Where Do Technology Shocks Come From? Public Funding and Private Ownership*

Andrea Gazzani   Joseba Martinez   Filippo Natoli   Paolo Surico

July 2026

Abstract

We construct postwar U.S. patent time series by funding source and ownership. In local projections with rich macroeconomic controls, disciplined by external instruments and information-set tests, unanticipated filings serve as source-specific technology-shock proxies. Government-funded but privately owned patents are two percent of filings, yet explain about one-fifth of medium-run fluctuations in TFP and GDP. Responses propagate through private R&D and investment, with model-implied social returns to public R&D about twice those to private R&D. Effects concentrate in science-based patents supported by NIH and NSF or assigned to universities, research institutes, and start-ups. Government-owned patents have weak average effects but a consequential upper tail.

JEL: E32, C32, O47

Keywords: technology shocks, local projections, patents, shock proxies, public R&D, productivity.

*This supersedes a previous version titled "The Public Origins of American Innovation". We are grateful to our discussant, Valery Ramey, as well as to Òscar Jordà, Maarten De Ridder, Regis Barnichon, Silvia Miranda-Agrippino, Marta Prato, John Van Reenen, Juan Antolin-Diaz, Pierre Azoulay, Katharina Bergant, Antonin Bergeaud, Gary Dushnitsky, Andrew Fieldhouse, Andrea Galeotti, Luca Gambetti, Yuriy Gorodnichenko, Pierre-Olivier Gourinchas, Dan Gross, Sergei Guriev, Karel Mertens, Stelios Michalopoulos, Fabio Panetta, Elias Papaioannou, Dimitris Papanikolaou, Lucrezia Reichlin, Jim Stock, Ludwig Straub, Chris Tonetti, Emil Verner, Jonathan Yiangou, and seminar participants at many conferences and seminars for very useful comments and suggestions. Paolo Surico gratefully acknowledges financial support from the Economic and Social Research Council (Grant ES/Y002490/1). The views expressed in this paper do not necessarily reflect those of the Bank of Italy or ESCB. Address for correspondence: Bank of Italy, andreagiovanni.gazzani@bancaditalia.it; London Business School and CEPR, jmartinez@london.edu; Bank of Italy, filippo.natoli@bancaditalia.it; London Business School and CEPR, psurico@london.edu.

1 Introduction

Technology shocks are central to empirical macroeconomics. They shape the behavior of productivity, output, investment, hours, asset prices, and living standards; they are a key input into business-cycle and growth models. A large literature identifies these shocks and traces their dynamic effects (Blanchard and Quah, 1989; Galí, 1999; Francis and Ramey, 2005; Fernald, 2012; Ramey, 2020), yet their economic origins remain largely unexplored. In most empirical applications, technology shocks arrive as primitive disturbances: the econometrician identifies their aggregate effects but observes little about the institutional source of their underlying innovation.

This paper asks whether the source of innovation helps explain the propagation of technology shocks. We construct postwar U.S. time series of patent filings by funding source and ownership, and use their unanticipated component as source-specific proxies for technology shocks. The distinction between funding and ownership is key. We separate privately funded and privately owned patents, government-funded patents owned by private entities, and patents both funded and owned by the government. This classification allows us to ask whether technologies supported by public funds but deployed outside the government generate different macroeconomic dynamics from technologies funded and owned entirely by the private sector or by the government.

Our main empirical object is the dynamic macroeconomic response to unanticipated source-specific arrivals of patented technologies. We ask whether innovations arriving from different funding–ownership regimes behave as distinct aggregate technology-shock proxies and whether they propagate differently through productivity, output, private R&D, investment, wages, and consumption. We use patent filings because they provide an observable and dated measure of the arrival of patented technologies. The data come from the Government Patent Register of Gross and Sampat (2025), which links U.S. patents to information on government interest and ownership. We time patents by filing date, following Miranda-Agrippino et al. (2025), since filings provide the earliest observable indication of patented innovative activity. We then estimate local projections (Jordà, 2005; Jordà and Taylor, 2025) in which the impulse variable is the log number of patent filings in a given funding–ownership category. The baseline specification conditions on a rich macroeconomic information set, including lags of the outcome, GDP, utilization-adjusted TFP, investment, stock prices, the T-bill rate, public R&D, private R&D, and patent filings in the other funding—ownership categories. The resulting residual movements in source-specific patent filings are our patent-based technology-shock proxies.

We discipline this shock interpretation in several ways. First, we show that the control set has predictive power for patenting activity, so the residualization removes a meaningful predictable component of patent filings. Second, using the sufficient-information tests of Forni and Gambetti (2014), we cannot reject the orthogonality of the implied patent shocks to professional forecasts or to standard identified macroeconomic shocks, including government spending, tax, and monetary policy shocks. Third, we show that patent filings are only weakly related to contemporaneous business-cycle conditions, unlike R&D expenditures. Fourth, we provide corroborating evidence using external-instrument local projections: privately funded patenting is instrumented with the patent-based measure of Miranda-Agrippino et al. (2025), while publicly funded patenting is instrumented with agency-level federal R&D appropriation shocks from Fieldhouse and Mertens (2023). Finally, the results are robust to alternative lag structures, outlier treatments, inference procedures, and time-series identification schemes. These exercises support the interpretation of residual source-specific patent filings as aggregate technology-shock proxies.

A main finding is that government-funded but privately owned innovations are a large source of aggregate fluctuations, accounting for about 20% of GDP and TFP medium-run forecast error variance. They yield larger and more persistent responses than privately funded patents, and are followed by sustained increases in private R&D, capital investment and living standards. By contrast, privately funded patents generate positive but smaller responses. Government-owned patents have weak average responses, but the upper tail of government-owned innovation contains highly consequential technologies. Finally, patents that are more important according to the textual-similarity measure of Kelly et al. (2021), or more reliant on science according to Marx and Fuegi (2020, 2022), generate much larger medium-run responses of GDP and TFP. This heterogeneity is especially pronounced among public–public patents: the most important and most science-based patents in this category generate large aggregate effects, while the remaining patents have little measurable macroeconomic impact. The weak average effect of government-owned innovation is consistent with a high-risk, high-reward distribution in which a small upper tail coexists with a large mass of patents with limited aggregate effects.

The propagation of public–private technology shocks points to a specific mechanism. Public funding appears most powerful when it supports upstream, science-based research whose resulting technologies are privately owned, and extended through further private research. The largest responses are associated with NIH and NSF, with healthcare and biotechnology, with universities and research institutes, and with start-ups. Public–private patent shocks in these parts of the innovation system generate large subsequent increases in private R&D, consistent with the classical view of Nelson (1959) and Arrow (1962): publicly supported research can generate spillovers that private firms do not fully internalize. It also sharpens the institutional vision of Bush (1945a): public support for basic research matters most for aggregate productivity when it creates technologies that can be deployed and extended outside the federal government.

We use a simple framework to map the estimated impulse responses into model-implied social returns to R&D expenditure. The shocks are measured in patent filings, whereas returns are measured relative to dollars of R&D spending. We therefore convert the present value of the output response into an output multiplier using the average R&D cost per measured patent. Because R&D expenditures are observed separately for publicly and privately funded research, but not separately for public–private and public–public patents, the return calculation compares public R&D and private R&D at the spending level. At the point estimates, the model-implied perpetuity-equivalent social return to publicly funded R&D is about twice that of privately funded R&D. Publicly funded R&D also generates substantially larger private-R&D crowd-in. The disaggregated patent impulse responses show that these aggregate public returns are primarily associated with the public–private component.

Related Literature. The paper relates to the empirical macroeconomic literature on technology shocks, which existing work identifies using long-run restrictions, news measures, external instruments, or other aggregate restrictions (Blanchard and Quah, 1989; Galí, 1999;

Fernald, 2012; Ramey, 2020). We contribute by measuring technology arrivals directly from patent data and by distinguishing their institutional source. The results show that the funding and ownership structure of innovation is informative about the propagation of aggregate technology shocks.

Second, our analysis supports work on shock proxies and applied time-series methods. Recent studies have used external information, narrative measures, or proxy variables to identify aggregate shocks (Stock and Watson, 2018; Ramey and Zubairy, 2018; Antolin-Diaz and Surico, 2025; Fieldhouse and Mertens, 2023). We construct proxies directly from observable measures of innovation activity and validate them using information-set tests, professional forecasts, policy shocks, external instruments, and alternative time-series specifications. In this respect, the paper is closest to Miranda-Agrippino et al. (2025), who study patent-based measures of technological news. We build on this approach by opening the innovation aggregate into funding–ownership components and showing that these components have sharply different macroeconomic effects.

Third, the paper contributes to the literature on endogenous growth and public R&D. A large body of work emphasizes that basic research may be underprovided by private markets because its social returns are diffuse, delayed, and difficult to appropriate (Nelson, 1959; Arrow, 1962; Bush, 1945a). We provide aggregate time-series evidence that publicly funded, privately owned, science-based innovation has been a central source of postwar U.S. productivity dynamics. Our results connect to Fieldhouse and Mertens (2023), who study the macroeconomic effects of public R&D spending on defence and non-defence. Our findings suggest that public funding and private deployment are complements: innovation policy has the largest aggregate effects when it supports upstream technologies that crowd in subsequent private research and investment.

A conservative interpretation is that our evidence provides robust and validated dynamic relationships between source-specific patent arrivals and aggregate outcomes. Under the shock-proxy interpretation, these relationships can be further viewed as impulse responses to source-specific technology shocks. Either interpretation implies technology shocks are not institutionally homogeneous. The source of innovation matters for medium-run productivity dynamics, for the crowd-in of private R&D, and for the macroeconomic returns to research.

Structure of the Paper. The remainder of the paper is organized as follows. Section 2 presents a simple framework linking source-specific technology arrivals to productivity, private R&D, and social returns. Section 3 describes the patent data, the funding–ownership classification, and the empirical framework. Section 4 estimates the aggregate effects of source-specific technology shocks, their variance contribution, and their model-implied returns. Section 5 presents sensitivity analysis and alternative time-series specifications. Section 6 studies heterogeneity by patent importance and reliance on science. Section 7 examines federal agencies, research fields, and private-R&D spillovers. Section 8 studies universities, research institutes, start-ups, and venture-capital-backed firms. Section 9 concludes.

2 Interpreting Source-Specific Technology Shocks

To study the macroeconomic effects of technical change, we use patent filings disaggregated by funding source and ownership as source-specific measures of the arrival of new technologies. This section provides a simple framework for interpreting this empirical exercise. In models of growth through ideas (Romer, 1990), new technologies are both outputs of past innovation and inputs into future innovation and productivity growth. Patent filings are then the empirical counterpart of additions to source-specific technology stocks. The empirical responses we estimate trace the effects of new technologies on productivity, R&D, and other variables.

Let $s \in \mathcal{S}$ denote the source of a technology and X_{st} the corresponding stock of technologies. Final output is produced using labor, L_{yt}, augmented by aggregate productivity, A_t:

$$Y_t = A_t L_{yt}. \tag{1}$$

Productivity growth depends on the stocks of technologies available for production:

$$\frac{\dot{A}_t}{A_t} = \sum_{s \in \mathcal{S}} \lambda_s \frac{X_{st}}{A_t}, \tag{2}$$

where the parameter λ_s summarizes the direct contribution of technologies from source s to aggregate productivity. Technologies can also matter indirectly by raising the productivity of subsequent innovation. Let R_{st} denote source-s R&D expenditure, which produces new technologies according to

$$\dot{X}_{st} \propto R_{s,t}^{\eta} \prod_{s' \in \mathcal{S}} X_{s't}^{\theta_{s'}}, \tag{3}$$

where $\eta > 0$ is the elasticity of patentable technologies to R&D expenditure. The parameters $\theta_{s'}$ capture spillovers from technologies of source s' to the productivity of source-s R&D. Sources may therefore differ in the extent to which technologies raise productivity directly and stimulate subsequent research. This formulation nests the classical view of Nelson (1959) and Arrow (1962), where publicly supported research is more basic in character and so upstream of commercial innovation.

The local projections estimate the response of TFP, R&D, GDP and other variables to unanticipated movements in source-specific patent filings. In Appendix A, we linearize the full model around its balanced growth path and show that these impulse responses are the empirical counterparts of the model dynamic responses to source-specific technology arrivals. The estimates can then be read as marginal, scale-invariant effects of unexpected changes in the arrival of patented technologies. The empirical object is not the patent-level causal effect of government funding, but rather the dynamic macroeconomic response to source-specific technology-shock proxies.

The framework also provides a transparent mapping from patent-based impulse responses to social returns to R&D expenditure. The shocks we estimate are measured in patent filings, whereas social returns are measured relative to dollars of R&D spending. In Appendix A, we show that the perpetuity-equivalent social return to source-s R&D expenditure can be written as

$$r_s = (r + \delta)\frac{\eta}{C_s} \sum_{h=0}^{T} \left(\frac{1}{1+r}\right)^h \hat{Y}_{h,s}, \tag{4}$$

where r is the discount rate, δ is the depreciation rate of R&D capital, so that $r + \delta$ is the user cost, $\hat{Y}_{h,s}$ is the dollar response of output at horizon h per additional source-s patent, constructed from the estimated impulse response to an unanticipated source-s patent shock, and C_s is the average R&D cost of an additional measured patent from source s. This implies that C_s/η is the marginal cost of generating an additional source-s patent. The output fiscal multipliers can then be computed by dividing each source-s social return by the user cost.

3 Data and Empirical Framework

This section describes the construction of the source-specific patent series and the empirical framework to estimate their macroeconomic effects. The framework in Section 2 interprets patent filings as observed arrivals of new technologies. Here we describe how we measure those arrivals by funding source and ownership, how we isolate their unanticipated component, and how we validate the resulting patent innovations as source-specific technology-shock proxies. We first present the patent classification and auxiliary data. We then document the main descriptive features of the three patent categories. Finally, we introduce the local-projection specification and the information-set tests that discipline the shock interpretation.

3.1 Patent Classification by Government Interest and Ownership

Our analysis draws on the *Government Patent Register* (GPR) compiled by Gross and Sampat (2025), which combines information on government interest and assignees for all U.S. patents granted since 1890. The GPR allows us to distinguish three mutually exclusive funding–ownership categories. The first consists of patents funded by the government but assigned to a private, non-federal entity, which we label *public–private* patents ("license" in the GPR).[1] The second consists of patents with no public interest, which we label *private–private* patents. The third consists of patents funded and owned by the government, which we label *public–public* patents ("title" in the GPR).[2] The GPR also identifies the federal agency associated with each government-interest patent. By combining multiple administrative sources, the GPR reduces measurement error arising from incomplete reporting of government interest (Gross and Sampat, 2025).[3]

The empirical analysis covers 1950–2015. Before 1950, *public–private* patenting is sparse and key macroeconomic controls are unavailable at quarterly frequency. After 2015, patent registration delays become severe.[4] Throughout the paper, we time patents by filing date, following Miranda-Agrippino et al. (2025). Filing dates provide the earliest observable indi-

[1] Non-federal assignees include firms, universities, non-profit research institutes, and other private organizations.

[2] "Title" patents list a government agency as assignee, implying agency ownership, sometimes alongside private entities.

[3] A residual "unknown" category arises from conflicting government-interest records, particularly after 1980, following the digitization of patent records and the Bayh–Dole Act. We classify these cases using a transparent rule: patents indicating government license rights or private development are treated as *public–private*, while the remainder are classified as *public–public*. Appendix C details this procedure and shows that excluding these cases does not affect our results.

[4] Although the GPR extends to 2020, post-2015 patent counts decline sharply because of reporting lags.

cation of patented innovative activity and therefore correspond most closely to the timing of technology arrivals in the framework above.

We complement the GPR patent data with several sources. Assignee characteristics come from *PatentsView*. Measures of patent *importance* and *reliance on science*, used in Section 6, are drawn from Kelly et al. (2021) and Marx and Fuegi (2020, 2022). We also incorporate market-valuation measures from Kogan et al. (2017), extended to non-listed firms following Kline et al. (2019). Sectoral analyses rely on USPTO CPC Master Classification Files. The classification of research entities uses *PatentsView* and USPTO historical files. Information on start-ups and venture-capital-backed firms comes from Ewens and Marx (2024). Aggregate macroeconomic variables are from Antolin-Diaz and Surico (2025). Appendix C reports the full list of data sources and variable construction details.

3.2 Descriptive Evidence

The GPR database covers over 7 million granted patents in our sample (Table C.1). Patents with no public interest account for about 97% of the total. Among the remaining patents, about 2% are government-funded but privately owned, and around 1% are both funded and owned by the government. The number of *public–public* patents has remained broadly stable over time, while *public–private* patents have trended upward, partly reflecting the Bayh–Dole Act of 1980, which allowed small businesses and non-profits to retain ownership of federally funded inventions (Figure C.1).

Publicly funded patents tend to be more disruptive and more reliant on basic science than privately funded patents, consistent with the view that public funding is more concentrated in upstream research.[5] Figure 1 shows that, although all publicly funded patents account for only about 3% of total patenting, their share rises sharply among patents that are more disruptive or more reliant on science. Among the top 0.5% and 0.1% most disruptive patents according to the importance measure of Kelly et al. (2021), the publicly funded share is about 9% and 14%, respectively. Among patents with the highest reliance on science according to Marx and Fuegi (2020, 2022), the corresponding shares are about 22% and 20%.[6] These facts

[5] Our results are robust to pooling all government-funded patents into a single category (Appendix D), and to controlling for major institutional changes such as TRIPS in 1995. Following Kelly et al. (2021), *importance* combines forward and backward textual similarity, while *reliance on science* measures citations to scientific publications (Marx and Fuegi, 2020, 2022).

[6] The publicly funded share is about 10% among patents narratively classified by Kelly et al. (2021) as historically important.

motivate the mechanism analysis below, where we show that more basic and science-based patent shocks are associated with larger subsequent gains in GDP and TFP.

Figure 1: Patent *Importance* and *Reliance on Science*

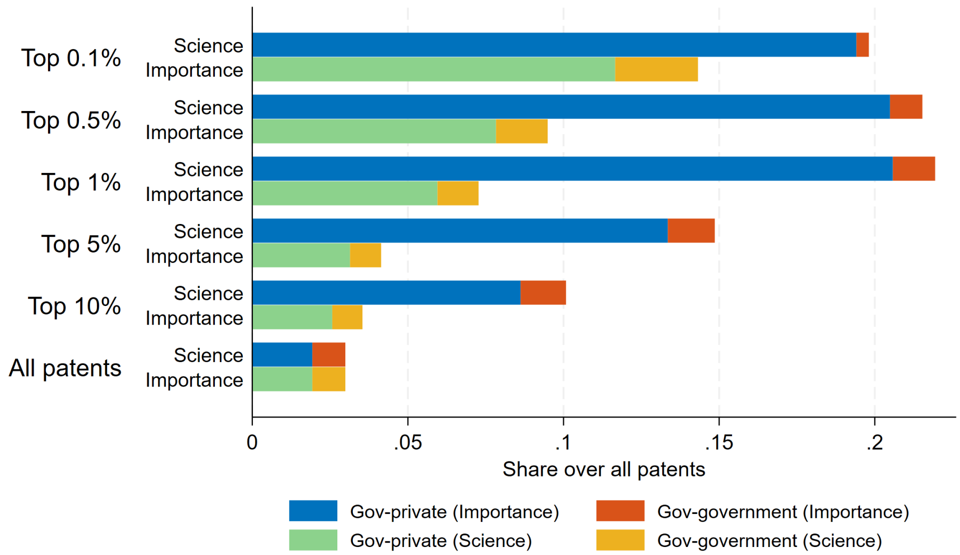

Note. The figure reports the share of public-interest patents for different levels of patent basicness from Kelly et al. (2021) and reliance on science, measured by the number of citations to scientific papers from Marx and Fuegi (2020, 2022). The bottom bar reports the unconditional share of public-interest patents among all patents. Sample: quarterly data, 1950:Q1–2015:Q4.

Federal agencies differ substantially in the volume and ownership structure of the patents they support. The Department of Defense accounts for the largest number of public-interest patents, followed by NIH, DoE, NSF, and NASA (Table C.1, Panel D). NIH and NSF primarily fund privately owned innovations, whereas NASA, DoD, and DoE retain larger public-ownership shares. DoD's public-ownership share declined from nearly 100% in 1950 to about 20% in 2015, while NIH and NSF became increasingly prominent after the 1970s (Figure C.1, Panel C). Appendices B and C report illustrative examples and additional statistics.

3.3 Time-Series Approach

Our empirical object is the dynamic macroeconomic response to unanticipated source-specific arrivals of patented technologies. Let $s \in \mathcal{S}$ index the three patent categories: *public–private*, *private–private*, and *public–public*. For each source s and horizon h, we estimate:

$$y_{t+h} = \alpha_{h,s} + \beta_{h,s} x_{s,t} + \boldsymbol{\gamma}'_{h,s} \mathbf{w}_{s,t-1} + \upsilon_{h,s,t}, \qquad h = 0, 1, \ldots, H, \tag{5}$$

as in (Jordà, 2005), where y_{t+h} is the outcome of interest, in log levels as per Montiel Olea et al. (2025), and $x_{s,t}$ is the log number of granted patents in category s filed in quarter t. The coefficient $\beta_{h,s}$ traces the impulse response of y_{t+h} to an unanticipated increase in source-s patent filings. We apply the small-sample bias correction for local projections proposed by Herbst and Johannsen (2024).

The vector $\mathbf{w}_{s,t-1}$ contains the information set used to isolate the unanticipated component of patent filings. In the baseline specification, it includes four lags of the dependent variable, four lags of real per-capita GDP, utilization-adjusted TFP, real per-capita investment, real stock prices, the T-bill rate, public R&D spending, private R&D spending, and patent filings in the other categories. These controls absorb predictable variation in patenting associated with the business cycle, financial conditions, monetary policy, past R&D, and innovation activity in other funding–ownership groups. Since the development and filing of new patents take time, conditioning on this rich lagged information set helps isolate movements in patent filings that were not predictable from prior macroeconomic conditions. Equivalently, by the Frisch–Waugh–Lovell theorem, the shock underlying equation (5) is the residual component of source-s patenting after projection on the control set:

$$\widehat{\varepsilon}_{s,t} = x_{s,t} - \widehat{\Pi}_s' \mathbf{w}_{s,t-1}. \tag{6}$$

We interpret $\widehat{\varepsilon}_{s,t}$ as a source-specific patent-based proxy for a technology shock.[7] The identifying assumption is that, after conditioning on the macroeconomic information set, residual movements in source-specific patent filings are valid proxies for unanticipated arrivals of source-specific technologies.

Table 1 reports diagnostics for this shock-proxy interpretation. Column (i) asks whether the control set has predictive power for patenting activity. The null is rejected for each funding–ownership group, confirming that residualization removes a meaningful predictable component of patenting. Columns (ii)–(v) then implement the sufficient-information tests of Forni and Gambetti (2014), regressing the implied patent shocks $\widehat{\varepsilon}_{s,t}$ on external information that should not predict a properly isolated technology-shock proxy: identified government spending shocks (Ramey and Zubairy, 2018), tax shocks (Romer and Romer, 2010; Mertens

[7] In the words of Montiel Olea et al. (2025): *"In an LP, we are estimating impulse responses with respect to a shock that is defined as the residual from projecting the impulse variable on the control variables."*

and Ravn, 2013), monetary policy shocks (Romer and Romer, 2004), and professional forecasts. Across the three patent categories, we cannot reject the null of orthogonality to these external information sets. Appendix D.1 further shows that, unlike R&D spending, patent filings are only weakly related to contemporaneous business-cycle conditions.[8]

Table 1: Sufficient-Information Diagnostics for Patent-Shock Proxies

Panel A: *p-values from F-test of joint significance — by patent type*					
	Test for controls' explanatory power	*Test for implied patent shocks' orthogonality*			
Variable	(i) LP specification	(ii) Military spending shocks	(iii) All shocks	(iv) SPF	(v) Shocks + SPF
Total patents	0.001	0.897	0.976	0.154	0.384
Private–private	0.001	0.834	0.976	0.164	0.402
Public–private	0.008	0.507	0.996	0.580	0.909
Public–public	0.038	0.787	0.564	0.766	0.614

Panel B: *p-values from F-test of joint significance — by agency*					
	Test for controls' explanatory power	*Test for implied patent shocks' orthogonality*			
Variable	(i) LP specification	(ii) Military spending shocks	(iii) All shocks	(iv) SPF	(v) Shocks + SPF
DoD	0.000	0.514	0.558	0.598	0.182
NASA	0.167	0.498	0.895	0.098	0.825
NIH	0.014	0.557	0.809	0.270	0.771
NSF	0.638	0.310	0.409	0.175	0.318
DoE	0.161	0.237	0.682	0.942	0.587

Note. Column (i) reports p-values for the F-test of joint significance on the coefficients associated with the set of controls from the baseline local projections, excluding patents' own lags and the constant. Columns (ii)–(v) report p-values from orthogonality tests in which the implied patent shocks $\widehat{\varepsilon}_{s,t}$ are regressed on external information. "Military spending shocks" are from Ramey and Zubairy (2018). "All shocks" also include the monetary policy shocks of Romer and Romer (2004), updated using Coibion et al. (2017), and personal and corporate income tax changes from Mertens and Ravn (2013), following Romer and Romer (2010). SPF stands for Survey of Professional Forecasters and includes one- and four-quarter-ahead forecasts for unemployment rate, GDP deflator, real non-residential investment, and real corporate net profits, as in Miranda-Agrippino et al. (2025).

Our empirical analysis combines two complementary approaches. We first use external-instrument local projections to provide corroborating evidence on the dynamic effects of innovation at the aggregate funding–ownership level. In Section 4.1, we instrument private innovation using the patent-based measure of Miranda-Agrippino et al. (2025), and publicly funded innovation using agency-level federal R&D appropriation shocks from Fieldhouse and Mertens (2023). This exercise helps discipline the interpretation of the patent series as innovation-shock proxies.

Because the available instruments do not provide sufficiently granular variation to study agencies, research fields, patent characteristics, assignee types, and innovator age, we then use

[8] We use the monetary policy series of Romer and Romer (2004), extended until 2007 by Coibion et al. (2017). We also cannot reject the orthogonality of $\widehat{\varepsilon}_{s,t}$ to factors extracted from the FRED-QD quarterly macroeconomic database. We use local projections rather than VARs as the baseline because we are interested in medium-run effects of shocks with delayed propagation. In this setting, extrapolating VAR responses from short-run autocovariances may lead to biases (Plagborg-Møller and Wolf, 2021; Montiel Olea et al., 2025). Section 5 verifies that the baseline ranking is not overturned using VAR-based alternatives. Baseline inference is based on heteroskedasticity-robust standard errors.

the conditional local-projection specification in equation (5) as the main empirical workhorse for the rest of the paper. In Section 4.1, we show that the conditional estimates are closely aligned with the IV responses in sign, timing, and relative magnitude. Section 5 and Appendix D show that the results are robust to a broad set of checks, including outlier-robust specifications, richer R&D lag structures, alternative inference procedures, alternative time-series identifications, and specifications that pool all government-funded patents to address mechanical changes in the public–private versus public–public split around the Bayh–Dole Act. Across these exercises, the central ranking across patent categories remains unchanged.

4 Aggregate Impact of Source-Specific Technology Shocks

This section estimates the macroeconomic effects of the source-specific technology-shock proxies constructed above. The empirical object is the dynamic response of aggregate variables to unanticipated movements in patent filings by funding source and ownership. We proceed in four steps. First, we estimate external-instrument Local Projections (LPs) and compare them to the corresponding conditional LP estimates. The two approaches yield similar responses, which motivates using the more flexible conditional specification as the benchmark given that external instruments are too coarse to study mechanisms. Second, we trace the propagation of these source-specific technology-shock proxies to private R&D, investment, wages, and consumption. Third, we quantify their contribution to medium-run movements in GDP and TFP using variance and historical decompositions. Finally, we use the mapping in Section 2 to translate the estimated impulse responses into model-implied social returns to public and private R&D expenditure.

The central result is that government-funded but privately owned patents are a quantitatively important source of medium-run technology shocks. Of the three sources we study in our baseline estimates, these patents are associated with the largest and most persistent responses of GDP and TFP, the largest responses of private R&D and capital formation, and a sizable share of the medium-run forecast variances of output and productivity. Privately funded patents also generate positive effects, but the responses are smaller and less persistent. Patents both funded and owned by the government have weak average effects, although the later sections show that this average masks substantial heterogeneity.

4.1 Results from External-Instrument Local Projections

We begin with an external-instrument exercise at the aggregate funding–ownership level. The goal is to discipline the interpretation of the patent series as technology-shock proxies using sources of variation that are external to the conditional LPs. The LP-IV specification follows equation (5), with the coefficient on source-specific patenting estimated by two-stage least squares at each horizon. Responses are normalized to a shock that raises total patent filings by 1% on impact through the shocked patent category. The control vector includes four lags of the shocked patent category, the dependent variable, real per-capita GDP, utilization-adjusted TFP, real per-capita investment, real stock prices, the T-bill rate, patenting activity in the other categories, and the aggregate R&D capital stock; the latter summarizes the cumulative effects of past public and private R&D flows and helps absorb persistent movements in the research environment.

The instruments differ across patent categories. For private–private patents, we use the orthogonalized patent-based instrument of Miranda-Agrippino et al. (2025). The instrument is constructed by residualizing patent-filing growth with respect to past patenting, identified macroeconomic shocks, and professional forecasts. Conditional on our control set, this residual variation provides an external source of private-sector innovation shocks that is plausibly orthogonal to contemporaneous macroeconomic conditions and anticipated future developments.[9] For the two government-funded patent categories, we use agency-specific federal R&D appropriation shocks from Fieldhouse and Mertens (2023). This choice reflects two features of the R&D-to-patent process. First, federal funding pipelines differ across agencies in mission orientation, scientific content, and patenting propensity, so a single aggregate R&D instrument would mask relevant heterogeneity across DoD, DoE, NASA, NIH, and NSF. Second, the lag between research funding and patent filings is long and heterogeneous: NIH-funded patents, for instance, peak at lags of 11–14 years in Azoulay et al. (2019). We therefore allow agency-specific appropriation shocks to affect patenting over a distributed lag

[9] In the Miranda-Agrippino et al. (2025) framework, this series is used as an instrument for stock prices in order to isolate the technological component of a forward-looking variable related to news. In contrast, our objective is to study technology shocks rather than technology news. We therefore use their orthogonalized patent series directly as an instrument for privately funded patenting, while including stock prices as controls in the LP-IV specification.

window of 0–80 quarters. To reduce dimensionality, we compress each agency-specific lag profile using a third-degree Almon polynomial basis, as in Barnichon and Mesters (2020). This reduces the instrument set to $6 \times 4 = 24$ instruments, while preserving cross-agency heterogeneity and imposing smoothness only across lags. In Appendix Figure D.2, we show that reducing the instrument set by considering only selected agencies does not materially affect the estimates.

