicitly controlled for. Empirical evidence shows public research funding generates larger innovation responses in regions with strong local venture capital supply, highlighting the “complementarity” between public R&D and VC (Samila and Sorenson 1349). In this context, failing to control for VC investment could lead to over-attribution of patenting responses to defense R&D when they are partly driven by concurrent private investment dynamics.

A related concern arises from differences in counties' underlying industrial structure. Regions with a large manufacturing or high-technology industrial base tend to exhibit persistently higher levels of inventive activity due to established firms, supplier networks, and accumulated technical expertise (Kenney and Burg 94). These same characteristics also increase the likelihood of receiving DoD R&D contracts in sectors like microelectronics which rely on specialized production capabilities and skilled labor. While county fixed effects absorb time-invariant structural differences, they cannot capture time-varying changes in industrial composition that may confound estimates of spillover effects. Prior research documents the intense geographic clustering of invention, finding that 63 percent of U.S. patents are developed in just 20 metropolitan areas, which house only 34 percent of the population (Rothwell et al. 12).

7.5 Directions for Future Research

Future research would benefit from moving beyond purely geographic measures of exposure. Network-based indicators, such as shared defense contractors, supplier linkages, or co-participation in specific procurement programs, may outperform distance-based metrics in predicting where spillovers occur. Combining USAspending contract data with firm-level supply-chain datasets or program-level participation records would allow researchers to test whether innovation diffuses along contractual and institutional networks rather than across physical space. This approach would help reconcile the weak distance-based effects observed here with theoretical expectations of knowledge diffusion in technologically complex, IP-intensive industries.

This research could also be extended in several directions. Incorporating measures of patent quality, such as forward citations, technological breadth, or originality scores, would allow assessment of whether DoD-related innovation concentrates in a small number of high-value inventions rather than in aggregate counts. Linking contract-level data to firm-level patenting outcomes, rather than county aggregates, would permit more precise distinction between recipient responses and external spillovers.

The evidence suggests that, within the limits of this county-year panel and patent-based outcome measure, Department of Defense microelectronics R&D does not produce sizable, geographically diffuse spillovers in semiconductor patenting among non-recipient counties. Instead, innovation in this sector appears to be governed more by path dependence and entrenched clustering than by local spatial diffusion from federal R&D, tempering expectations about the ability of defense spending alone to rebalance the geography of semiconductor innovation.

8. Conclusion

This study examined whether the mechanisms of innovation that characterized the early semiconductor industry - localized spillovers, regional clustering, and the broad diffusion of defense-funded knowledge - still structure the contemporary era. The results suggest that the mechanisms linking public investment to private innovation in semiconductors have shifted from regionally embedded spillovers to institutional ones. While defense R&D continues to stimulate private invention in the semiconductor industry, its effects are narrowly concentrated among actors with direct contractual exposure rather than diffusing outward through geographic proximity. This pattern

contrasts with the "triangular alliance" model of the Cold War, not because public investment has ceased to matter, but because the organizational structure of innovation has changed, where innovation spillovers in mature, capital-intensive industries are conditional on institutional access like funding relationships and supplier networks (Decarolis et al. 1997) rather than spatial proximity.

The empirical analysis yields two complementary conclusions. First, the Temporal Causal Hypothesis (H1) receives limited and specification-sensitive support. Lagged defense microelectronics R&D obligations are positively associated with subsequent private semiconductor patenting, with effects concentrated at a one-year horizon and dissipating thereafter. The point estimates are consistently positive but statistically imprecise, suggesting that federal R&D operates less as a long-term builder of new regional innovation capacity than as a short-run stimulus that activates existing capabilities where institutional linkages already exist. Second, the Geographic Spillover Hypothesis (H2) is rejected - once direct funding exposure is accounted for, physical proximity to a treated county confers no increase in patent output. Being 25 kilometers from a major defense innovation hub is, statistically, no different than being 500 kilometers away.

These findings suggest that the relationship between government, industry, and academia has evolved from a regionally embedded ecosystem into a series of vertical, exclusive silos. The islands of innovation pattern observed in this paper implies that the structural barriers to entry in modern semiconductors (Platzer and Sargent Jr. 1999) have weakened the proximity-based spillovers that characterized earlier technological booms. Innovation gains are concentrated primarily where defense contracts directly land, and even neighboring counties fail to exhibit measurable spillover effects.

Semiconductor clusters thus appear less as engines of broad regional diffusion and more as collections of institutionally connected actors whose innovative activity is tied to access rather than location.

The results also cast doubt on the spatial ambitions of recent industrial policy. The CHIPS and Science Act explicitly aims to broaden the spatial distribution of U.S. innovation beyond established innovation centers (136 Stat. 1643). Yet this research suggests that federal contracting, as currently structured, is unlikely to achieve this goal on its own. By generating patenting activity primarily in directly funded locations - and failing to induce activity even in immediately adjacent counties - defense R&D obligations will continue to reinforce existing centers of dominance rather than flattening regional disparities. In a setting where innovation responds more to funding intensity than to geographic proximity, dispersing small awards across many regions is unlikely to generate the meaningful spillovers policymakers anticipate.

Finally, the null results regarding geographic proximity invite a reconsideration of how political scientists study technology transfer. The failure of the distance-band models in this research suggests that physical proximity may be an obsolete proxy for knowledge diffusion in highly specialized, high-barrier sectors. Future research on industrial policy and federally linked innovation would benefit from shifting attention toward institutional and supply-chain networks as the primary units of analysis. In an era of digital collaboration and specialized global supply chains, the relevant neighbor for a semiconductor firm is not the county next door, but the specific supplier, contractor, or research partner embedded in its innovation chain. Understanding the geography of twenty-first-century innovation may therefore require moving beyond the map and toward the network.

