

Local Returns to Federal Innovation Investment: Evidence from the National Laboratories

Susan Helper, Resem Makan, and Daniel Shoag*

Case Western Reserve University

June 30th, 2026

Abstract

Federal innovation investments generate large social returns, but whether they can seed new innovation ecosystems far from the technological frontier remains an open question. The U.S. national laboratories offer an unusually informative setting: their locations were chosen for security and isolation rather than economic potential, and they combine large federal investments with decades of outcome data. Drawing on newly digitized historical records, we find that host counties experienced large and persistent gains in patenting by non-laboratory inventors, retail sales, income, and educational attainment relative to counterfactual counties. These effects grew over time; two decades after founding, every dollar of laboratory operating budgets was associated with over five dollars of additional local economic activity. The returns varied widely across laboratories: university-affiliated labs generated roughly 2.6 times the per-dollar spillovers of privately operated facilities. This variation suggests the importance of governance, publication norms, and personnel mobility in spreading knowledge to local firms. We argue these gains reflected genuine productivity improvements rather than spatial reallocation: pre-existing residents saw significant wage gains (about 25 percent for college-educated workers between 1940 and 1950), and the former collaborators of relocating scientists (perhaps the most likely group to be negatively impacted) show no detectable decline in productivity. Because the laboratories were not designed to maximize local development, these estimates likely represent a lower bound on what deliberately-designed investments could achieve. The findings speak directly to current debates over place-based innovation policy, as governments worldwide commit tens of billions to regional clusters even as the U.S. innovation system faces deep proposed cuts.

*We thank Josh Lerner, Ben Jones, Andrew Fieldhouse, Jeffrey Alexander, Nick Barendt, and Andrew Schrank for their helpful comments. We are grateful to the National Science Foundation Technology, Innovation, and Partnerships Directorate for financial support. Lastly, we thank Junyang Wang and Isma'il Seddon for providing excellent research assistance.

1 Introduction

It is well established that the social returns to innovation in general are very high. Knowledge spillovers and other appropriability problems mean that private markets underinvest in research ([Arrow, 1962](#); [Nelson, 1959](#)); quantitative estimates confirm large gaps between social and private returns ([Jones and Summers, 2020](#)). But it's less clear whether federal innovation investments can reliably spark new ideas and economic activity, especially in places far from the existing technological frontier.

This question is timely. Catalyzing innovation-based regional ecosystems that deliver broad economic and societal benefits has been a major focus of bipartisan policy and academic discussions in recent years. For instance, [Gruber and Johnson \(2019\)](#) proposed to “jump-start America” by establishing competitive federal grants to help municipalities become new technology hubs. Partially in response to this work, the U.S. Congress in 2021 and 2022 authorized \$80 billion for regions to “invest in clusters of technology, innovation, and competitiveness.”¹ These initiatives include Tech Hubs, RECOMPETE, and the Build Back Better Regional Challenge. Similarly, the NSF Engines program “aims to fund regional coalitions of partnering organizations to catalyze technology and science-based regional innovation ecosystems.”

This activity is by no means limited to the United States. For example, the European Union has developed a methodology for “Regional Innovation Strategies for Smart Specialisation” over the last decade and continues to invest in “Innovation for place-based transformations”. A significant part of the ongoing 7-year, 94 billion euro “Horizon Europe” program is devoted to building “excellence hubs” and a “European Innovation Ecosystems (EIE) programme that strengthens innovation ecosystems. Individual nations continue to invest in such programs, such as the Catapult network in the United Kingdom, and the 800 clusters that make up the European Cluster Alliance.

The question remains: Do these efforts to foster innovation actually result in meaningful and sustained economic gains for the regions they target? Answering this question is empirically difficult because government (and other) R&D dollars typically flow to

¹[Breaking down an \\$80 billion surge in place-based industrial policy | Brookings](#)

elite research universities and existing clusters — places where complementary human capital and institutions are already present and where innovation would likely have increased even absent the investment (Howell, 2024). Moreover, public funding can crowd out private effort, political incentives can distort targeting, and resources directed to one institution or place may come at the expense of another (Bloom et al., 2019).

This paper uses the case of the U.S. national laboratories to make progress and suggest future research on this question. The laboratories combine three features rarely observed together: large ongoing federal research investments, plausibly exogenous location decisions, and long-run data spanning several decades. The national laboratories provide a particularly informative setting because they sidestep the usual problem of “identification”, that is, of separating correlation and causation. Laboratory locations were not chosen to maximize economic spillovers — they were chosen to protect security and secrecy, thus pushing facilities toward remote areas with very little prior research capacity. Evidence from newly digitized historical data shows that counties in which laboratories were located experienced large and persistent increases in innovation and economic activity — including patenting by non-laboratory-affiliated inventors, shifts in local invention toward laboratory research fields, and sustained growth in retail sales, income, and educational attainment (Helper et al., 2026).

Given that substantial innovation ecosystems emerged even under these unpromising conditions, the resulting effects likely represent a lower bound on the local spillover returns to federal research infrastructure. Returns to programs that have been deliberately designed to leverage local institutions and human capital would likely be much higher, as we discuss below. Moreover, because knowledge is nonrival and diffuses beyond the county, the local economic gains we measure capture only part of the total social return.

In this chapter, we review the evidence that federal investment generated large local spillovers, even starting from near-zero innovative capacity, drawing on our work on national labs (see Helper et al. (2026)). Next, we examine what predicts the largest returns across laboratories. The most consistent finding is that university connections and institutional embeddedness amplify spillovers. Finally, we provide suggestive evidence that local gains are not merely offset by losses elsewhere. Wage gains for residents who

lived in treated counties before laboratory establishment indicate that incumbents' productivity genuinely increased. Unlike wage increases driven by in-migration — which can bid up local prices and draw workers from other locations — productivity increases for people already in a place are more likely to be due to net output gains rather than to spatial reallocation. These wage gains are the component of local effects that, in the spatial equilibrium framework of [Kline and Moretti \(2014\)](#), most reliably translates into aggregate welfare improvement. Similarly, we do not find a decline in the research productivity of scientists' former collaborators. We discuss one possible mechanism for these results, a "big push" interpretation in which federal investment triggered self-sustaining innovation trajectories rather than merely reallocating activity across space.

These findings connect to a growing body of evidence on federal R&D and innovation policy. We know that public research spending can generate large social returns, that wartime and biomedical programs can crowd in private innovation, and that the design of intermediary institutions matters for whether knowledge diffuses broadly ([Gross and Sampat, 2023](#); [Azoulay et al., 2019](#); [Moretti et al., 2025](#); [Howell, 2024](#)). Universities, in part due to federal funding, play a key role, as shown by [Liu \(2015\)](#), [Lee \(2019\)](#) and [Andrews \(2023\)](#). We also know that national laboratories occupy an important but comparatively understudied place in that policy landscape, especially relative to university funding and firm-level subsidies ([Siegel et al., 2023](#); [Guzman et al., 2024](#)). What we know much less well is whether large federal research facilities can seed durable local innovation ecosystems in places with little preexisting innovative capacity, which institutional features make that outcome more likely, and whether local gains mainly reflect new productivity rather than spatial reallocation.

The national laboratories help fill these gaps because they combine unusually large federal investments, historically idiosyncratic siting decisions, and meaningful variation in organizational form. Their history therefore lets us examine not only whether federal research infrastructure can generate spillovers, but also what kinds of institutional design and other features tend to increase those spillovers into the surrounding economy.

The paper proceeds as follows. Section 2 lays out a nontechnical conceptual framework. Section 3 introduces the national laboratories. Section 4 summarizes our main

empirical findings. Section 5 examines heterogeneity. Section 6 discusses crowd-out and equilibrium dynamics. Section 7 draws implications for current policy and Section 8 provides suggestions for future research and concludes.

2 Conceptual Framework

To organize the analysis, we present a simple, nontechnical framework for thinking about how federal research investments translate into innovation and economic outcomes.

Innovation production. Consider a location that receives a federal research investment, for example in a national lab. The change in local innovation has two components: the laboratory's own output and innovation by non-lab entities that is induced by the activities of the lab. This second, or spillover component, depends importantly on presence of complementary investments nearby. The same federal investment will generate larger spillovers when effective local complements are present or are developed.

Three features of the local environment are especially important in determining local complementarity. Human capital determines the extent to which local actors can absorb frontier knowledge. Nearby institutions, including universities, supplier networks, and entrepreneurial firms, give local actors places to use and extend that knowledge. And information transmission, shaped by management quality, collaboration mechanisms, and the degree to which research is classified or openly shared, determines whether ideas diffuse widely or remain inside the laboratory.

This complementarity is central to our lower-bound interpretation of the national-laboratory evidence. Many laboratories were located in places with few complements: i.e., they had limited human capital, thin local institutions, and sometimes significant secrecy, yet they still generated large innovation gains. Similar investments designed from the outset to leverage stronger local ecosystems would plausibly generate even larger spillovers.