A possible concern is that federal R&D appropriations may affect aggregate outcomes through channels other than patented innovation, including the employment of researchers, graduate training, or research infrastructure. The LP-IV identifying assumption is that, conditional on the aggregate R&D capital stock, lagged macroeconomic conditions and the other controls, the excluded variation in agency-level appropriations affects future GDP and TFP mainly through its effect on patented innovation. This assumption is stronger than the shock-proxy interpretation in the conditional LPs, and for this reason we use the LP-IV results here as corroborating evidence rather than as the sole basis for our findings. As the instrument structure differs across patent categories, inference also differs across specifications. For fully private patents, which are instrumented with a single external series, confidence intervals are based on heteroskedasticity- and autocorrelation-robust standard errors (Newey and West, 1987). For public patents, the high-dimensional distributed-lag instrument set and the long, heterogeneous transmission from R&D appropriations to patent filings raise the possibility of weak first stages. Hence, we construct weak-IV robust confidence bands (Anderson and Rubin, 1949), which do not require a strong first stage.[10] Weak instruments reduce precision, resulting in wider and potentially unbounded confidence sets, but do not invalidate inference in this framework.

Figure 2 reports the external-instrument LP estimates together with the corresponding conditional LP estimates. The first column shows the responses of GDP and TFP to a public–private patent shock, the second column reports the responses to a private–private patent shock, and the third column reports the responses to a public–public patent shock. The blue dash-dotted lines are the LP-IV estimates; the red lines are the corresponding conditional LP

[10] We test candidate values of the impulse-response coefficient using an Anderson–Rubin statistic based on the covariance of the instrument-score moments. The covariance matrix of these moments is estimated with a Newey–West-style HAC correction. The confidence band is the set of coefficient values that cannot be rejected at the chosen significance level.

Figure 2: The Effects of Technology Shocks on GDP and TFP

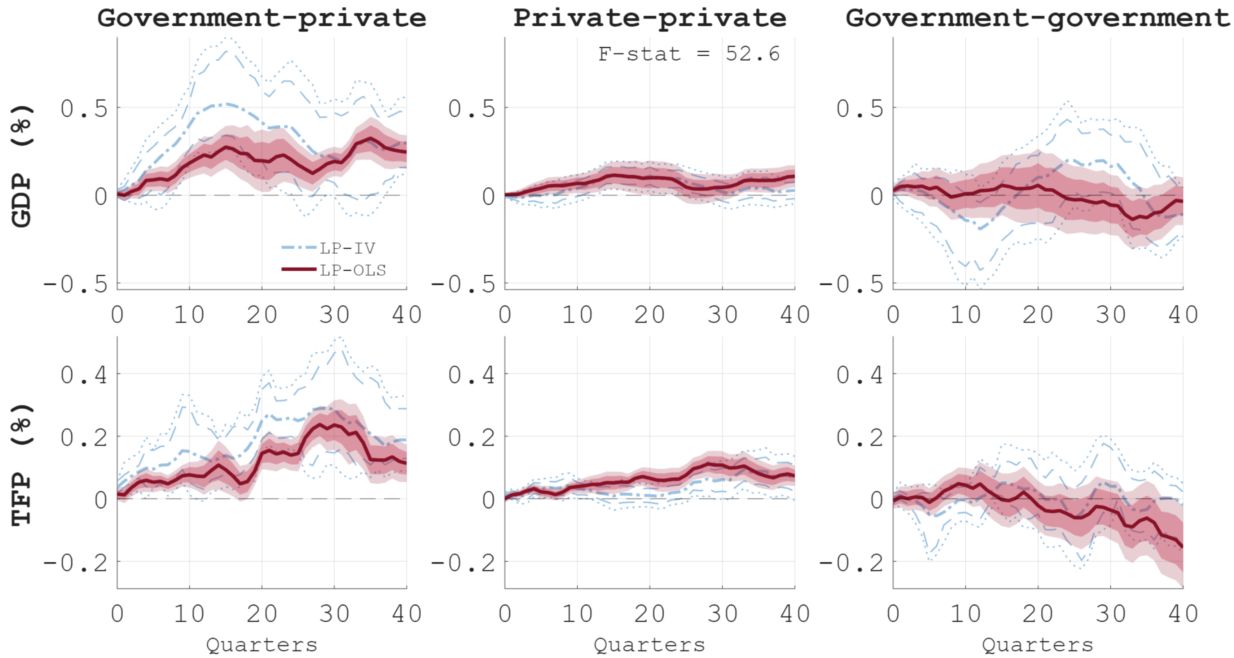

Note. The figure reports the dynamic effects of innovation shocks in each patent category—public–private, private–private, and public–public—on log real per-capita private GDP and log TFP. Red solid lines denote LP-OLS point estimates; red shaded areas report 68% and 90% confidence bands. Blue dash-dotted lines report the corresponding conditional LP-IV point estimates; blue dotted lines denote 90% confidence bands. For public–private and public–public patents, the instruments are agency-specific federal R&D appropriation shocks from Fieldhouse and Mertens (2023). The agency-level distributed lags span 0–80 quarters and are compressed using a third-degree Almon polynomial basis, yielding 24 instruments. Confidence bands for these two categories are weak-instrument-robust Anderson–Rubin bands constructed with Newey–West-style HAC covariance matrices. For private–private patents, the instrument comes from Miranda-Agrippino et al. (2025), and confidence intervals are based on standard inference with Newey–West HAC standard errors. The first-stage F-statistic is 52.6. Shocks are normalized so that total patents increase by 1% on impact. The IV control set includes four lags of the shocked patent group, the dependent variable, real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill rate, patenting activity in the other groups, and the aggregate R&D capital stock. The conditional local-projection overlay uses the baseline control set described in Section 3. Sample: quarterly data, 1950:Q1–2015:Q4.

estimates. A clear ranking emerges. Public–private patent shocks generate the largest and most persistent increases in GDP and TFP. The responses build gradually over several years and peak at medium-run horizons, consistent with the delayed diffusion of new technologies into aggregate production. Private–private patent shocks also raise GDP and TFP, but the responses are smaller and less persistent. By contrast, public–public patent shocks display weak and imprecisely estimated average effects. As shown below, this muted average response masks substantial heterogeneity: some government-owned patents are highly consequential, but they coexist with a larger mass of patents with little measurable aggregate effect.

The LP-IV estimates provide important evidence on identification at the aggregate funding–ownership level. Their scope, however, is necessarily limited: the available external instruments are designed for aggregate private patenting and agency-level public R&D appropriations, and do not provide sufficiently granular variation to study research fields, assignee types, patent characteristics, and other margins that are central to the transmission mechanism. Since the conditional local-projection estimates are broadly aligned with the IV responses in sign, timing, and relative magnitude, we use the conditional LP specification as the main econometric workhorse for the remainder of the paper.

4.2 Propagation to Firms and Households

We now use the local projections in equation (5) to study the broader propagation of source-specific technology shocks. The previous subsection showed that the conditional estimates for GDP and TFP are closely aligned with the LP-IV responses, so we do not report a separate GDP and TFP figure here. Instead, we examine whether the same shocks propagate through private R&D, capital investment, real wages, and consumption. Figure 3 shows that the broader propagation of the shocks mirrors the ranking documented for GDP and TFP. In Panel A, an unanticipated rise in public–private patenting is followed by a sustained expansion in private R&D expenditure and private capital investment. The response of private R&D is particularly important: it indicates that public–private technology arrivals are not merely associated with contemporaneous patenting activity, but are followed by additional private research effort. This pattern is consistent with the spillover channel in equation (3), in which some source-specific technologies raise the productivity of subsequent private innovation.

The responses to private–private patent shocks are positive but smaller. Public–public patent shocks generate some movement in private R&D and investment, but these responses do not translate into comparably large average gains in GDP, TFP, real wages, or consumption. Panel B shows that public–private patent shocks are followed by persistent increases in real wages and consumption, consistent with a technology-driven expansion in living standards. The corresponding responses to private–private patents are positive but more muted, while public–public patents again display weak average effects. Overall, the evidence suggests that the large GDP and TFP responses associated with public–private innovation are accompanied by broader propagation through firms and households.

Figure 3: The Dynamic Effects of Technology Shocks on Firms and Households

(a) Private R&D and Investment

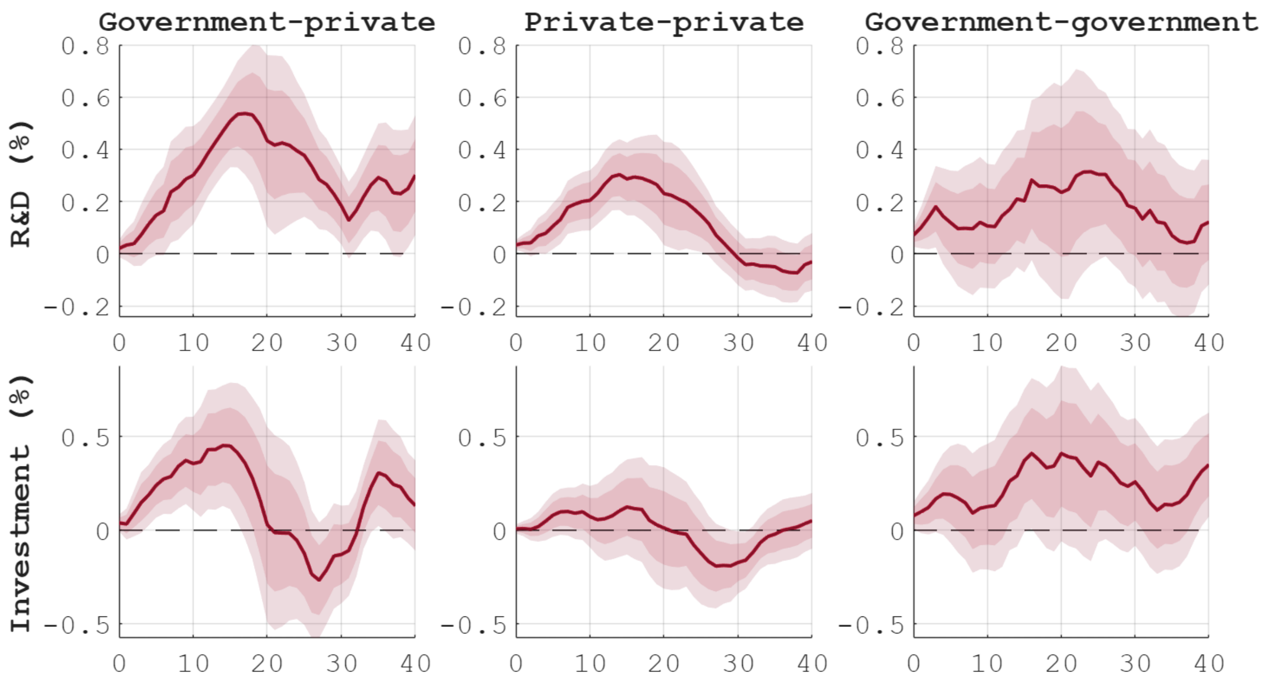

(b) Real Wages and Consumption

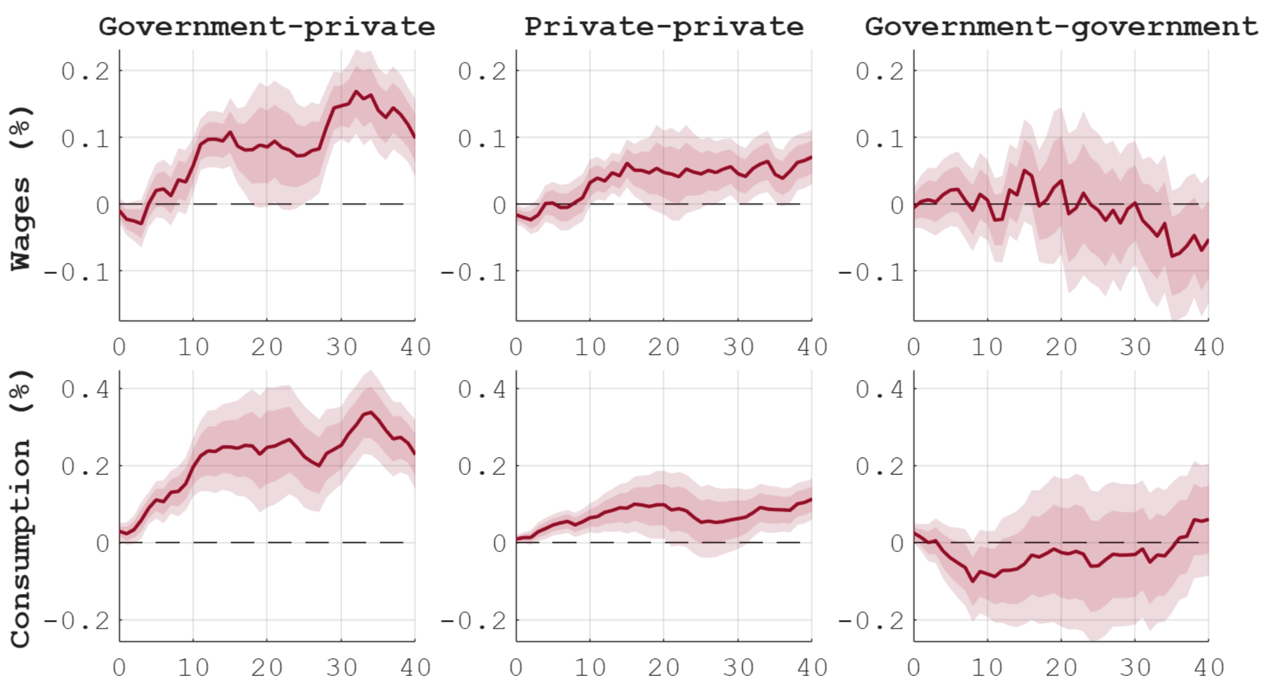

Note. The figure displays the dynamic effects of innovation shocks in each patent category—public–private, private–private, and public–public—on log real per-capita private R&D and log real per-capita investment in Panel A, and on log real wages and log real per-capita consumption in Panel B. The estimation follows equation (5). Shocks are normalized so that total patents increase by 1% on impact. The control set includes four lags of the shocked patent group, the dependent variable, real per-capita GDP, utilization-adjusted TFP, real per-capita investment, real stock prices, the T-bill rate, real per-capita R&D expenditure, and patenting activity in the other groups. All variables except the T-bill rate are expressed in logs. Solid lines denote local-projection point estimates after applying the sequential bias correction of Herbst and Johannsen (2024). Shaded areas report 68% and 90% confidence intervals based on heteroskedasticity-robust standard errors, centered on the bias-corrected estimates. Sample: quarterly data, 1950:Q1–2015:Q4.

4.3 Contribution to Medium-Run Fluctuations

The previous results show that public–private patent shocks are associated with the largest medium-run responses of productivity and output. We now ask how much of aggregate fluctuations these shocks account for. We use the R^2 decomposition for local projections proposed by Gorodnichenko and Lee (2020) to conduct two exercises. The first is a forecast error variance contribution that measures the share of GDP and TFP forecast variance explained by each source-specific patent shock at each horizon. The second exercise is a historical decomposition of TFP growth at an eight-year horizon, $h = 32$, which isolates the contribution of public–private patent shocks to medium-run productivity movements beyond business-cycle frequencies.[11]

Figure 4 Panel A reports the shares of GDP and TFP forecast error variance explained by each source-specific patent shock. Public–private patent shocks account for about 20% of the medium-run variance, while contributing little at short horizons. Fully private patent shocks play a smaller role and peak at business-cycle frequencies, while fully public patent shocks explain a negligible average share at all horizons. Panel B shows that public–private patent shocks also account for important historical movements in medium-run productivity growth. The historical decomposition at $h = 32$ tracks several major swings in postwar TFP growth, including the strong growth of the 1950s, the 'roaring' 1990s, and the slowdown of the 2000s.[12] These results reinforce the interpretation that public–private patent shocks are not primarily short-run business-cycle disturbances. Their contribution is concentrated at the medium-run horizons at which new technologies diffuse into production.

Counterfactual exercise. Appendix E.2 uses the historical decomposition to quantify the implications of the decline in the public–private shock contribution after the 1990s. If the medium-run contribution of public–private innovation shocks in the 2000s had remained at its 1990s peak, the level of TFP in 2007 would have been 5.4% to 10% higher.[13] This counterfactual is not a policy experiment; it is a way of summarizing the quantitative importance of the measured shock contribution for medium-run productivity dynamics.

[11] Eight years is a standard choice in the empirical macro literature when focusing on frequencies beyond the business cycle (Christiano and Fitzgerald, 2003). The historical decomposition combines shocks implied by the local-projection model with estimated impulse responses following Gorodnichenko and Lee (2020). Appendix E provides the formal details of this procedure.

[12] By contrast, public–private patent shocks explain little of the short-run component of TFP growth at $h = 6$ (Appendix E).

[13] These results complement existing explanations for the post-2005 productivity slowdown (e.g. De Ridder, 2024).

Figure 4: Forecast Error Variance Contribution and Historical Decomposition

(a) Forecast Error Variance Contribution

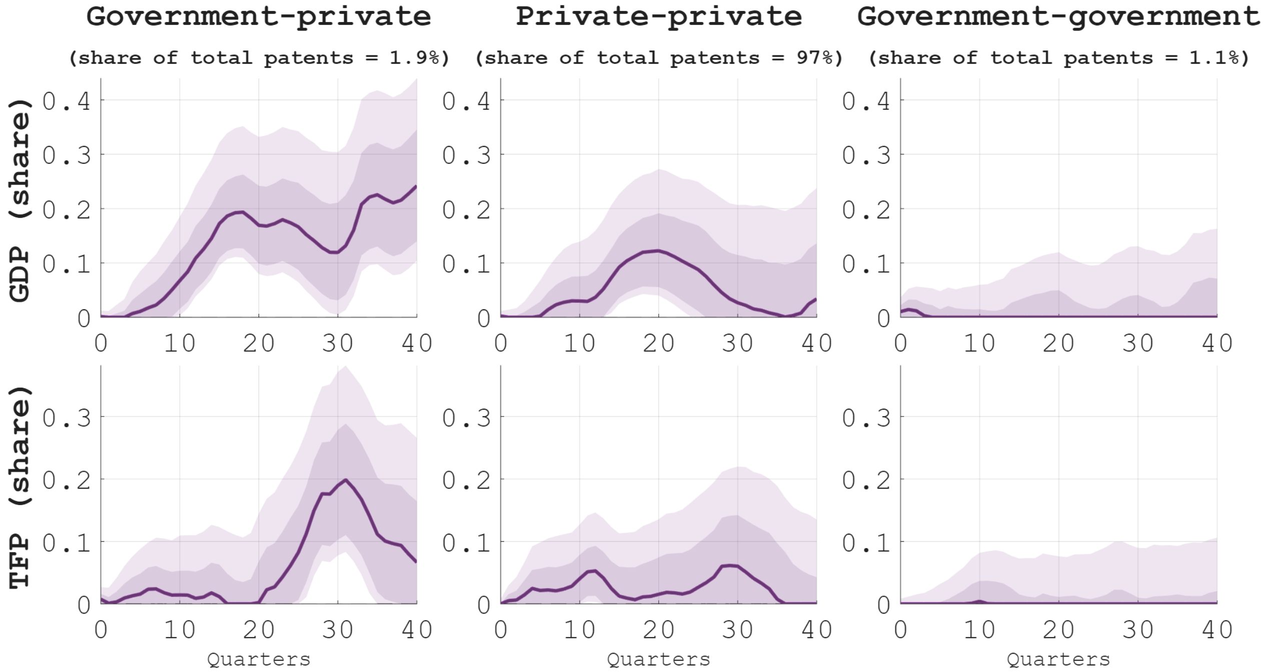

(b) Historical Decomposition of the Medium-Run Component of TFP Growth

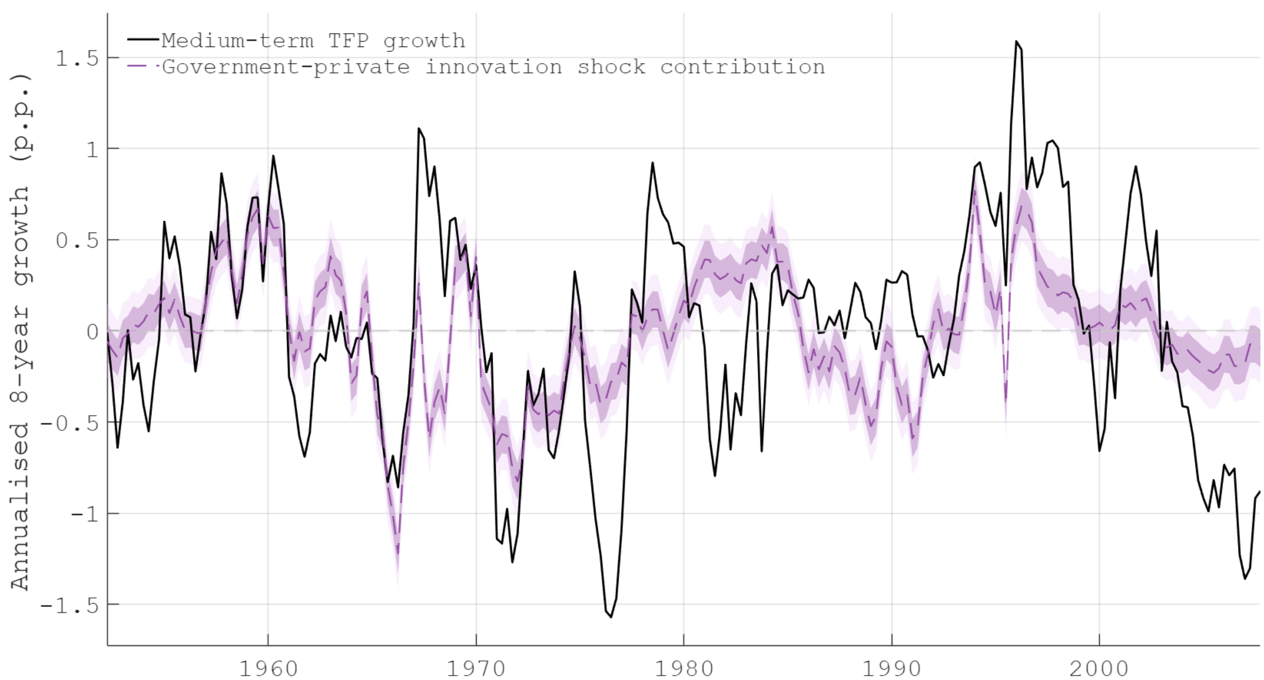

Note. *Panel A displays the forecast error variance contribution of innovation shocks in each patent category—public–private, private–private, and public–public—to log real per-capita GDP and log TFP. The estimation is based on the R^2 method for local projections in Gorodnichenko and Lee (2020). Solid lines denote point estimates; shaded areas report 68% and 90% confidence intervals. Panel B displays the historical contribution of public–private innovation shocks to the medium-run component of TFP growth. The solid black line is the medium-run TFP growth component; the purple line and shaded bands report the point estimate and 68% and 90% confidence intervals for the contribution of public–private innovation shocks. Inference is based on 2,000 bootstrap replications with small-sample adjustment. Sample: quarterly data, 1950:Q1–2015:Q4.*

4.4 Model-Implied Social Returns to R&D by Funding Source

The results so far are expressed in units of patent shocks. We now use the mapping in Section 2 to convert the estimated impulse responses into model-implied social returns to R&D expenditure. A measurement issue arises as R&D data distinguish privately funded and publicly funded research, but do not separately allocate public R&D expenditures to government-funded patents owned by private entities and patents owned by the government. We therefore compute returns at the public-versus-private R&D spending level. For privately funded R&D, the benefit is based on the estimated impulse response to private–private patent shocks. For publicly funded R&D, the benefit is based on the response to pooled government-funded patent shocks, reported in Appendix Figure D.7. However, the disaggregated impulse responses of the previous sections suggest that the public return is primarily associated with the public–private component.

The benefit of a source-s shock is the present value of the dollar GDP response it generates. The cost is the present value of the dollar R&D expenditure required to produce the patents it induces. Both objects are computed over the same 40-quarter window and discounted back to the shock date. Because the estimated patent shocks are persistent, they raise filings over several quarters rather than by a single patent on impact. The denominator in equation (4) therefore values the entire induced patent sequence at the sector-average cost per patent. Sector-average costs are constructed using historical R&D expenditure data from the NSF NCSES National Patterns of R&D Resources tables and patent counts over 1953–2015. Privately and publicly funded patents cost on average \$1.22 million and \$25.90 million, respectively, in 2017 dollars.

Three parameters are needed to compute multipliers and returns. Following Fieldhouse and Mertens (2023), we set $r = 7\%$ and $\delta = 16\%$ per year, corresponding to the average return on capital in NIPA data and the average BEA depreciation rate for R&D capital. We set the elasticity of patenting with respect to R&D expenditure to $\eta = 0.40$, in the middle of the range of estimates in the literature (Griliches, 1990; Danguy et al., 2014). This parameter affects the level of the multipliers and returns. The ratio of public to private returns, however, is independent of η under the maintained assumption that the elasticity of patenting to R&D expenditure is the same for publicly and privately funded patents.

Table 2: Returns to Public and Private R&D Expenditure

funding source	*cost per patent*	ESTIMATED RETURN	OUTPUT MULTIPLIER	R&D MULTIPLIER
private	1.22M$	0.88	4.01	0.16
		[0.55, 1.17]	[2.50, 5.30]	[0.01, 0.27]
public	25.90M$	1.91	8.67	0.52
		[1.16, 2.93]	[5.28, 13.30]	[0.29, 0.82]

Note. Estimates based on equation (4) with $\eta = 0.40$. The numerator is the discounted GDP response over 40 quarters, valued at sample-mean GDP. The denominator is the present value of the induced own-category patent sequence, valued at the average cost per measured patent. All values are in 2017 dollars. The estimated return multiplies the output multiplier by $r + \delta$, with $r = 0.07$ and $\delta = 0.16$. The R&D multiplier is computed by replacing the GDP response with the private R&D response, valued at sample-mean private R&D. Brackets report 90% wild-bootstrap bands.

Table 2 reports the resulting estimates, with 90% wild-bootstrap bands. The model-implied perpetuity-equivalent social return to publicly funded R&D is 1.91, more than twice the return of 0.88 for privately funded R&D. The relative return magnitudes are consistent with the estimates of Dyèvre (2024). The corresponding 90% confidence intervals are $[1.16, 2.93]$ for public R&D and $[0.55, 1.17]$ for private R&D, with the latter containing the values estimated by Bloom et al. (2013) and Jones and Summers (2021). In multiplier terms, publicly funded R&D generates a present-value output multiplier of 8.67, compared with 4.01 for privately funded R&D. The confidence set for the public R&D output multiplier, $[5.28, 13.30]$, contains the range of estimates in Fieldhouse and Mertens (2023) and is in line with other recent estimates based on different data and identification strategies (De Lipsis et al., 2023; CBO, 2025; DSIT, 2025).

The last column reports the present-value private R&D crowd-in per dollar of source-s R&D, computed by substituting the estimated private R&D response for the GDP response. Publicly funded patents generate roughly three times the private R&D crowd-in of privately funded patents at the 40-quarter horizon, with multipliers of 0.52 versus 0.16 dollars of induced private R&D per dollar of R&D expenditure, consistent with the estimates of Bergeaud et al. (2025). This is consistent with the mechanism emphasized above: public R&D has its largest aggregate effects when it supports technologies that are subsequently deployed and extended by the private sector. Since the public-spending calculation pools all government-funded patents, the return estimates should be interpreted as returns to publicly funded R&D

as a whole. The disaggregated impulse responses indicate that the public–private component is likely the main source of these aggregate returns.

5 Sensitivity Analysis

The evidence above implies a stable ranking across source-specific technology shocks: public–private patents generate the largest medium-run responses of GDP and TFP, private–private patents generate smaller positive responses, and public–public patents have weak average effects. This section examines two threats to this ranking that are especially relevant for interpreting residual patent filings as technology-shock proxies. The first is that patent filings contain occasional spikes, some of which may reflect institutional, regulatory, or reporting changes in the patent system rather than economically meaningful technology arrivals. The second challenge is that patent filings may respond with long and heterogeneous lags to past R&D expenditure, so that a short-lag specification for R&D controls could leave predictable research-driven variation in the patent shocks. In Figure 5, we show that neither concern overturns the main results.

Panel A studies whether the baseline estimates are driven by exceptional movements in patent filings. We define spike episodes as quarters in which the absolute quarterly growth rate of the seasonally adjusted patent series exceeds one within-sample standard deviation, grouping contiguous quarters into a single episode. We then compare the baseline local-projection estimates with three alternative treatments of these episodes. The first dummies out all spike episodes. The second dummies out one episode at a time. The third removes the identifying variation from one episode at a time by setting the residual patent shock to zero during that episode.[14]

The ranking across categories is essentially unchanged. Public–private patents continue to generate positive and persistent medium-run responses of GDP and TFP, which are close to the baseline across all spike treatments. Private–private patents have positive but smaller impact. Public–public patents still display weak average effects. Thus, the main results are

[14] In the latter exercise, we first residualize patent filings with respect to the baseline information set, $\widehat{\varepsilon}_{s,t} = x_{s,t} - \widehat{\Pi}'_s \mathbf{w}_{s,t-1}$, and then re-estimate the responses after setting $\widehat{\varepsilon}_{s,t} = 0$ during each spike episode in turn. This exercise does not discard observations; it asks whether any single spike episode is responsible for the estimated impulse responses.