Appendix A. Data Sources and Technical Notes

This appendix provides the full data provenance and technical details supporting the county-year dataset used in the analysis. It lists each data source, explains why it was chosen, and notes the key processing steps necessary to harmonize these sources into a unified spatial panel. These details appear in the appendix rather than the main text to preserve readability while ensuring transparency and replicability.

All data cleaning, variable construction, spatial processing, and statistical analysis scripts used in this study are publicly available in a GitHub repository. The repository includes Jupyter notebooks for data preprocessing and spatial analysis, Stata scripts for regression estimation, and documentation describing the execution order of all files. The full replication archive is available at:

<https://github.com/thedanielju/semiconductors-patent-analysis>

A1. PatentsView (USPTO Office of the Chief Economist)

Use: Dependent variable construction; patent stock.

Tables used: g_patent, g_cpc_at_issue, g_inventor_disambiguated, g_location_disambiguated, g_assignee_disambiguated.

PatentsView provides the only nationally consistent, disambiguated dataset linking patents to inventor locations at county-level precision. It includes standardized Cooperative Patent Classification (CPC) codes, enabling the identification of semiconductor technologies through subclasses H01L, H10, and G11C. Unlike annualized summary tables, the bulk data files provide granular metadata,

including full CPC hierarchies, which are necessary to distinguish semiconductor technologies from broader electronics innovation.

Exclusion of Software and AI: The restriction to these CPCs ensures the dependent variable is in line with the treatment (DoD microelectronics R&D spending), as other technologies like AI or software cannot be cleanly mapped to the DoD’s procurement taxonomies. There is no corresponding PSC structure for CPC classes in the G-series such as G06F (general computing) or G06N (AI systems) to filter for semiconductors.

Isolation of Private Sector Innovation: Patents with DoD assignees were removed using high-precision text and type-code filters in Python to isolate private, non-recipient inventive output. This was achieved by combining two screening methods: flagging assignee records where the assignee_type equals 6 ("U.S. federal government") and applying a high-precision keyword screen for terms such as "Army," "Navy," "DARPA," or "Department of Defense. PatentsView’s CPC and inventor-location coverage is stable beginning in 2013, which determines the earliest year for dependent-variable inclusion. The USPTO Bulk Data Directory provides patent-level XML files (PTGRXML), but these were deemed infeasible for research use. The raw corpus consists of over 700 weekly files (2011-2024), totaling several hundred gigabytes, with no server-side filtering by technology class or geography. Because each file embeds complete patent documents, extracting the limited fields relevant to this study would require specialized XML handling and substantial storage. Therefore, this project relies on PatentsView, which provides compact, relational tables with pre-computed disambiguation for inventors and locations.

Data Cleaning and Geolocation: The raw PatentsView tables contain 9.26 million granted patents.

Missingness appears only in the geographic elements; roughly 70 percent of location records lack valid county FIPS codes. This is primarily driven by foreign inventors (who make up ~65-70 percent of semiconductor patents) and U.S. addresses that do not resolve to the county level. However, this is broadly consistent with fractional inventor residence shares derived from the same underlying NSF patent tabulations, which place the U.S. share of semiconductor inventive activity in the ~22% range in 2022 (Statistics (NCSES) 49). In addition, around 0.7% of inventor-patent links lack a location_idm, which similarly reflect foreign inventors or unparseable addresses. Structural geographic missingness regarding U.S. territories and administrative changes is detailed in Appendix A10.

Summary of Filtering Impact:

- ❖ Total Semiconductor Patents: (US-Locatable): 94,910 unique patents.
- ❖ Multi-county patents: 50.38% (mean 1.77 counties per patent).
- ❖ DoD Exclusion: 2,476 patents were removed (2.60% of US-locatable semiconductor patents).

This includes 428 DoD-owned patents and 2,142 DoD-funded patents identified via government interest tags.

A2. USAspending Prime Award Summaries (Department of Treasury)

Use: Independent variable (DoD microelectronics R&D obligations).

Tables used: prime_award_summaries

USAspending provides authoritative federal procurement data with contract-level detail on obligations, product/service codes (PSCs), and place-of-performance geography. DoD

microelectronics contracts were isolated using a targeted set of PSCs representing the microelectronics lifecycle.

Included Product Service Codes (PSCs):

- ❖ **AC61** (Basic Research): Captures upstream microelectronics science, material science, and nanoelectronics fundamentals.
- ❖ **AC62** (Applied Research / Exploratory Development): Covers semiconductor R&D, including early ASIC design and rad-hard prototyping.
- ❖ **AC63** (Advanced Development): Includes midstage electronics development, subsystem electronics, prototypes, and FPGA subsystems.
- ❖ **AC64** (Engineering Development): Captures late-stage engineering development.
- ❖ **5961** (Semiconductor Devices): Covers discrete chips and device-level semiconductors.
- ❖ **5962** (Microcircuits): Covers hybrid and monolithic ICs, ASICs, logic chips, memory chips, and rad-hard microcircuits.
- ❖ **3670** (Specialized Manufacturing Machinery): Covers capital goods like semiconductor fabrication tools and wafer processing equipment. This code was included to capture DoD investment into manufacturing capacity and strengthen the theoretical link between federal spending and private patenting spillovers.