One might worry about the opposite — that laboratories generated outsized local effects precisely because they were the only game in town, and that the same investment in a denser research environment would get lost in the noise. We address that concern

empirically in Section 5: we find that laboratories with stronger university ties and richer local institutional environments generated larger spillovers per dollar, not smaller ones. This cross-laboratory evidence points to increasing returns from complementarity rather than diminishing returns from saturation.

From innovation to economic activity. Local economic gains can arise through at least two pathways. One is direct federal spending for payroll, procurement, and construction expenses related to the labs. The other is spillover innovation that raises productivity for local firms and workers.

The distinction matters because direct spending effects resemble a standard fiscal multiplier, while innovation spillovers can generate much larger and more persistent returns. The empirical challenge is to separate short-run demand effects from longer-run changes in local productive capacity.

Dynamics. If federal investment raises local innovation, which in turn trains workers, attracts firms, and deepens supporting institutions, later spillovers can become larger than initial ones. Under those conditions, a sufficiently large federal investment can serve as a “big push” that moves a region from a low-innovation to a high-innovation trajectory. We return to this mechanism in Section 6.

This framework highlights three questions that guide the empirical discussion:

- (i) Does federal research infrastructure increase local innovation outside the laboratory itself?
- (ii) Through what channels does that innovation translate into broader economic activity?
- (iii) Which complementary institutions and local conditions amplify or limit these effects?

In Sections 4-6 we use these concepts as an organizing device, not a structural model. Laboratory operating budgets and facilities capture the scale of federal research investment. We use changes in education to indicate changes in local human capital. University ties and surrounding industrial or research infrastructure capture complementary institutions. Differences in openness, classification, and collaboration norms help indicate how easily knowledge diffuses. This mapping helps identify which local complements raise

the returns to federal research capital.

3 The National Laboratories

The U.S. national laboratory system emerged from wartime mobilization and postwar expansion of federal scientific institutions. The Manhattan Project created large federally financed sites where scientists, engineers, production facilities, and specialized equipment were co-located to perform an urgent national mission. Rather than dismantling that model after World War II, policymakers adapted and expanded it during the early Cold War for atomic energy, weapons design, reactors, materials science, and eventually a broader portfolio of mission-oriented research ([Gross and Sampat, 2023](#); [Howell, 2024](#)). Within Howell's taxonomy of innovation policy arenas, laboratories therefore represent a fairly direct form of government intervention: public provision of research infrastructure at a large scale. All the labs but one became government-owned, contractor-operated facilities: publicly owned laboratories run day to day by a university, university consortium, nonprofit, or industrial manager. The GOCO model was meant to combine a public mission with long-term funding with managerial flexibility in hiring, collaboration, and research administration ([Jaffe and Lerner, 2001](#)). As we discuss below, the nature of the contractor is not a minor administrative detail; it helps determine how open the laboratory is to local universities and firms, how easily students and researchers move through it, and how readily ideas spill over into the surrounding economy.

Their wartime and early Cold War origins also explain why laboratories were located where they were. Sites were chosen primarily for security, isolation, access to land and power, and political feasibility rather than because they sat inside existing innovation hubs ([Helper et al., 2026](#)). Los Alamos was placed in a county with roughly 200 residents; Oak Ridge was built as a secret city in rural Tennessee; the Idaho site was selected in part because it was at least 50 miles from any city of 50,000 people. Even Brookhaven, though closer to major universities, occupied a former military site on Long Island chosen as a compromise among federal officials and the universities involved, not as an effort to maximize local spillovers. Because these locations were not selected based on existing

regional innovation clusters, they provide an unusually credible setting for studying the causal impact of federal research infrastructure.

For a subset of laboratories, archival sources identify runner-up locations that narrowly lost the siting competition — sometimes for arbitrary reasons. These provide close counterfactuals for causal inference, and our companion paper (Helper et al., 2026) uses them alongside synthetic control methods to estimate the effects of laboratory establishment.

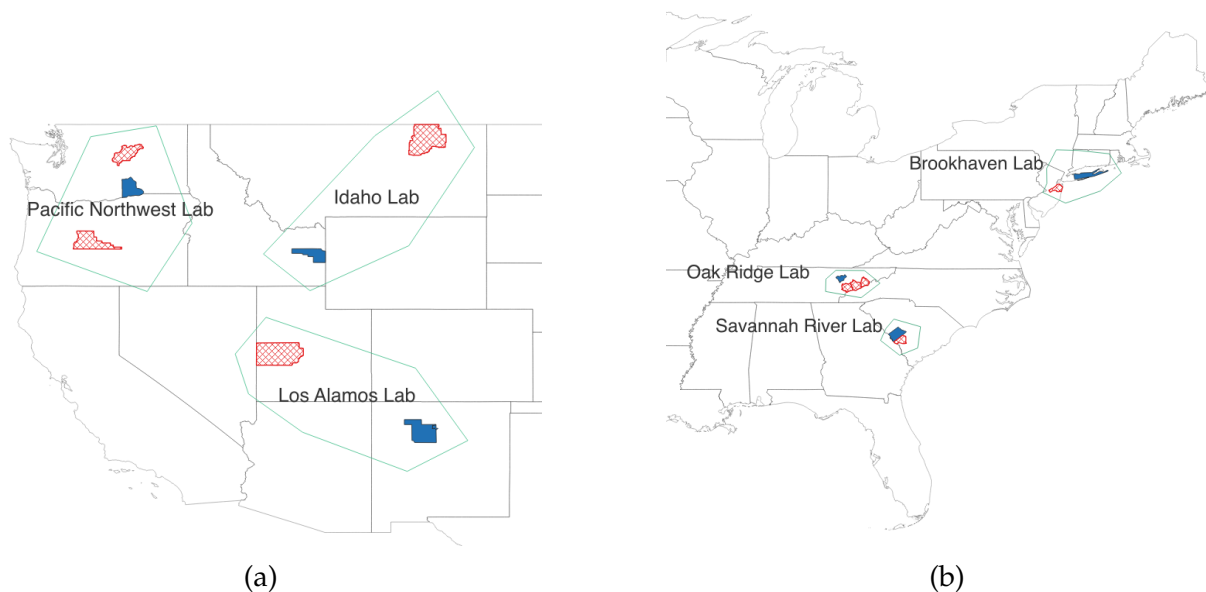


Figure 1: Map of National Laboratory Locations and Runner-Up Sites. *Source: Helper et al. (2026).*

Alt text: Map of the United States showing the locations of National Laboratories alongside the locations of their respective runner-up candidate sites.

Despite these seemingly unpromising beginnings, robust regional ecosystems did develop around most of the labs. One example is Argonne National Laboratory, located in Lemont, Illinois. Established in 1946 as the first national lab in the U.S., Argonne evolved from the University of Chicago’s Metallurgical Laboratory, which was instrumental in the Manhattan Project. Initially, its mission centered on developing nuclear reactors for both peaceful and defense purposes. Argonne’s rural location was specifically chosen to facilitate safe testing of nuclear technologies (previously the testing was done underneath the University of Chicago’s football stadium, which was located in a populated area). Argonne had 2,727 employees in 1957, 2.5% of the total workforce of DuPage County. In

2024, it had 3,836 full-time employees and a budget of \$1.2 billion.² Over the years, Argonne has expanded its research scope to include a broad array of scientific disciplines, such as physics, chemistry, biology, and engineering. Its collaborations with universities, industries, and government agencies have been instrumental in solving complex scientific challenges and driving innovations in a variety of fields, including energy storage, supercomputing, and quantum technologies.

4 Evidence: Local Returns to Federal Research Infrastructure

We summarize findings from historical patent records, newly digitized county-level economic data from the Survey of Buying Power (1936–1971)³, and linked 1940–1950 Census records, drawing on multiple identification strategies: synthetic controls, runner-up comparisons, and pooled event-study estimators.⁴ We organize the evidence around the framework’s three questions.

4.1 Does federal research infrastructure increase local innovation outside the laboratory?

Counties hosting national laboratories experienced dramatic increases in patent activity. Anderson County, Tennessee — site of Oak Ridge National Laboratory — rose from the bottom quintile of the U.S. patent distribution in 1940 to near the top 5 percent by 1975. Crucially, these increases extend well beyond laboratory-affiliated patents. Patenting by non-lab inventors rises sharply, with a lag of roughly five to ten years — consistent with gradual knowledge diffusion rather than an immediate demand shock (see Figure

²Source: <https://www.anl.gov/reference/argonne-by-the-numbers>.

³See [Helper et al., 2026](#) for details on this recently discovered data.

⁴Due to data limitations, we restrict our sample to 12 labs. One (NETL) was founded before our data range (1910). Two of them—NREL and Thomas Jefferson Lab—were founded in the 1970s and later, outside of our data range. Another two were founded in the same counties as existing laboratories but in later years. Fermi Lab was founded in DuPage County, Illinois, home to Argonne National Laboratory, and Lawrence Livermore National Laboratory was founded in Alameda County, California, home to Lawrence Berkeley National Laboratory. See [Helper et al. \(2026\)](#) for more details.