Figure 5: Outliers and R&D Lag Structure

(a) The Role of Outliers

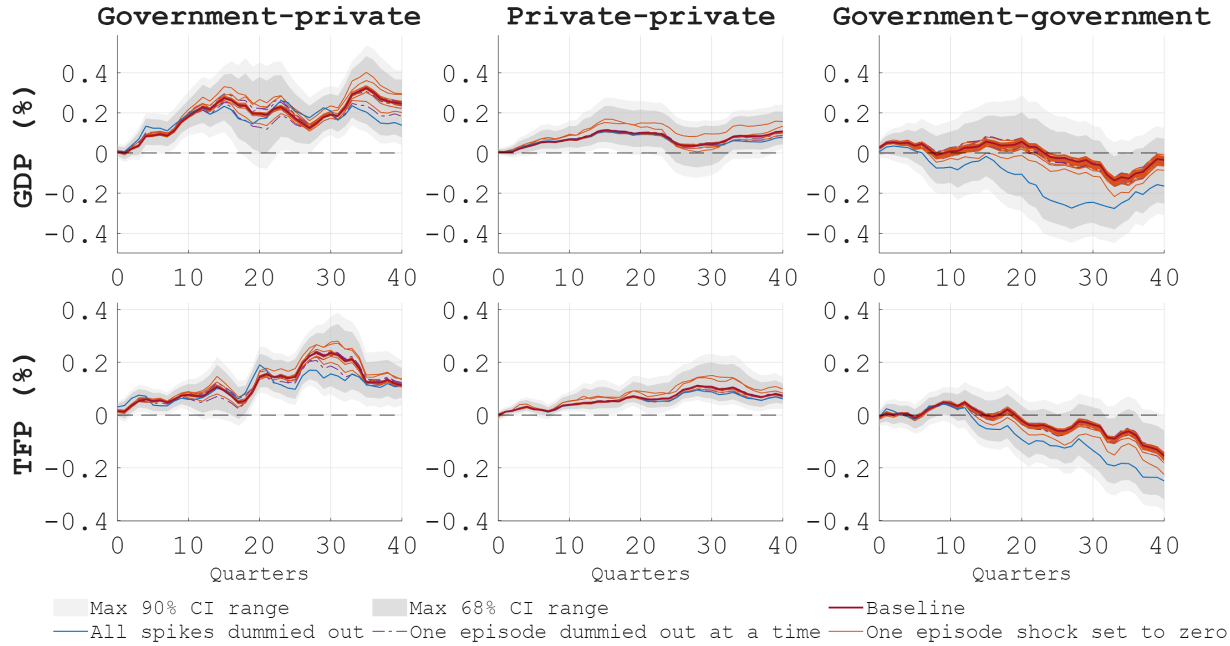

(b) The Role of R&D Lag Structure

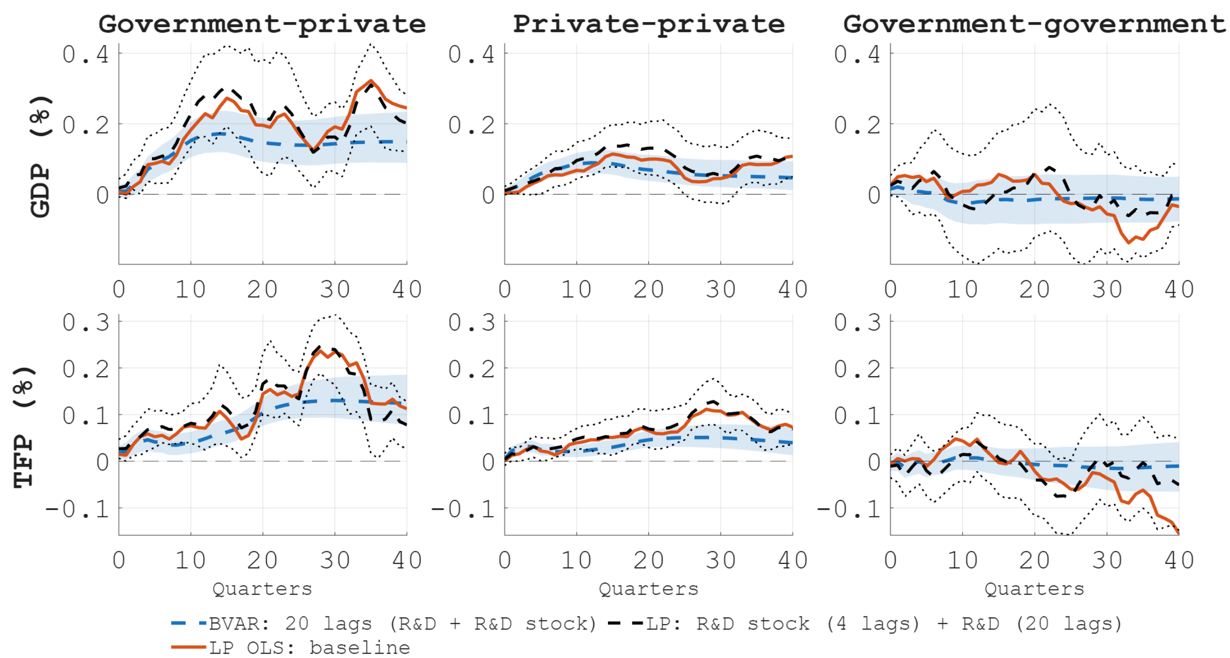

Note. *Panel A reports outlier-robustness exercises for the dynamic effects of technology shocks in each patent category—public–private, private–private, and public–public—on log real per-capita GDP and log TFP. Spike episodes are quarters in which absolute quarterly growth in the seasonally adjusted patent series exceeds one within-sample standard deviation, with contiguous quarters grouped into a single episode. The figure compares the baseline estimate with specifications that dummy out all spike episodes, dummy out one episode at a time, or set the residual patent shock to zero during one episode at a time. Panel B compares the baseline local-projection specification with alternative treatments of R&D dynamics. The orange solid line is the baseline specification, which includes four lags of all baseline variables and excludes the aggregate R&D capital stock. The black dashed line adds the aggregate R&D capital stock with four lags and extends the lag structure of total R&D expenditure to 20 quarters; dotted black lines denote 90% confidence bands. The blue dashed line and shaded area report the Antolin-Diaz and Surico (2025) BVAR benchmark and its 90% credible interval. Shocks are normalized so that total patents increase by 1% on impact. All variables except the T-bill rate are expressed in logs. Sample: quarterly data, 1950:Q1–2015:Q4.*

not driven by a small number of exceptional patent-filing episodes or by discrete changes in the patent environment.

Panel B examines the treatment of R&D dynamics. The baseline specification controls for four lags of real per-capita R&D expenditure, together with the rest of the macroeconomic information set. We compare this specification with a richer treatment that adds the aggregate R&D capital stock and extends the lag structure of total R&D expenditure to 20 quarters. This alternative is deliberately demanding: it allows patent filings to depend on a substantially longer history of research expenditures and on the accumulated stock of past R&D. We also report a Bayesian VAR in the spirit of Antolin-Diaz and Surico (2025), using 20 lags of the endogenous variables. The point estimates remain close to the baseline.

Taken together, these exercises support the interpretation of the residual patent series as source-specific technology-shock proxies. The dominant role of public–private innovation in medium-run output and productivity dynamics is not an artifact of patent-filing spikes, nor of a particular short-lag treatment of R&D controls.

Additional results and robustness. Appendix D reports a broader set of checks. First, we show that patent filings are only weakly related to contemporaneous business-cycle conditions, whereas shocks to total patenting generate responses consistent with standard technology disturbances: output and productivity rise, while consumer prices fall. Second, we verify that the LP-IV estimates for publicly funded patents are not driven by the particular set of agency-level R&D appropriation instruments used in the baseline by progressively reducing the set of agencies included in the first stage. Third, we directly test whether the responses to public–private and private–private patent shocks differ and reject the null of equality for both GDP and TFP. Fourth, we show that inference based on the equal-weighted cosine estimator of Lazarus et al. (2018) delivers results close to the baseline robust-inference estimates. Fifth, we compare the baseline specification with alternative time-series approaches, including VAR-based maximum-share identification, recursive VARs, factor-augmented local projections, and identification through non-Gaussianity or time-varying volatility. Finally, we address the possibility that the Bayh–Dole Act mechanically affects the split between public–private and public–public patents by pooling all government-funded patents into a single category. Across all these sets of exercises, the main ranking is unchanged: public–private patent shocks, or pooled public patent shocks in the Bayh–Dole exercise, generate the largest medium-run responses of GDP and TFP; private–private patent shocks generate smaller positive responses; and public–public patent shocks display weak average effects.

6 The Role of Basic Research

The baseline results estimate average responses to source-specific patent shocks. These are informative about the aggregate contribution of each funding–ownership category, but they may mask substantial heterogeneity within categories. This is especially important for interpreting the weak average response of public–public patents. If government-owned innovation has a high-risk, high-reward structure, then a small upper tail of consequential technologies may coexist with a large mass of patents with little aggregate impact. This section asks whether the macroeconomic effects of source-specific technology shocks are concentrated among patents that are more disruptive, more science-based, and thus closer to the notion of upstream or basic research.

The concept of basic research is central to Bush (1945a) and to the economics of innovation, but it is difficult to measure directly in aggregate R&D expenditure data. We use two complementary patent-level measures. The first is the textual-similarity of patent importance by Kelly et al. (2021). The second is the reliance on science by Marx and Fuegi (2020, 2022). These measures capture different aspects of basicness. The importance measure identifies patents that depart from previous technologies while influencing subsequent innovation. The reliance-on-science measure identifies patents that cite scientific publications and are therefore more directly connected to formal research. In the next sections, we complement these patent-characteristic cuts by studying federal agencies, research fields, universities, start-ups, VC-backed firms and incumbents.

As discussed in Section 3, the importance measure of Kelly et al. (2021) combines forward and backward textual similarity. A patent receives a high importance score when later patents are textually close to it, while it is less similar to earlier patents. The intuition is that more disruptive technologies both move away from the existing knowledge base and shape subsequent innovation. We use this measure because it is available for the patent universe, including government-owned patents and patents assigned to universities, research institutes, and young firms.[15] Because both patent importance and reliance on science are ex-post characteristics, we do not use them as instruments or as real-time information variables. Instead, we use them to ask where, within each source-specific patent series, the aggregate responses are concentrated.

Figure 6 Panel A splits each funding–ownership category into patents in the top quartile of the Kelly et al. (2021) importance distribution and all remaining patents.[16] Three results stand out. First, within every funding–ownership category, shocks to more important patents generate larger medium-run responses of GDP and TFP than shocks to the remaining patents. The difference is small at short horizons and grows over time, consistent with the gradual diffusion of more fundamental technologies into aggregate production. Second, the strong average effects of public–private patent shocks documented above are concentrated among the more important patents in that category. Public–private shocks in the top quartile of the importance distribution generate large and persistent gains in both output and productivity, whereas the remaining public–private patents produce smaller responses. This result connects the average public–private effect to the mechanism in Section 2: government-funded technologies owned and deployed outside the federal government appear especially powerful when they are more disruptive and therefore more likely to raise the productivity of subsequent private innovation. Third, public–public patents display the sharpest within-category heterogeneity. On average, public–public patent shocks have weak effects, as shown in Section 4. But the top quartile of public–public patents is followed by some of the largest medium-run responses of GDP and TFP in the figure. By contrast, shocks to the remaining public–public patents have small or negative responses. This pattern reconciles the weak average effect of government-owned patents with the existence of highly consequential public-sector inventions. The public–public portfolio appears to contain a small upper tail of technologies with large aggregate effects and a larger mass of patents with little measurable

[15] A main reason to favor the measure of Kelly et al. (2021) over the market-value measure of Kogan et al. (2017) is that the latter is available only for listed firms and thus excludes many patents central to our analysis, such as government-owned or assigned to research institutes, universities, start-ups, and VC-backed firms. In Appendix F.3, we show that similar results are obtained using the patent value metric of Kogan et al. (2017) extended beyond listed firms following Kline et al. (2019).

[16] We rank patents within each funding–ownership category over the full sample using the five-year textual-similarity measure of Kelly et al. (2021). Results are similar when using the ten-year window.

Figure 6: Dynamic Effects of More Important/More Science-Based Technology Shocks

(a) Importance

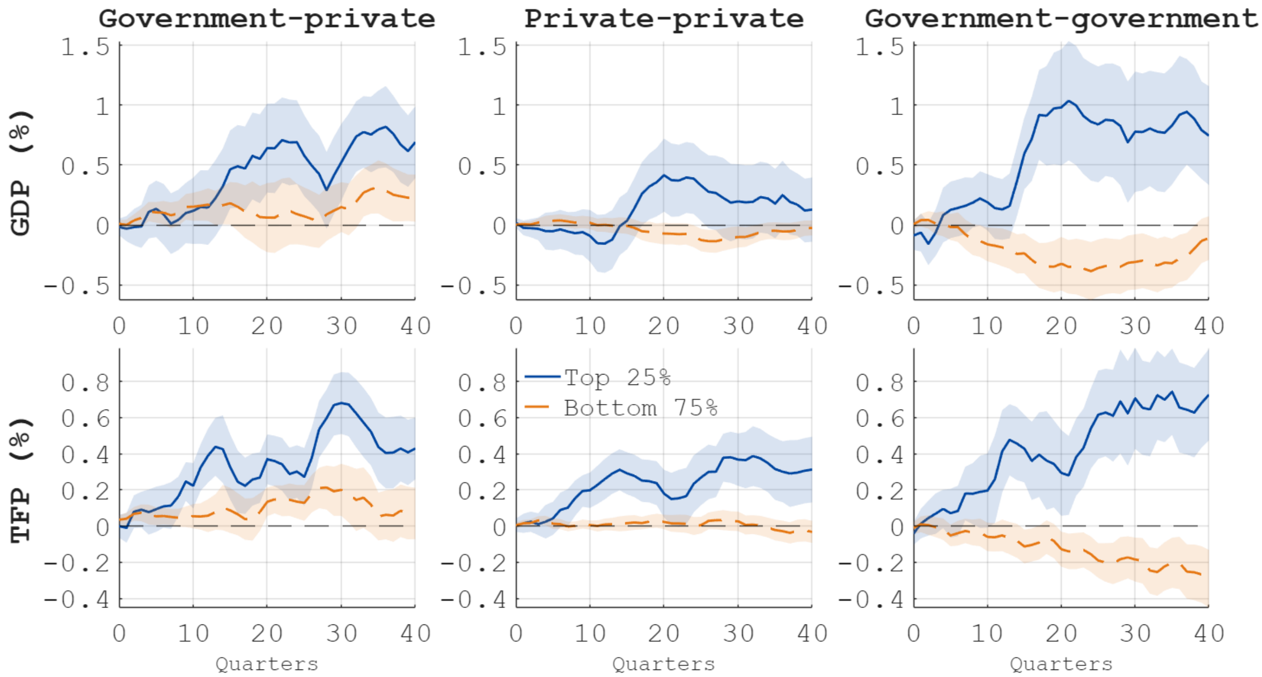

(b) Reliance on Science

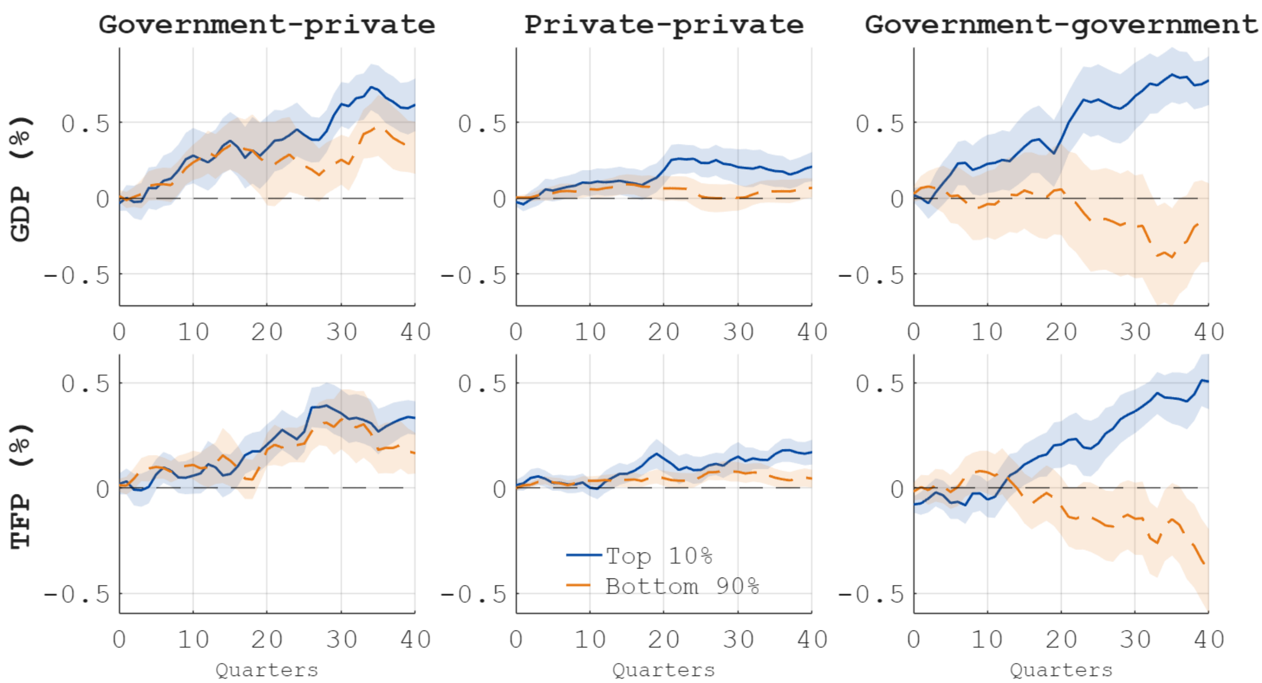

Note. Panel A reports the dynamic effects of technology shocks to patents in the top quartile of the Kelly et al. (2021) five-year importance measure, compared with shocks to the remaining patents in the same funding–ownership category. Panel B reports the same exercise using the reliance-on-science measure of Marx and Fuegi (2020, 2022), comparing patents in the top decile with the remaining patents in the same category. Columns correspond to public–private, private–private, and public–public patents. Rows report responses of log real per-capita GDP and log utilization-adjusted TFP. The estimation follows equation (5). Shocks are normalized so that total patents increase by 1% on impact. The control set includes four lags of the shocked patent group, the dependent variable, real per-capita GDP, utilization-adjusted TFP, real per-capita investment, real stock prices, the T-bill rate, real per-capita R&D expenditure, and patent filings in the other groups. All variables except the T-bill rate are in logs. Solid blue lines denote Herbst and Johannsen (2024) bias-corrected point estimates for top-importance or top-science patents; dashed orange lines denote point estimates for the remaining patents. Shaded areas report 90% confidence intervals based on heteroskedasticity-robust standard errors. Sample: quarterly data, 1950:Q1–2015:Q4.

macroeconomic impact.

Panel B uses reliance on science as an alternative proxy for basicness. The measure of Marx and Fuegi (2020, 2022) counts scientific publications cited by each patent. Patents with more citations to scientific articles are therefore more directly connected to formal scientific research. We split each funding–ownership category into patents in the top decile of the reliance-on-science distribution and all remaining patents.

The reliance-on-science results reinforce the evidence from the importance measure. In public–private and private–private patenting, shocks to more science-based patents generate somewhat stronger responses of GDP and TFP than shocks to less science-based patents. The contrast is again most pronounced for public–public patents. Government-owned patents with high reliance on science generate large positive medium-run responses, while the remaining public–public patents are largely uninformative for GDP and TFP at longer horizons. The result suggests that the aggregate value of government-owned innovation is concentrated in the science-based upper tail.

Appendix Figure F.4 shows that the strong response to science-based public–public patents is robust to using a more balanced 25–75 percentile split of the reliance-on-science measure. With this cutoff, the heterogeneity within public–private and private–private patents becomes less pronounced, consistent with the fact that reliance on science is more evenly distributed within those categories than the importance measure. Appendix Figure F.5 further shows that similar results are obtained when we classify patents according to whether they cite at least one scientific publication.

The results in this section sharpen the interpretation of the baseline findings. The average dominance of public–private innovation is concentrated among patents that are more disruptive and more connected to science. At the same time, the weak average response of fully public patents masks substantial upper-tail heterogeneity. Source-specific technology shocks thus differ not only by funding and ownership, but also by the basicness of the technologies they contain. This motivates the agency, field, and assignee analysis that follows.

7 Mechanisms: Agencies, Fields, and R&D Spillovers

The results above show that government-funded but privately owned patent shocks are associated with the largest medium-run responses of GDP and TFP, and that these effects are concentrated among more important and more science-based patents. This section studies where those effects come from. We examine heterogeneity across federal agencies, matched private-sector industries, research fields, and private-R&D spillovers. The goal is to locate the parts of the innovation system that account for the aggregate responses documented above. The mechanism suggested by Section 6 is that public funding has its largest macroeconomic effects when it supports upstream, science-based technologies that are subsequently owned, deployed, and extended outside the federal government. The evidence below is consistent with this interpretation. Public–private patent shocks have the largest effects when they are funded by NSF and NIH, when they are concentrated in fields such as healthcare and biotechnology, and when they generate large subsequent increases in private R&D. In the next section, we study the institutional players behind these responses, focusing on universities, research institutes, start-ups, VC-backed firms, and incumbents.

7.1 Federal Agencies and Industries

The postwar U.S. innovation system was built around a division of labor between federal funding and decentralized research performance. The architecture described by Bush (1945a) assigned a central role to federal support for basic research in colleges, universities, medical schools, and research institutes. We use the agency information in the Government Patent Register to ask whether the aggregate effects of publicly funded technologies differ across federal funders. We group government-funded patents into five main agencies or departments: NIH, NSF, NASA, the Department of Energy (DoE), and the Department of Defense (DoD). According to Table C.1, these agencies account for 98% of government-funded patents.[17] To construct private-sector counterparts to these agencies, we assign private patents to industries using CPC codes that capture the main technological domains relevant to each agency. The

[17]The Department of Education Organization Act of 1979 established the U.S. Department of Health and Human Services as the primary federal agency responsible for public health, social services, medical research, and Medicare and Medicaid administration. Since 1980, the vast majority of HHS patents in our data are funded by NIH. We therefore refer to patents funded by federal health agencies over the full sample as NIH patents.

exception is the education sector, which is identified using university assignee names. In Appendix G, we describe the mapping between federal agencies or departments, industries, research fields, and CPC codes.

Figure 7 reports the dynamic responses of GDP and TFP to source-specific technology shocks by agency, matched private industry, and research field. Panel A compares federal agencies and their private-sector counterparts. The first column reports public–private patents by funding agency; the second reports private–private patents in matched industries; and the third reports public–public patents by owning or funding agency. Panel B repeats the exercise by research field. Since some narrow groups display a negative impact response of total patents, the shocks in this figure are normalized so that the peak response of total patent filings over the forecast horizon is 1%, following Fieldhouse and Mertens (2023). Colored lines denote point estimates that are statistically significant at the 68% level; gray lines denote estimates that are not.

Three results emerge from Panel A. First, among public–private patent shocks, NSF and NIH generate the largest medium-run responses. Their effects peak at roughly 0.5% for GDP and 0.3% for TFP. The other agencies also generate positive responses, but their effects are smaller. This ranking is consistent with the view that the largest aggregate effects come from federal funding that supports upstream research while leaving ownership and deployment outside the federal government. Second, the matched private-sector industries display less heterogeneity, especially for TFP. Fully private patent shocks in health, education, aerospace, energy, and manufacturing generate positive responses in some cases, but the variation across industries is smaller than the variation across federal agencies. This suggests that the agency results are not simply a reflection of the technological domain in which each agency operates. Third, public–public patents display the strongest within-category heterogeneity. The average effect of public–public patent shocks is weak, as shown in Section 4, but Panel A reveals that this average effect masks sharply different responses across agencies. NIH-funded fully public patents generate large medium-run responses, whereas DoD public–public patents generate weak responses. This pattern is consistent with the high-risk, high-reward interpretation developed in Section 6: government-owned patents contain a small upper tail of highly consequential technologies, but also many patents with little

Figure 7: The Effects of Technology Shocks by Federal Agencies and Research Fields

(a) Federal Agencies and Industries

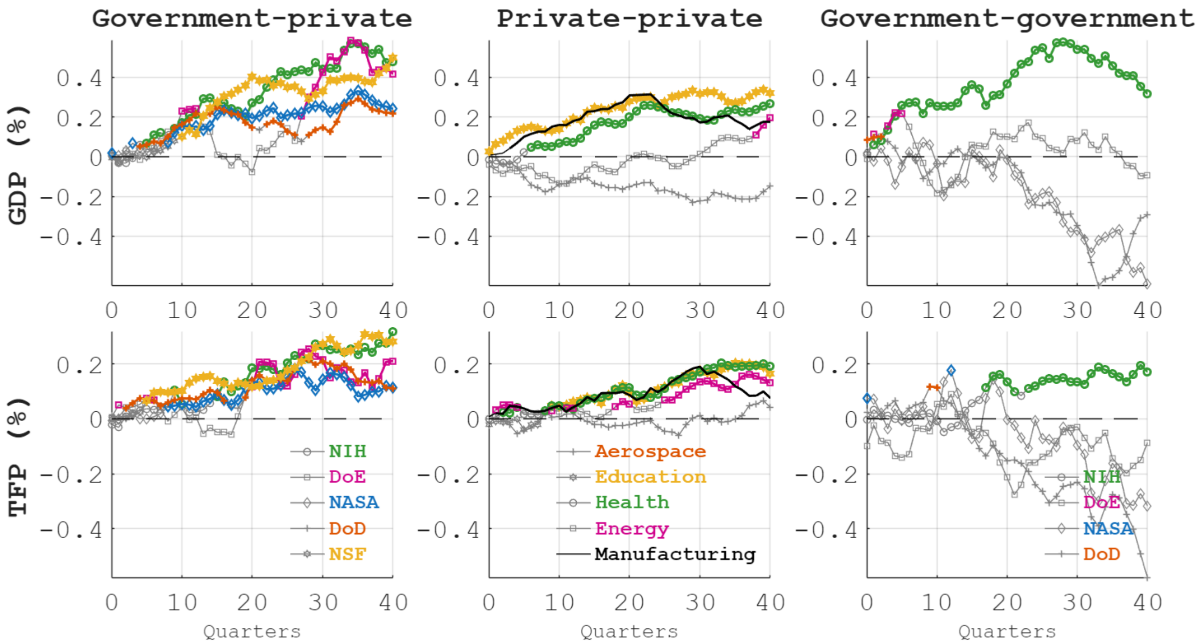

(b) Research Fields

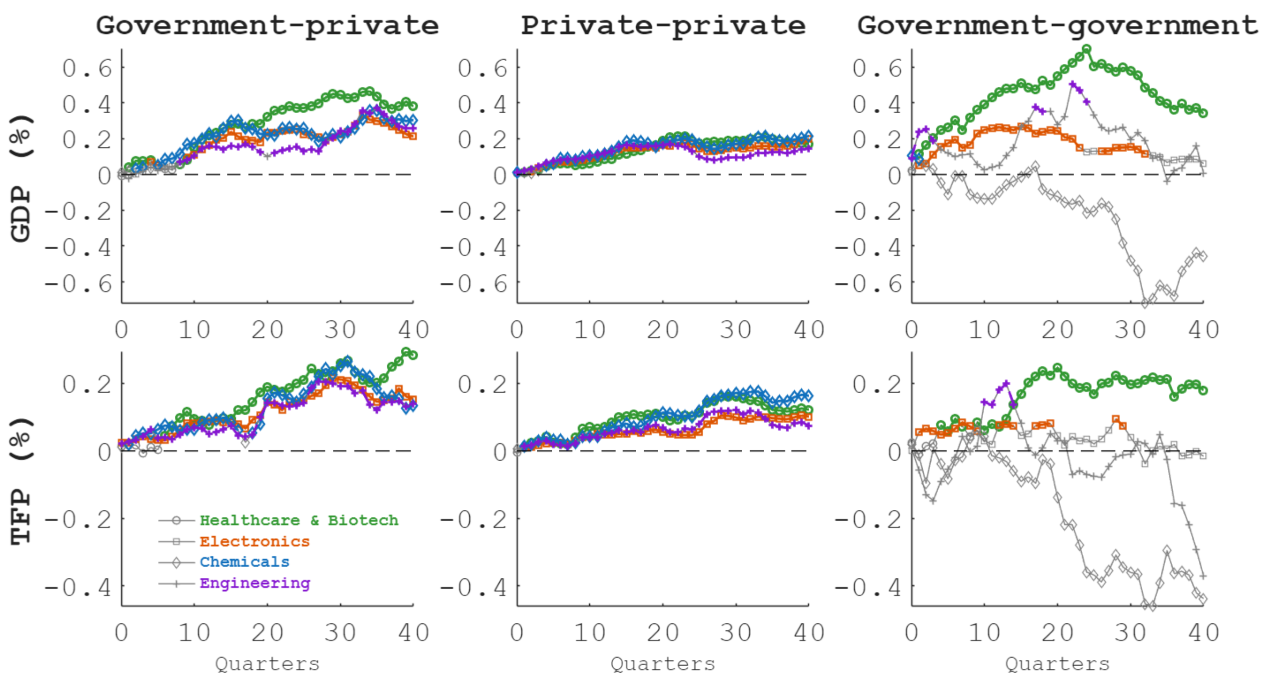

Note. The figure displays the dynamic effects of technology shocks in each patent category—public–private, private–private, and public–public—on log real per-capita GDP and log utilization-adjusted TFP. Panel A reports responses by federal agency and matched private-sector industry. Panel B reports responses by research field. The estimation follows equation (5), using the Herbst and Johannsen (2024) small-sample bias correction. The shock is normalized so that the peak response of total patent filings is 1% over the forecast horizon. The control set includes four lags of the shocked patent group, the dependent variable, real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill rate, real per-capita R&D expenditure, and patent filings outside the shocked group. All variables except the T-bill rate are in logs. Colored lines denote point estimates significant at the 68% level using heteroskedasticity-robust standard errors; gray lines denote estimates that are not significant at that level. Sample: quarterly data, 1950:Q1–2015:Q4.

measurable aggregate effect.

Agency composition helps rationalize these differences. Appendix Table I.3 shows that the agencies with larger effects fund a larger share of public–private patents by universities and research institutes. They also fund a larger share of start-ups among for-profit participants. This compositional evidence anticipates Section 8, where we show that public–private patents assigned to research entities and start-ups generate especially large macroeconomic responses.