Exclusions and Filtering:

- ❖ **Non-R&D PSCs:** Codes AC66 (management support) and AC67 (commercialized/post-development sustainment) were explicitly excluded as they represent unrelated R&D activities.

- ❖ **Grants:** Grant awards were excluded because DoD rarely uses them for technical semiconductor work; they typically fund non-technical activities like STEM outreach. Furthermore, grants are reported using recipient addresses rather than the site of performance, which would degrade spatial precision. Indefinite Delivery Vehicles (IDVs), such as IDIQ or GWAC framework contracts, were excluded because they are contracting vehicles rather than realized expenditures, and their inclusion would introduce measurement error through zero-obligation records and potential double counting. These restrictions ensure that the independent variable captures geographically anchored defense R&D activity rather than general innovation assistance.
- ❖ **Indefinite Delivery Vehicles (IDVs):** Framework contracts (IDIQ, GWAC, FSS Schedule, BPA parent vehicles, and BOAs) were excluded because they often carry zero obligations and can lead to double-counting. Only obligated transactions, such as Delivery Orders (DOs) and definitive contracts, were included.

Interpretation of Negative Obligations: The USAspending data contain negative obligation values, which represent deobligations or administrative corrections (e.g., reclaiming unspent funds). These values were retained rather than zeroed out as they reflect real financial processes and reduce the net obligation for that county-year without distorting the research's identification strategy, which relies on positive R&D injections.

A3. Bureau of Economic Analysis (BEA), Real GDP by County (CAGDP9)

Use: Economic control variable (real GDP).

BEA's annual county-level GDP estimates provide a standardized measure of regional economic capacity and productivity. The analysis uses Real GDP in Chained 2017 Dollars, as provided by BEA. BEA explicitly recommends chained-dollar GDP for longitudinal analyses because it reflects inflation-adjusted production volume rather than price changes. Real GDP is calculated using chain-type quantity indexes, ensuring that every county-year value is comparable over time and not confounded by shifts in the price level.

Reshaping: Following standard panel-data practice, the BEA wide-format series (where each year is a separate column) was reshaped into a long county-year structure using a melt transformation. This ensures consistency across all variables in the dataset without altering the underlying economic values.

A4. Bureau of Labor Statistics (BLS), Local Area Unemployment Statistics (LAUS)

Use: Economic control variable (unemployment rate).

The LAUS program publishes nationally consistent unemployment statistics for all U.S. counties. Annual averages (period "M13") from the not-seasonally-adjusted unemployment rate series (Measure Code 03) were extracted from the large `la.data.64.County` file using dictionary-based filtering on `area_type`, measure codes, and county-area identifiers embedded in the `series_id` string. These unemployment rates are used as a soft proxy for labor market conditions that may shape both the placement of research funding and patenting capacity.

A5. National Science Foundation (NSF), Higher Education R&D Survey (HERD)

Use: Innovation control variable (university R&D expenditures).

HERD provides institution-level R&D expenditure data disaggregated by source of funding. Total academic R&D expenditures were aggregated to the county level using the HUD ZIP-to-county crosswalk. Federal data suppression codes (D, S, X, NA) were universally converted to treated as NA, and institutions lacking valid ZIP codes were excluded. This variable captures the academic research environment, an important predictor of regional innovative capacity and a potential attractor of federal R&D contracts.

Geographic Assignment: The HERD survey reports institution-level R&D expenditures but relies on ZIP codes rather than county FIPS for location. Institutions were mapped to counties using the HUD ZIP-to-County crosswalk. Approximately ~98 percent of institutions were successfully mapped.

Unmatched Institutions: A small number of institutions could not be mapped due to these ZIP code issues:

- ❖ **PO Box-Only ZIPs:** (e.g., campus mailcodes) which do not correspond to census geographies.
- ❖ **Campus-Only ZIPs:** Internal mail codes not present in USPS delivery data.
- ❖ **U.S. Territory ZIPs:** (e.g., Puerto Rico) which fall outside the county analysis structure.
- ❖ **Malformed Formats:** Truncated or incorrect digit sequences.
- ❖ **Missing ZIPs:** Encoded as "00000".

Aggregation: Because these unmatched cases do not systematically correlate with R&D intensity and represent a negligible share of national expenditures (~0.2 percent), they were dropped. For all mapped

institutions, total R&D expenditures were aggregated to the county-year level to capture the local academic research environment.

A6. Census Bureau Population Estimates (Intercensal and Postcensal Series)

Use: Population normalization; denominator for the dependent variable.

County-level population estimates for 2000-2024 were harmonized from three Census series with differing schemas. These data allow patenting to be expressed per 100,000 residents, reducing the influence of population size on measured innovation rates. County name normalization and FIPS assignment were validated using the 2024 TIGER/Line county shapefile.

Harmonization: To construct a complete county-year population panel from 2000-2024, three distinct Census population series were harmonized. The 2000-2010 series provided clean FIPS codes and required only reshaping. The 2010-2020 and 2020-2024 series, provided in non-standard formats, required substantial restructuring: metadata rows were removed, and county names were parsed from the combined “Geographic Area” field.

Normalization: Because later files lacked FIPS codes, county and state names were normalized to a standardized format (stripping suffixes like “Parish” or “Borough”) and merged to a modern county boundary shapefile to assign five-digit FIPS identifiers. Observations failing to match modern county definitions were dropped to ensure full mergeability across the master panel.

A7. Census TIGER/Line County Boundaries and Centroid Construction

Use: Spatial processing; centroid construction; map generation.