2). Non-laboratory patents shift toward the technological fields in which laboratories are active, and citation patterns show that nearby inventors disproportionately build on laboratory research, consistent with the classic finding that knowledge flows are meaningfully localized (Jaffe et al., 1993).




Figure 2: Pooled Event Study — Non-Lab Patenting in Laboratory vs. Control Counties. Source: [Helper et al. \(2026\)](#).

Alt text: Event study plot showing the difference in non-laboratory patenting rates between laboratory counties and synthetic control counties before and after the establishment of the laboratories.

4.2 Through what channels does innovation translate into broader economic outcomes?

Innovation spillovers translated into broader economic gains. We find sustained increases in retail sales and household income relative to synthetic controls. The most striking pattern is the trajectory of the relationship between laboratory budgets and local activity. In 1948, every dollar of the lab operating budget was associated with roughly 48 cents in lo-

cal retail sales impact — a ratio below one, consistent with direct spending effects alone. By 1970, every budget dollar was associated with more than five dollars of additional local economic activity relative to control counties. We emphasize that this is a descriptive ratio, not a spatial general equilibrium–adjusted fiscal multiplier. In a pure spending model,⁵ this ratio would remain roughly constant; its dramatic growth suggests that laboratories catalyzed broader innovation ecosystems whose economic effects compound over time.

These growing ratios substantially exceed the typical local fiscal multiplier of roughly 1.8 (Chodorow-Reich, 2019). They also contrast with the experience of NASA’s Space Race spending, where Kantor and Whalley (2025) estimate a fiscal multiplier of about 0.3 and find limited spillovers beyond firms who received government contracts directly. The two estimates are not directly comparable — Kantor and Whalley’s is a spatial GE-corrected estimate while ours is a simple ratio — but the contrast in trajectories is striking: NASA’s effects are flat over time while the laboratories’ grow. The difference is consistent with evidence that sustained research infrastructure generates larger and more persistent spillovers than does application-oriented R&D and procurement (Akcigit, Hanley, and Serrano-Velarde 2020; Fieldhouse and Mertens 2024).

The human capital channel further explains these growing returns. Using linked Census records for 1940 and 1950, we found that cohorts exposed to laboratory establishment during school-age years attained more education and were substantially more likely to complete high school, with effects appearing within four to seven years of laboratory founding. In the language of the framework above, laboratories increase the local stock of trained workers and the region’s absorptive capacity for future spillovers. That creates a feedback loop in which research investment generates innovation, innovation encourages education and skill formation, and those added capabilities support still more innovation.

Using the same census data, we also find that male residents who lived in treated counties before laboratory establishment experienced wage gains of roughly 5 percent relative to comparable residents of runner-up counties, with substantially larger effects for

⁵That is, a model that takes into account only current spending and not impacts on building up human and other forms of innovation capital.




Figure 3: Local Activity-to-Budget Ratio Over Time — Lab Operating Budget vs. Retail Sales. *Source: Helper et al. (2026).*

Alt text: Line graph plotting the local activity-to-budget ratio over time, comparing the total laboratory operating budget against local retail sales.

college-educated workers (approximately 25 percent) and professionals (approximately 14 percent). We discuss the aggregate implications of these incumbent wage gains in Section 6.

5 What Predicts High Returns?

The cross-laboratory variation in spillovers is substantial — and understanding what drives it is a key question for policy. We examine heterogeneity using patent data matched to laboratory operating budgets, computing adjusted patents above the synthetic control per million dollars of budget.

The heterogeneity results align with the framework’s main mechanisms:

University connections and managing entity. The most consistent predictor of high returns is whether a laboratory has an institutional connection to a university, specifically whether its GOCO managing entity is a university or university consortium as opposed

Table 1: Innovation and Economic Returns by Lab Contractor Type.

(a) Adjusted patents per \$1 million of budget

Year	Actual	Synthetic	Diff	Budget (\$1 million)	Raw Patents per \$1M	Adjusted ^a
<i>Panel A: University-Operated Laboratory Counties^b</i>						
1957	447.0	412.6	34.4	60.706	7.36	0.57
1965	756.0	560.7	195.3	161.721	4.67	1.21
1970	717.0	501.6	215.4	213.481	3.36	1.01
<i>Panel B: Privately Operated Laboratory Counties^c</i>						
1957	52.0	10.1	41.9	78.959	0.66	0.53
1965	58.0	8.9	49.1	121.953	0.48	0.40
1970	63.0	13.5	49.5	175.656	0.36	0.28

(b) Adjusted retail sales per \$1 million of budget

Year	Actual	Synthetic	Diff	Budget (\$1 million)	Raw Retail per \$1M	Adjusted ^d
<i>Panel A: University-Operated Laboratory Counties^b</i>						
1957	842,160.0	449,356.0	392,804.0	60.706	13.87	6.47
1965	2,935,190.0	1,465,093.0	1,470,097.0	161.721	18.15	9.09
1970	4,527,141.0	1,999,439.0	2,527,702.0	213.481	21.21	11.84
<i>Panel B: Privately Operated Laboratory Counties^c</i>						
1957	162,721.0	65,057.0	97,664.0	78.959	2.06	1.24
1965	245,634.0	95,082.0	150,552.0	121.953	2.01	1.23
1970	262,432.0	119,811.0	142,621.0	175.656	1.49	0.81

Notes: ^aAdjusted patents per \$1M = [(Actual patents – Synthetic patents) × 1,000,000]/Budget. ^dAdjusted retail per \$1M = [(Actual retail – Synthetic retail) × 1,000,000]/Budget. ^bUniversity-operated laboratories: Ames Laboratory, Argonne National Laboratory, Brookhaven National Laboratory, Lawrence Berkeley National Laboratory, and SLAC National Accelerator Laboratory. ^cPrivately operated laboratories: Oak Ridge National Laboratory, Pacific Northwest National Laboratory, and Savannah River National Laboratory. Data from [Helper et al. \(2026\)](#).

to a private entity.⁶ The prime contractor shapes publication norms, collaboration rules, student pipelines, joint appointments, and the ease with which local researchers and firms can access the laboratory’s knowledge base. University-affiliated labs — such as Argonne (University of Chicago), Brookhaven (Associated Universities), and Ames (Iowa State)⁷ — generated about 1.21 adjusted patents per million dollars of budget in 1965, compared

⁶We thank Josh Lerner for suggesting we look into the role of contractors as determinants of lab performance.

⁷The Appendix table 2 shows our classification of labs on the dimensions of rural and/or university-connected, and Appendix Figure 7 shows names and active dates of contractors for each lab.

to about 0.4 for labs without university ties. Universities supply skilled researchers, create collaborative networks, and reinforce norms of open publication rather than classification. [Jaffe and Lerner \(2001\)](#) find a similar pattern: among national laboratories, those managed by universities produced more patents and maintained their quality even as volumes increased after the Bayh-Dole reforms. This evidence also addresses the concern raised in Section 2 that isolated places might benefit simply because they are empty: laboratories embedded in richer institutional settings generated larger spillovers per dollar, not smaller ones.

The *absolute* volume of additional non-lab patents is concentrated overwhelmingly at university-affiliated laboratories. Argonne alone generated 182 excess patents⁸ by 1970, compared to a combined total of roughly 59 across all four privately-run laboratories in our sample (Ames, Oak Ridge, Pacific Northwest, and Savannah River).⁹ Both facts are policy-relevant: federal research investments can stimulate measurable innovation even in places far from the frontier, but complementary institutions multiply the total output.

Laboratory scale. Medium-budget laboratories generate the highest adjusted patent returns — about 1.1 adjusted patents per million dollars, compared to 0.7 for the largest facilities — outperforming both large and small labs. Very large facilities may function as self-contained campuses with less interaction with surrounding firms and inventors; very small ones lack critical mass. The inverted-U suggests a tension between scale and embeddedness that is relevant for designing new investments.

Collaboration effectiveness and openness. We do not yet have direct evidence from our study on every organizational mechanism that affects how well laboratories collaborate with outside partners, though related research suggests these institutional features are important ([Lerner et al., 2024](#)). Several aspects of the laboratories may hinder diffusion, e.g. their focus on research that is often highly classified, dangerous, or organizationally inward-looking. [Tartari and Stern \(2021\)](#) find that the more cloistered nature of national laboratories is associated with much weaker entrepreneurship spillovers than at universities. That is, governance structures influence how effectively research is trans-

⁸That is, patents compared to those from control counties.

⁹Note that Ames is both rural and university -run.

lated into surrounding innovative activity.

Economic outcomes: Retail sales data tell a similar story as the innovation data: university-affiliated laboratory counties generated roughly \$8-9 of local retail activity per dollar of laboratory budget, compared to about \$1 per dollar for privately-run laboratory counties (see Figure 4). The consistency across innovation and economic outcomes strengthens confidence that governance and embeddedness matter economically, not just for patent counts.