7.2 Research Fields

The agency results raise a natural question: do NSF and NIH matter because these agencies are special, or because they fund research in fields whose technologies have larger aggregate effects? We examine this question by splitting patents into research fields. Appendix Figures G.1 and G.2 report the field composition of patents for federal agencies, private-sector innovators, research institutes, and universities. Although each group patents across many fields, four fields account for at least 40% of all patents for each group in the postwar sample: Healthcare & Biotechnology, Electronics, Chemicals, and Engineering. We therefore focus on these four fields.[18]

Panel B of Figure 7 focuses on research field. The main result is that, among government-funded patents, healthcare and biotechnology generate the largest medium-run responses of GDP and TFP. The result is especially pronounced for public–public patents: while average effects are weak, fully public patents in healthcare and biotechnology generate large positive responses. This finding helps explain why NIH stands out in the agency analysis and reinforces the evidence from Section 6 that science-based public innovation contains a consequential upper tail.

The contrast with private–private healthcare and biotechnology patents is informative. Among private patents, healthcare and biotechnology do not generate markedly larger GDP or TFP responses than other fields. This does not imply that private biomedical innovation has low social value. Many health innovations improve welfare through longevity, quality of life, and morbidity reductions that may not be fully captured by GDP or TFP over the medium-term. The narrower macroeconomic point is that, in measured productivity and

[18] Appendix Table G.2 reports the mapping between CPC codes and research fields.

output, government-funded healthcare and biotechnology shocks propagate more strongly than their fully private counterparts.

One interpretation is that private and social incentives differ especially sharply in biomedical innovation. Strategic patenting, including patent thickening and evergreening, can extend market exclusivity without commensurate productivity gains (Frakes and Wasserman, 2025; Dwivedi et al., 2010). Private pharmaceutical R&D may also be directed toward treatments that improve longevity or quality of life for older patients, whose health benefits are valuable but less directly reflected in labor productivity (Benmelech et al., 2021). Finally, private firms may target larger or more profitable markets rather than the areas with the largest aggregate productivity spillovers (Boldrin and Levine, 2008; Acemoglu and Linn, 2004; Moen-Vorum, 2025). These channels are not separately identified by our time-series design, but they provide plausible reasons why public and private patent shocks in the same broad field can have different aggregate propagation.

Taken together, the agency and field results support the mechanism. The large public–private effects come from parts of the federal innovation system that support upstream, science-based technologies and that are connected to subsequent private deployment.

7.3 Innovation Spillovers

The baseline results show that public–private patent shocks generate large responses of private R&D and investment. We now examine whether these spillovers are also visible within the agency and field cuts studied above. This exercise is directly linked to the model equation (3): source-specific technology arrivals may matter not only through their direct effect on productivity, but also by raising the productivity of subsequent private research. Figure 8 displays the response of business-sector R&D spending to shocks that raise patenting in a single agency, matched industry, or research field. The first row compares NSF, NIH, and NASA public–private patents with their closest private-sector counterparts: education, healthcare and biotechnology, and aerospace. The second row refers to DoE and DoD public–private patents alongside the energy and manufacturing sectors. The remaining panels compare public–private and private–private shocks within Chemicals, Healthcare & Biotechnology, Electronics, and Engineering.

Figure 8: The Effects of Technology Shocks on R&D by Agencies, Industries, and Fields

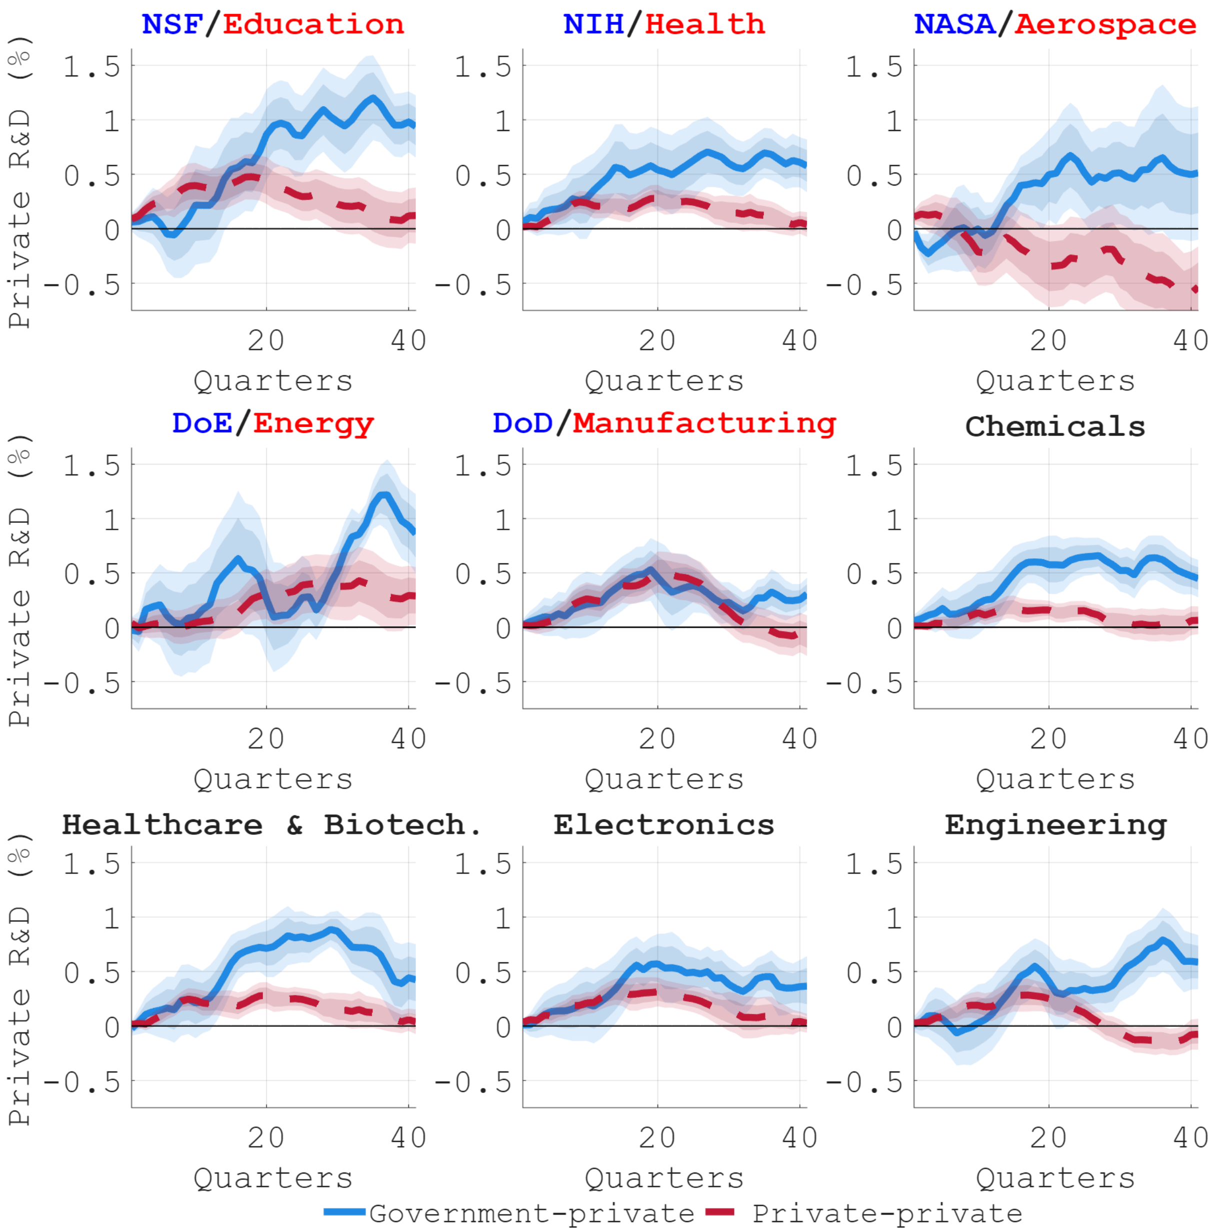

Note. *The figure displays the dynamic effects of public–private technology shocks and private–private technology shocks on log real private R&D per capita, by agency, matched industry, and research field. Solid blue lines report public–private patent shocks; dashed red lines report private–private patent shocks. The estimation follows equation (5), using the Herbst and Johannsen (2024) small-sample bias correction. The shock is normalized so that the peak response of total patent filings is 1% over the forecast horizon. The control set includes four lags of the shocked patent group, the dependent variable, real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill rate, real per-capita R&D expenditure, and patent filings outside the shocked group. All variables except the T-bill rate are in logs. Shaded areas report 68% and 90% confidence intervals based on heteroskedasticity-robust standard errors, centered on the bias-corrected estimates. Sample: quarterly data, 1950:Q1–2015:Q4.*

Three findings stand out. First, public–private patent shocks generate larger private-R&D spillovers than private–private shocks in most agencies, industries, and fields. This pattern is consistent with the model's spillover channel: publicly funded but privately owned technologies appear to raise the expected return to subsequent private research. Second, the largest differences occur in education, health, and energy, corresponding to NSF, NIH, and DoE. NSF-funded public–private patents generate especially large increases in private R&D, consistent with the view that university-based and science-based research can open new lines of private experimentation and product development.[19] NIH and DoE also generate sizable private-R&D responses, reinforcing the evidence that health and energy technologies are important channels through which public funding affects later private research. Third, the difference between public–private and private–private spillovers is smaller, though still visible, in Chemicals and Engineering, and more modest in Electronics, DoD/Manufacturing, and NASA/Aerospace. This pattern is consistent with evidence that the postwar aggregate effects of public R&D were larger outside NASA and defence (Kantor and Whalley, 2025; Fieldhouse and Mertens, 2023).[20]

Overall, the spillover evidence connects the agency and field results to the aggregate propagation documented in Section 4.2. Public–private technology shocks are followed by larger private-R&D responses not only on average, but also within the agencies and research fields that account for the strongest GDP and TFP effects. This suggests a transmission mechanism in which the largest effects arise when public funding generates technologies that are privately deployable and that crowd in subsequent private research and investment.

8 Institutional Sources of Technology Shocks

The previous sections show that the macroeconomic effects of public–private patent shocks are concentrated in more basic, more science-based technologies and in the agencies and fields most closely associated with upstream research. This section asks which institutions

[19]The finding that NSF-funded patents crowd in private R&D, while privately funded university patents may crowd it out, suggests that the effect of university research depends on the funding and ownership environment. This is consistent with evidence on commercial spillovers from university research in Lerner et al. (2025), while also qualifying the concern in Arora et al. (2023) that some forms of public science can be rivalrous and excludable. The large private-R&D response to DoE patents is consistent with Myers and Lanahan (2022).

[20]The private-R&D spillovers in Figure 8 are also consistent with Fleming et al. (2019), who document a sharp increase since the mid-1960s in corporate patents citing government-funded innovations.

produce these source-specific technology shocks. More specifically, we study two margins. First, we distinguish research entities—universities and research institutes—from other private assignees. This split captures an institutional dimension of basic research that is complementary to the patent-text and science-citation measures used in Section 6. Second, within the for-profit sector, we distinguish start-ups and venture-capital-backed firms from other (incumbent) firms. The data on universities and research institutes come from *PatentsView* since 1976 and are extended back to 1950 using historical USPTO assignee records. The start-up and venture-capital classifications are from Ewens and Marx (2024).

8.1 Research Entities and Other Organizations

The postwar U.S. research system placed universities and research institutes at the center of publicly supported basic science. Bush (1945a) defined basic research as research aimed at expanding general knowledge rather than immediate commercial application, and Bush (1945b) emphasized the distinctive role of colleges, universities, and research institutes in creating new scientific knowledge. This institutional distinction is useful for our purposes because the basicness of research is difficult to measure directly in aggregate R&D expenditure data or in patent text alone.

Appendix Figure H.1 documents the aggregate relevance of this margin. Most postwar U.S. R&D spending has been directed toward applied research and experimental development, largely performed by the private sector. At the same time, the share of basic research has risen over time, reaching about 20% of total R&D by 2010. The federal government has financed roughly 60% of annual basic R&D spending since the 1950s, and non-business sources account for about 70% once universities' own funds are included. Within federally funded basic R&D, universities and research institutes account for about 70% of expenditures. Accordingly, we treat research entities as an institutional proxy for environments specialized in upstream knowledge production.[21]

Figure 9 reports the dynamic effects of patent shocks by assignee type and funding source. Research entities include universities and research institutes, while other organizations in-

[21] This classification complements the patent-based measures of importance and reliance on science used in Section 6. It may understate the role of basic research, since universities and research institutes also perform applied research.

Figure 9: The Effects of Technology Shocks by Research Entities and Other Organizations

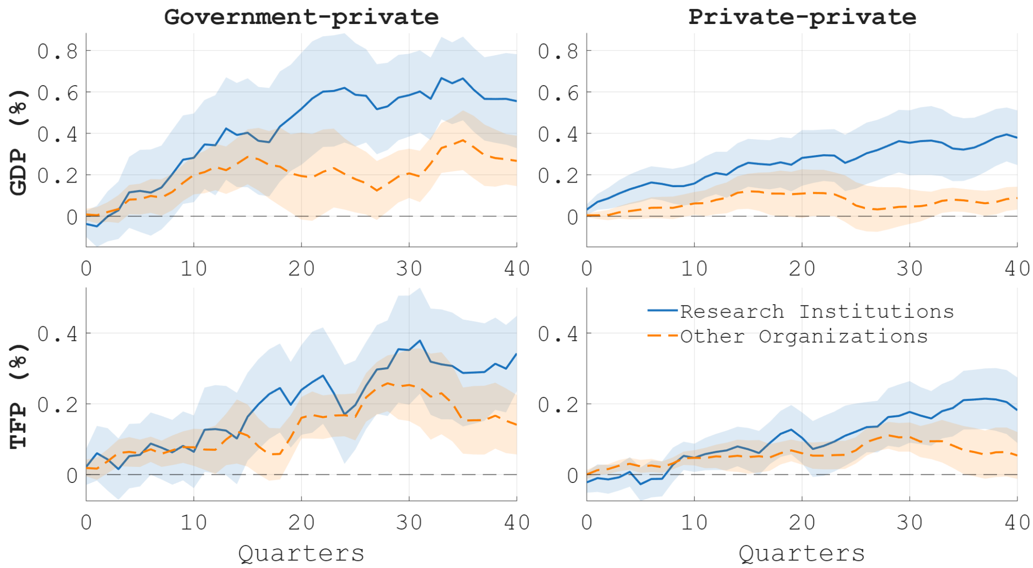

Note. The figure compares the dynamic effects of technology shocks for public–private and private–private patents across research entities and other organizations on log real per-capita GDP and log utilization-adjusted TFP. Research entities include universities and research institutes; other organizations include all remaining private assignees. The estimation follows equation (5), using the Herbst and Johannsen (2024) small-sample bias correction. Shocks are normalized so that total patents increase by 1% on impact. The control set includes four lags of the shocked patent group, the dependent variable, real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill rate, real per-capita R&D expenditure, and patent filings in the other groups. All variables except the T-bill rate are in logs. The solid blue line reports the point estimate for research entities; the dashed orange line reports the point estimate for other organizations. Shaded areas report 90% confidence intervals based on heteroskedasticity-robust standard errors, centered on the bias-corrected estimates. Sample: quarterly data, 1950:Q1–2015:Q4.

clude all other private assignees. Since Gross and Sampat (2025) classify universities and research institutes as private entities, both columns refer to privately owned patents: the left captures technologies that are publicly funded and the right column captures those that are privately funded. The comparison thus isolates whether the same broad class of private assignees generates different macroeconomic responses depending on whether the underlying research is publicly or privately funded.

Three results emerge. First, among public–private patents, shocks to research-entity patenting generate substantially larger and more persistent responses of GDP and TFP than shocks to patents assigned to other organizations. The latter also generate positive responses, especially for TFP at longer horizons, but the effects are smaller and less persistent. This result connects the average public–private effect to the institutional channel emphasized by Bush (1945a): publicly funded research performed in universities and research institutes appears to be a major source of medium-run technology shocks. Second, research entities also generate larger responses than other organizations among fully private patents, but the magnitudes are smaller than for publicly funded research entities. This comparison is important because it suggests that the research-entity effect is not purely an assignee-composition effect. The largest responses occur when the research-entity environment is combined with public funding. Third, the results help rationalize the agency evidence in Section 7.1. NIH and NSF, the agencies with the largest public–private responses, also fund a high share of patents assigned to research entities. The institutional and agency evidence therefore point to the same mechanism: public funding has its largest aggregate effects when it supports upstream research in institutions specialized in knowledge production and when the resulting technologies are privately owned and deployable.

8.2 Start-ups and Venture-Capital-Backed Firms

We now turn to for-profit assignees. Start-ups and venture-capital-backed firms may be especially important for the propagation of new technologies because they are more likely to commercialize frontier ideas, enter new markets, and reallocate resources toward new technological opportunities. At the same time, young firms may face financing constraints (as shown by Cloyne et al., 2023) and may be unable to undertake socially valuable research without public support. We therefore ask whether the macroeconomic effects of public–private patent shocks are concentrated among start-ups or venture-capital-backed firms. The classification follows Ewens and Marx (2024). We compare start-ups with other firms in Panel A of Figure 10, and venture-capital-backed firms with non-venture-capital-backed firms in Panel B.[22] The sample begins in 1976, when the assignee information needed for this classification is available.

Panel A shows that the distinction between start-ups and other firms matters primarily for publicly funded innovation. Among public–private patents, start-up patent shocks generate large and persistent responses of GDP and TFP, with peak effects around 0.6% for GDP and 0.3% for TFP. The corresponding responses for other firms are much smaller. In contrast, among private–private patents, the difference between start-ups and other firms is limited. Thus, start-ups are not uniformly associated with larger macroeconomic effects. Their dis-

[22] Start-up status and venture-capital backing can overlap. "Other firms" denotes non-start-ups in Panel A and non-venture-capital-backed firms in Panel B, and should be interpreted as separate comparisons rather than as a single decomposition.

Figure 10: The Effects of Technology Shocks by Firm Characteristics

(a) Start-ups Versus Other Firms

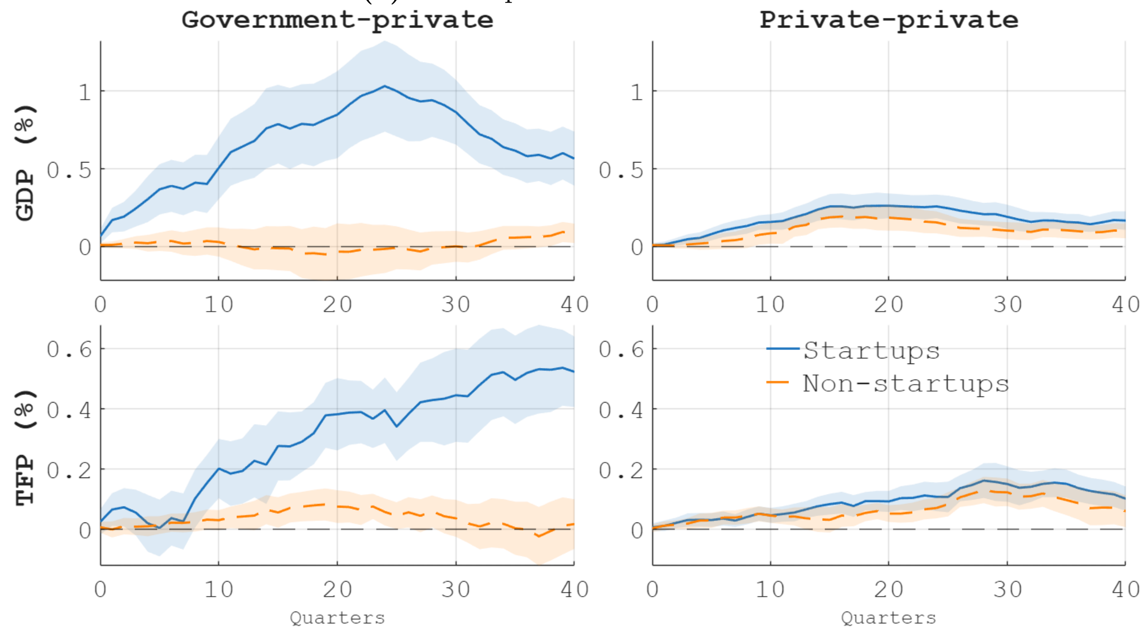

(b) Venture-Capital-Backed Firms Versus Non-VC-Backed Firms

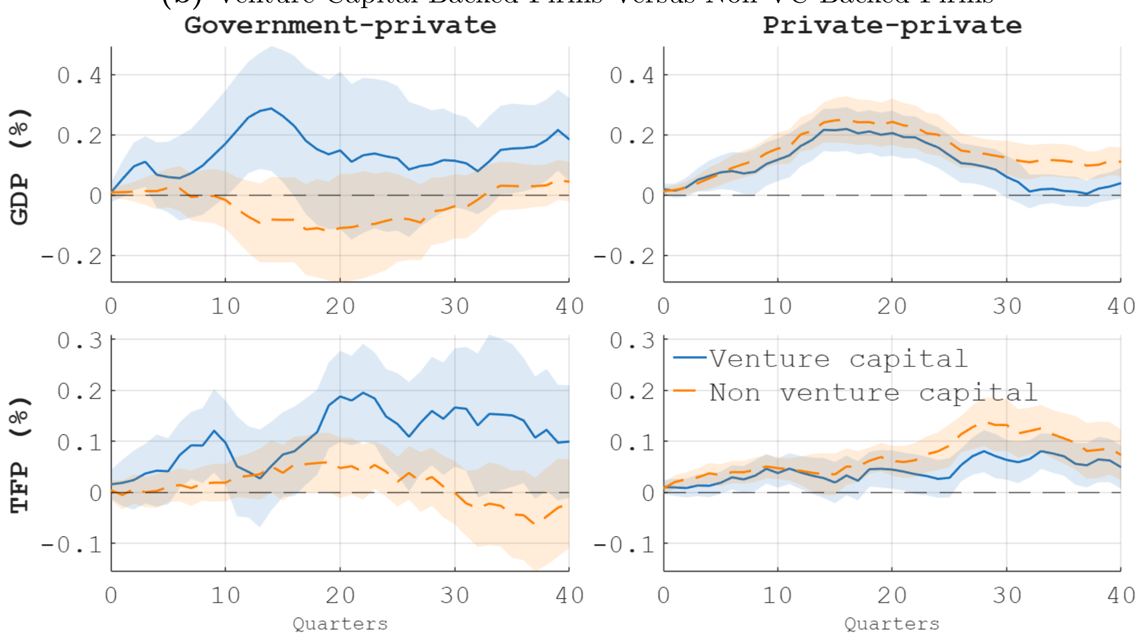

Note. The figure compares the dynamic effects of technology shocks for public–private and private–private patents across firm types on log real per-capita GDP and log utilization-adjusted TFP. Panel A compares start-ups with other firms. Panel B compares venture-capital-backed firms with non-venture-capital-backed firms. The estimation follows equation (5), using the Herbst and Johannsen (2024) small-sample bias correction. Shocks are normalized so that total patents increase by 1% on impact. The control set includes four lags of the shocked patent group, the dependent variable, real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill rate, real per-capita R&D expenditure, and patent filings in the other groups. All variables except the T-bill rate are in logs. In Panel A, the solid blue line reports start-ups and the dashed orange line reports other firms. In Panel B, the solid blue line reports venture-capital-backed firms and the dashed orange line reports non-VC-backed firms. Shaded areas report 90% confidence intervals based on heteroskedasticity-robust standard errors, centered on the bias-corrected estimates. Sample: quarterly data, 1976:Q1–2015:Q4.

tinctive contribution appears when their technologies are supported by public funding. This is consistent with a complementarity between public funding and young-firm experimentation. Public support may help start-ups commercialize frontier technologies whose social value is high but whose private financing is difficult because of uncertainty, long horizons, or spillovers. The result also connects back to the agency evidence: Appendix Table I.3 shows that NSF and NIH, the agencies associated with the largest public–private effects, fund relatively high shares of young firms among for-profit public–private participants.

Panel B repeats the exercise for venture-capital-backed firms. There is some evidence that, among public–private patents, venture-capital-backed firms generate larger responses than other firms. The magnitudes are smaller than for start-ups, however, and the confidence sets overlap. Among private–private patents, there is little systematic heterogeneity by venture-capital status. Publicly funded VC-backed patents also tend to yield larger responses than privately funded VC-backed patents, consistent with evidence on the role of venture capital in public–private innovation (Li, 2025) and with broader evidence on state-directed innovation and young firms (Beraja et al., 2024).

Taken together, the institutional evidence helps paint a fuller picture about the mechanism behind the public–private effect. The largest source-specific technology shocks arise when public funding supports technologies that are developed in research-focused institutions or young firms and are then owned, deployed, and extended outside the federal government. This is the institutional counterpart of the spillover channel in Section 7: public funding is most macroeconomically powerful when it creates privately deployable technologies that crowd in subsequent private research, investment, and diffusion.

9 Conclusions

This paper opens the black box of aggregate technology shocks. We construct postwar U.S. time series of patent filings by funding source and ownership to isolate their unanticipated movements as source-specific technology-shock proxies. Government-funded patents owned by private entities are technologies supported by public funds but deployed outside the federal government. We find that this institutional margin is quantitatively important for the transmission of technological change.

Our main finding is that government-funded but privately owned innovations are a large source of medium-run macroeconomic fluctuations, explaining roughly one-fifth of GDP and TFP forecast error variance at longer horizons. They generate larger and more persistent responses to output and productivity than privately-funded patents and are followed by sizable increases in private R&D, capital investment, real wages, and consumption. Model-implied return calculations point in the same direction: publicly funded R&D has output returns about twice those of privately funded R&D, with substantially larger crowd-in of subsequent private research. The response to government-funded but privately owned innovations shocks are strongest for patents associated with NIH and NSF, healthcare and biotechnology, universities and research institutes, and start-ups. This is consistent with a propagation mechanism in which public funding supports basic research that raises the productivity of subsequent private R&D and investment. Turning to patents funded and owned by the government, we find that they have weak average responses. But this average masks substantial heterogeneity: the upper tail of those patents, especially more important and more science-based, generates large aggregate effects, with the largest responses concentrated in upstream, science-based, privately deployable innovation.

The results speak first to empirical macroeconomics. Technology shocks are typically identified from restrictions on aggregate time series or from external proxies, while their economic origins remain largely unobserved. By measuring technology arrivals by funding source and ownership, this paper shows that the institutional source of innovation is informative about both the size and the propagation of aggregate technology shocks. The findings also speak to growth and innovation. They support the classical view of Nelson (1959), Arrow (1962), and Bush (1945a) that upstream research can generate social returns that are not fully internalized by private firms: the largest aggregate effects arise when public funding supports science-based technologies that are privately owned, privately deployed, and extended through further private research.

The broader implication is that the technology frontier should not be modeled as a single aggregate residual. It is shaped by institutions that determine which ideas are funded, who owns them, and how they are deployed and extended. The empirical moments documented in this paper—the disproportionate role of public–private patent shocks, the crowd-in of private R&D, the importance of science-based research, and the upper-tail nature of government-owned innovation—provide guidance for endogenous growth models in which technologies differ not only by sector or quality, but also by funding and ownership regimes. Incorporating these margins can bring models of growth and fluctuations closer to the institutional architecture that produces new ideas. The origins of technology shocks are not incidental to macroeconomic growth: they are part of the very mechanism through which growth occurs.

References

Acemoglu, D. and J. Linn (2004): "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry*," *The Quarterly Journal of Economics*, 119, 1049–1090.

Anderson, T. W. and H. Rubin (1949): "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *Annals of Mathematical Statistics*, 20, 46–63.

Antolin-Diaz, J. and P. Surico (2025): "The Long-Run Effects of Government Spending," *American Economic Review*, 115, 2376–2413.

Arora, A., S. Belenzon, L. C. Cioaca, L. Sheer, and H. Zhang (2023): "The Effect of Public Science on Corporate R&D," Working Paper 31899, NBER.

Arrow, K. (1962): "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, National Bureau of Economic Research, Inc, 609–626.

Azoulay, P., J. S. Graff Zivin, D. Li, and B. N. Sampat (2019): "Public R&D Investments and Private-sector Patenting: Evidence from NIH Funding Rules," *The Review of Economic Studies*, 86, 117–152.

Barnichon, R. and G. Mesters (2020): "Identifying modern macro equations with old shocks," *The Quarterly Journal of Economics*, 135, 2255–2298.

Benmelech, E., J. Eberly, D. Papanikolaou, and J. Krieger (2021): "Private and Social Returns to RD: Drug Development and Demographics," *AEA Papers and Proceedings*, 111, 336–40.

Beraja, M., W. Peng, D. Y. Yang, and N. Yuchtman (2024): *Government as Venture Capitalists in Artificial Intelligence*, University of Chicago Press, 81–102.

Bergeaud, A., A. Guillouzouic, E. Henry, and C. Malgouyres (2025): "From public labs to private firms: magnitude and channels of R&D spillovers," *The Quarterly Journal of Economics, forthcoming*.

Blanchard, O. J. and D. Quah (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review*, 79, 655–673.

Bloom, N., M. Schankerman, and J. Van Reenen (2013): "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 81, 1347–1393.

Boldrin, M. and D. K. Levine (2008): *Against Intellectual Monopoly*, no. 9780521879286 in Cambridge Books, Cambridge University Press.

Bush, V. (1945a): *Science, the Endless Frontier*, Princeton University Press.

——— (1945b): "Transmittal Letter and Summary of the Report 'Science: The Endless Frontier'," A report to the president, United States Government Printing Office.

CBO (2025): "Preliminary Analysis of How Federal Investment in Nondefense Research and Development Affects the Economy and the Federal Budget," www.cbo.gov/publication/61375, Congressional Budget Office.

Christiano, L. and T. Fitzgerald (2003): "The Band Pass Filter," *International Economic Review*, 44, 435–465.

Cloyne, J., C. Ferreira, M. Froemel, and P. Surico (2023): "Monetary Policy, Corporate Finance, and Investment," *Journal of the European Economic Association*, 21, 2586–2634.

Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2017): "Innocent Bystanders? Monetary policy and inequality," *Journal of Monetary Economics*, 88, 70–89.