County polygons from the 2024 TIGER/Line cartographic boundary files were projected to EPSG:5070 (Conus Albers) to compute accurate centroids and reprojected to EPSG:4326 for great-circle distance calculations, following Duke University Libraries' geospatial research guide (Center for Data and Visualization Sciences). These centroids serve as the basis for all spatial measurements, including minimum-distance-to-treated computations.

Centroid Projection Bias: Distance calculations rely on geographic centroids. In large counties (e.g., San Bernardino, Los Angeles), the geometric centroid can be somewhat distant from the population or economic center. This introduces a known measurement error where the distance to treatment may differ from the effective distance to the labor market or source of innovative activity. This error is assumed to be random and strictly increases distance estimates in rural areas, making the spillover estimates conservative.

A8. HUD USPS ZIP-to-County Crosswalk

Use: University R&D geographic assignment.

The ZIP-to-county crosswalk provides allocation ratios that map each ZIP code to its most likely county of residence or activity. The county with the highest allocation ratio was chosen for each HERD institution. This method is standard practice for assigning institutions without explicit county identifiers.

A9. Panel Harmonization and Lag Construction

The master panel was assembled through sequential left joins onto a county-year skeleton. The skeleton was constructed as the Cartesian product of all unique FIPS codes and the union of all years

observed across constituent datasets (1986-2025). The distance-to-treated frame - derived from the spatial analysis of DoD funding locations - provided the initial skeleton with centroid coordinates. Variables were merged in the following order: (1) dependent variable (semiconductor patents); (2) patent stock; (3) independent variable (DoD obligations); (4) county GDP; (5) unemployment rate; (6) university R&D. After all merges, the panel was sorted by county and year, and within-county lag variables (1-, 2-, and 3-year) were generated for the DV, IV, and patent stock using grouped shifts. Population was merged last, and per-capita normalizations (patents per 100,000, DoD funding per 100,000) were computed as final derived variables.

Temporal Coverage: Because constituent datasets span different ranges, the master panel contains structural missingness in early years. The DV is undefined before 2013; university R&D is undefined before 2008; GDP is undefined before 2001. These missing values are inherent to source data availability and are handled via listwise deletion in regression specifications. The full temporal span is retained to maximize coverage for lagged IV specifications, where funding in year $t - 3$ may predict patenting in year t .

A10. Technical Notes on Data Completeness and Missing Locations

Geographic missingness in the master panel arises from structural features of federal statistical programs rather than random data loss. This section documents each category of missing locations and explains why these gaps do not threaten the validity of the analysis.

FIPS Prefi x	Territory	Counties	GDP	Unemployment	Population	DoD Funding (2013-23)
--------------------	-----------	----------	-----	--------------	------------	-----------------------

60	American Samoa	5	X	X	X	\$65k
66	Guam	1	X	X	X	\$84k
69	N. Mariana Islands	4	X	X	X	\$0
78	U.S. Virgin Islands	3	X	X	X	\$0

U.S Territories: BEA county GDP, BLS LAUS, and Census intercensal estimates do not cover U.S.

territories (Knoedl and Leubsdorf). These areas exhibit no semiconductor patenting activity and trivial DoD microelectronics funding (~\$149k total). Exclusion has no effect on results.

Puerto Rico (FIPS 72xxx): Puerto Rico is excluded from BEA county GDP tables and uses different BLS reporting (Knoedl and Leubsdorf). Census population estimates exist but were not merged due to different geographic identifiers (municipalities vs. counties). Puerto Rico has no semiconductor innovation cluster or semiconductor patents.

❖ **Scope:** 78 municipalities missing GDP, unemployment, and population.

❖ **DoD funding:** \$32k total (2013-2023), negligible.

Alaska Boroughs (FIPS 02063, 02066): Chugach (02063), Copper River (02066) were created from prior Census areas after 2019; BEA has not yet updated coverage. These boroughs receive no DoD funding and lack any semiconductor activity.

Hawaii Counties (FIPS 15005): Kalawao County is the former Kalaupapa leper colony and is frequently suppressed.

Connecticut Planning Regions (FIPS 09110-09190): The 2024 TIGER/Line shapefile uses newly assigned Census planning region FIPS codes (09110-09190), but BEA, BLS, and Census historical data use legacy county FIPS (09001-09015). No crosswalk was applied.

- ❖ **Scope:** 9 planning regions replaced Connecticut's 8 counties in 2022 Census geographic reorganization (United States Census Bureau).
- ❖ **DoD funding:** Although Connecticut received approximately \$239.9 million in microelectronics obligations (1986-2025), this represents a negligible share of national R&D and occurs in a state with no significant semiconductor fabrication or patenting clusters.

Virginia Independent Cities (38 FIPS codes): Virginia independent cities are legally separate from surrounding counties and receive distinct FIPS codes. BEA reports GDP for independent cities combined with adjacent counties, creating FIPS mismatches. This is the most consequential gap - several Virginia independent cities (particularly in Northern Virginia) receive substantial DoD microelectronics contracts.

- ❖ **DoD Funding:** Calculations from USAspending using individual FIPS codes put independent city spending at \$3.13B over 2022-2023, which is substantial.

Distance Variable Missingness: County-year observations with missing distance values (NaN) represent locations where spatial proximity could not be calculated (e.g., mismatched geographies). These are explicitly excluded from distance-based analyses to avoid treating unknown locations as "untreated" or "distant."