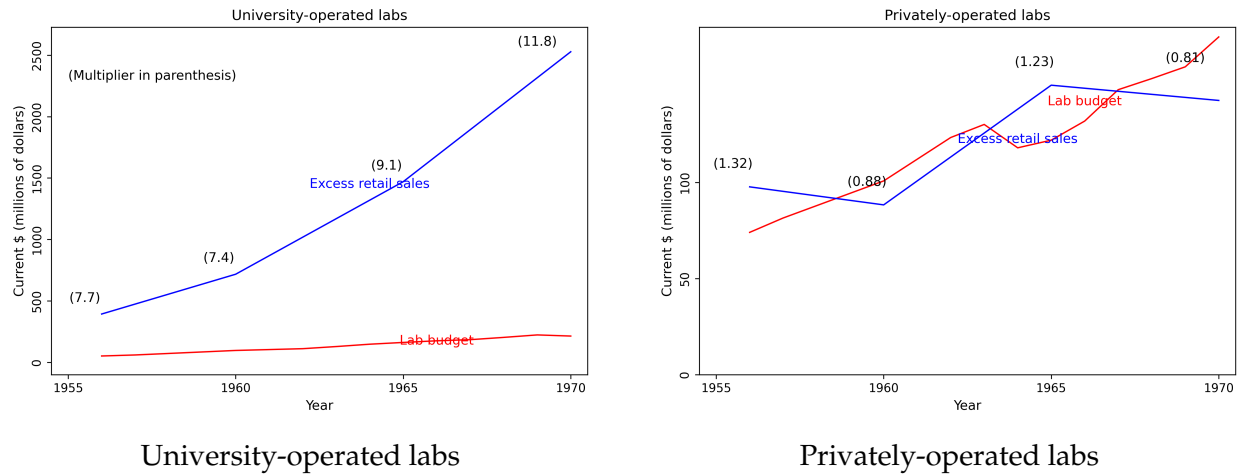


Figure 4: Gaps in retail sales between lab counties and synthetic control counties in blue. Lab budget in red. Multipliers in parenthesis.

Alt text: Line graphs comparing the aggregate lab budget and excess retail sales for University-operated and Privately-operated National Laboratories from 1955 to 1970. The university labs show a steadily growing excess retail sales gap that consistently outpaces their budget, while privately operated labs show excess retail sales that closely track their budget.

Why do University-Operated Labs generate higher returns on average?

a) Research Culture and Publication Norms

The most consistent difference between university and corporate contractors is their treatment of publication. University-managed labs were shaped by academic norms of peer review and open dissemination. They actively encouraged scientists to publish in journals, attend conferences, and maintain ties to the broader scientific community. Brookhaven under the Associated Universities Inc. (AUI, a consortium of northeastern

universities), explicitly chartered in 1947 for peaceful uses of the atom and basic research¹⁰ exemplifies this: its Physics Department became one of the most productive basic science groups in the country, with high publication rates.¹¹ Argonne under the University of Chicago similarly maintained a culture of publication and academic collaboration, hosting regular symposia and supporting faculty joint appointments.¹²

Corporate contractors imposed different norms. Union Carbide at Oak Ridge was primarily a chemical engineering company whose Oak Ridge work was physically adjacent to its commercial operations in isotope chemistry and uranium processing. The company's industrial culture prioritized deliverables over publications, and much of the knowledge generated about gaseous diffusion, isotope separation, reactor fuel chemistry, was either classified or treated as proprietary. A 1962 Union Carbide Nuclear Division publication celebrating twenty years at Oak Ridge reads as an industrial achievement report, emphasizing throughput metrics and cost reductions rather than scientific contributions.

b) Talent Recruitment and Personnel Mobility

The National Research Council (2005) identifies personnel exchange as one of the most important mechanisms of knowledge transfer between national labs and universities, and notes that university-managed labs had systematically stronger exchange programs.¹³ This finding is consistent with our findings of greater spillovers, in terms both of innovation (measured by patents) and by economic output (measured by retail sales) for university-managed labs compared to privately-affiliated ones.

University contractors leveraged academic networks to recruit scientists who might not have accepted corporate employment. Faculty dual appointments, graduate student

¹⁰See <https://www.bnl.gov/esh/env/ser/05ser/chapter.1.pdf> and Brookhaven National Lab 2005 Site Environmental Report, pg 1.

¹¹See <https://www.aps.org/funding-recognition/historic-sites/brookhaven> for an excerpt from the American Physical Society hailing Brookhaven Lab as “one of the world’s foremost research institutions. . .” and that “its many contributions to myriad branches of physics, chemistry, materials science, biology, and environmental science are renowned, and have resulted in seven Nobel Prizes (so far).”

¹²For a range of collaborative activities between Argonne Lab and U Chicago, see <https://www.uchicagoargonnellc.org/stewardship/collaboration>.

¹³National Research Council (2005). National Laboratories and Universities: Building New Ways to Work Together: Report of a Workshop. Washington, DC: National Academies Press

research programs, and postdoctoral pipelines created ongoing flows of talent between labs and universities. This had important consequences for knowledge spillover: scientists trained at university-managed labs were more likely to move into academic positions, carrying lab-acquired tacit knowledge into university departments and from there into the broader scientific community (National Research Council and others, 2014, pp 155-176).¹⁴

Not every laboratory generated large spillovers. Lawrence Berkeley and SLAC show negative adjusted patent returns by 1970, though in both cases this reflects the extraordinary growth of their control regions — the Bay Area and the Stanford corridor — rather than laboratory failure. These were among the most innovative places in the country and would have been regardless of laboratory presence. Similarly, several small and remote laboratories generated modest effects (eg., Los Alamos and Savannah River Lab).

To get a more precise sense of this heterogeneity, in Figure 5, we plot excess patents per million dollars of lab budget (comparing treated counties with their synthetic counties) against 1940 human capital levels (log of population aged 25+ with a high school diploma).¹⁵ The pattern suggests an inverted-U relationship: as we move right, the impact of the labs gradually rises, then starts to decline at the higher end (see Figure 5).

Growing returns over time. The strongest laboratories show growing effects: Argonne’s treatment effect rises from 52 excess patents in 1957 to 182 by 1970; Brookhaven moves from a negative difference of 60 to a positive gap of 75; Oak Ridge’s excess patents grow from 25 to 40. This trajectory is consistent with dynamic feedback: laboratories build local human capital, attract complementary firms and institutions, and thereby make the ecosystem progressively more productive over time.

The heterogeneity results reinforce the lower-bound interpretation. The features that predict the largest returns — university ties, embeddedness, time to develop feedback loops — are precisely the features the national laboratories were not designed to maxi-

¹⁴Another related mechanism is the cross-appointments between the labs and their affiliated universities. For example, more than 120 personnel have joint appointments between the University of Chicago and Argonne (<https://news.uchicago.edu/story/university-seed-grants-spur-collaboration-argonne-and-fermilab-spark-doe-support>).

¹⁵The analysis is restricted to the nine labs for which we have budget data in the early years after lab founding.




Figure 5: Relationship between excess patents and county education in 1940.

Alt text: Scatter plot with a fitted curve showing an inverted U-shaped relationship between a county’s excess patents and its 1940 high school education level.

mize. Investments that deliberately select on these margins should generate local spillover returns at least as large.

6 Crowd-Out, Displacement, and the Big Push

A natural concern about place-based innovation policy is that local gains may come at the expense of activity elsewhere. If laboratories attract scientists who would have been at least equally productive in existing clusters, local spillovers could overstate aggregate returns. We examine this possibility below.

Incumbent productivity gains and aggregate welfare. Residents who lived in treated counties before laboratory establishment experienced significant wage increases — roughly 25 percent for college-educated workers and 14 percent for professionals. For non-college educated workers the gains are about 1.5%.

Using the framework of [Kline and Moretti \(2014\)](#), we examine two kinds of local effects: a *reallocation* effect (workers and firms move toward treated locations, with offsetting losses elsewhere) and a *net productivity* effect (improvements in the treated location lead to aggregate productivity gains). Wage increases for pre-existing residents who do not relocate are evidence that there is a significant net productivity component. These patterns are difficult to reconcile with a pure composition effect from in-migration: the laboratories made incumbent workers more productive, particularly those best positioned to absorb knowledge spillovers. Unlike wage gains driven by in-migration — which can bid up local prices — productivity increases for people already in place are more naturally interpreted as net output gains rather than spatial reallocation.