Danguy, J., G. de Rassenfosse, and B. van Pottelsberghe de la Potterie (2014): "On the origins of the worldwide surge in patenting: an industry perspective on the R&D–patent relationship," *Industrial and Corporate Change*, 23, 535–572.

De Lipsis, V., M. Deleidi, M. Mazzucato, and P. Agnolucci (2023): "Macroeconomic Effects of Public R&D," .

De Ridder, M. (2024): "Market Power and Innovation in the Intangible Economy," *American Economic Review*, 114, 199–251.

DSIT (2025): "The value of public R&D: Research and Analysis," https://www.gov.uk/government/publications/the-value-of-public-rd, Department for Science, Innovation and Technology of the UK government.

Dwivedi, G., S. Hallihosur, and L. Rangan (2010): "Evergreening: A deceptive device in patent rights," *Technology in Society*, 32, 324–330.

Dyèvre, A. (2024): "Public R&D Spillovers and Productivity Growth," Mimeographed, london School of Economics.

Ewens, M. and M. Marx (2024): "Firm Age and Invention: An Open-Access Dataset," .

Fernald, J. G. (2012): "A quarterly, utilization-adjusted series on total factor productivity," Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.

Fieldhouse, A. J. and K. Mertens (2023): "The Returns to Government R&D: Evidence from U.S. Appropriations Shocks," Working Papers 2305, Dallas Fed.

Fleming, L., H. Greene, G. Li, M. Marx, and D. Yao (2019): "Government-funded research increasingly fuels innovation," *Science*, 364, 1139–1141.

Forni, M. and L. Gambetti (2014): "Sufficient information in structural VARs," *Journal of Monetary Economics*, 66, 124–136.

Frakes, M. D. and M. F. Wasserman (2025): "Strategic Patenting: Evidence from the Biopharmaceutical Industry," Working Paper 34024, NBER.

Francis, N. and V. A. Ramey (2005): "Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited," *Journal of Monetary Economics*, 52, 1379–1399.

Galí, J. (1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89, 249–271.

Gorodnichenko, Y. and B. Lee (2020): "Forecast error variance decompositions with local projections," *Journal of Business & Economic Statistics*, 38, 921–933.

Gouriéroux, C., A. Monfort, and J.-P. Renne (2020): "Identification and estimation in non-fundamental structural VARMA models," *The Review of Economic Studies*, 87, 1915–1953.

Griliches, Z. (1990): "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 28, 1661–1707.

Gross, D. P. and B. N. Sampat (2025): "The Government Patent Register: A new resource for measuring U.S. government-funded patenting," *Research Policy*, 54, 105142.

Herbst, E. P. and B. K. Johannsen (2024): "Bias in Local Projections," *Journal of Econometrics*, 240, 105655.

Himberg, J., A. Hyvarinen, and F. Esposito (2004): "Validating the independent components of neuroimaging time series via clustering and visualization," *NeuroImage*, 22, 1214 – 1222.

Hyvarinen, A. (1999): "Fast and robust fixed-point algorithms for independent component analysis," *IEEE Transactions on Neural Networks*, 10, 626–634.

Jones, B. F. and L. H. Summers (2021): *1 A Calculation of the Social Returns to Innovation*, Chicago: University of Chicago Press, 13–60.

Jones, C. I. (1995): "R&D-Based Models of Economic Growth," *Journal of Political Economy*, 103, 759–784.

Jordà, Ò. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161–182.

Jordà, Ò. and A. M. Taylor (2025): "Local Projections," *Journal of Economic Literature*, 63, 59–110.

Kantor, S. and A. T. Whalley (2025): "Moonshot: Public R&D and Growth," *American Economic Review*, 115, 2891–2925.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021): "Measuring Technological Innovation over the Long Run," *American Economic Review: Insights*, 3, 303–20.

Keweloh, S. A. (2021): "A generalized method of moments estimator for structural vector autoregressions based on higher moments," *Journal of Business & Economic Statistics*, 39, 772–782.

Kline, P., N. Petkova, H. Williams, and O. Zidar (2019): "Who Profits from Patents? Rent-Sharing at Innovative Firms," *The Quarterly Journal of Economics*, 134, 1343–1404.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017): "Technological Innovation, Resource Allocation, and Growth*," *The Quarterly Journal of Economics*, 132, 665–712.

Lazarus, E., D. J. Lewis, J. H. Stock, and M. W. Watson (2018): "HAR Inference: Recommendations for Practice," *Journal of Business & Economic Statistics*, 36, 541–559.

Lerner, J., H. J. Manley, C. Stein, and H. L. Williams (2025): "The Wandering Scholars: Understanding the Heterogeneity of University Commercialization," *Econometrica*, forthcoming.

Lewis, D. J. (2021): "Identifying shocks via time-varying volatility," *The Review of Economic Studies*, 88, 3086–3124.

Li, X. (2025): "Complementarity in RD Investment: Evidence from the Space Race," *mimeo*.

Martinez, J. and J. Moen-Vorum (2025): "Task Relevance: A New Measure of Technological Change," Working Paper.

Marx, M. and A. Fuegi (2020): "Reliance on science: Worldwide front-page patent citations to scientific articles," *Strategic Management Journal*, 41, 1572–1594.

——— (2022): "Reliance on science by inventors: Hybrid extraction of in-text patent-to-article citations," *Journal of Economics & Management Strategy*, 31, 369–392.

McCracken, M. and S. Ng (2020): "FRED-QD: A quarterly database for macroeconomic research," Tech. rep., National Bureau of Economic Research.

Mertens, K. and M. O. Ravn (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103, 1212–47.

Miranda-Agrippino, S., S. Hacioğlu-Hoke, and K. Bluwstein (2025): "Patents, News, and Business Cycles," *Review of Economic Studies, forthcoming.*

Moen-Vorum, J. (2025): "Chronic Underinvestment: How Patents Change the Direction of Pharmaceutical R&D," Working paper.

Montiel Olea, J. L., M. Plagborg-Møller, E. Qian, and C. K. Wolf (2025): "Local Projections or VARs? A Primer for Macroeconomists," Working Paper 33871, National Bureau of Economic Research.

Myers, K. R. and L. Lanahan (2022): "Estimating Spillovers from Publicly Funded RD: Evidence from the US Department of Energy," *American Economic Review*, 112, 2393–2423.

Nelson, R. (1959): "The Simple Economics of Basic Scientific Research," *Journal of Political Economy*, 67.

Newey, W. K. and K. D. West (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Plagborg-Møller, M. and C. K. Wolf (2021): "Local projections and VARs estimate the same impulse responses," *Econometrica*, 89, 955–980.

Ramey, V. A. (2020): "Secular Stagnation or Technological Lull?" *Journal of Policy Modeling*, presented at the AEA session "The United States Economy: Growth, Stagnation or New Financial Crisis?".

Ramey, V. A. and S. Zubairy (2018): "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 126.

Rigobon, R. (2003): "Identification through heteroskedasticity," *Review of Economics and Statistics*, 85, 777–792.

Romer, C. D. and D. H. Romer (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94, 1055–1084.

——— (2010): "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, 100, 763–801.

Romer, P. M. (1990): "Endogenous Technological Change," *Journal of Political Economy*, 98, S71–S102.

Stock, J. H. and M. W. Watson (2018): "Identification and estimation of dynamic causal effects in macroeconomics using external instruments," *The Economic Journal*, 128, 917–948.

Uhlig, H. (2003): "What moves real GNP?" *Unpublished manuscript.*

Appendix

"Where Do Technology Shocks Come From? Public Funding and Private Ownership"

by Andrea Gazzani (Bank of Italy), Joseba Martinez (London Business School), Filippo Natoli (Bank of Italy) and Paolo Surico (London Business School)

A A Model of Technology by Funding Source

This appendix presents the model summarized in the main text and derives all the results used there. Technologies are indexed by their funding source $s \in \mathcal{S}$, matching the source categories we track in the data.

The model serves three purposes. First, it separates the two channels through which source-s technologies affect the economy: a direct contribution to TFP, and an indirect contribution through the productivity of every source's subsequent research. The classical view of Nelson (1959) and Arrow (1962), with publicly supported research upstream of commercial innovation, is the nested case in which the direct effect of public innovation is small, but the indirect effect is large. Second, it establishes that the log-linear impulse responses around the balanced growth path depend only on elasticities, balanced-growth rates, and balanced-growth shares, not on the levels of the idea stocks. Through the lens of our model, around a balanced growth path, our pooled time-series estimates recover a stable object. Third, it delivers a transparent mapping from impulse responses measured per patent to multipliers and rates of return measured per dollar of R&D.

A.1 Model

Final output is produced with TFP A_t and production labor L_{yt}. TFP accumulates from the stocks of patented ideas of all sources, with the direct contribution of source s governed by λ_s. New source-s patents arrive as draws from a Poisson process with arrival rate $\mu_{s,t}$ that depends on source-s R&D input $R_{s,t}$, on the idea stocks of all sources through the spillover elasticities $\theta_{s'}$, and on the aggregate technology level A_t through the knowledge spillover elasticity ϕ. Each idea stock evolves as the running count of its Poisson arrivals net of depreciation. Sources differ in how their R&D is funded. Privately funded sources, collected in $\mathcal{S}_m$, hire research labor in the market; publicly funded sources, collected in $\mathcal{S}_g$, receive R&D set by a policy rule. Population grows at rate n and is allocated between production and research in constant shares. Table A.1 collects the equations.

Table A.1: The Multi-Source Model

	Equation	
Final output	$Y_t = A_t L_{yt}$	(M1)
TFP accumulation	$\dot{A}_t / A_t = \sum_{s\in\mathcal{S}} \lambda_s X_{st}/A_t - \delta$	(M2)
Idea arrival rates	$\mu_{s,t} = \chi_s R_{s,t}^{\eta} A_t^{\phi} \prod_{s'\in\mathcal{S}} X_{s't}^{\theta_{s'}}$	(M3)
Idea stocks	$\dot{X}_{st} = \mu_{s,t} - \delta X_{st}$	(M4)
Private R&D input	$R_{s,t} = A_t L_{r,s,t}, \quad s \in \mathcal{S}_m$	(M5)
Public R&D input	$R_{s,t} = R_{s,0} e^{g_y t}, \quad s \in \mathcal{S}_g$	(M6)
Population	$L_{yt} + \sum_{s\in\mathcal{S}_m} L_{r,s,t} = L_0 e^{nt}$	(M7)

Equation (M2) restates equation (2) in the main text, adding depreciation for the reasons discussed below. Equation (M3) generalizes equation (3): the normalization constants $\chi_s > 0$ absorb the proportionality in equation (3), and the term A_t^{ϕ} adds a knowledge spillover from the aggregate technology level to research productivity.

Four remarks on the environment. First, the Poisson layer makes the connection to the data direct. An empirical 1-patent shock is one realization of (M3) and moves the corresponding stock X_{st} up by one unit. Equation (M4) describes the deterministic-mean dynamics of X_{st} around the shock, which is what enters the log-linearization below. The arrival functions are isoelastic in own R&D with elasticity η, the central patent-production elasticity in the model. The spillover elasticity $\theta_{s'}$ is indexed by the source of origin: technologies of source s' raise the productivity of every source's research with the same elasticity. We write $\theta \equiv \sum_{s'\in\mathcal{S}} \theta_{s'}$ for the total spillover elasticity. The term A_t^{ϕ} makes research productivity depend on the aggregate level of technology itself. The elasticity ϕ is the intertemporal knowledge spillover of semi-endogenous growth models (Jones, 1995).

Second, because the marginal product of production labor in (M1) is A_t, the competitive wage is A_t and the research input $R_{s,t} = A_t L_{r,s,t}$ of a market source equals its R&D wage bill. Measured in units of output, $R_{s,t}$ is source-s R&D spending, for market and policy-funded sources alike. Third, the policy rule (M6) holds public R&D at a constant share of GDP. This choice reflects the idea that government R&D budgets are set by policymakers and therefore do not respond to market forces. On the balanced growth path derived below, aggregate output grows at rate g_y, so each public R&D share of GDP is constant, and g_y is the unique growth rate consistent with this policy. Production labor L_{yt} and the research labor of each market source are allocated in constant shares, so both grow at rate n. Fourth, the model in the main text does not include depreciation of patented ideas. We add it here in (M2) and (M4) for two reasons. Empirical work consistently finds that R&D capital depreciates at rates in the range of 10–20 percent per year, and the U.S. national accounts treat R&D as a depreciating capital input on those same grounds. Since the social returns we compute are intended to be compared with estimates from the R&D-returns literature, we adopt the same convention here. A single rate δ applies to TFP in (M2), to the patent stocks in (M4), and to the user cost of R&D capital below; this is a parsimony choice, and the derivations go through with distinct rates. The depreciation rate δ affects the timing of the impulse responses, but it does not change any of the qualitative conclusions of the model.

A.2 Balanced Growth Path

Let g denote the balanced-growth rate of GDP per capita and g_y the balanced-growth rate of aggregate output. With population growing at rate n, the two are related by

$$g_y = g + n. \tag{1}$$

We use the model equations one at a time to identify the balanced-growth rate of each remaining variable.

From equation (M1), $Y_t = A_t L_{yt}$, with production labor growing at n, TFP A_t grows at g. From equation (M5), market R&D inputs grow at $g + n = g_y$. The policy rule (M6) sets public R&D to grow at g_y so that its share of GDP is constant on the BGP. All R&D inputs therefore grow at the common rate g_y.

Equation (M2) requires each ratio X_{st}/A_t to be constant on the BGP, so every idea stock grows at g, the growth rate of TFP, and the BGP levels satisfy

$$\sum_{s\in\mathcal{S}} \lambda_s \frac{\bar{X}_s}{\bar{A}} = g + \delta, \tag{2}$$

where bars denote detrended BGP levels. Equation (M4) in turn requires each arrival rate $\mu_{s,t}$ to grow at g as well, with BGP ratio

$$\frac{\bar{\mu}_s}{\bar{X}_s} = g + \delta \qquad \text{for every } s \in \mathcal{S}. \tag{3}$$

Equation (M3) says each arrival rate grows at $\eta g_y + (\phi + \theta)g$, since TFP and every idea stock grow at g, where $\theta = \sum_{s'} \theta_{s'}$ is the total spillover elasticity. Consistency with the growth rate g of the arrival rates requires

$$\eta g_y + (\phi + \theta)g = g. \tag{4}$$

Using $g_y = g + n$ and solving for g,

$$g = \frac{\eta}{1 - \eta - \phi - \theta}\, n, \qquad 1 - \eta - \phi - \theta > 0. \tag{5}$$

Growth is semi-endogenous. It is pinned down by population growth and the elasticities, while the scale of research effort, given by the research labor shares and the public R&D levels $R_{s,0}$, affects only the levels of the stocks. Both spillover margins amplify the effect of population growth on g: the knowledge spillover ϕ and the patent spillovers θ enter the denominator symmetrically.

A.3 Impulse Responses Around the Balanced Growth Path

We model the arrival of an unusually large cohort of patents as a one-time innovation to the relevant stock at $t = 0$, and we characterize the resulting transition back to the BGP. Let $\hat{x}_t$ denote the log deviation of a detrended variable from its BGP path. Because the system below is linear, responses scale with the size of the innovation, and we normalize to a unit log impulse. Log-linearizing (M2) and (M3)–(M4) around the BGP gives

$$\dot{\hat{A}}_t = \sum_{s \in \mathcal{S}} \omega_s \big(\hat{X}_{st} - \hat{A}_t \big), \tag{6}$$

$$\dot{\hat{X}}_{st} = (g + \delta) \Big(\eta \hat{R}_{st} + \phi \hat{A}_t + \sum_{s' \in \mathcal{S}} \theta_{s'} \hat{X}_{s't} - \hat{X}_{st} \Big), \tag{7}$$

where

$$\omega_s \equiv \lambda_s \frac{\bar{X}_s}{\bar{A}}, \qquad \sum_{s \in \mathcal{S}} \omega_s = g + \delta, \tag{8}$$

is the BGP contribution of source-s technologies to TFP growth plus depreciation, from equation (2), and the coefficient $g + \delta$ in equation (7) is the BGP arrival-to-stock ratio from equation (3). For market sources, constant labor shares imply that research spending moves with the wage, $\hat{R}_{st} = \hat{A}_t$; for policy-funded sources the rule (M6) implies $\hat{R}_{st} = 0$. TFP therefore feeds back into patent production with total elasticity $\eta + \phi$ for market sources, through the R&D input and the knowledge spillover, and ϕ for policy-funded sources, through the knowledge spillover alone. Each source transmits a shock through two channels: the direct channel, with strength ω_s in equation (6), and the spillover channel, with strength θ_s in every stock equation (7).

Invariance of the IRFs along the BGP. Every coefficient in equations (6)–(7) is a combination of the elasticities η, ϕ, and $\theta_{s'}$, the BGP growth rate g, the depreciation rate δ, and the BGP shares ω_s. The shares are constant along the BGP because every idea stock grows at the same rate as TFP. The levels of A_t and X_{st} do not appear, and the normalization constants χ_s in (M3) enter only through the BGP shares ω_s. A one percent innovation to any idea stock therefore produces the same percent response of TFP and output wherever the economy sits on its growth path, even though the stocks grow over time and a single idea moves its stock by an ever smaller percentage. Pooled estimation over a long sample therefore recovers a stable object rather than an average of drifting responses, provided the economy stays close to its BGP. Diminishing returns to research determine the BGP growth rate and the levels of the stocks, but they leave the log-linear impulse responses unchanged. Two properties of the IRF follow from this construction: the response at any horizon scales one-for-one with the size of the shock, so it is the marginal effect of a small unanticipated innovation evaluated at the BGP; and it is the response to a date-zero patent surprise. These are the same two properties the local projections identify in the data.

A.4 Multiplier and Return Formulas

This section derives the dollar multiplier on R&D in two steps. First, we convert the per-patent TFP IRF from Section A.3 into a dollar present value per patent. Second, we translate the per-patent value into the marginal GDP multiplier per dollar of R&D, which multiplies the per-patent value by the idea-production elasticity η and divides it by the average dollar cost per patent. In what follows, t and the horizon h are measured in quarters, while the discount rate r, the growth rate g, and the depreciation rate δ are expressed at annual rates. The date of the shock is normalized to $t = 0$, so the detrended BGP levels below, such as $\bar{X}_s$, $\bar{Y}$, and $\bar{N}_s$, equal the corresponding levels at the date of the shock.

From TFP IRF to dollar GDP dividend per patent. The IRF $\hat{A}_t^{(s)}$ is the log response of TFP to a unit log innovation in the source-s patent stock $X_{s,t}$. One additional source-s patent moves the stock by $1/\bar{X}_s$ in logs, where $\bar{X}_s$ is the BGP source-s patent stock at the date of the shock. The per-patent TFP IRF is therefore

$$\widetilde{A}_t^{(s)} \equiv \frac{\hat{A}_t^{(s)}}{\bar{X}_s}. \tag{9}$$

With $Y_t = A_t L_{yt}$ and L_{yt} on its BGP path, the dollar GDP change at horizon t per patent equals BGP dollar GDP at t scaled by $\widetilde{A}_t^{(s)}$,

$$D_s(t) = \bar{Y}\,(1+g_y)^{t/4}\,\widetilde{A}_t^{(s)} = \frac{\bar{Y}}{\bar{X}_s}\,(1+g_y)^{t/4}\,\hat{A}_t^{(s)}, \tag{10}$$

where $\bar{Y}$ is current annual real GDP. The ratio $\bar{Y}/\bar{X}_s$ is BGP dollar GDP per source-s patent in the stock. It converts the unit-free TFP IRF into a dollar flow. The object $D_s(t)$, whose quarterly flow $\Delta D_s(t)$ appears in the main text under the notation $\hat{Y}_{h,s}$, is measured in real dollars of GDP per additional source-s patent at horizon h.

Present value per patent. Discounting at the annual rate r with quarterly increment $\Delta = 1/4$, the present value per source-s patent over horizon H is

$$V_s(H;r) = \sum_{t=0}^{H} \Delta(1+r)^{-t/4} D_s(t) = \frac{\bar{Y}}{\bar{X}_s} \sum_{t=0}^{H} \Delta \left(\frac{1+g_y}{1+r}\right)^{t/4} \hat{A}_t^{(s)}. \tag{11}$$

Equation (11) is the model's one-patent dollar multiplier, the present-value GDP attributable to a single source-s patent. Differences in V_s across sources inherit the differences between the TFP IRFs $\hat{A}_t^{(s)}$ generated by the direct and spillover channels in Section A.3.

Average dollar cost per patent. The R&D input $R_{s,t}$ is source-s R&D spending in dollars, and the Poisson arrival rates in (M3) are the patent arrival flows. The average dollar R&D cost per source-s patent is

$$C_s^{\text{pat}} \equiv \frac{R_{s,t}}{\mu_{s,t}}. \tag{12}$$

C_s^{pat} is observable as the ratio of source-s R&D spending to source-s patent counts. On the BGP it is not constant. Source-s R&D spending grows at g_y while the source-s patent flow grows at g, so the real cost per patent rises at the population growth rate n for every source. We use this BGP rate when costs need to be projected forward below.

Marginal arrival product and the role of η. All arrival rates are isoelastic in own R&D with elasticity η, so

$$\frac{\partial \mu_{s,t}}{\partial R_{s,t}} = \eta \frac{\mu_{s,t}}{R_{s,t}} = \frac{\eta}{C_s^{\text{pat}}}. \tag{13}$$

A marginal dollar of source-s R&D buys η / C_s^{pat} extra patents. The elasticity $\eta < 1$ captures the diminishing returns to R&D in idea production.

Marginal GDP multiplier per dollar of R&D. The marginal GDP multiplier on source-s R&D is the marginal patent arrival product multiplied by the dollar present value per patent,

$$M_s(H;r) = \frac{\partial \mu_{s,t}}{\partial R_{s,t}} V_s(H;r) = \eta \frac{V_s(H;r)}{C_s^{\text{pat}}}. \tag{14}$$

This is the central formula. The marginal GDP multiplier is the ratio of per-patent value to per-patent cost, scaled by the elasticity η. The two ingredients are observable in principle, V_s from the per-patent IRF in equation (11) and C_s^{pat} from R&D and patent count data.

From the multiplier to a flow social rate of return. The marginal multiplier $M_s(H; r)$ is the present-value dollar gain per dollar of source-s R&D evaluated at the discount rate r. The rate-of-return concept used in the empirical R&D literature (Fieldhouse and Mertens, 2023) is instead a flow social rate of return, and a user-cost argument links the two.

Treat one dollar of source-s R&D as a unit of depreciable knowledge capital with depreciation rate δ and discount rate r. A flow of r_s dollars of GDP per year per dollar of remaining capital has present value $r_s/(r+\delta)$. Setting this present value equal to M_s and solving for the flow,

$$r_s = (r+\delta)\, M_s, \tag{15}$$

which is the annualized per-period GDP gain per dollar of source-s R&D that justifies the multiplier M_s under perpetuity accounting. Equation (15) is the model's flow social rate of return on source-s R&D, gross of depreciation.

Combining equations (14) and (15) yields equation (4) in the main text,

$$r_s = (r+\delta)\frac{\eta}{C_s}\sum_{h=0}^{T}\left(\frac{1}{1+r}\right)^{h}\hat{Y}_{h,s},$$

with $\hat{Y}_{h,s} = \Delta D_s(h)$ the quarterly dollar GDP dividend per source-s patent implied by equation (10), the discount factor $(1+r)^{-h}$ evaluated at the quarterly equivalent of the annual rate r while the user cost $r+\delta$ remains at annual rates, $C_s = C_s^{\text{pat}}$ the average dollar R&D cost per source-s patent, and $T = H$ the valuation horizon. The output multiplier reported in the main text, $r_s/(r+\delta)$, is $M_s(H; r)$. The empirical implementation below replaces a single per-patent value with the value of the entire patent sequence induced by the shock.

Empirical implementation with patent sequences. The empirical shock generates a sequence of additional patents rather than a single patent on impact, so we value the GDP response of the shock against the R&D cost of the entire induced patent sequence. The empirical marginal GDP multiplier is

$$\widehat{M}_s(H;r) = \eta \, \frac{V_s^{\mathrm{GDP}}(H;r)}{C_s^{\mathrm{seq}}(H;r)}. \tag{16}$$

This is the empirical counterpart of equation (14), with V_s^{GDP} in place of V_s and C_s^{seq} in place of C_s^{pat}, and with η entering identically. Note that because the numerator and the denominator are responses to the same shock, any normalization of the shock, for instance per impact patent, cancels from the ratio, so the multiplier is invariant to the scale of the shock.

B The Anatomy of Public Innovation

This Appendix provides further details on government-funded innovations. We report the top government-private and private-private assignees in our dataset, provide details on government-funded historically important patents, and offer insights from key studies.

B.1 Top assignees in public-private partnerships and examples of important public-interest patents

Table B.1 reports entities that own most patents in the government-private and private-private categories. The University of California, MIT, and General Electric are the top three assignees in the first category. IBM, Samsung and Canon are the top three in the privately funded patents.

Table B.2 reports government-funded patents in the historically important group identified by Kelly et al. (2021).[1] As the list ended in 2002, we have extended it till 2015 by employing *Chat-GPT Deep Research* and manually verifying all information.

[1] The list was based on a list of 250 historically important patents formerly available from the USPTO.

Table B.1: Top Assignees by Public and Private Interest

Public-Private		*Private-Private*	
Name	**Count**	**Name**	**Count**
University of California	6537	IBM Corporation	113608
Massachusetts Institute of Technology	3301	Samsung Electronics	87658
General Electric Company	2547	Canon	65872
California Institute of Technology	2266	Fujitsu	48771
Wisconsin Alumni Research Foundation	2158	Sony Group	41398
Stanford University	1523	General Electric Company	40694
University of Texas System	1425	Toshiba	39858
Johns Hopkins University	1256	Hitachi	37172
University of Michigan	1157	Intel Corporation	34992
Harvard University	1079	Mitsubishi Electric	32575
Columbia University	1013	Sumitomo Electric Industries	30523
Boeing Company	1000	NEC Corporation	28600
UT-Battelle, LLC	1000	Siemens AG	27206
Northwestern University	966	Microsoft Corporation	26115
Raytheon Company	885	Micron Technology	24477
Honeywell International Inc.	869	LG Electronics	23918
University of Pennsylvania	841	Seiko Epson	22734

Table B.2: Historically important patents with public interest

Patent	Year	Inventor	Invention	Govt. Interest	Private actor	Agency
1980972	1934	Small Lyndon Frederick	Morphine	pub–pub		-
2206634	1940	Enrico Fermi et al	Radioactive Isotopes	pub–pub		DOE
2329074	1943	Paul Muller	DDT – Insecticide	pub–priv	Novartis	DOD
2404334	1946	Frank Whittle	Jet Engine	pub–priv	Power Jets Limited	DOD
2682050	1954	Andrew Alford	Radio Navigation System	pub–priv	Andrew Alford	DOD
2708656	1955	Enrico Fermi et al	Neutronic reactor	pub–pub		DOE
2708722	1955	An Wang	Magnetic Core Memory	pub–priv	An Wang	DOD
2816721	1957	R. J. Taylor	Rocket Engine	pub–pub		DOD
2835548	1958	Robert C. Baumann	Satellite	pub–pub		DOD
2879439	1959	Charles H. Townes	Maser	pub–priv	Charles H. Townes	DOD
3093346	1963	M. A. Faget et al	First Manned Space Capsule	pub–pub		NASA
3156523	1964	G. T. Seaborg	Americium (Element 95)	pub–pub		DOE
3478216	1969	G. Carruthers	Far-Ultraviolet Camera	pub–pub		DOD
4237224	1980	Boyer & Cohen	Molecular chimeras	pub–priv		NSF
4363877	1982	H. M. Goodman et al	Human Growth Hormone	pub–priv	UCSD	NIH
4399216	1983	Richard Axel et al	Genetic transformation	pub–priv	Columbia University	NIH
4468464	1984	Boyer & Cohen	Molecular chimeras	pub–priv	Stanford	NSF
4634665	1987	Richard Axel et al	Genetic transformation	pub–priv	Columbia University	NIH
4838644	1989	Ellen Ochoa et al	Recognizing Method	pub–pub		DOE
5149636	1992	Richard Axel	Genetic transformation	pub–priv	Columbia University	NIH
6285999	2001	Larry Page	Google PageRank	pub–priv	Stanford	NSF
6677082	2001	Michael Thackeray et al	Lithium ion batteries	pub–priv	UChicago Argonne	DoE
6413802	2002	Chenming Hu et al	3D-transistor geometry	pub–priv	UCSD	DoD
6506559	2003	Andrew Fire et al	RNA interference (dsRNA)	pub–priv	Carnegie Institution, UMass Boston	NIH
6794534	2004	Robert H. Grubbs	Olefin-metathesis Ru catalyst	pub–priv	Caltech	NIH
8278036	2009	Katalin Kariko and Drew Weissman	mRNA	pub–priv	U Penn	NIH
7797367	2010	David Gelvin et al	Wireless Integrated Network Sensors	pub–priv	Sensoria Corporation	DoD
8367991	2011	Timothy Bradley	Modulation device for a mobile tracking device	pub–pub		DoD
8183038	2012	James A. Thomson	Induced-pluripotent stem cells	pub–priv	U of Wisconsin	NIH
8185551	2012	Bradley Kuzmaul et al	Disk-resident streaming dictionary	pub–priv	Rutgers University et al	NSF
8399645	2013	Dario Campana et al	CAR-T cell therapy (cancer treatment)	pub–priv	St Jude Childrens Research Hospital	NIH
8697359	2014	Feng Zhang	CRISPR–Cas9 genome editing	pub–priv	MIT, Broad Institute	NIH

Note. The table reports public-interest patents within the list of historically important patents defined by Kelly et al. (2021), extended to 2015 using ChatGPT Deep Research.

B.2 Case Studies

In this Appendix, we present eight case studies – six success stories and two notable failures – of patents that exemplify the main categories of interest in this paper: patents funded by the government but developed and owned by the private sector; patents funded and owned by the public sector; patents funded by the government and developed by private startups.