A11. Technical Notes on Patent-Dating, Fiscal-Calendar Adjustments

Patent Dating Convention: The analysis relies on patent grant years rather than application years to define the dependent variable. While this distinction does introduce a natural administrative lag between the conceptualization of the invention and the official grant, it aligns the dependent variable with the administrative availability of the data. This lag is mechanically absorbed by the lag structure of the independent variable (DoD obligations), which is explicitly tested at $t - 1$ through $t - 3$. Therefore, it may be more precise to say this approach ensures that the model captures the downstream effect of R&D investment on finalized patents.

Calendar Harmonization: A structural mismatch exists between federal spending data (reported by Fiscal Year, Oct 1-Sep 30) and all other datasets (reported by Calendar Year). To harmonize these series, DoD obligations had their precise calendar year from the `award_base_action_date` of each contract extracted. This transformation is critical because approximately 20 percent of federal contracts (those signed between October 1 and December 31) are assigned to the subsequent fiscal year in raw USAspending data. Reassigning these awards to their actual calendar year of execution, ensures that a contract signed in December 2015 is correctly treated as a 2015 economic injection, aligning it with the native calendar-year reporting of USPTO patent grants, BEA GDP, and BLS unemployment data. The only variable remaining in a fiscal-year format is University R&D (HERD), which is reported by academic fiscal year (typically ending June 30). This variable was not converted because monthly expenditure data was unavailable. However, because university research infrastructure is a slow-moving stock variable with low yearly variability, this minor offset is unlikely to introduce meaningful bias into the control structure.


Appendix B

Appendix B presents visual and statistical materials that support the empirical strategy and contextualize the main results. The appendix documents the spatial distribution of semiconductor patenting and DoD microelectronics R&D exposure, details the choices made in choropleth design, and provides the full regression output underlying the main specifications discussed in the Results section.


B1. Intensity of Private Semiconductor Patenting, 2017-2020.

The construction of Map 1 intentionally excludes Alaska, Hawaii, and other U.S territories. This decision is grounded in both substantive and technical considerations. Substantively, none of these regions exhibit measurable engagement in the semiconductor innovation ecosystem or DoD microelectronics R&D obligations during the study period, and no treated counties are located within these geographies. Technically, these areas rely on distinct administrative structures (e.g., Alaska's borough-census area system), and datasets such as BEA county GDP and BLS county unemployment frequently provide incomplete or suppressed coverage for them (A11). Including these regions would introduce artificial spatial discontinuity into a map whose purpose is to communicate the contiguous U.S. spillover environment.


County-level Semiconductor Patents per 100,000 Residents, 2017




County-level Semiconductor Patents per 100,000 Residents, 2018



County-level Semiconductor Patents per 100,000 Residents, 2019




County-level Semiconductor Patents per 100,000 Residents, 2020




B2: Geographic Distribution of DoD Microelectronics R&D Obligations


Distance-to-Treated Map (Year 2012)




Distance-to-Treated Map (Year 2016)



Distance-to-Treated Map (Year 2020)



Distance-to-Treated Map (Year 2024)



Map 2 visualizes the geographic "treatment" structure across four representative years: 2012, 2016, 2020, and 2024. Treated counties (defined as those receiving positive DoD microelectronics R&D obligations) are highlighted, surrounded by concentric distance bands (0-25 km, 25-50 km, 50-100 km). As detailed in the appendix, Alaska, Hawaii, and U.S territories are excluded to focus on the contiguous spatial spillover environment.

Map 2 does not match Map 1's year, because the purpose of this choropleth is not to show parallel visualization series but to show treatment geographies, which does not vary meaningfully in a year-to-year context, as treatments tend to cluster and persist. Mapping the same bands in 2017-2020 would produce nearly identical visual patterns with minor noise, because DoD R&D awards are not randomly redistributed every year. Per the time-series (Fig. 5), these four periods cover a wide range of policy-relevant periods.

B3: Full Regression Output

This subsection reports the full regression output corresponding to Table 1 of the main text. All coefficients, including year fixed effects and control variables omitted from the main table for readability, are reported here. The purpose of this appendix is to provide full transparency regarding the estimation results without distracting from the primary hypotheses tested in the main analysis. Across all specifications, county fixed effects and year fixed effects are included, and standard errors are clustered at the county level. The inclusion of year fixed effects generates large and statistically significant coefficients for the post-2014 period, reflecting the nationwide surge in semiconductor