Co-author networks. We examine whether relocating scientists to laboratory locations disrupted productive research networks. If relocation destroys valuable collaborations, local innovation gains could be offset by losses in origin locations. The impact of a single scientist’s departure is too small, relative to baseline noise, to detect in aggregate patterns. We therefore turn to direct co-authors under the assumption that any negative impact would be easiest to detect in this narrow group. To get a better understanding of this we first identify the complete list of lab scientists up to the year 1975.¹⁶ We then identify their co-authors in the years preceding their respective lab founding years (we restrict to authors who had at least 5 publications total in their career). This subset of co-authors is our treated group. The control group is a list of scientists who never co-authored with a lab scientist but who published in the same set of journals and fields as the former lab co-authors. Specifically, we use a nearest neighbor matching method based on year of first publication and pre-lab founding outcome means.¹⁷ Using the number of publications in select top publishers as the outcome variable, Figure 6 plots time series trends of average annual publications for the lab co-authors (blue) and control authors (in red). Focusing on the two most active research fields among laboratory scientists—physics and chemistry—

¹⁶We use the set of 12 labs from [Helper et al., 2026](#). The data for identifying publications was scraped from the OpenAlex website (<https://openalex.org/>).

¹⁷We use the complete set of journals from the American Institute of Physics and Elsevier publishers since these are the two most frequent sources that lab scientists publish in during our treatment years. Examples of journals operated by these publishers include *Applied Physical Letters*, *The Journal of Chemical Physics*, *Nuclear Physics B*, etc.

we find no evidence of a decline in the productivity of treated scientists following the relocation of their collaborators.

The absence of negative effects on former co-authors suggests that the relocation of scientists did not sever collaborations in a way that reduced research output at origin locations. Taken together, this evidence alleviates a key concern for interpreting the local innovation effects of laboratories. Rather than simply redistributing existing knowledge production across space or crowding out research activity elsewhere, the laboratories appear to have expanded overall research capacity without harming the productivity of pre-existing collaborators.

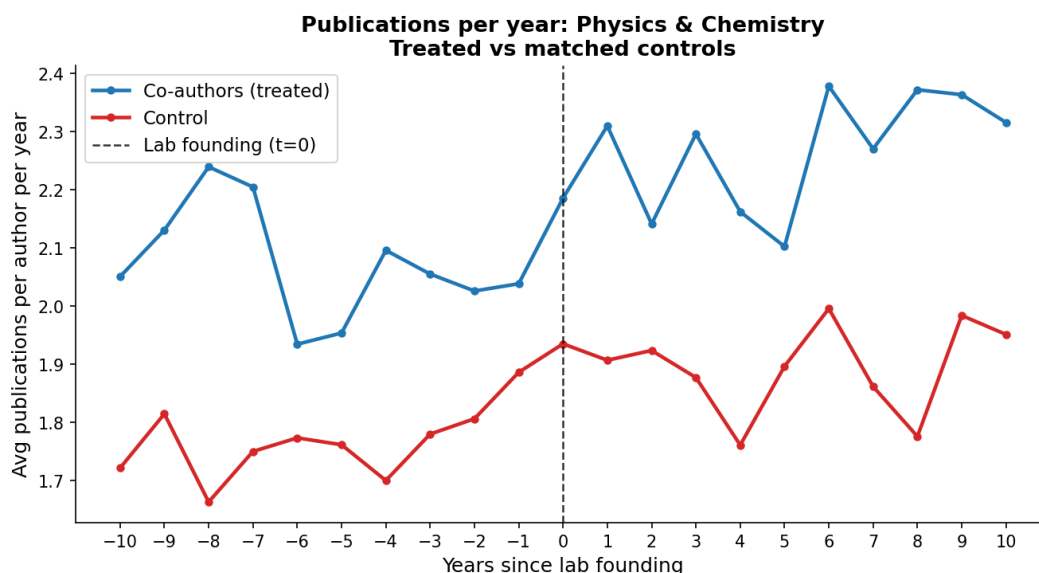


Figure 6: Co-Author Productivity Before and After Scientists Departure to the Labs. Plots shown for publications in chemistry and physics, the two most frequent categories the labs publish.

Alt text: Event study plots for chemistry and physics showing the productivity of co-authors before and after a focal scientist departs for a National Laboratory. The plots indicate changes in publication rates surrounding the departure year.

The big push interpretation. One possibility is that the sizable injection of new demand and innovation infrastructure from lab investments helped their regions jump out of low- or moderate-income equilibrium traps (Rosenstein-Rodan, 1943; Sachs and Warner, 1999; Jaramillo and Kim, 2025). The temporal pattern of laboratory effects - gradual emergence, growing activity-to-budget ratios, long-run persistence — is more

consistent with a transition to a new equilibrium than with a simple resource transfer. Before a laboratory arrives, a remote county may have too little research investment, too little human capital, and too few complementary institutions to generate spillovers. Federal establishment of a laboratory changes several of those conditions at once: it brings research funding, trains workers, attracts firms, and creates new institutional linkages. [Garg \(2025\)](#) provides a formal framework for distinguishing fundamental improvements from equilibrium selection in settings with strong complementarities.

This concern about crowd-out is less compelling for the postwar period specifically. The G.I. Bill and expanding federal support for graduate education produced a rapid surge in the number of trained scientists through the 1950s and 1960s. With the total workforce growing quickly, laboratories could recruit new personnel without meaningfully depleting the pool available to universities and private firms elsewhere.

These findings suggest that counties that shifted from a low-innovation to a high-innovation trajectory expanded the total stock of innovation-supporting infrastructure. Taken together, the wage gains of pre-existing residents, the co-author null, the elastic supply context, and the big push dynamics suggest that these local gains did not come at the expense of other regions.

7 Implications

Catalyzing innovation-based regional ecosystems that deliver broad economic and societal benefits has been a major focus of academic and policy discussion in recent years. Two concerns motivate this interest. One is that the United States has often struggled to translate its historic strength in basic research into products and processes that are globally competitive. The other is that innovation-led growth has become highly concentrated geographically, enriching a small set of superstar regions while leaving much of the country behind. Scholars such as [Atkinson et al. \(2019\)](#), [Guzman et al. \(2024\)](#) and [Gruber and Johnson \(2019\)](#) have proposed that investments in new regional innovation ecosystems can address both commercialization and regional inequality issues.

As mentioned above, in part as a result, the U.S. federal government has significantly

expanded its place-based innovation spending, such as through the CHIPS and Science Act. [Fieldhouse and Mertens \(2025\)](#) estimate that if the R&D provisions of the CHIPS Act were fully funded, the resulting expansion in nondefense R&D would boost U.S. productivity by 0.2–0.4 percent within several years — with direct effects eventually exceeding total outlays. Our results provide complementary, microeconomic evidence on the local mechanisms through which such investments generate returns. To put the magnitudes in context: the roughly \$5.2 million required to generate an additional non-lab patent at university-affiliated laboratories compares to approximately \$4.5 million per private-sector patent induced by NIH funding ([Azoulay et al., 2019](#)), approximately \$195,000 per additional patent from DOE SBIR grants ([Howell, 2017](#)), and substantial cluster-level effects from World War II OSRD contracts, where a doubling of wartime R&D intensity was associated with 30 percent higher patenting by 1970 ([Gross and Sampat, 2023](#)). These figures are not directly comparable — they measure different patent types, at different margins, over different time horizons, and using different identification strategies. But they suggest that the per-dollar innovation returns to national laboratory infrastructure fall within the range of other well-studied federal programs, despite the laboratories’ unfavorable siting conditions.

The U.S. national system of innovation faces deep cuts under the Trump administration. The administration’s 2027 budget request includes deep cuts to federal research spending,¹⁸ including double-digit proposed cuts to the national lab budget.¹⁹ While Congress last year limited these proposed cuts, agencies have laid off staff anyway. Above we showed that increased federal spending on innovation had significant spillover effects; it is likely that these effects are symmetric, meaning cuts in innovation spending would lead to significant declines in overall R&D, schooling, and wages; an American University study ([Gonzalez Garcia et al., 2025](#)) found that 25% cut to non-defense public R&D would reduce GDP by about 3.8%.

The Trump Administration has proposed re-orienting some of the spending that is left. For example, they have created a new “NSF X-Labs” program, which would fund or-

¹⁸[Slasher sequel: Trump again proposes major cuts to U.S. science spending | Science | AAAS.](#)

¹⁹[One Year into Trump II: Budget Cuts and the Path Forward](#)

ganizations “operating outside of existing academic, start-up, and industry constraints,” according to NSF’s accompanying [request for information](#). The X-Labs Initiative will support teams to “move beyond traditional research outputs,” such as publications and datasets, to focus instead on transitioning critical technology from early concept or prototypes to commercially viable platforms.

Our research suggests some considerations for the design of new programs such as this one.

Federal investments have the greatest local spillovers in places with a moderate level of absorptive capacity. Figure 5 shows an inverted-U relationship between a lab county’s number of additional patents per million dollars budget (compared to control counties) and its absorptive capacity (as measured by its number of high-school graduates). And, as Figure 6 suggests, federal investments in moving researchers to new areas appear to have very little effect on existing research networks (consistent with the argument of [Gruber and Johnson \(2019\)](#)).