I) Public-private patents funded by the National Institutes of Health (NIH). There are many notable examples of successful patents funded by the government but owned by the non-federal entities that have laid the foundation for profound technological acceleration across various fields. Among others, breakthroughs in medical research conducted together with the private sector improved disease treatment planning and enabled earlier and more accurate diagnosis. One of them is the groundbreaking research on RNA interference (RNAi) using double-stranded RNA, published by Andrew Fire and Craig Mello in 1998. A patent on that discovery was filed by the two researchers under US Patent No. 6,506,559, granted in 2003. This discovery, funded in part by the NIH, revolutionized molecular biology by providing a powerful tool to selectively inhibit genes, enabling advances in drug development and therapeutic interventions. RNAi technology has had vast implications in biotechnology and medicine, leading to new treatments for diseases such as viral infections and cancers and spurring a thriving industry centered on gene-silencing technologies. For such achievement, Andrew Fire and Craig Mello received the Nobel Prize in Physiology or Medicine in 2006.

II) Public-private patents funded by the National Science Foundation (NSF). Other examples pertain to the role of the National Science Foundation in developing foundational research across manufacturing, networking, and computer technologies. In the field of computer science, the development of the PageRank algorithm at Stanford University under NSF's Digital Library Initiative formed the backbone of Google's search engine. This foundational algorithm that considers the quality of links and captures the 'collective intelligence' of the Web revolutionized Internet search and catalyzed the growth of the digital economy. The assignee of the patent was Stanford University, as Larry Page was a Ph.D.

student there when, together with Sergey Brin, he developed the algorithm. The patent was filed in January 1998 (U.S. Patent No. 6,285,999) and preceded Google's foundation by a few months.

III) Public-private patents involving a startup. Publicly funded research has been particularly effective when collaborating with private startups, which have greatly benefited from such partnerships. A prime example is U.S. Patent No. 7,148,040, filed in 2002 and titled "Method of Rapid Production of Hybridomas Expressing Monoclonal Antibodies on the Cell Surface," which was developed shortly after the founding of a pioneering biotech startup, Abeome Corporation. This patent, which includes a government interest statement acknowledging direct public-sector involvement in its development, describes innovative methods for rapidly generating hybridoma cells that display monoclonal antibodies on their surfaces, greatly accelerating antibody discovery and production. The underlying research was firmly rooted in publicly funded work in cellular biology and immunology, supported by the National Institutes of Health (NIH). The startup utilized this foundational knowledge to create practical, scalable technologies that streamlined monoclonal antibody development for therapeutic and diagnostic applications. Patents like this have advanced biotechnology by enabling faster drug discovery and personalized medicine, reinforcing the vital role of translating government-sponsored research into impactful commercial innovations.

IV) Highly-disruptive public patents funded by the Department of Defense. Some patents with public interest and public ownership have also been extraordinarily disruptive. One of the most consequential examples is U.S. Patent No. 3,789,409, titled "Navigation System Using Satellites and Passive Ranging Techniques," issued in 1974 and attributed to Roger L. Easton, a key figure in the development of the Global Positioning System (GPS). This patent laid the groundwork for satellite-based positioning by describing a system that used precise timing signals from orbiting satellites to determine the location of a receiver on Earth. Crucially, the invention introduced concepts such as passive ranging, time synchronization via atomic clocks, and orbital tracking—all of which became central to modern GPS systems. Developed at the Naval Research Laboratory (NRL) and fully funded by the U.S. Department of Defense, the invention was assigned to the U.S. government, exemplifying a case of direct public ownership of high-impact intellectual property. Initially intended for military navigation and targeting, the technology was eventually declassified and released for civilian use, catalyzing a transformation in global navigation, logistics, agriculture, and everyday consumer behavior. The patent's core principles have since been embedded in virtually every GPS-enabled device—from smartphones and car navigation systems to precision-guided farming equipment and autonomous vehicles—demonstrating the enormous long-term value of publicly funded innovation.

V) Highly disruptive public patents funded by the Department of Energy and the Department of Defense. Other examples of important patents developed through federal agency initiatives come from the energy sector. One of the most consequential is U.S. Patent No. 2,708,656, titled "Neutronic Reactor," filed on December 19, 1944, by Enrico Fermi and Leo Szilard, and granted on May 17, 1955. This patent describes a nuclear reactor design using graphite as a moderator and natural uranium arranged in a geometric lattice to sustain a controlled nuclear chain reaction. The innovation solved critical challenges in reactor stability and safety, laying the foundation for nearly all modern nuclear reactor designs. The research was conducted under the Manhattan Project, a top-secret World War II initiative managed by the Department of War (a precursor to today's Department of Defense), and later transferred to the Department of Energy, which now oversees many of the resulting technologies and intellectual property. The University of Chicago, through its Metallurgical Laboratory (Met Lab), was a central site of this work and where Fermi's team achieved the world's first controlled nuclear chain reaction in 1942. The patent's long-term impact has been profound: it enabled the development of both civilian nuclear power plants—now a major source of carbon-free energy—and numerous advances in nuclear safety, reactor design, and scientific instrumentation. The reactor principles outlined in this patent continue to influence nuclear innovation across energy, medicine, and defense.

VI) The role of the public in developing GPS technologies. Some patents with public interest and public ownership have also been extraordinarily disruptive. A prime example is U.S. Patent No. 3,126,545, titled "Satellite Hyperbolic Navigation System," issued on March 24, 1964, to I. D. Smith Jr. and assigned to the United States government. This patent encapsulated the core principles of the Transit navigation system, the world's first operational satellite-based navigation network. Developed in the late 1950s and early 1960s by Johns Hopkins University's Applied Physics Laboratory (APL) in close collaboration with and under the sponsorship of the U.S. Navy, the project was fully funded by the Department of Defense (DoD). Transit eliminated the need for fixed ground-based infrastructure and delivered global positioning capabilities long before the advent of GPS. Operational from 1964 through the 1990s, the system introduced precursor technologies to the architecture of the Global Positioning System (GPS).

NASA subsequently played a critical role in extending these capabilities into space-based applications. One such contribution is captured in U.S. Patent No. 7,548,199, titled "Radiation Hardened Fast Acquisition Weak Signal Tracking System and Method," issued in 2009 and assigned to NASA. This invention enabled GPS receivers to rapidly acquire and track weak positioning signals in the challenging environment of Low Earth Orbit, allowing for precise autonomous navigation of satellites and spacecraft. By hardening GPS technology for use in space, NASA extended the utility of satellite navigation beyond Earth, reinforcing the broad and enduring impact of federally funded innovation.

VII) Public failure (a). A notable example of government-funded research that is regarded as a failure is U.S. Patent No. 2,992,981, titled "Neutronic reactor core", granted on July 08, 1961, and assigned to the U.S. Atomic Energy Commission (AEC), a predecessor to the Department of Energy. This patent outlined a design for a nuclear reactor that could be used to power a jet aircraft, promising virtually unlimited flight endurance. Funded by the AEC and the Department of Defense through the Aircraft Nuclear Propulsion (ANP) program in the 1950s, the initiative included the NB-36H "Crusader" testbed and plans for the Convair X-6 nuclear bomber. Despite over $1 billion of funding and years of high-profile research, reactor flights never moved beyond shielding tests, and no aircraft ever flew using nuclear power. The entire program was canceled in 1961, with the X-6 never built. Though the scientific work informed later nuclear and shielding technologies, the core promise—an operational nuclear-powered airplane—was never realized.

VIII) Public failure (b). Another example is that of U.S. Patent No. 7,394,016, titled "Bifacial Elongated Solar Cell Devices with Internal Reflectors," granted on July 15, 2008, to Benyamin Buller et al. and assigned to Solyndra LLC. The patent outlines the company's signature cylindrical CIGS (copper indium gallium selenide) solar cell design, arranged in tubular casings with internal reflectors, intended to enhance efficiency. The technology was backed by a large federal commitment—a $535 million loan guarantee from the U.S. Department of Energy (DOE) under the 2009 stimulus program, making Solyndra a high-profile centerpiece of clean-energy investment. Yet, despite its patented innovation, Solyndra could not compete with rapidly declining costs of silicon-based panels from overseas, and by August 2011, the company filed for bankruptcy, leaving taxpayers with a substantial loss.

These cases stand as cautionary examples: even with robust federal funding, dramatic technological breakthroughs are not guaranteed.

C Data Sources and Definitions

This Appendix provides further details on the data, its sourcing and the transformations of the key variables used in the paper's empirical analysis.

C.1 Descriptive Statistics

Figure C.1: The Evolution of Patenting Activities by Funding Entities and Ownership

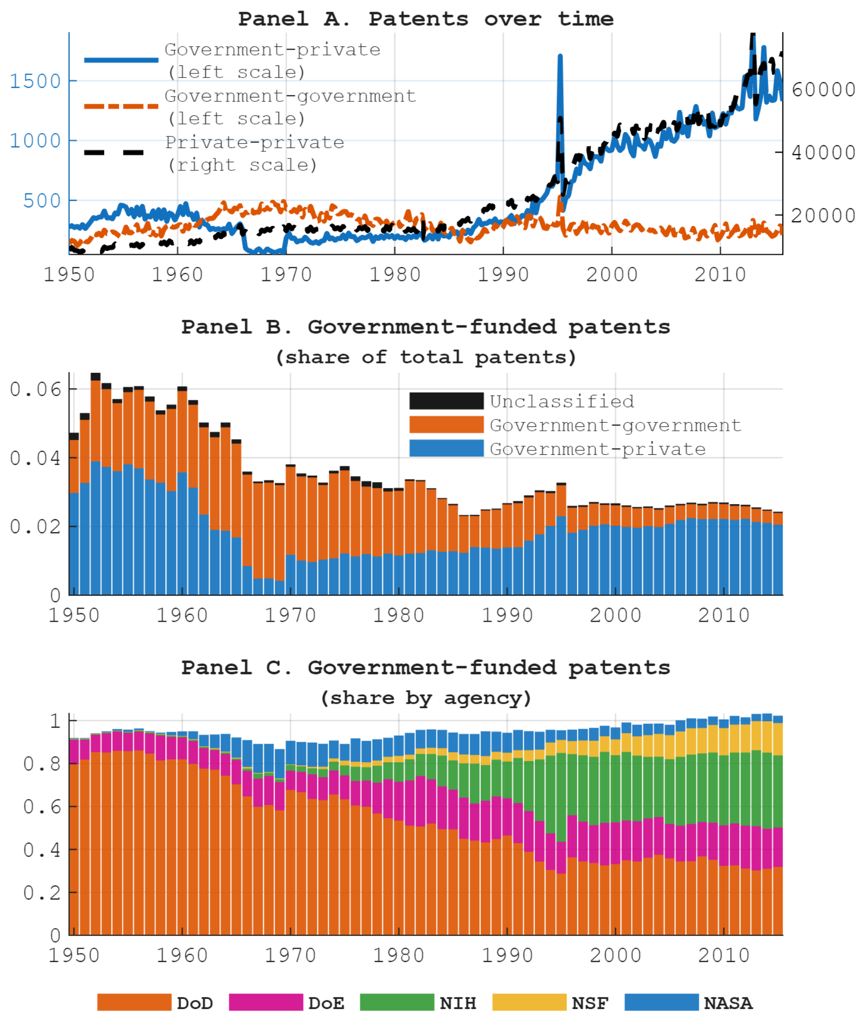

Note. Panel A reports the quarterly total number of patents by government interest over time (GPR database - Gross and Sampat (2025). Panel B displays public interest patents over time as share of the total number of patents. Panel C represents the breakdown of public interest patents over time by federal agencies. The shares may sum to more than 1 because several agencies may fund the same patent.

Table C.1: Descriptive Summary of Patent Types and Importance

Panel A: Publication Type Breakdown

	Priv–priv	Pub–pub	Pub–priv
Number of Patents	7,000,953	76,601	139,191
Share of Total	97.0%	1.1%	1.9%

Panel B: Importance by Public Interest Type

Statistic	Priv–priv	Pub–pub	Pub–priv
Median	0.17	0.15	0.18
Skewness	0.91	1.59	1.50
Kurtosis	7.49	11.78	11.56

Panel C: Reliance on Science

Statistic	Priv–priv	Pub–pub	Pub–priv
Mean	2.67	3.85	23.15
Median	0.00	0.00	3.00
75th percentile	1.00	1.00	25.00
90th percentile	4.00	7.00	70.00
95th percentile	10.00	19.00	107.00

Panel D: Agency Breakdown

Agency	Patents	Pub–pub share	Pub–priv share
Department of Defense	106,574	43.0%	55.1%
National Institutes of Health	47,830	10.4%	84.0%
Department of Energy	35,439	37.6%	60.5%
National Science Foundation	14,670	≈0%	93.8%
NASA	12,869	56.1%	36.8%

Panel E: Government Interest by Assignee Type

Assignee type	Patents	Pub-share	Share of govt-funded
Research institutes	153,956	8.7%	8.4%
Universities	90,785	52.9%	30.1%
Start-ups	131,572	1.5%	1.2%
VC-backed assignees	410,959	1.6%	4.1%
Federal agencies	42,147	100.0%	26.4%

Note: Panel A is the patent distribution of the government interest classification in the GPR database of Gross and Sampat (2025). Panel B shows the importance distribution in Kelly et al. (2021) by patent category. Panel C shows the reliance on science distribution in Marx and Fuegi (2020, 2022) by patent category. Panel D reports the distribution of public interest patents across main federal agencies; the last two columns do not sum to 100 due to a small group of patents that Gross and Sampat (2025) leave 'unclassified'. Panel E reports for the sample 1975-2015 the total number of patents produced by entity types, together with the shares of government-funded patents within each assignee type (Pub-share) and their (%) contribution to the number of total government-funded patents (Share of govt-funded; they do not sum to 100 because not all possible actors are included). Government-funded patents in Panel E include both pub-pub and pub-priv. Data on startups and venture capital is available since 1975 from Ewens and Marx (2024).

Additional details on patent dynamics and composition. The patents assigned to *public-private* or solely private display an upward time trend, which partly reflects legislative changes. Enacted to encourage the commercialization of inventions arising from federally funded research, the Bayh–Dole Act of 1980 introduced a uniform patent licensing policy across federal funding agencies under which small businesses and non-profit institutions could retain ownership of patents that benefited from public funding while the government retained a license for use. This primarily affected innovations developed by research institutes and universities, some of which may have been classified as public-public rather than public-private without the Act. Moving to shares, we observe that the incidence of solely public patents has been progressively falling over time (Figure C.1, Panel B), while the proportion of *public-private* patents has stayed roughly constant with only mild variations over the years. In 1995, the number of patent applications surged due to the Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS).

Importance and reliance on science distributions. In Panel B of Table C.1, we report that *public-private* patents tend to be more disruptive than *private-private* ones, followed by *public-public* patents. However, their distribution reveals that the importance of public interest patents is heavily skewed to the right and is leptokurtic. In Panel C, we report summary statistics for the number of scientific publications cited in each patent group. Publicly funded innovations emerge as the most exposed to academic publications, especially at the top end of the distribution. These patterns are consistent with the notion that the government typically funds research that is more basic or fundamental relative to the research financed and developed solely by the private sector.

Agency heterogeneity. The innovations of most federal agencies tend to be associated with a majority of private ownerships with government funds. NASA, DoD and—to a lesser extent—DoE display a much higher share of *public-public* patents, while the NIH and the NSF mainly fund innovations that are eventually developed outside the federal government. As revealed by Panel C of Figure C.1, patents funded by NASA increased between the 1960s and 1980s, as a result of the Moonshot race with the USSR; since the 1990s, however, they have started to play an increasingly minor role.

Assignee type breakdown. Finally, in Panel E of Table C.1, we describe the data employed in Section 8. We combine information on government funding, both for private and public ownerships, with the patents' assignee type, including research institutes, universities, startups, and venture capital-backed companies. Within each assignee type (under the heading 'Pub-share'), more than half of university patents are funded by the government; this share drops to about 9% for research institutes and sits around 1.5% for startups and venture capital companies. On the other hand, universities and federal agencies are the main contributors to the total number of government-funded patents (in the 'Share of govt-funded' column), which is far higher than the shares for research institutes, VC-backed firms, and startups.

C.2 Government Patent Registry Database

Gross and Sampat (2025) constructs the Government Patent Register database by combining historical administrative records with several modern data sources to create a comprehensive measure of U.S. government-funded patents. The core of the database is the digitization of the historical U.S. Patent and Trademark Office (USPTO) Register of Government Interest in Patents ("historical GPR"), which the authors digitized and cross-validated against Google Patents. To supplement the historical GPR and extend the data through 2020, the authors integrated several modern sources:

USPTO Patent Assignment Dataset (UPAD). They identified government-interest patents by searching for conveyance text referencing "Executive Order 9424" or "Confirmatory License" and by systematically identifying all federal agencies listed as assignees in the dataset. The authors then manually classified these transactions as conveying either *title* or *license* to the government.

PatentsView Government Interest Statements. The authors used the government interest statement data from PatentsView. Finding occasional imprecision in the existing agency identification, they developed a new approach using a large language model (GPT-4) to extract the specific funding agencies from the text of the interest statements.

Government Assignee Data. They identified government-assigned patents using assignee data from Fleming et al. (2019) for the pre-1976 period and from PatentsView for 1976 onwards.

Gap-filling using Fleming et al. (2019). For the pre-1976 period, the authors identified a set of patents that appeared in the Fleming et al. (2019) data as having an interest statement but were not in the historical GPR. Using a semi-automated process involving GPT-4 and manual review, they isolated the true positives from this sample and extracted the funding agency, adding 818 additional patents to their data.

The final dataset is the union of all patents identified through these sources. The unique feature of the GPR is the classification of government-funded patents into *title* and *license*. In the former, the federal government is the owner of the patent, whereas in the latter, the government retains licensed use of the patent because it funded its development. The method creates a more complete accounting of government-funded patents than previously available, as many patents, particularly the *license* ones from the mid-twentieth century, can only be identified via the historical GPR.

Imputing Unclassified Patents in GPR

The GPR contains a residual category marked as "unknown", primarily arising from multiple records with discordant government-interest information. The size of this group increased after 1980, due to the introduction of digital records and the adoption of the Bayh–Dole Act. Although our main results are unaffected if we exclude these ambiguous cases from the analysis, for completeness, we categorize unknown patents into public-private and public-public. To do so, we apply a straightforward strategy: any indication of a license right for the government or any involvement of private parties in the innovation's development signals joint *public-private* efforts; the remaining cases are classified as *public-public*.

Specifically, to classify patents with unclassified government interest (code 3) in the GPR database, we distinguish between public-public patents (code 1), which are patents developed entirely within government agencies with government retention of title, and public-private patents (code 2), which are patents developed with private sector involvement where contractors/grantees retain title with government license.

We first classify patents based on the `assignee_type` field from PatentsView. Patents with `assignee_type = 6` (U.S. federal government agency) are classified as public-public, while those with `assignee_type = 2` (U.S. company organization) or `assignee_type = 3` (foreign company organization) are classified as public-private.

For remaining unclassified patents, we apply regular expression searches to the government interest statement text (`gi_statement`), converted to lowercase. Patents are classified as public-public if the government interest statement contains any of the following phrases: `"assigned to the united states"`, `"assigned to the uṡȧ"`, `"title to this invention (is|has been) vested"`, or `"owned by the (u?s?|government)"` where `?` allows for optional periods (in Stata). Patents are instead classified as public-private if the statement contains `"nonexclusive"` (indicating nonexclusive license to government) or `"license"` (indicating contractor/grantee retention of title).

Finally, when a patent has multiple government interest records with conflicting classifications, we identify duplicate patent IDs with different government interest codes and retain the public-private record when both public-private and public-public records exist. This prioritization reflects the principle that any private sector involvement classifies the patent as public-private.

C.3 Additional Patent Data Sources

Importance. We employ the measure of patents' importance from Kelly et al. (2021), which is the ratio of forward to backward textual similarity, computed using natural language processing techniques. We use the variable `lqsim05` (importance based on a 5-year window) in our baseline analysis and `lqsim010` (importance based on a 10-year window) in our robustness analysis. In the baseline, we define dummy groups based on percentiles computed within each category. In Appendix F, we report that similar results hold when computing percentiles over the whole distribution of patents.

Reliance on Science. We employ the measure of reliance on science from Marx and Fuegi (2020, 2022), which counts scientific citations in patents both in the title and in the body of the patent. We rely on the measure that aggregates papers' citations across title and text.

In Appendix F, we report that similar results hold when computing different percentiles or considering papers' citations as a dichotomous variable.

Research Entities. The construction of the research-entity series given by the classification of universities and research institutes and its extension to the full 1950–2015 sample using historical USPTO assignee records, is described in Appendix I.

Startups and VC. The startup and venture capital-backed flags for patents' assignees are based on data constructed by Ewens and Marx (2024).[2] The patent-level dataset provides assignees' founding year. Following Ewens and Marx (2024), we assign a patent-level startup flag in cases when the patent filing date is 3 or fewer years from the foundation. The VC-backed flag is directly provided by Ewens and Marx (2024) as a dummy equal to 1 if the assignee has received venture capital at any point in its life, and 0 otherwise.

C.4 Macroeconomic variables

Our primary source for macroeconomic variables is Antolin-Diaz and Surico (2025). We measure economic activity using log Real GDP per capita, and its main private components also in per-capita logs: Private Consumption, sourced from BEA NIPA data (1947-2015), and Private Investment, which is based on unpublished BEA estimates from 1901 onwards and interpolated to a quarterly frequency. Research and Development (Total R&D) comes from FRED `Y694RX1Q020SBEA` and is expressed in (log) per capita terms by dividing it by the FRED series `CNP16OV` (civilian noninstitutional population). We similarly transform the FRED series `Y006RC1Q027SBEA` for Private R&D. The GDP deflator is the FRED series `GDPDEF`. For productivity, we use Total Factor Productivity (TFP), calculated as the Solow residual from a Cobb-Douglas production function (with a capital share $\alpha = 0.28$) and subsequently adjusted for capacity utilization. The aggregate price level is measured by the logarithm of the GDP Deflator, and we include Wages using the FRED series `COMPRNFB`, also in logs. Finally, the model includes two financial variables: the short-term interest rate series from Ramey and Zubairy (2018) and stock prices, sourced from R.J. Shiller website (series

[2] The dataset used in Ewens and Marx (2024) is shared through foundingpatents.com

`Real price`) and expressed in logs.

D Additional Results

This Appendix collects additional evidence and robustness checks supporting the baseline results. We first document that patent filings are only weakly cyclical and that aggregate patent shocks display the standard properties of technology shocks. We then assess the robustness of the LP-IV estimates to alternative sets of agency-level public R&D appropriation instruments. Next, we test directly whether the responses to public–private and private–private patent shocks differ, evaluate the sensitivity of inference to the long-run variance estimator proposed by Lazarus et al. (2018), and report the definition of outlier episodes used in the spike-robustness analysis. We also compare the baseline local-projection estimates with several alternative time-series approaches, including recursive and maximum-share VARs, factor-augmented local projections, and VARs identified through non-Gaussianity, moment restrictions, pseudo-maximum likelihood, and time-varying volatility. Finally, we address the role of the Bayh–Dole Act by pooling all government-funded patents into a single category and showing that the central results are unchanged.

D.1 Patent Filing Cyclicality and Effects of Total Patents

In this section, we show that, unlike R&D expenditure, patent filings tend to not exhibit much business-cycle cyclicality, neither economically nor statistically. This holds for all patents as well as government-interest patents. In Figure D.1, we report the contemporaneous correlation between these variables and real GDP. The only possible exception are Government-private patents in the third panel. However, while patenting activity in this group is somewhat negatively correlated with GDP, driven by the 1970 recession, the coefficient is small and the significance marginal. Figure D.3 plots the effects on key macroeconomic aggregates of estimating the empirical model described in Section 3 using total patents as the right-hand side variable. In line with previous literature (Miranda-Agrippino et al., 2025), an innovation shock measured from total patents has the properties of a standard technology shock, increasing GDP and TFP while lowering prices.

Figure D.1: Patent filing cyclicality

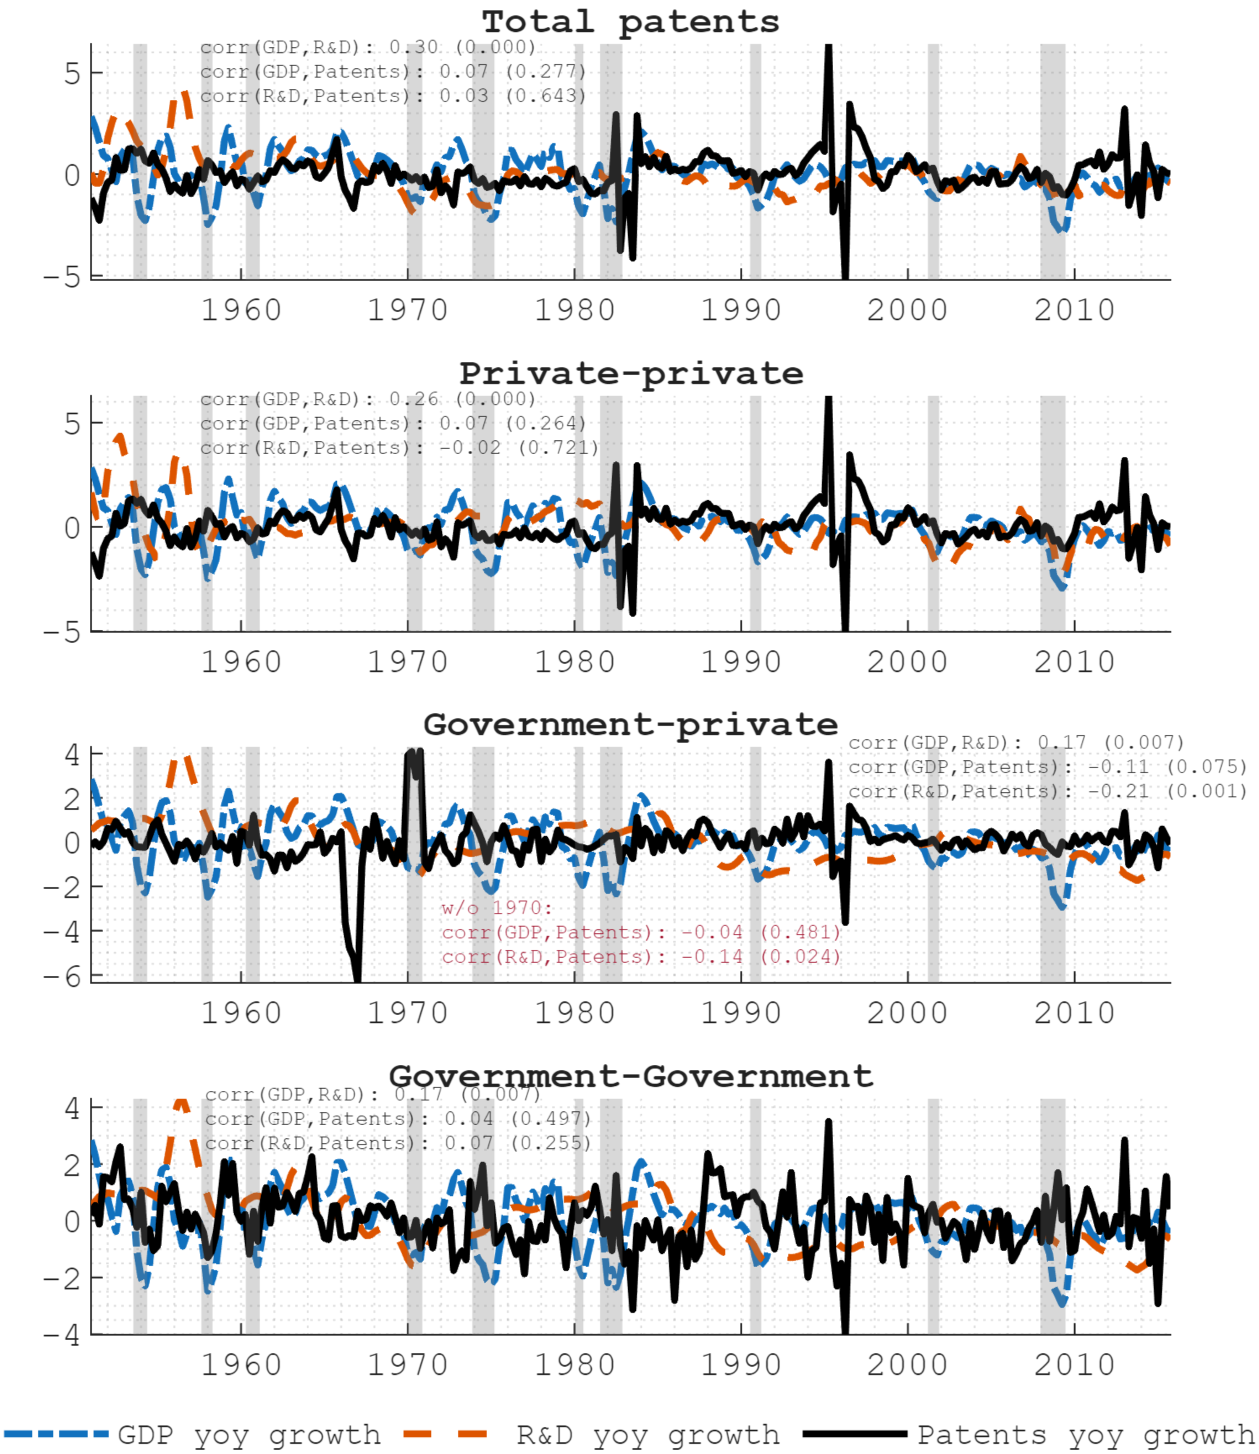

Note. The figure displays the growth rates of real per-capita GDP, R&D expenditure, and the number of patents filed in each category (total, public-private, private-private, and public-public), along with their contemporaneous correlations (p-values in parentheses). The figure shows that, unlike R&D expenditure, patent filings do not exhibit strong cyclical behavior.

D.2 LP-IV Estimates with Fewer Agency Instruments

The baseline LP-IV estimates for public–private and public–public patents use agency-specific federal R&D appropriation shocks from Fieldhouse and Mertens (2023). The instrument set includes distributed lags of R&D appropriations for DoD, DoE defense, DoE nondefense, NASA, NIH, and NSF, compressed using a third-degree Almon polynomial basis. This specification preserves cross-agency heterogeneity and allows for long gestation lags between public R&D funding and patent filings. At the same time, one may worry that the results depend on the dimensionality of the instrument set.

To assess this concern, Figure D.2 re-estimates the LP-IV specification using progressively smaller sets of agency instruments. The first specification corresponds to the baseline six-agency set. The second excludes DoD, thereby removing the largest defense agency from the first stage. The third further reduces the instrument set by excluding DoE appropriations. In all cases, the distributed lag window, Almon compression, control set, shock normalization, and weak-instrument-robust inference are kept unchanged.