Table 1: Effect of DoD Microelectronics R&D on Private Semiconductor Patenting

	Model (1)	Model (2)	Model (3)	Model (4)
	Distance Bands (<i>Spatial Exposure</i>)	Funding Intensity (<i>t - 1, Two-Way FE</i>)	Distributed Lag (<i>Two-Way FE</i>)	Funding Intensity (<i>Excl. Top 1%</i>)
0_25	0.020 [-2.009,2.050]			
25_50	2.190 [-2.525,6.906]			
50_100	-1.608 [-6.119,2.904]			
100_plus	-1.187 [-4.393,2.019]			
Patent Stock	0.019*** [0.010,0.029]	0.023*** [0.009,0.036]	0.014 [-0.006,0.033]	0.022*** [0.007,0.036]
GDP	-0.000** [-0.000,-0.000]	-0.000* [-0.000,0.000]	-0.000 [-0.000,0.000]	-0.000 [-0.000,0.000]
Unemployment Rate	0.740* [-0.037,1.518]	0.857 [-0.629,2.342]	0.634 [-0.765,2.033]	0.896 [-1.005,2.796]
University R&D	-0.000 [-0.000,0.000]	-0.000 [-0.000,0.000]	-0.000 [-0.000,0.000]	-0.000 [-0.000,0.000]
year=2013	0.000 [0.000,0.000]	0.000 [0.000,0.000]	0.000 [0.000,0.000]	0.000 [0.000,0.000]
year=2014	30.104*** [9.476,50.733]	27.006*** [8.775,45.237]	19.589*** [4.833,34.346]	42.116*** [15.288,68.945]
year=2015	38.316*** [14.181,62.451]	34.409*** [12.629,56.189]	25.623*** [7.184,44.062]	50.806*** [20.286,81.327]
year=2016	40.343*** [14.550,66.137]	35.191*** [12.193,58.189]	27.438*** [7.500,47.376]	52.079*** [20.046,84.113]
year=2017	40.953*** [14.117,67.790]	35.444*** [11.823,59.064]	27.657*** [6.960,48.354]	52.546*** [19.750,85.341]
year=2018	40.359*** [13.508,67.210]	34.683*** [11.522,57.845]	26.113*** [6.591,45.635]	51.592*** [19.258,83.926]
year=2019	40.709*** [13.669,67.748]	36.424*** [11.628,61.219]	27.366*** [7.333,47.398]	53.568*** [19.489,87.646]
year=2020	36.786*** [12.673,60.898]	30.342*** [10.643,50.040]	23.938*** [7.720,40.156]	47.411*** [18.480,76.342]
year=2021	36.708*** [12.632,60.784]	32.772*** [10.806,54.738]	23.761*** [5.928,41.595]	50.124*** [18.830,81.418]

Table 1: Effect of DoD Microelectronics R&D on Private Semiconductor Patenting

	Model (1)	Model (2)	Model (3)	Model (4)
	Distance Bands (<i>Spatial Exposure</i>)	Funding Intensity (<i>t - 1, Two-Way FE</i>)	Distributed Lag (<i>Two-Way FE</i>)	Funding Intensity (<i>Excl. Top 1%</i>)
year=2022	34.205*** [11.604,56.806]	31.237*** [10.280,52.194]	23.586** [4.009,43.163]	47.547*** [17.913,77.182]
year=2023	34.504*** [11.841,57.167]	33.242*** [10.529,55.954]	23.755** [3.660,43.850]	49.531*** [18.144,80.917]
iv_millions_l1		0.106* [-0.005,0.217]	0.043* [-0.005,0.091]	0.158** [0.021,0.295]
iv_millions_l2			-0.007 [-0.038,0.025]	
iv_millions_l3			0.013 [-0.018,0.043]	
Constant	-14.872 [-37.960,8.216]	2.463 [-26.978,31.904]	3.094 [-33.518,39.706]	-2.384 [-57.750,52.983]
Observations	1974.000	1126.000	883.000	927.000
Counties	295.000	202.000	138.000	181.000

95% Confidence Intervals Reported In brackets.

DV: semiconductor patents per 100k residents.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

patenting documented in the descriptive time series (Fig. 5). These coefficients are reported here for completeness but are not interpreted substantively in the main text, as they absorb common national shocks rather than county-level treatment effects.

Control Variable Estimates: Several controls, most notably population and university R&D

expenditures, are imprecisely estimated once county fixed effects are included. This reflects the highly persistent and structural nature of these variables rather than a lack of substantive relevance. On the other hand, GDP and university R&D primarily capture long-run cross-sectional differences in economic scale and research capacity across counties; because identification in two-way fixed effects

models relies on within-county variation over time, these slow-moving variables contribute limited identifying variation after fixed effects and accumulated patent stock are accounted for. For this reason, the main text focuses on treatment coefficients and key theoretically relevant controls, while full control estimates are reported here for completeness and transparency.

Alternative Scaling of the Treatment Variable: To improve interpretability of coefficient magnitudes, the Department of Defense microelectronics R&D treatment variable is scaled in millions of dollars per 100,000 residents in the main specifications. This normalization aligns the scale of the independent variable with the dependent variable, which is measured as semiconductor patents per 100,000 residents, and facilitates a more readable interpretation of effect sizes.

Works Cited

- Acosta, Manuel, et al. "Patents and Dual-Use Technology: An Empirical Study of the World's Largest Defence Companies." *MPRA Paper*, MPRA Paper, no. 123085, 2018. *ideas.repec.org*, <https://ideas.repec.org/p/pra/mprapa/123085.html>.
- Bendig, David, et al. "A Multi-Level Analysis of the Relationship between Formal Institutions and Ambitious Entrepreneurship: The Moderating Role of Economic Development." *International Business Review*, vol. 33, no. 5, Oct. 2024, p. 102303. *ScienceDirect*, <https://doi.org/10.1016/j.ibusrev.2024.102303>.
- Bureau of Economic Analysis. "Chained-dollar estimates." U.S. Bureau of Economic Analysis (BEA), 25 Apr. 2018, <https://www.bea.gov/help/glossary/chained-dollar-estimates>.
- Bureau of Economic Analysis. "GDP by County, Metro, and Other Areas." U.S. Bureau of Economic Analysis (BEA), 29 Aug. 2025, <https://www.bea.gov/data/gdp/gdp-county-metro-and-other-areas>.
- Bureau of Labor Statistics. "download.bls.gov - /pub/time.series/la/." U.S. Bureau of Labor Statistics, n.d., <https://download.bls.gov/pub/time.series/la/>.
- Bureau of Labor Statistics. "Local Area Unemployment Statistics Overview." U.S. Bureau of Labor Statistics, 17 Mar. 2025, <https://www.bls.gov/lau/lauov.htm>.
- Center for Data and Visualization Sciences. "LibGuides: R: Mapping and Geospatial: Projections / CRS." *Duke University Libraries*, 23 Sept. 2025, <https://guides.library.duke.edu/r-geospatial/CRS>.