As noted above, we also found a second inverted-U relationship, between spillovers and the size of the lab. A small facility may have little impact, but a too-large facility may insulate itself from the larger community.

Governance is a key design feature. Because most national laboratories are government-owned but contractor-operated, the prime contractor helps determine whether a facility behaves like a porous regional institution or a self-contained enclave. Managers rooted in universities or broad research consortia are often better positioned to take actions that promote spillovers of information and output to the broader community, such as supporting joint appointments, graduate training, shared facilities, open publication norms, and repeated interaction with local firms. Managers chosen mainly for internal mission execution may deliver on the laboratory’s core task while generating fewer spillovers to the surrounding community.

Indeed, University-governed laboratories generated about 1.2 adjusted patents per million dollars of budget, compared to 0.4 for labs without university ties (Table 1). This distinction suggests a potential tradeoff between goals of success at a well-defined mission and more diffuse and longer-term goals of accelerating regional ecosystems.

Research has shown that high-impact knowledge production increasingly depends on team-based and cross-institution collaboration (Jones et al., 2008). Explicit attention to building institutional connections such as co-location with universities, formal collaboration agreements, joint appointments, and student pipelines are important determinants of whether an investment generates spillovers or operates as an enclave (Guzman et al., 2024).

Consider long-term and dynamic impacts. The five-to-ten-year lag before non-lab patenting takes off, the growing activity-to-budget ratio over decades, and the education effects in exposed cohorts all indicate that the full returns to federal research investments materialize over long horizons. Evaluation frameworks that judge success within a few years will systematically undervalue investments whose returns operate through ecosystem dynamics. Many U.S. economic development and technology programs are designed with the idea that only catalytic funding is needed, such as the Manufacturing USA Institutes in 2014 and the currently-proposed X-labs; in each case winners were initially slated to receive only 3-5 years of federal funding. The case of the national labs shows that federal funding can indeed “crowd in” private investment (Boushey, 2022) – but that such crowding in is likely increased with on-going investment.²⁰

Short-run metrics — jobs created within two years, immediate private co-funding — will not capture ecosystem effects. These effects can be substantial. As our results above show, the labs had significant effects in building local eco-systems, as measured by indicators such as non-affiliated patenting, technological field-shifting, researcher mobility, and wages and human capital formation over five-to-fifteen-year horizons. The national laboratories would have looked like expensive remote facilities on any short-run cost-benefit metric; their true returns became visible only over decades. However, labs’ own reporting of their local impacts focus on jobs created due to direct spending, and do not mention these indicators of eco-system impacts, hindering understanding of their overall impacts.²¹

Summarizing the above points, the development of new regional research hubs would

²⁰See <https://www.nationalacademies.org/publications/18448>.

²¹See an example here for Idaho National Lab: <https://inl.gov/impacts/economic/>.

be facilitated by policies such as the following: Build formal bridges to universities and workforce pipelines — joint appointments, shared facilities, graduate fellowships — rather than excluding them, or relying on informal proximity. Researchers have found that such bridges increase patenting (Adams et al., 2001). Institutions can be designed for diffusion, with open user facilities, publishing norms, and collaborative IP arrangements that lower the cost of local firms accessing frontier knowledge (Bryan and Ozcan, 2021; Helmers and Overman, 2017). Regions may insularity by aligning procurement and vendor policies with local capability-building and innovation-system objectives (Lember et al., 2014; Kähkönen et al., 2025). Such goals are promoted by program evaluations that track innovation diffusion and human-capital indicators over long horizons, not only on near-term job counts. Of course regional economic development is not the only goal of federal research programs; they may also seek quick returns, and/or to limit diffusion for national security reasons.

8 Conclusion and Suggestions for Future Research

The national laboratories were created to advance national security through a distinctive form of publicly-funded research infrastructure, not to promote regional development. Early research was often classified, many sites were chosen for secrecy and isolation, and surrounding communities frequently had little prior innovative capacity. Yet large and persistent innovation ecosystems emerged — ecosystems whose economic effects grew over decades and raised wages for residents who were already there before the laboratories arrived. Not every laboratory generated large measurable spillovers, and the effects varied substantially with institutional context. But the overall pattern is striking, and the historical conditions under which it emerged make it especially informative.

We draw three conclusions. First, the local spillover returns to national laboratory investment are substantial and likely represent a lower bound on what the returns would be to public investments that were deliberately designed to leverage complementary institutions and human capital. Federal research institutions likely underestimate their impact on their local economies, because they tend to report economic impacts using models

that tally only their demand-side impact. Better models are needed to capture the catalyzing effects we find. Second, across laboratories, the strongest predictor of high returns is institutional connectivity and governance — particularly university affiliations and contractor arrangements conducive to openness — not simply the dollar amount of investment. Third, evidence from incumbent wage gains, co-author networks, and the temporal dynamics of spillovers points toward creation of new higher-innovation local equilibria rather than mere resource reallocation.

As policymakers design the next generation of innovation investments, the historical lesson is clear: federal R&D can generate large spillovers, but the magnitude of those returns depends critically on institutional connections, long-run commitment, and the ability of investments to trigger self-reinforcing innovation dynamics. The national laboratories demonstrate that this is possible even when starting from nothing.

The possibilities for future research on the local returns to public innovation investment are rich. The topic remains important, as U.S. R&D spending is under threat, and its distribution remains highly skewed across regions ([Howard and Liebersohn, 2025](#)). We make several (non-exhaustive) suggestions below. These start from the observation that while there is a fair amount of research investigating individual determinants of scientific productivity, there is relatively little on the role of organizational structures.

1. Investigate local spillovers from other anchor institutions.

Scholars might apply the tools of organizational economics to understand the strengths and weaknesses of these different institutions ([Powell et al., 2026](#)). That is, how do organizations structure incentives, governance, and decision-making to align individual behavior with organizational and social goals? What promotes productive interaction among institutions within an ecosystem?

There are many dimensions of institutional design (duration of the organization, intellectual property regime, who receives funding (people, projects, missions, organization?), performance metrics (papers, patents, trainees) – enough to keep scholars and policy makers busy for years, drawing on seminal work on such institutions as the Defense Advanced Research Projects Agency (DARPA) ([Fuchs, 2010](#)); universities ([Andrews, 2023](#); [Glaeser and Cutler, 2026](#)), and regional innovation engines ([Guzman et al., 2024](#)).

As noted above, the national labs were designed as GOCOs (government-owned, contractor-operated) institutions, on the theory that these would be more flexible than if run directly by the government. We have shown that the nature of the contractor (university vs. firm) matters for a variety of outcomes. But what would happen in practice if most of those who work at these institutions were civil servants, as at the U.S. National Institutes of Health.

2. *What mix of structures and relationships is most productive for different kinds of ecosystems?* Stuart Buck and Aishwarya Khanduja (2026)²² think of the innovation eco-system as a garden, in which the government's role should be like a gardener cultivating a rich and diverse ecosystem, by seeding some efforts and weeding out others.

What promotes productive and fair interaction *among* institutions within an ecosystem? Do efforts that seek explicitly to foster innovation ecosystems (such as NSF Engines, Tech Hubs, etc) actually result in meaningful and sustained economic gains for the regions they target? In contrast to the national labs, which have hundreds (sometimes thousands) of employees dedicated to a single technology all under the direction of a single contractor, these programs typically aim to build coalitions of existing institutions, such as firms (large and small), financial institutions, universities and community colleges, labor and community organization.

For example, what is it that makes universities such effective partners in amplifying the effects of national labs? A key factor may well be that the two institutions have complementary incentives – academic freedom promotes the creative control valuable for early stage research, while labs (and private sector research) have an ability to direct focused resources that is valuable for later-stage research (Aghion et al., 2008). Audretsch et al. (2024) provide empirical support for this model, in finding that increased academic freedom in a nation leads to significantly greater quantity and quality of patents.

3. *Estimate functional form of spillovers with respect to ecosystem density.*

The influential work of Gruber and Johnson (2019) argued that rather than devote almost all innovation resources to regions that are already highly dense, nations should identify and fund regions that could “jump-start” their innovation productivity at little

²²Stuart Buck and Aishwarya Khanduja. “Seeing Like a Gardener”, Liberal Currents, forthcoming.

or no loss of efficiency. Despite relatively little quantitative evidence, the book inspired programs that invested billions of dollars worth in regions (such as RECOMPETE and TechHubs, mentioned above). Our results showing inverted-U shaped relationship between innovation density and productivity suggest support for this idea—but is based on only a dozen labs. Research should more clearly identify the conditions under which this happy coexistence of equity and efficiency is likely.