The results show that the main LP-IV evidence is not driven by a single agency or by the dimensionality of the instrument set. Public–private patent shocks continue to generate positive and persistent responses of GDP and TFP across all instrument sets. The point estimates are close to the baseline and the weak-IV robust confidence bands overlap substantially. Public–public patent shocks remain weaker and more imprecisely estimated, consistent with the average results reported in the main text. These findings support the interpretation that the aggregate LP-IV estimates reflect broad variation in public R&D-driven patenting rather than the influence of a specific agency-level instrument.

Figure D.2: LP-IV reducing the number of instruments/agencies

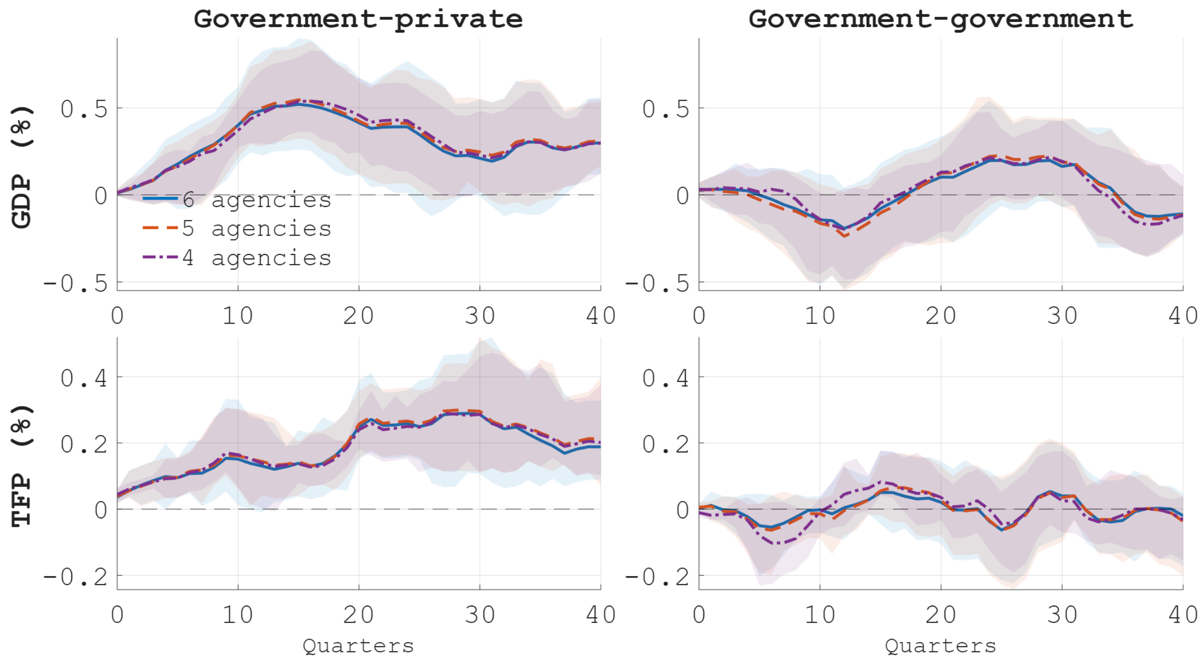

Note. The figure compares LP-IV estimates obtained using alternative sets of agency-level public R&D appropriation instruments. Columns refer to government-private and government-government patents; rows report responses of log real per-capita private GDP and log TFP. Solid lines denote LP-IV point estimates, normalized to a one-percent increase in the corresponding public patent category. Shaded areas report 90% weak-instrument-robust confidence bands. The 6-agency specification uses DOD, DOE Defense, DOE Nondefense, NASA, NIH, and NSF; the 5-agency specification excludes DOD; and the 4-agency specification further excludes DOE Nondefense. In all cases, instruments are based on Almon-compressed distributed lags of real agency-level R&D appropriations. The LPs include four lags and the same macroeconomic and patent controls as in the baseline specification. Sample: quarterly data, 1950:Q1–2015:Q4.

D.3 Macroeconomic Impact of Total Patenting Activity

In Figure D.3, we show that the dynamic effects of our shocks to patenting activity generate response of main macroeconomic indicators that align with a technology shock interpretation, especially the negative conditional correlation between prices and quantities in the first row.

D.4 Are the effects of government-private and private-private patents different?

To explicitly assess the statistical difference between the effects of patenting in government-private and private-private groups, we rely on wild bootstrap with the same multiplier across the three local projection equations. Both the Sup-Wald test and the Cramér-von Mises test

Figure D.3: Macroeconomic effects of a patent shock - all patents

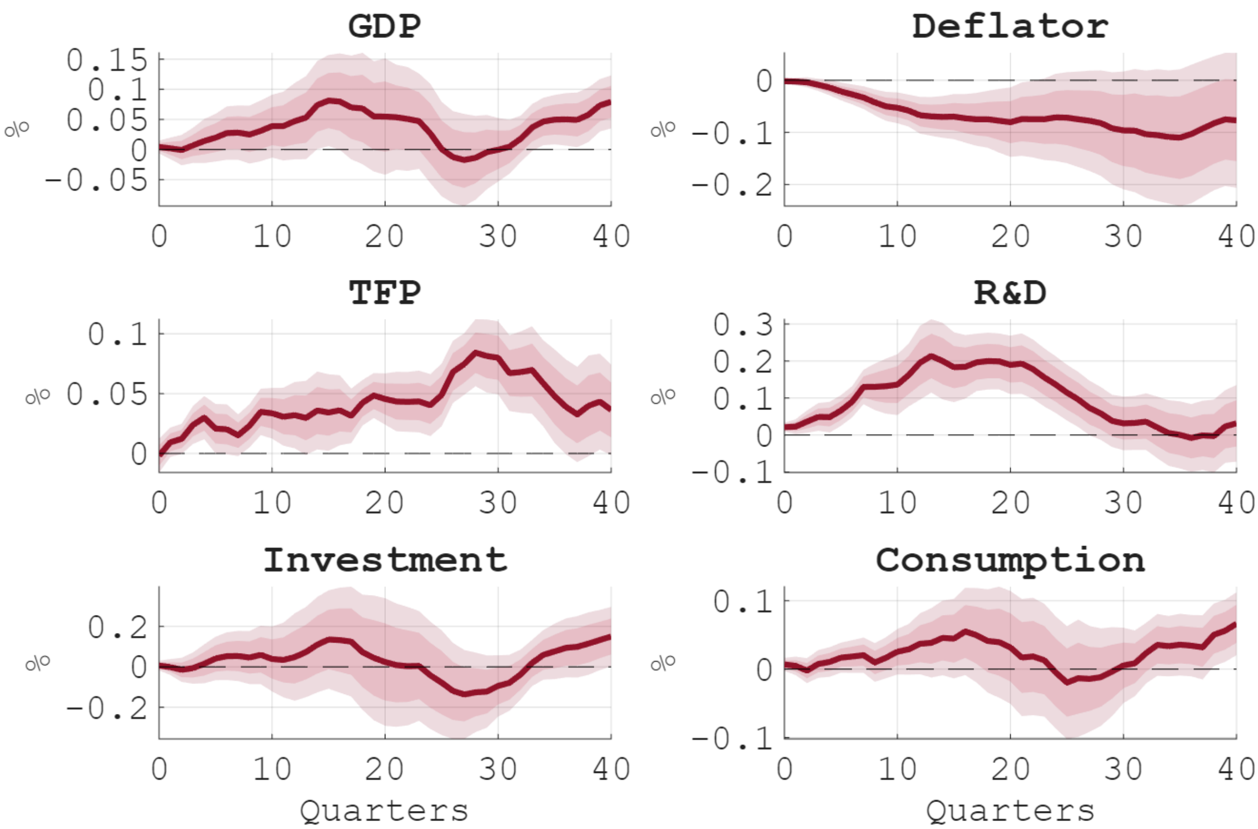

Note. The figure displays the dynamic effects of an aggregate innovation shock on (log) real per-capita GDP, (log) utilization-adjusted TFP, (log) real per-capita private R&D expenditure, (log) real per-capita private investment, (log) real wages, and (log) real per-capita consumption. The shock is an unanticipated increase in the total number of patents, normalized to increase total patents by 1% on impact. The estimation by local projections follows eq.(5). The set of controls includes 4 lags of total patents, real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill, and real per-capita R&D expenditure. All variables except the T-bill are in logs. The solid line represents the point estimate, while the shaded areas report 68% and 90% confidence intervals computed from heteroskedasticity-robust standard errors. Sample: quarterly data 1950:Q1-2015:Q4

reject the equality of the estimated response of GDP and TFP. Figure D.4 explicitly reports the difference between the estimated response across the two patent groups. The delta is economically sizable and highly statistically significant.

D.5 Lazarus et al. Inference

This section assesses the sensitivity of the baseline local-projection results to the choice of standard errors. We re-estimate the baseline OLS local projections using the same sample, control set, lag structure, transformation of the dependent variables, and shock normalization as in the main specification. The only change concerns inference. Instead of using the standard heteroskedasticity robust inference, we compute standard errors following the rec-

Figure D.4: Difference between Government-private and fully private effects

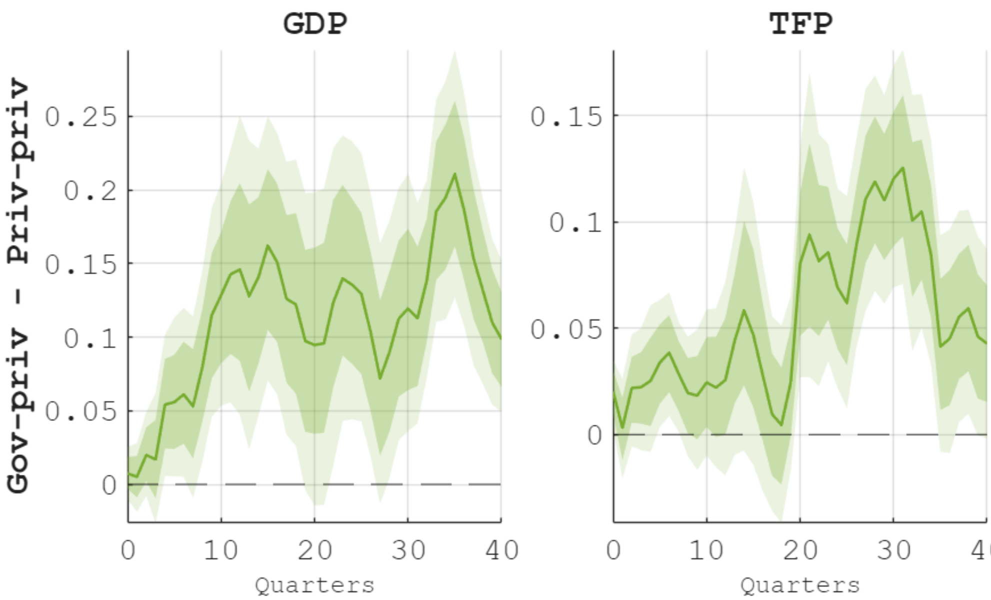

Note. The figure displays the differential effects of government-private and private-private shocks on (log) real per-capita GDP and (log) utilization-adjusted TFP. The shock is an unanticipated increase in the number of patents by category, normalized to increase total patents by 1% on impact. The estimation by local projections follows eq.(5). The set of controls includes 4 lags of the patent group shocked and real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill, real per-capita R&D expenditure, and the number of patents in other groups. All variables except the T-bill are in logs. The solid line represents the point estimate, while the shaded areas report 68% and 90% confidence intervals computed from wild bootstrap standard errors. Sample: quarterly data 1950:Q1-2015:Q4

ommendation of Lazarus et al. (2018) based on the equal-weighted cosine estimator. Specifically, the number of cosine terms is set to $B = \lceil 0.4T^{2/3} \rceil$, and confidence intervals are constructed using the corresponding fixed-b t_B critical values. The exercise therefore isolates whether the statistical significance of the estimated dynamic responses is sensitive to using a more conservative long-run variance estimator designed for highly persistent and serially correlated projection residuals. For comparison, the figure overlays the impulse responses and 90 percent confidence intervals obtained under the standard robust inference used in the baseline estimates. Results are very similar to the baseline.

Figure D.5: IRFs under Lazarus et al. (2018) inference

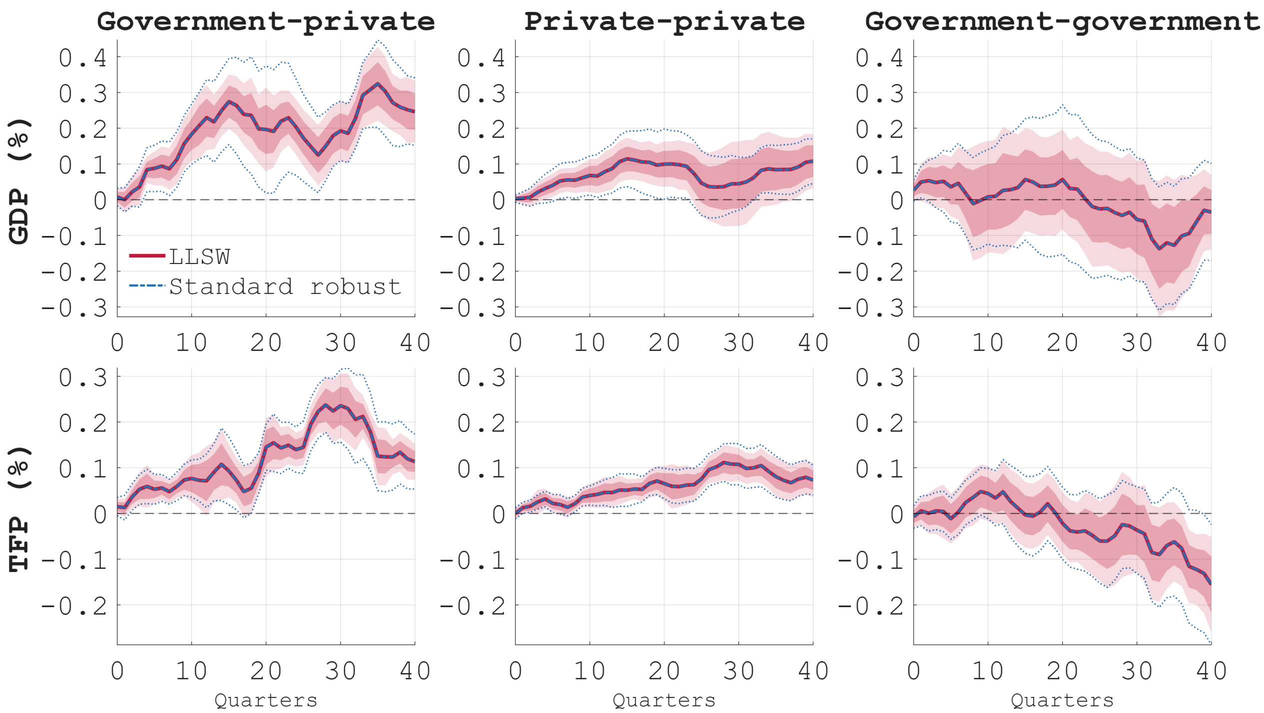

Note. The figure reports impulse responses of real GDP and TFP to a patenting shock estimated with the baseline OLS local-projection specification. The shock is normalized to increase total patenting by 1 percent on impact. Red solid lines denote point estimates obtained using the Lazarus et al. (2018) inference procedure; red shaded areas report the associated 68 and 90 percent confidence intervals based on the equal-weighted cosine estimator with $B = \lceil 0.4T^{2/3} \rceil$ cosine terms and fixed-b critical values. Blue dash-dotted lines report the corresponding point estimates from the standard robust-inference specification, while blue dotted lines denote the associated 90 percent confidence intervals. The set of controls includes four lags of the shocked patent category, the dependent variable, real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill rate, real per-capita R&D expenditure, and patenting activity in the other categories. All variables except interest rates are in logs. Sample: quarterly data, 1950:Q1–2015:Q4.

D.6 Definition of Outliers in Patent Filings

We assess whether the baseline results are driven by unusually large quarterly movements in patent filings. For each funding–ownership category, we compute quarterly log changes in the seasonally adjusted patent series and classify as outliers those quarters in which the absolute standardized change exceeds two standard deviations. The baseline estimates keep all observations; the outlier classification is used only for robustness.

D.7 Alternative Time-Series Approaches

This section assesses the robustness of the baseline results to alternative time-series specifications and identification schemes. The exercises are designed to address whether the estimated

Table D.1: Outlier episodes in patent filings

Patent category	Threshold	Outlier episodes
Government-private	17.54	1959:Q4–1960:Q1; 1962:Q1; 1966:Q2–1966:Q3; 1967:Q1–1967:Q4; 1969:Q3–1970:Q1; 1972:Q2; 1973:Q4; 1977:Q1; 1978:Q1; 1982:Q4–1983:Q1; 1983:Q4; 1992:Q4–1993:Q1; 1995:Q2–1995:Q4; 1996:Q4; 2013:Q1–2013:Q2; 2014:Q1–2014:Q2; 2015:Q2.
Private-private	7.23	1966:Q1; 1972:Q1; 1982:Q3–1983:Q1; 1995:Q2–1995:Q4; 1996:Q3; 2013:Q1–2013:Q3.
Government-government	11.58	1950:Q4; 1952:Q4–1953:Q1; 1954:Q2–1954:Q3; 1958:Q1; 1958:Q4; 1959:Q3; 1960:Q4–1961:Q2; 1963:Q3; 1964:Q3; 1966:Q1; 1972:Q4; 1973:Q4; 1974:Q2; 1974:Q4; 1976:Q2; 1977:Q1; 1978:Q3; 1982:Q1–1982:Q4; 1986:Q1–1986:Q2; 1986:Q4–1987:Q1; 1988:Q1; 1992:Q2; 1992:Q4–1993:Q2; 1994:Q4; 1995:Q2–1995:Q3; 1997:Q3; 1998:Q1; 1998:Q4–1999:Q1; 2000:Q1; 2003:Q4–2004:Q1; 2005:Q1–2005:Q2; 2008:Q1–2008:Q4; 2009:Q4; 2010:Q2; 2012:Q1; 2013:Q1–2013:Q3; 2014:Q1–2015:Q2.

Note. The table reports outlier episodes in the seasonally adjusted quarterly patent series. Outliers are defined as quarters in which the absolute quarterly growth rate of patent filings exceeds one within-sample standard deviation. The threshold column reports the category-specific cutoff, computed as the standard deviation of quarterly log growth multiplied by 100. Consecutive outlier quarters are grouped into a single episode. The baseline estimates retain all observations; the outlier classification is used only for robustness exercises that either include indicators for the outlier episodes or set the patent shock to zero during one episode at a time.

dynamic effects depend on the local-projection framework, the recursive timing assumptions implicit in the baseline specification, or the information set used to isolate innovation shocks. Figure D.6 reports impulse responses for the baseline specification and seven alternatives: a VAR max-share identification, a standard recursive VAR, a factor-augmented local projection, and four VAR-based approaches that exploit non-Gaussianity, alternative moment restrictions, pseudo-maximum likelihood, and time-varying volatility. In all cases, shocks are normalized to increase total patents by 1% on impact.

The first specification is the baseline OLS local projection described in equation (5). It identifies innovation shocks as residual variation in each patent category after conditioning on four lags of the shocked patent group, the dependent variable, real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill rate, real per-capita R&D expenditure, and patenting activity in the other groups.

The second specification uses a VAR-based maximum-share identification, in the spirit of Uhlig (2003). This approach selects the structural innovation shock as the disturbance that explains the largest share of the forecast error variance of the relevant patent series over the medium-run horizon. Relative to the baseline, this procedure is more agnostic in terms of short-run assumptions as it replaces a recursive impact restriction with a medium-run restriction on forecast-error variance shares.

The third specification is a standard recursive VAR. The VAR includes the same core variables as the baseline specification, estimated with a longer lag structure, and identifies patent shocks through a Cholesky ordering that places the patent series before the macroeconomic aggregates.

The fourth specification augments the local projection with factors extracted from the FRED-QD macroeconomic dataset (McCracken and Ng, 2020). Specifically, it includes lags of the first two principal components, which summarize a large set of quarterly macroeconomic indicators. This exercise addresses the possibility that the baseline control set omits aggregate information relevant for aggregate outcomes.

The remaining specifications identify structural shocks within VAR models using alternative statistical restrictions. VAR-ICA exploits non-Gaussianity in reduced-form residuals to recover statistically independent structural innovations (Hyvarinen, 1999; Himberg et al., 2004). VAR-GMM uses moment restrictions implied by higher-order departures from Gaussianity Keweloh (2021). VAR-PML implements a related pseudo-maximum-likelihood criterion for non-Gaussian structural shocks (Gouriéroux et al., 2020). VAR-TVV identifies structural innovations from changes in residual volatility over time, following the logic of identification through heteroskedasticity (Rigobon, 2003; Lewis, 2021). These approaches provide a demanding robustness check because they rely on identifying information that is distinct from the recursive timing assumptions of the baseline specification.[3]

Across specifications, the main ranking of patent categories is stable. Public–private patent shocks are associated with the largest and most persistent medium-run responses of GDP and TFP. Private–private patent shocks generate smaller positive responses, while public–public patent shocks display weaker average effects. Although the alternatives differ in short-run dynamics and in the smoothness of the estimated responses, they preserve the central conclusion that government-funded but privately owned innovation is most strongly associated with medium-run aggregate fluctuations.

[3] Due to computation constraints, this latter approach is based on a smaller VAR system compared to the baseline.

Figure D.6: Alternative Time-series Strategies

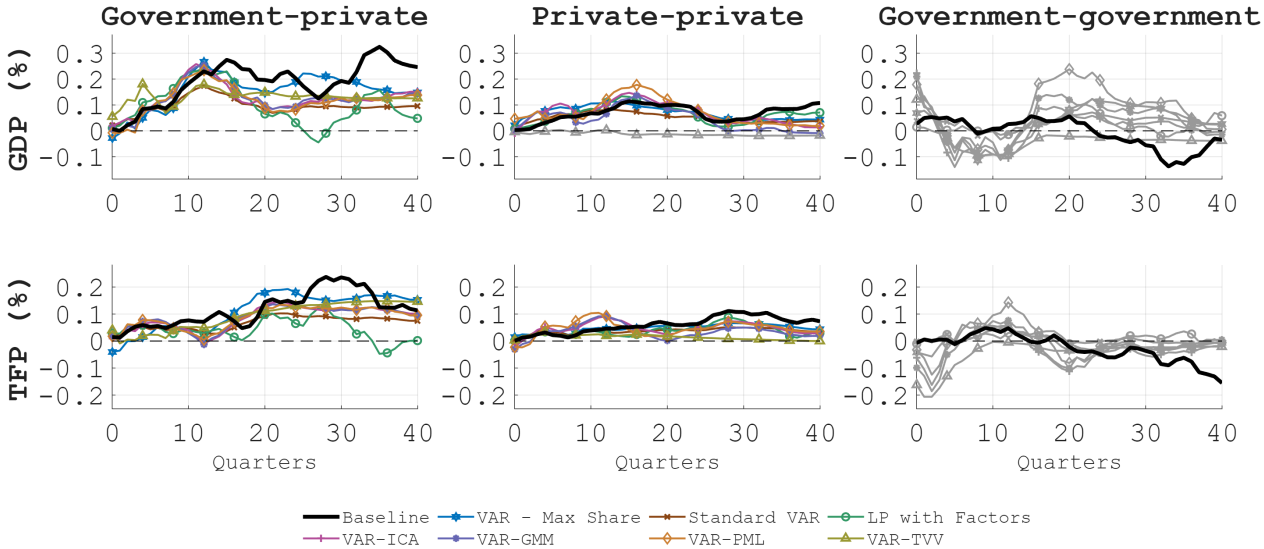

Note. *Panel (a) displays the dynamic effects of innovation shocks identified through a LP-IV approach in each category of patents (government-private, private-private, government-government) on (log) real per-capita private GDP and (log) TFP. The solid line represents the local-projection point estimate after applying the sequential bias correction of Herbst and Johannsen (2024). The shaded areas report 68% and 90% confidence intervals based on heteroskedasticity-robust standard errors, centered on the bias-corrected estimates. Sample: quarterly data, 1950:Q1–2015:Q4. Panel (b) reports point estimates from the baseline OLS local-projection specification and seven alternative time-series identification strategies. These alternatives include a VAR-based maximum-share identification, a standard recursive VAR, a factor-augmented local projection specification, and four VAR-based approaches relying on alternative identifying restrictions or objective functions: ICA, GMM, PML, and time-varying volatility. The shock is normalized to increase total patents by 1% on impact. The baseline control set includes four lags of the shocked patent group, real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill, real per-capita R&D expenditure, and the number of patents in other groups. The factor-augmented local projection further includes four lags of the first two FRED-QD factors. All variables except the T-bill and the factors are expressed in logs. Sample: quarterly data, 1950:Q1–2015:Q4. Grey lines indicate impulse responses that are statistically insignificant at conventional levels for the vast majority of forecast horizons.*

D.8 The Bayh–Dole Act

A significant legislative change in patenting occurred in 1980 with the Bayh–Dole Act. The Act changed the rules for patents funded by the federal government and harmonized regulations across agencies. Before the Bayh-Dole Act, the federal government typically retained ownership of inventions arising from its funded research. However, the policies were fragmented across agencies. The Act was intended to increase the incentives for universities and research centers to participate in federally funded projects and patent the resulting research, with the government retaining a license to the patent and acknowledgments of funding. Our time series approach captures the increased patenting incentives introduced by the Act.

Additionally, the Act also led to the reclassification of the ownership of government-funded patents from the government to contractors (universities, firms, etc.). To assess whether classification affects our results, we consolidate the two types of government-funded patents into a single category. Figure D.7 provides very similar insights to our baseline specification: the boost that government-funded patents provide to GDP and TFP is twice as large as the stimulus coming from privately-funded patents. Figure D.8 - which corresponds to Figure 7a) - is also consistent with the results in the main text: NIH, HHS, and DoE shocks lead to large gains to both GDP and TFP, while the contribution of DoD and NASA is smaller or not statistically significant.

Figure D.7: Effects on GDP and TFP - grouping government-funded patents

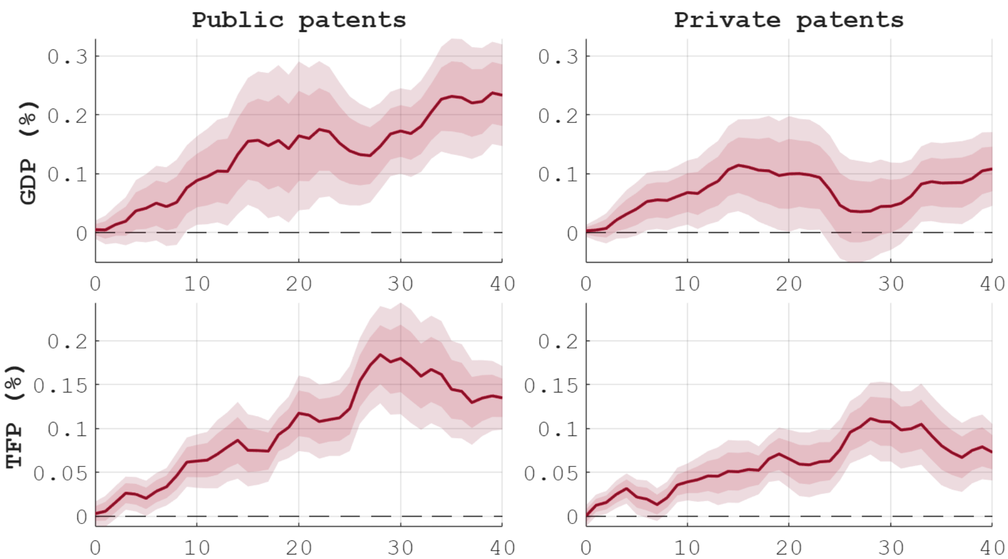

Note. The figure displays the dynamic effects of innovation shocks in each category of patents on (log) real per-capita GDP and (log) utilization-adjusted TFP. This is a robustness check of the baseline results. The estimation by local projections follows eq.(5). The size of the shock is normalized to increase total patents by 1% on impact. The set of controls is the same as in the baseline specification. The solid line represents the point estimate, while the shaded areas report 68% and 90% confidence intervals computed from Newey-West standard errors. Sample: quarterly data 1950:Q1-2015:Q4.

Figure D.8: Federal Agencies breakdown - unique government-funded category

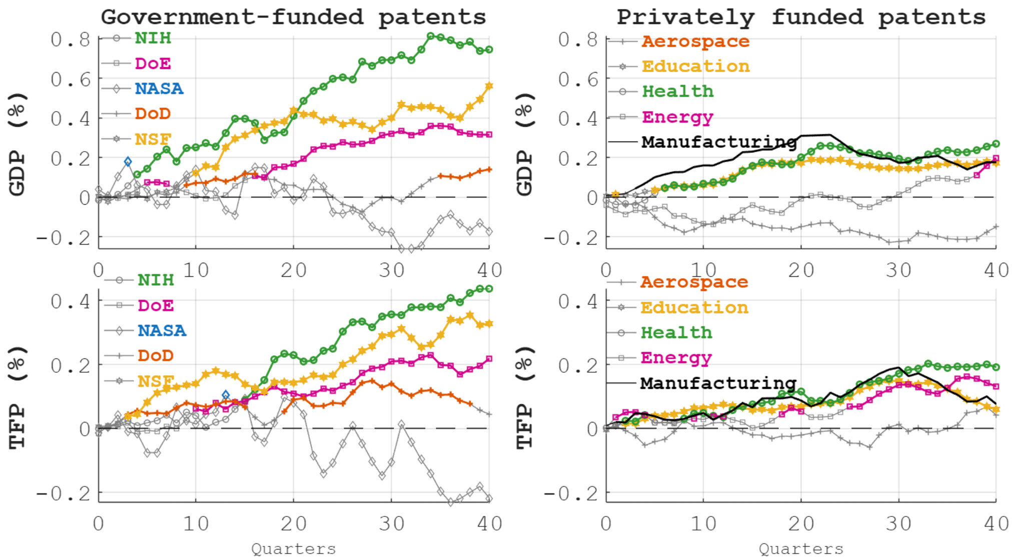

Note. The figure displays the dynamic effects of innovation shocks in each category of patents (government-funded, privately-funded; by column) on (log) real per-capita GDP and (log) utilization-adjusted TFP (by row) by agency-sectoral breakdown. The estimation by local projections follows eq.(5). The size of the shock is normalized such that the peak response of total patents is 1%. The set of controls includes 4 lags of the patent group shocked and real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill, real per-capita R&D expenditure, and the number of patents in other groups. All variables except the T-bill are in logs. Colored (gray) lines denote (no) significance at the 68% level according to heteroskedasticity-robust standard errors. Sample: quarterly data 1950:Q1-2015:Q4 (except for 'University' in the second column, 1975q1-2015:Q4).