- Chowdhury, Farhat, et al. "Public Support for Research in Artificial Intelligence: A Descriptive Study of U.S. Department of Defense SBIR Projects." *The Journal of Technology Transfer*, vol. 47, no. 3, June 2022, pp. 762–74. *Springer Link*, <https://doi.org/10.1007/s10961-022-09943-z>.
- Corrocher, Nicoletta, et al. *Innovation and the Labor Market: Theory, Evidence and Challenges*. IZA - Institute of Labor Economics, 2023. *JSTOR*, <https://www.jstor.org/stable/resrep67525>.
- Council, National Research. *An Assessment of the SBIR Program*. Edited by Charles W. Wessner, The National Academies Press, 2008, <https://doi.org/10.17226/11989>.
- Decarolis, Francesco, et al. "Buyers' Role in Innovation Procurement: Evidence from US Military R&D Contracts." *Journal of Economics & Management Strategy*, vol. 30, no. 4, 2021, pp. 697–720. *Wiley Online Library*, <https://doi.org/10.1111/jems.12430>.
- Johnson, Sandra, and David Armstrong. "Methodology, Limitations and Applications of the 2010-2020 Intercensal Population and Housing Unit Estimates." *Research Matters* (U.S. Census Bureau), 7 Nov. 2024, <https://www.census.gov/newsroom/blogs/research-matters/2024/11/2010-2020-intercensal-population-and-housing-unit-estimates.html>.
- Fuchs, Erica R. H. "Rethinking the Role of the State in Technology Development: DARPA and the Case for Embedded Network Governance." SSRN Scholarly Paper no. 1545155, Social Science Research Network, 7 June 2010. *papers.ssrn.com*, <https://papers.ssrn.com/abstract=1545155>.
- Goodman, Matthew P., and Dylan Gerstel. *Investing in Innovation*. Center for Strategic and International Studies (CSIS), 2020, pp. 5–15. Sharpening America's Innovative Edge. *JSTOR*, <https://www.jstor.org/stable/resrep26537.5>.

- Hall, Bronwyn H., et al. "Market Value and Patent Citations." *The RAND Journal of Economics*, vol. 36, no. 1, 2005, pp. 16–38. *JSTOR*, <https://www.jstor.org/stable/1593752>.
- Hausman, Jerry, et al. "Econometric Models for Count Data with an Application to the Patents-R & D Relationship." *Econometrica*, vol. 52, no. 4, 1984, pp. 909–38. *JSTOR*, <https://doi.org/10.2307/1911191>.
- Howell, Sabrina T., et al. "Opening Up Military Innovation: Causal Effects of Reforms to US Defense Research." *Journal of Political Economy*, vol. 133, no. 11, Nov. 2025, pp. 3605–51. *journals.uchicago.edu (Atypon)*, <https://doi.org/10.1086/737235>.
- Imai, Kosuke, and In Song Kim. "On the Use of Two-Way Fixed Effects Regression Models for Causal Inference with Panel Data." *Political Analysis*, vol. 29, no. 3, July 2021, pp. 405–15. *Cambridge University Press*, <https://doi.org/10.1017/pan.2020.33>.
- Kenney, Martin, and Urs Burg. "Technology, Entrepreneurship and Path Dependence: Industrial Clustering in Silicon Valley and Route 128." *Industrial and Corporate Change*, vol. 8, Feb. 1999, pp. 67–103. *ResearchGate*, <https://doi.org/10.1093/icc/8.1.67>.
- Knoedl, Taylor R, and Ben Leubsdorf. *Federal Statistical Data for U.S. Territories: Issues and Resources*. Legislation. <https://www.congress.gov/crs-product/R48522>.
- Linville, John G., and C. Lester Hogan. "Intellectual and Economic Fuel for the Electronics Revolution." *Science*, vol. 195, no. 4283, 1977, pp. 1107–13. *JSTOR*, <https://www.jstor.org/stable/1743557>.

- Maliranta, Mika, et al. “Is Inter-Firm Labor Mobility a Channel of Knowledge Spillovers? Evidence from a Linked Employer–Employee Panel.” *Industrial and Corporate Change*, vol. 18, no. 6, Dec. 2009, pp. 1161–91. *Silverchair*, <https://doi.org/10.1093/icc/dtp031>.
- Moretti, Enrico, et al. *The Intellectual Spoils of War? Defense R&D, Productivity and International Spillovers*. IZA - Institute of Labor Economics, 2019. *JSTOR*, <https://www.jstor.org/stable/resrep59444>.
- Mowery, David C. “Federal Policy and the Development of Semiconductors, Computer Hardware, and Computer Software: A Policy Model for Climate Change R&D?” *NBER Chapters*, 2011, pp. 159–88. *ideas.repec.org*, <https://ideas.repec.org/h/nbr/nberch/11753.html>.
- Mowery, David C. “Military R&D and Innovation.” *Handbook of the Economics of Innovation*, vol. 2, Elsevier, 2010, pp. 1219–56. *DOI.org (Crossref)*, [https://doi.org/10.1016/S0169-7218\(10\)02013-7](https://doi.org/10.1016/S0169-7218(10)02013-7).
- National Center for Science and Engineering Statistics. “Higher Education Research and Development (HERD) Survey: FY 2023.” National Science Foundation, 2024, <https://nces.nsf.gov/surveys/higher-education-research-development/2023>.
- Nicastro, Luke A. *The U.S. Defense Industrial Base: Background and Issues for Congress*. Legislation. 23 Sept. 2024, <https://www.congress.gov/crs-product/R47751>.
- Pallante, Gianluca, et al. “Does Public R&D Funding Crowd-in Private R&D Investment? Evidence from Military R&D Expenditures for US States.” *Research Policy*, vol. 52, no. 8, Oct. 2023, p. 104807. *DOI.org (Crossref)*, <https://doi.org/10.1016/j.respol.2023.104807>.