Researchers might integrate the two flavors of “place-based” research, one focused on poverty alleviation (e.g., [Neumark and Simpson, 2015](#); [Bartik, 2020](#)) and the other focused on building regional innovation ecosystems (e.g., [Guzman et al., 2024](#)). The national labs appear to have improved both sets of outcomes, but high levels of inequality in clusters such as Silicon Valley ([Hendrickson et al., 2018](#)) suggest that this is not always the case. It would be valuable to know the types of agglomerations that promote both equity and efficiency. For example, while some metros with world class universities (like Raleigh and Austin) have seen average income grow significantly relative to the national average over the past fifty years, others like New Haven and Cleveland have seen local incomes decline relatively. Understanding when research institutions have significant effects on their immediate neighbors is important for place-based policy.

4. *Build data infrastructure.* Studies of programs that aim to promote regional innovation ecosystems would be facilitated by agency efforts to preserve and make available to researchers data needed to construct counterfactuals, such as information on runners-up for grant programs.

It would also be helpful to develop additional measures of key concepts; two examples are innovation and information flow.

Innovation. Most studies (such as ours) use patent data to measure innovation. Patents are useful because they can be traced to a narrow geographic area, and represent a signal that is costly to agents, adding credibility. However, not all patents represent significant innovation, and many significant innovations are not patented ([Nagaoka et al., 2010](#)). And, “because patents create an excludable right to use an innovation, they may hinder follow-on research” ([Williams, 2022](#)). Thus, patents may be less helpful in building innovation ecosystems than are other kinds of activity, such as publishing papers, developing

non-patented products or processes (which are measured in Europe by the Eurostat *Community Innovation Survey (CIS)*).²³

Information flow. An important aspect of knowledge spillovers is how easily they circulate. A key lesson for the first author from her time in the federal government was the value of convenings in overcoming information barriers that were surprisingly severe. For example, one aspect of the “Investing in Manufacturing Communities Partnership” (an Obama-era cluster-based competition that included both innovation and workforce criteria) was that economic development and workforce development folks were incentivized to meet each other for the first time. Applicants to the program (even those who didn’t win) cited this feature of the program as highly beneficial.²⁴

Relatively few economic studies look directly at the nature of information flow (or “ideas in the air” (Marshall, 1890; Helper and Stanley, 2010). It would be useful to add more, both qualitative studies of what makes networks effective such as (Safford, 2009; Feldman, 2013; Feldman and Langford, 2021), and studies using cell phone and other types of data now more readily available, such as Atkin et al. (2022), who study these interactions in Silicon Valley.

The topic of local returns to national R&D spending seems poised to grow in importance. Within the U.S., the Trump Administration is proposing the greatest reshaping of the federal role in innovation in 80 years. In addition, the rapidly improving capability of artificial intelligence will likely have large impacts on both regional innovation and regional inequality.

²³<https://ec.europa.eu/eurostat/web/microdata/community-innovation-survey>.

²⁴<https://tcf.org/content/facts/future-new-defense-manufacturing-communities-support-program/>

References

- Adams, J. D., E. P. Chiang, and K. Starkey (2001). Industry-University Cooperative Research Centers. *The Journal of Technology Transfer* 26(1), 73–86.
- Aghion, P., M. Dewatripont, and J. C. Stein (2008). Academic Freedom, Private-Sector Focus, and the Process of Innovation. *The RAND Journal of Economics* 39(3), 617–635.
- Akcigit, U., D. Hanley, and N. Serrano-Velarde (2020). Back to Basics: Basic Research Spillovers, Innovation Policy, and Growth. *Review of Economic Studies* 87(4), 1743–1778.
- Andrews, M. (2023). The Role of Universities in Regional Innovation. *Journal of Economic Perspectives*.
- Arrow, K. J. (1962). Economic Welfare and the Allocation of Resources for Invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, pp. 609–626. Princeton, NJ: Princeton University Press.
- Atkin, D., M. K. Chen, and A. Popov (2022). The Returns to Face-to-Face Interactions: Knowledge Spillovers in Silicon Valley. Working Paper 30147, National Bureau of Economic Research.
- Atkinson, R. D., M. Muro, and J. Whiton (2019). The Case for Growth Centers: How to Spread Tech Innovation Across America. Technical report, Brookings Institution, Washington, DC.
- Audretsch, D. B., A. Fiedler, B. Fath, and M.-L. Verreyne (2024). The Dawn of Geographically Unbounded Entrepreneurial Ecosystems. *Journal of Business Venturing Insights* 22, e00487.
- Azoulay, P., J. S. G. Zivin, D. Li, and B. N. Sampat (2019). Public R&D Investments and Private-Sector Patenting: Evidence From NIH Funding Rules. *Review of Economic Studies* 86(1), 117–152.

- Bartik, T. J. (2020). Using Place-Based Jobs Policies to Help Distressed Communities. *Journal of Economic Perspectives* 34(3), 99–127.
- Bloom, N., J. Van Reenen, and H. Williams (2019). A Toolkit of Policies to Promote Innovation. *Journal of Economic Perspectives* 33(3), 163–184.
- Boushey, H. (2022). *Unbound: How Inequality Constricts Our Economy and What We Can Do About It*. Cambridge, MA: Harvard University Press.
- Bryan, K. A. and Y. Ozcan (2021). The Impact of Open Access Mandates on Invention. *Review of Economics and Statistics* 103(5), 954–967.
- Buck, S. and A. Khanduja (2026). Seeing Like a Gardener. *Liberal Currents*. Forthcoming.
- Chodorow-Reich, G. (2019). Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned? *American Economic Journal: Economic Policy* 11(2), 1–34.
- Feldman, M. P. (2013). What Makes Networks Effective.
- Feldman, M. P. and S. W. Langford (2021). Knowledge Spillovers Informed by Network Theory and Social Network Analysis. In *Handbook of Regional Science*, pp. 957–970. Berlin: Springer.
- Fieldhouse, A. J. and K. Mertens (2024). The Returns to Government R&D: Evidence From U.S. Appropriations Shocks. Working Paper 2305, Federal Reserve Bank of Dallas.
- Fieldhouse, A. J. and K. Mertens (2025). The Social Returns to Public R&D. Working Paper 33780, National Bureau of Economic Research.
- Fuchs, E. R. H. (2010). Rethinking the Role of the State in Technology Development: DARPA and the Case for Embedded Network Governance. *Research Policy* 39(9), 1133–1147.
- Ganong, P. and D. Shoag (2017). Why Has Regional Income Convergence in the U.S. Declined? *Journal of Urban Economics* 102, 76–90.

- Garg, T. (2025). Can Industrial Policy Overcome Coordination Failures? Theory and Evidence. Job market paper.
- Glaeser, E. and D. Cutler (2026). *Survival of the City: Living and Thriving in an Age of Isolation*. New York: Penguin Press.
- Gonzalez Garcia, I., J. Montecino, and V. Ramaswamy (2025). Preliminary Estimates of the Macroeconomic Costs of Cutting Federal Funding for Scientific Research. Technical report, Institute for Macroeconomic and Policy Analysis.
- Gross, D. P. and B. N. Sampat (2023). World War II R&D and the Takeoff of the U.S. Innovation System. *American Economic Review* 113(12), 3323–3356.
- Gruber, J. and S. Johnson (2019). *Jump-Starting America: How Breakthrough Science Can Revive Economic Growth and the American Dream*. New York: PublicAffairs.
- Guzman, J., F. Murray, S. Stern, and H. Williams (2024). Accelerating Innovation Ecosystems: The Promise and Challenges of Regional Innovation Engines. *Entrepreneurship and Innovation Policy and the Economy* 3(1), 9–75.
- Helmets, C. and H. G. Overman (2017). My Precious! The Location and Diffusion of Scientific Research: Evidence From the Synchrotron Diamond Light Source. *The Economic Journal* 127(604), 2006–2040.
- Helper, S., R. Makan, and D. Shoag (2026). Knowledge Spillovers and Local Outcomes: An Existence Proof From the Establishment of the National Labs. Working Paper 35011, National Bureau of Economic Research.
- Helper, S. and Stanley (2010). Ideas in the Air.
- Hendrickson, C., M. Muro, and W. A. Galston (2018). Countering the Geography of Discontent: Strategies for Left-Behind Places. Technical report, Brookings Institution, Washington, DC.
- Howard, G. and J. Liebersohn (2025). How Regional Inequality and Migration Drive Housing Prices and Rents. *Journal of Economic Perspectives* 39(3), 3–26.

- Howell, S. T. (2017). Financing Innovation: Evidence From R&D Grants. *American Economic Review* 107(4), 1136–1164.
- Howell, S. T. (2024). Government Intervention in Innovation. *Annual Review of Financial Economics* 16, 367–390.
- Jaffe, A. B. and J. Lerner (2001). Reinventing Public R&D: Patent Policy and the Commercialization of National Laboratory Technologies. *RAND Journal of Economics* 32(1), 167–198.
- Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993). Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *Quarterly Journal of Economics* 108(3), 577–598.
- Jaramillo, L. F. and C. Kim (2025). Innovation Spurred: Evidence From South Korea’s Big R&D Push. Working paper.
- Jones, B. F. and L. H. Summers (2020). A Calculation of the Social Returns to Innovation. Working Paper 27863, National Bureau of Economic Research.
- Jones, B. F., S. Wuchty, and B. Uzzi (2008). Multi-University Research Teams: Shifting Impact, Geography, and Stratification in Science. *Science* 322(5905), 1259–1262.
- Kähkönen, A.-K., A. Jääskeläinen, E. Karttunen, C. B. Pedroso, and K. Lintukangas (2025). Promoting Supply Market Development by the Dynamic Capabilities of Innovative Public Procurement. *Journal of Purchasing and Supply Management* 31(4), 101056.
- Kantor, S. and A. T. Whalley (2025). Moonshot: Public R&D and Growth. *American Economic Review* 115(9), 2891–2925.
- Kline, P. and E. Moretti (2014). Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence From the Tennessee Valley Authority. *Quarterly Journal of Economics* 129(1), 275–331.
- Lee, J. (2019). Universities and Regional Innovation Ecosystems. *Regional Studies*.

- Lember, V., R. Kattel, and T. Kalvet (2014). *Public Procurement, Innovation and Policy: International Perspectives*. Berlin: Springer.
- Lerner, J., H. J. Manley, C. Stein, and H. L. Williams (2024). The Wandering Scholars: Understanding the Heterogeneity of University Commercialization. Working Paper 32069, National Bureau of Economic Research.
- Liu, S. (2015). Spillovers From Universities: Evidence From the Land-grant Program. *Journal of Urban Economics* 87, 25–41.
- Marshall, A. (1890). *Principles of Economics*. London: Macmillan.
- Moretti, E., C. Steinwender, and J. Van Reenen (2025). The Intellectual Spoils of War? Defense R&D, Productivity, and International Spillovers. *Review of Economics and Statistics* 107(1), 14–27.
- Nagaoka, S., K. Motohashi, and A. Goto (2010). Patent Statistics as an Innovation Indicator. In *Handbook of the Economics of Innovation*, Volume 2, pp. 1083–1127. North-Holland.
- National Research Council and others (2014). *Furthering America’s Research Enterprise*. Washington, DC: National Academies Press.
- Nelson, R. R. (1959). The Simple Economics of Basic Scientific Research. *Journal of Political Economy* 67(3), 297–306.
- Neumark, D. and H. Simpson (2015). Place-Based Policies. In *Handbook of Regional and Urban Economics*, Volume 5, pp. 1197–1287. Amsterdam: North-Holland.
- Powell, M., R. Gibbons, and D. Barron (2026). *Organizational Economics: Foundations and Applications*. Princeton University Press.
- Rosenstein-Rodan, P. N. (1943). Problems of Industrialisation of Eastern and South-Eastern Europe. *Economic Journal* 53(210/211), 202–211.
- Sachs, J. D. and A. M. Warner (1999). The Big Push, Natural Resource Booms and Growth. *Journal of Development Economics* 59(1), 43–76.

- Safford, S. (2009). *Why the Garden Club Couldn't Save Youngstown: The Transformation of the Rust Belt*. Cambridge, MA: Harvard University Press.
- Schrank, A. (2015). Green capitalists in a purple state: Sandia National Laboratories and the renewable energy industry in New Mexico. In *State of Innovation*, pp. 96–108. Routledge.
- Siegel, D., M. Bogers, P. D. Jennings, and L. Xue (2023). Technology Transfer From National/Federal Labs and Public Research Institutes. *Research Policy* 52(1), 104645.
- Tartari, V. and S. Stern (2021). More Than an Ivory Tower: The Impact of Research Institutions on the Quantity and Quality of Entrepreneurship. Working Paper 28846, National Bureau of Economic Research.
- Williams, H. (2022). Innovation Policies, Including the US Patent System. *Journal of Economic Perspectives*.

Appendix

A Timeline of Lab Contractors

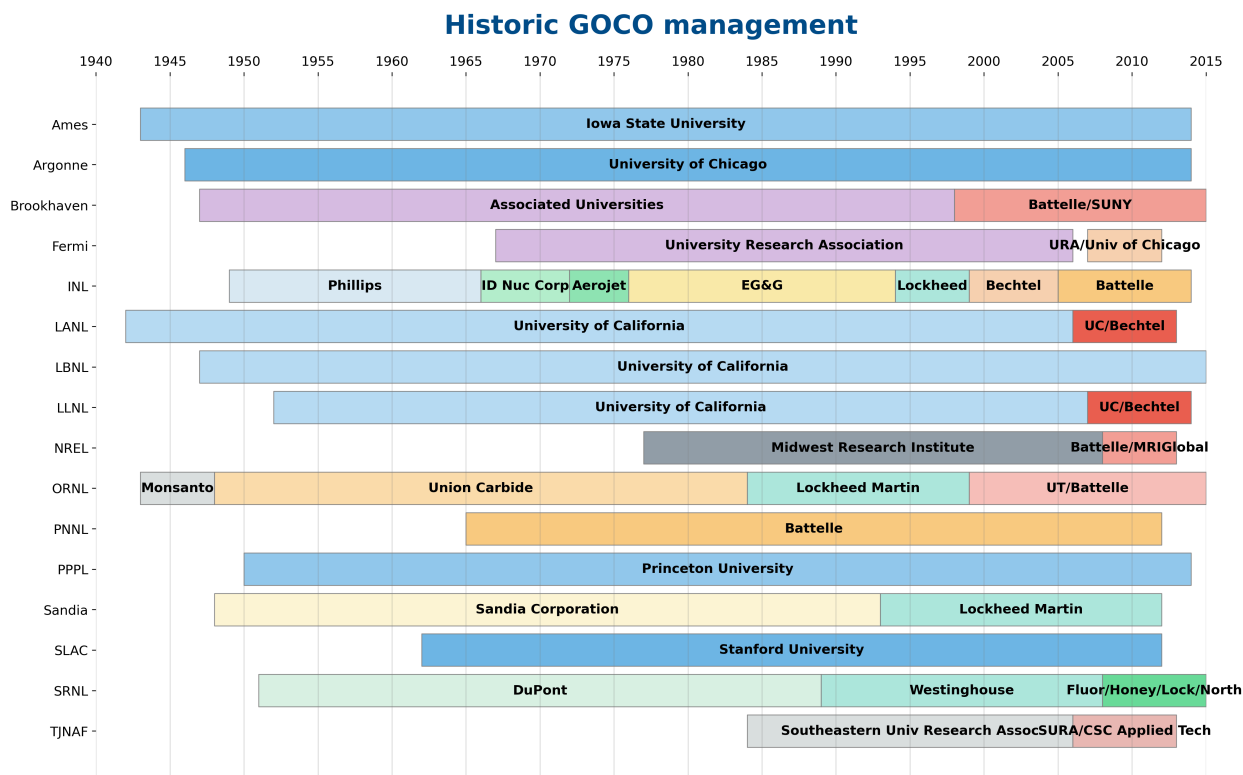


Figure A.1: Notes: INL = Idaho National Laboratory; LANL = Los Alamos National Laboratory; LBNL = Lawrence Berkeley National Laboratory; LLNL = Lawrence Livermore National Laboratory; NREL = National Renewable Energy Laboratory; ORNL = Oak Ridge National Laboratory; PNNL = Pacific Northwest National Laboratory; PPPL = Princeton Plasma Physics Laboratory; SLAC = Stanford Linear Accelerator Center; SRNL = Savannah River National Laboratory; TJNAF = Thomas Jefferson National Accelerator Facility. *Source:* <https://www.energy.gov/sites/prod/files/2014/08/f18/July%202018%20Kusnezov%20FINAL.pdf>.

Alt text: Gantt chart tracking the operating timelines of 14 U.S. National Laboratories from 1943 to 2014, grouped into University-operated and Privately-operated categories. University-operated labs show mostly continuous operation, while privately-operated labs show multiple operating periods broken up by different corporate contractors.

B Classification of lab categories

Table A1: Labs by Operator Affiliation and Remoteness

Lab Name	Rural	University-Affiliated
Ames Laboratory	1	1
Argonne National Laboratory	0	1
Brookhaven National Laboratory	0	1
Fermi National Accelerator Laboratory	0	1
Lawrence Berkeley National Laboratory	0	1
Oak Ridge National Laboratory	1	0
Pacific Northwest National Laboratory	1	0
Princeton Plasma Physics Laboratory	0	1
SLAC National Accelerator Laboratory	0	1
Lawrence Livermore National Laboratory	0	1
Los Alamos National Laboratory	1	1
Sandia National Laboratories	0	0
Idaho National Laboratory	1	0
Savannah River National Laboratory	1	0

A laboratory is classified as ‘Rural’ if it was deliberately established in a geographically isolated location away from major population centers (often for security or safety reasons during the Manhattan Project or early Cold War), or if its host county is situated outside of a major traditional Metropolitan Statistical Area (MSA). To classify university versus privately-run labs, we use the source from Figure A.1.