- PatentsView. “Data Download Tables.” PatentsView, n.d.,
<https://patentsview.org/download/data-download-tables>.
- Paunov, Caroline, et al. “Cross-Country Evidence on the Contributions of Research Institutions to Innovation.” *OECD Science, Technology and Industry Policy Papers*, Sept. 2019.
www.oecd.org, <https://doi.org/10.1787/d52d6176-en>.
- Platzer, Michaela D., and John F. Sargent Jr. “U.S. Semiconductor Manufacturing: Industry Trends, Global Competition, Federal Policy.” Legislation. *Congress.Gov*,
<https://www.congress.gov/crs-product/R44544>.
- Ponds, Roderik, et al. “Innovation, Spillovers and University—Industry Collaboration: An Extended Knowledge Production Function Approach.” *Journal of Economic Geography*, vol. 10, no. 2, 2010, pp. 231–55. *JSTOR*, <https://www.jstor.org/stable/26161397>.
- Reinsch, William Alan, et al. *Taking Stock of Government Involvement in Research and Development: Assessing Public R&D Effects*. Center for Strategic and International Studies (CSIS), 2020. *JSTOR*, <https://www.jstor.org/stable/resrep24776>.
- Rep. Ryan, Tim [D-OH-13. H.R.4346 - 117th Congress (2021-2022): CHIPS and Science Act. nos. 117–167, 9 Aug. 2022, 2021-07-01, <https://www.congress.gov/bill/117th-congress/house-bill/4346>.
- Robbins, Carol, and B Cosby. *One-Fifth of USPTO Semiconductor Utility Patents Granted to Residents of Santa Clara County, California between 2000 and 2020*. NSF 23-342, 5 June 2023, <https://nces.nsf.gov/pubs/nsf23342/>.

- Rothwell, Jonathan, et al. "Patenting Prosperity: Invention and Economic Performance in the United States and Its Metropolitan Areas." *Brookings Institute*. Zotero, <https://www.brookings.edu/articles/patenting-prosperity-invention-and-economic-performance-in-the-united-states-and-its-metropolitan-areas/>.
- Samila, Sampsa, and Olav Sorenson. "Venture Capital as a Catalyst to Commercialization." *Research Policy*, vol. 39, no. 10, Dec. 2010, pp. 1348–60. *ScienceDirect*, <https://doi.org/10.1016/j.respol.2010.08.006>.
- Statistics (NCSES), National Center for Science and Engineering. "Invention, Knowledge Transfer, and Innovation." *Invention Indicators: Protecting Useful Ideas Contributors to U.S. Patent Activity*, NSB-2024-1, Feb. 2024, <https://ncses.nsf.gov/pubs/nsb20241/invention-indicators-protecting-useful-ideas>.
- Suzuki, Kohei, and Mehmet Akif Demircioglu. "The Association Between Administrative Characteristics and National Level Innovative Activity: Findings from a Cross-National Study." *Public Performance & Management Review*, vol. 42, no. 4, July 2019, pp. 755–82. *Taylor and Francis+NEJM*, <https://doi.org/10.1080/15309576.2018.1519449>.
- Tomayko, James E. "Computers Take Flight: A History of NASA's Pioneering Digital Fly-By-Wire Project." 1 Jan. 2000. *NASA NTRS*, <https://ntrs.nasa.gov/citations/20050157919>.
- U.S. Census Bureau. "Cartographic Boundary Files." U.S. Census Bureau, 10 Apr. 2025, <https://www.census.gov/geographies/mapping-files/time-series/geo/cartographic-boundary.html>.

United States Census Bureau. “Change to County-Equivalents in the State of Connecticut.”

Department of Commerce Census Bureau, 6 June 2022,

<https://www.federalregister.gov/documents/2022/06/06/2022-12063/change-to-county-equivalents-in-the-state-of-connecticut>

U.S. Census Bureau. “County Intercensal Datasets: 2000-2010.” U.S. Census Bureau, 16 Dec.

2021, <https://www.census.gov/data/datasets/time-series/demo/popest/intercensal-2000-2010-counties.html>.

U.S. Census Bureau. “TIGER/Line Shapefiles.” U.S. Census Bureau, n.d.,

<https://www.census.gov/geographies/mapping-files/time-series/geo/tiger-line-file.html>.

U.S. Department of Housing and Urban Development, Office of Policy Development and

Research. “HUD USPS ZIP Code Crosswalk Files.” HUD User, n.d.,

https://www.huduser.gov/portal/datasets/usps_crosswalk.html.

U.S. General Services Administration. Product and Service Codes (PSC) Manual.

Acquisition.gov, 30 Apr. 2024, <https://www.acquisition.gov/psc-manual>.

USAspending.gov. “Award Data Archive.” USAspending.gov, n.d.,

https://www.usaspending.gov/download_center/award_data_archive.

USAspending.gov. Federal Spending Guide. USAspending.gov, 26 Apr. 2023,

<https://www.usaspending.gov/data/Federal-Spending-Guide.pdf>.