E Historical Decomposition and Counterfactuals

In this Appendix, we outline how we compute historical decompositions across multiple forecast horizons and construct counterfactual scenarios using our baseline LP model.

E.1 Historical Decomposition

Given the set of observables $\mathbf{y_t}$ that constitute our information set (or set of controls), consider pat_t as the patent group of interest from which we want to extract an innovation shock. In this discussion, we consider total factor productivity (TFP) as the variable of interest (but the same holds for any other variables, e.g., GDP). Let h^* denote the forecast horizon.

Following Gorodnichenko and Lee (2020), we extract our innovation shock ε_t by regressing pat_t on the information set at $t-1$:

$$\text{pat}_t = \sum_{j=1}^{4} \boldsymbol{\phi_j}\, \mathbf{y}_{t-j} + \varepsilon_t \tag{1}$$

Then, define the h^* ahead TFP_t level, that isolates low-frequency movements as $h^* = 32$ ($h^* = 6$ for a cross check reported in Figure E.1):

Regress TFP_{t+h^*} on lagged $\mathbf{y_t^*} = [\mathbf{y_t}\, \varepsilon_t]$ that contains the observables and the shock themselves to construct the forecast error of the baseline model:

$$\text{TFP}_{t+h^*} = \alpha + \sum_{j=1}^{4} \boldsymbol{\phi_j^*} \mathbf{y^*}_{t-j} + e_{t+h^*} \tag{2}$$

Then we regress the forecast error on the shocks realized between t and the forecast horizon h^*:

$$e_{t+h^*} = \sum_{k=0}^{h^*} \theta_{h^*,k}\, \varepsilon_{t+k} + u_{t+h^*} \tag{3}$$

Equation 3 is estimated separately for each forecast horizon $h = 0, \ldots, H$ by setting $h^* = h$. The R^2 of this regression for each forecast horizon represents the variance contribution (i.e. the average over the sample). For the historical decomposition, we fix $h^* \in \{6, 32\}$ and the fitted values from (3) deliver the historical contribution of the shock of interest. Inference is based on a VAR-based bootstrap as in Gorodnichenko and Lee (2020).

E.2 Counterfactual Scenario

We consider a counterfactual scenario for the 2000s to quantify the implications of the decline in the government-private innovation shock contribution after its peak. The exercise aims to answer the question: what would the level of medium-term TFP have been if the shock's contribution had remained at its 1996 peak level? We use the bias-corrected median historical series for the 32-quarter annualized TFP forecast error growth (g_t) and the corresponding median contribution of the innovation shock (s_t).

First, we identify the peak median shock contribution that occurred in 1996, which we denote s^*. We then construct a counterfactual series for the TFP medium-term growth, $\tilde{g}_t$, starting from 1996 and extending to 2007. For all data points after the 1996 peak, the counterfactual is defined as:

$$\tilde{g}_t = g_t - s_t + s^*$$

Figure E.1: Historical Decomposition of the Short-Term Component of TFP Growth

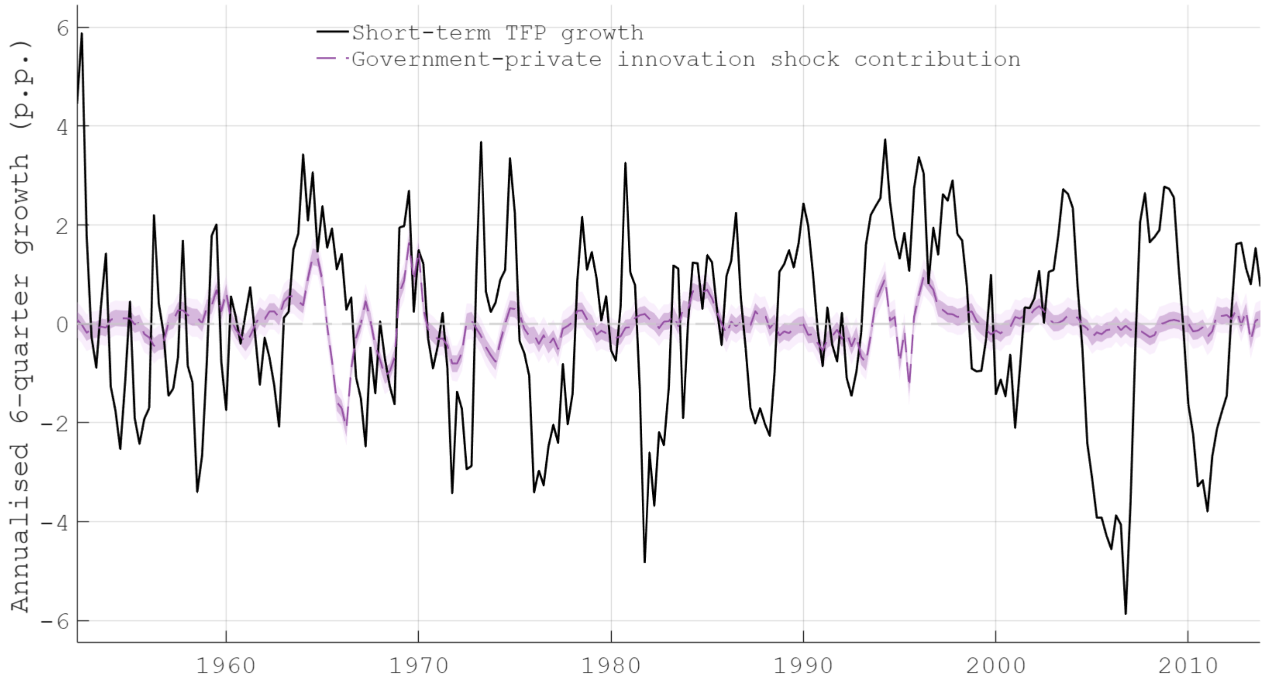

Note. The figure displays the historical contribution of public-private innovation shocks to the forecast errors of the 6-quarter TFP growth rate. The estimation by local projections is based on the method in Gorodnichenko and Lee (2020). The solid black line represents the stochastic component of TFP at the 6-quarter horizon, while the purple line (bands) represents the point estimate (68% and 90%) contribution of the public-private innovation shocks. Inference is based on 2000 bootstrap replication with small-sample adjustment. Sample: quarterly data 1950:Q1-2015:Q4.

This procedure effectively replaces the actual shock contribution from the peak onwards with the constant peak value, while leaving the baseline growth and the contributions of all other shocks unchanged.

Both the actual series (g_t) and the counterfactual series ($\tilde{g}_t$) are given in annualized percentage points. We convert these to equivalent quarter-on-quarter growth rates (g_t^q and $\tilde{g}_t^q$) using the standard compounding formula: $g_t^q = (1 + g_t/100)^{1/4} - 1$.

We then construct the level of TFP by setting it to a normalized value of 100 at the quarter of the 1996 peak. The time path of the TFP level is computed by compounding these quarterly rates. To construct a 90% confidence interval for our counterfactual estimate, we repeat the entire exercise using the 5th and 95th percentile levels of the shock contribution at the 1996 peak.

Figure E.2: Counterfactual Scenario

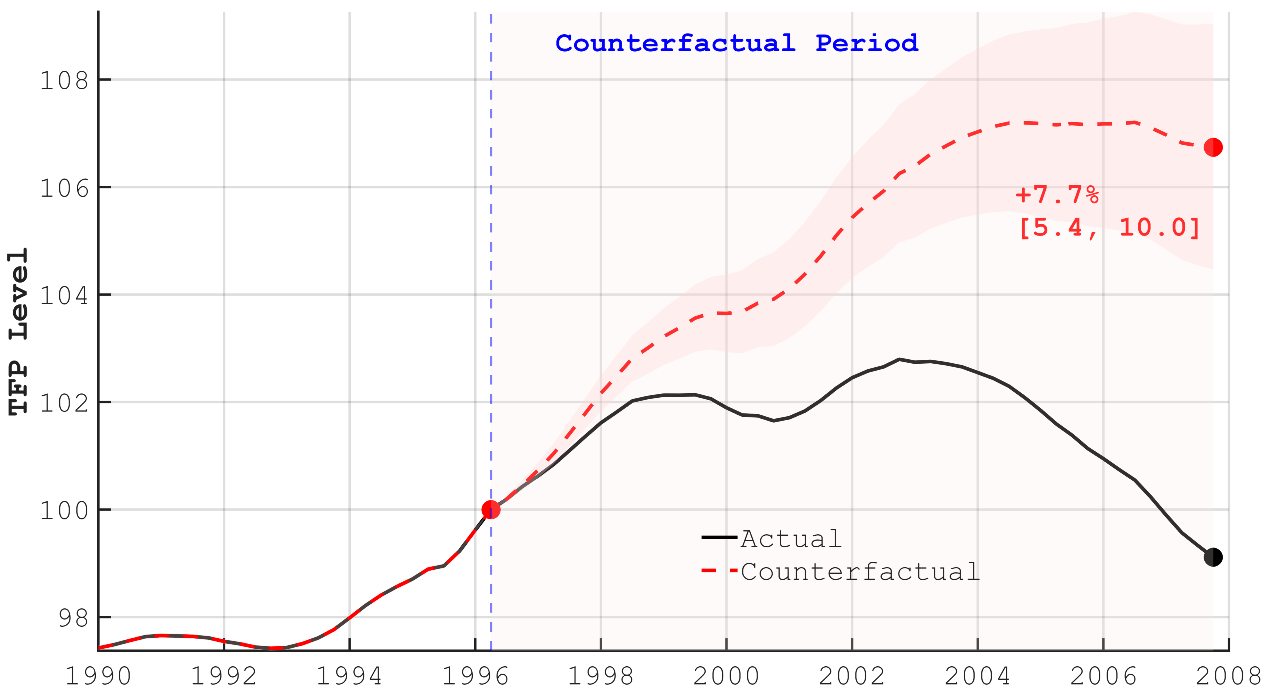

Note. The figure displays the level of realized medium-term TFP (black solid line) and a counterfactual scenario (red dashed line; bands denote 90% confidence bands) where we hold the government-private innovation shock contribution at the 1996 peak level till 2007. The numbers displayed in red denote the difference (with confidence interval in square brackets) between the counterfactual TFP level at the end of the sample and the realized value.

F Alternative Measures of Basicness

In this Appendix, we show that our main findings of: (i) a larger average association between GDP/TFP and innovations funded by the government and owned by the private sector; (ii) a leading role for fully public patents among the most fundamental, basic and scientific innovations, are robust to several variants of the measures of basicness developed by Kelly et al. (2021); Kogan et al. (2017); Marx and Fuegi (2020) respectively.

F.1 Importance

We consider an alternative definition of the *importance* measure from Kelly et al. (2021). In the paper, we use the 5-year version and construct dummies for percentiles that are computed within each group of patents (i.e., government-private, private-private, and government-government).

Figure F.1 reports the corresponding results obtained by using the 10-year version of the importance measure. This measure provides a longer-term perspective, but due to the truncation of the 10-year forward similarity, it ends in 2010. The picture is very similar to our baseline.

We also construct alternative dummies for above versus below the median (Figure F.2) and across the entire distribution of patents (Figure F.3). In both cases results are consistent with our baseline.

F.2 Reliance on Science

We also consider an alternative definition of top/bottom reliant on science patents based on the top 25% percentile (Figure F.4) and another one that exploits exclusively the extensive margin of papers' citations (Figure F.5). These two exercises yield overall consistent results with our baseline, although in this case, only government-government patents display marked statistical significance, possibly reflecting a more even distribution of this measure within each group. The exercise based on the extensive margin is important because it indicates that the trend in citations within the sample does not influence our conclusions about the role of science-based patents.

Figure F.1: Importance - 10 year measure

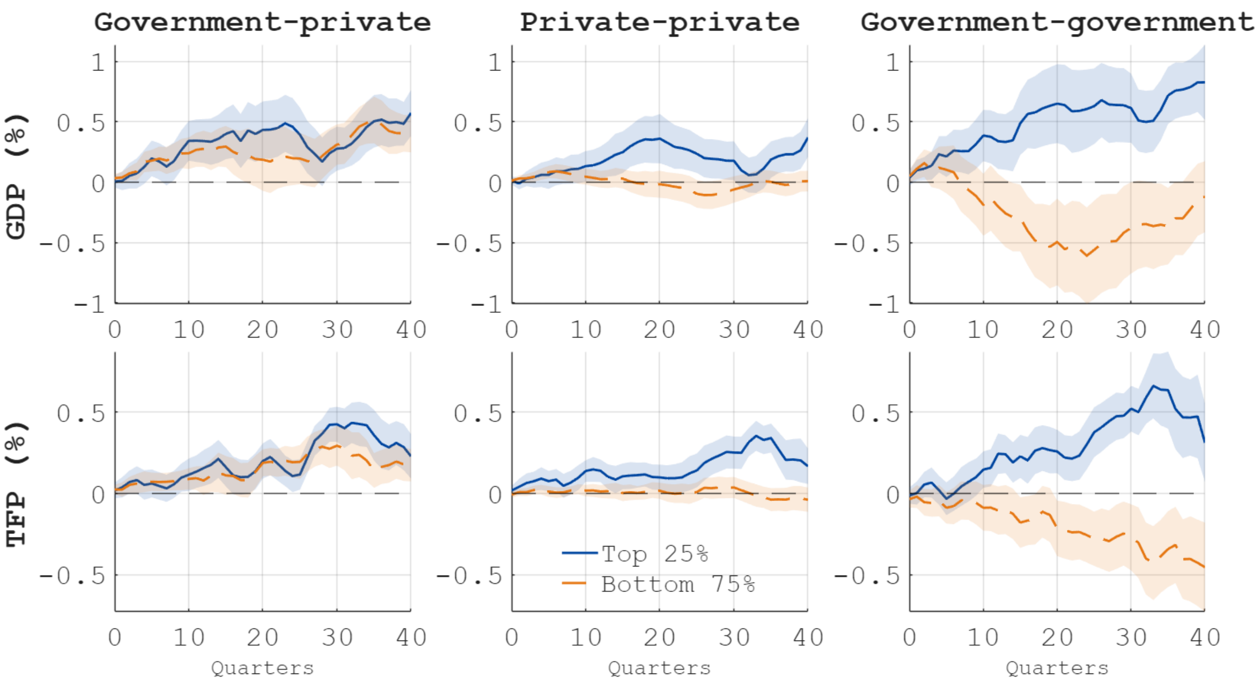

Note. The figure replicates the analysis in Figure 6, comparing the dynamic effects of innovation shocks from the top quartile versus the bottom three quartiles of patents ranked by importance. As a robustness check, this figure uses the 10-year patent similarity measure from Kelly et al. (2021) instead of the 5-year measure used in the main text. The estimation by local projections follows eq.(5). The size of the shock is normalized to increase total patents by 1% on impact. The solid blue (dashed orange) line represents the point estimate for top 25% (bottom 75%) important patents, while the corresponding shaded areas report 90% confidence intervals. Sample: 1950:Q1-2010:Q4 (due to the 10-year window).

Figure F.2: Importance - above/below median

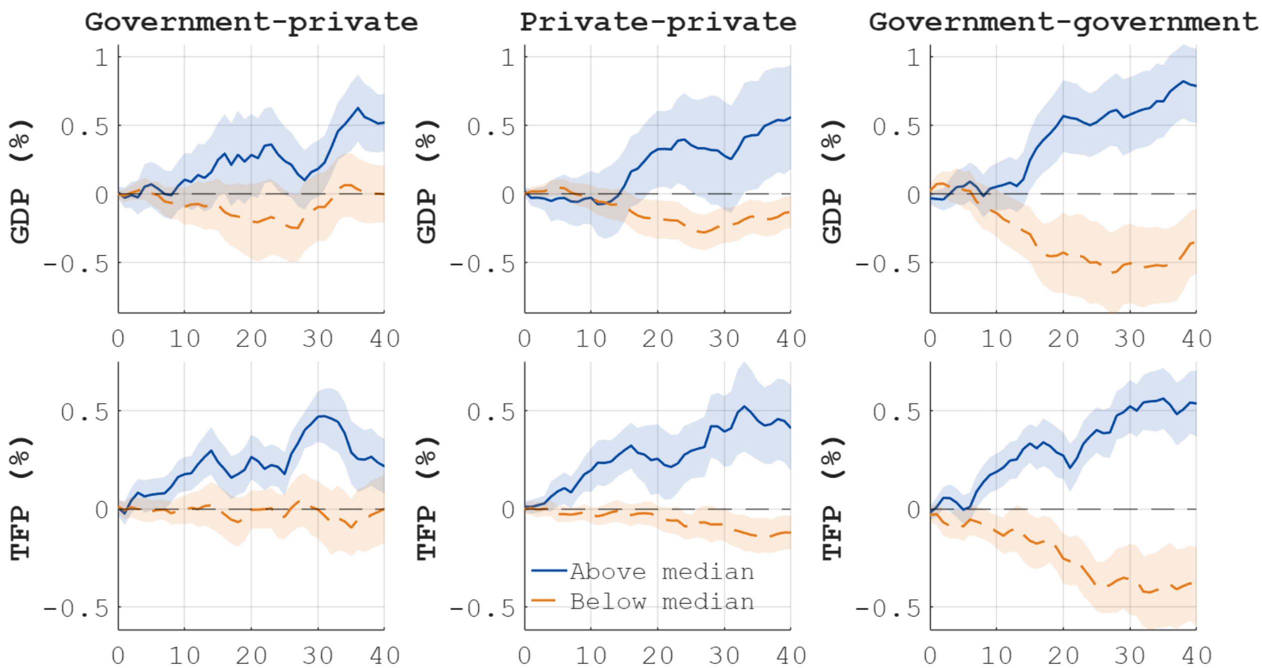

Note. The figure replicates the analysis in Figure 6, but splits patents at the median of the importance distribution (top 50% vs. bottom 50%) instead of the top quartile. The 5-year patent similarity measure from Kelly et al. (2021) is used. The estimation by local projections follows eq.(5). The size of the shock is normalized to increase total patents by 1% on impact. The solid blue (dashed orange) line represents the point estimate for top 50% (bottom 50%) important patents, while the corresponding shaded areas report 90% confidence intervals. Sample: quarterly data 1950:Q1-2015:Q4.

F.3 Results Based on an Extended Kogan et al. (2017) Measure

As a further robustness exercise, we construct an alternative measure of patent importance based on the approach of Kline et al. (2019) to extend the Kogan et al. (2017) patent value measure to the universe of patents. The objective is to impute a value proxy for patents that lack observed market-based valuations, using only information available at the time of filing.

We estimate a predictive model of the Kogan et al. (2017) value measure. The explanatory variables include (following Kline et al. (2019)) patent family size and the number of claims, which we extend to pre-1976 patents using text analysis of the Google Patents public corpus. We also include four measures of patent technological relevance, constructed by Martinez and Moen-Vorum (2025), who demonstrate that these indicators are strong predictors of the Kogan et al. (2017) patent values, as well as CPC 3-digit technology classes and filing-year fixed effects. Since our objective is to cover the patent universe, we cannot include firm-level characteristics such as revenue and employment as predictors, which Kline et al. (2019) show have high explanatory power for patent values.

Figure F.3: Importance - percentiles for each category

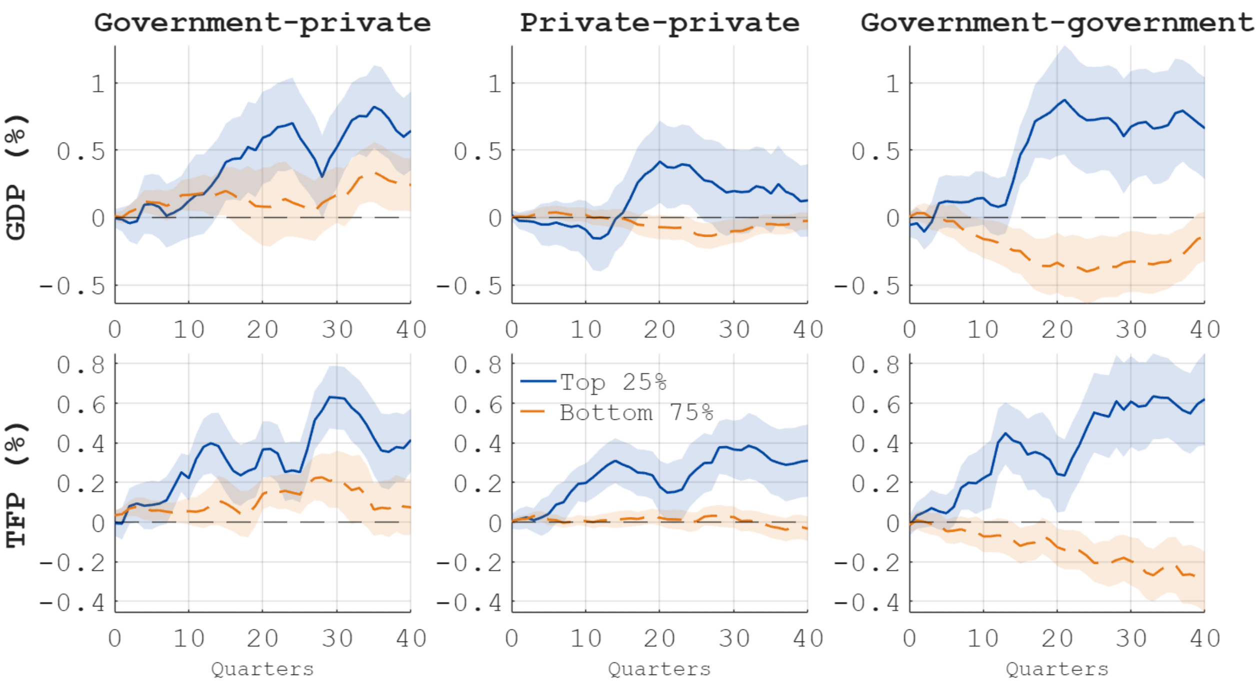

Note. The figure replicates the analysis in Figure 6, but the importance ranking is done across all patents rather than within each patent category (public-private, private-private, and public-public). This tests whether the top quartile of patents in each funding category is more impactful than the bottom three quartiles of the same category. The estimation by local projections follows eq.(5). The size of the shock is normalized to increase total patents by 1% on impact. The solid blue (dashed orange) line represents the point estimate for top 25% (bottom 75%) important patents within each category, while the corresponding shaded areas report 90% confidence intervals. Sample: quarterly data 1950:Q1-2015:Q4.

Figure F.4: Reliance on Science - Alternative percentiles

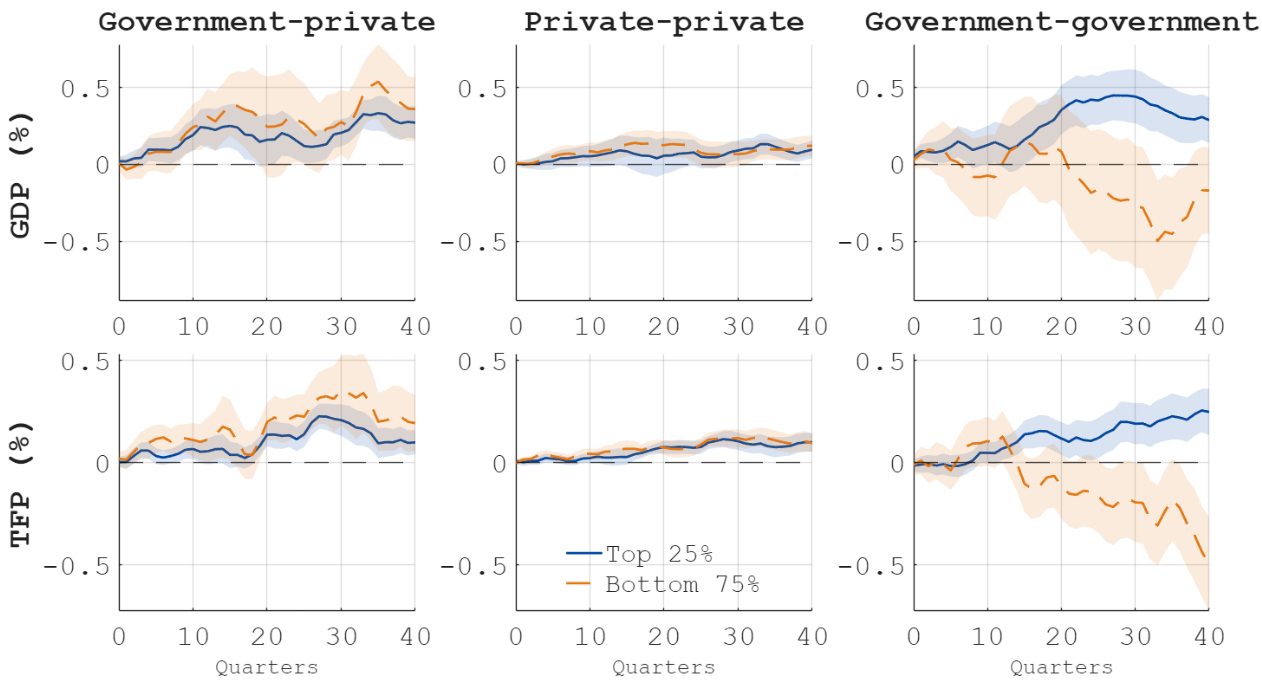

Note. The figure replicates the analysis in Figure 6b, but compares the top 25% with the bottom 75% percentiles. The estimation by local projections follows eq.(5). The size of the shock is normalized to increase total patents by 1% on impact. The solid blue (orange dashed) line represents the point estimate for top 25% (bottom 75%) most reliant on science patents within each category, while the corresponding shaded areas report 90% confidence intervals. Sample: quarterly data 1950:Q1-2015:Q4.

The fitted coefficients from this regression are then used to predict patent values for all observations, including those without market-based valuations.

We employ this measure as a robustness check alongside the Kelly et al. (2021) importance metric used in the main analysis. Following the approach described in Section 6, in Figure F.6 we divide each patent category into two groups—top 25 percent (blue solid lines) and bottom 75 percent (orange dashed lines)—using the extrapolated Kogan et al. (2017) value measure. Two main results emerge. First, the top 25 percent of patents are associated with larger effects on GDP and TFP across all categories. Second, the bottom 75 percent of patents by value exhibit mostly insignificant effects for private-private and government-government patents. Overall, these results confirm that the heterogeneity observed in the main analysis is closely linked to underlying differences in patent value: patents with higher predicted values account for a disproportionate share of the macroeconomic impact of innovation.

Figure F.5: Reliance on Science - Extensive Margin

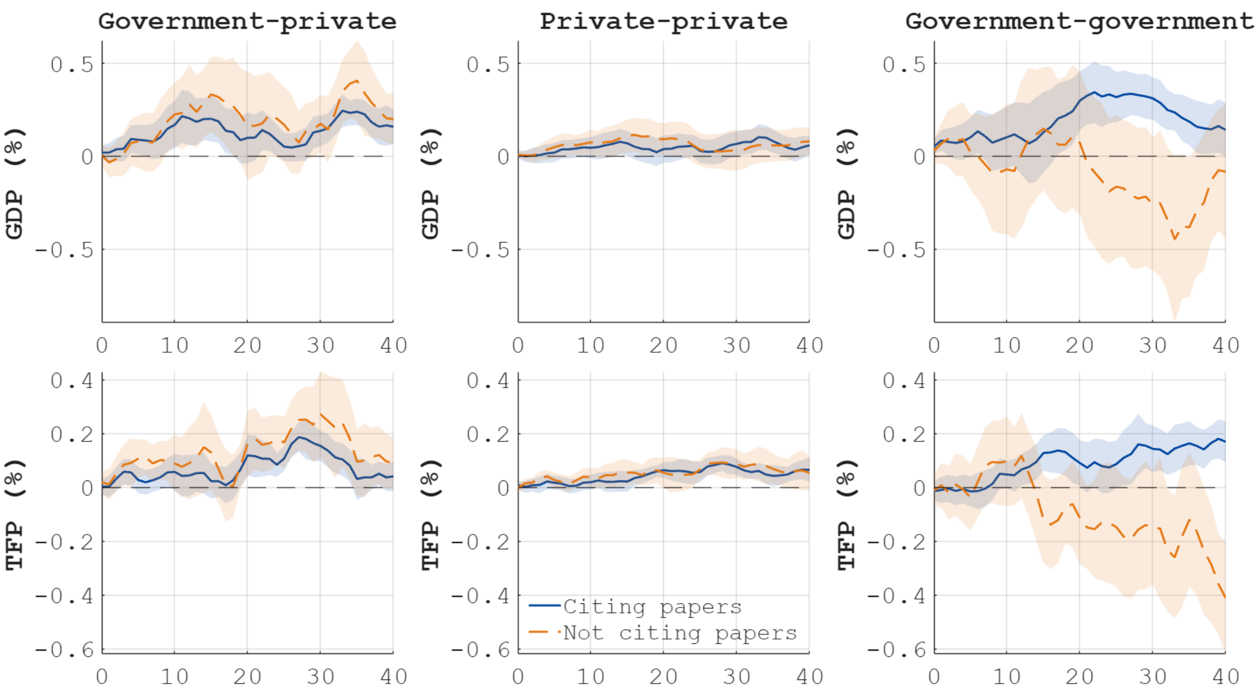

Note. The figure replicates the analysis in Figure 6b, but the reliance on science groups are based on citing at least one paper versus not citing any paper. The estimation by local projections follows eq.(5). The size of the shock is normalized to increase total patents by 1% on impact. The solid blue (dashed orange) line represents the point estimate for patents citing at least one paper (not citing any papers) within each category, while the corresponding shaded areas report 90% confidence intervals. Sample: quarterly data 1950:Q1-2015:Q4.

Figure F.6: The Dynamic Effects of the Most Valuable Innovations

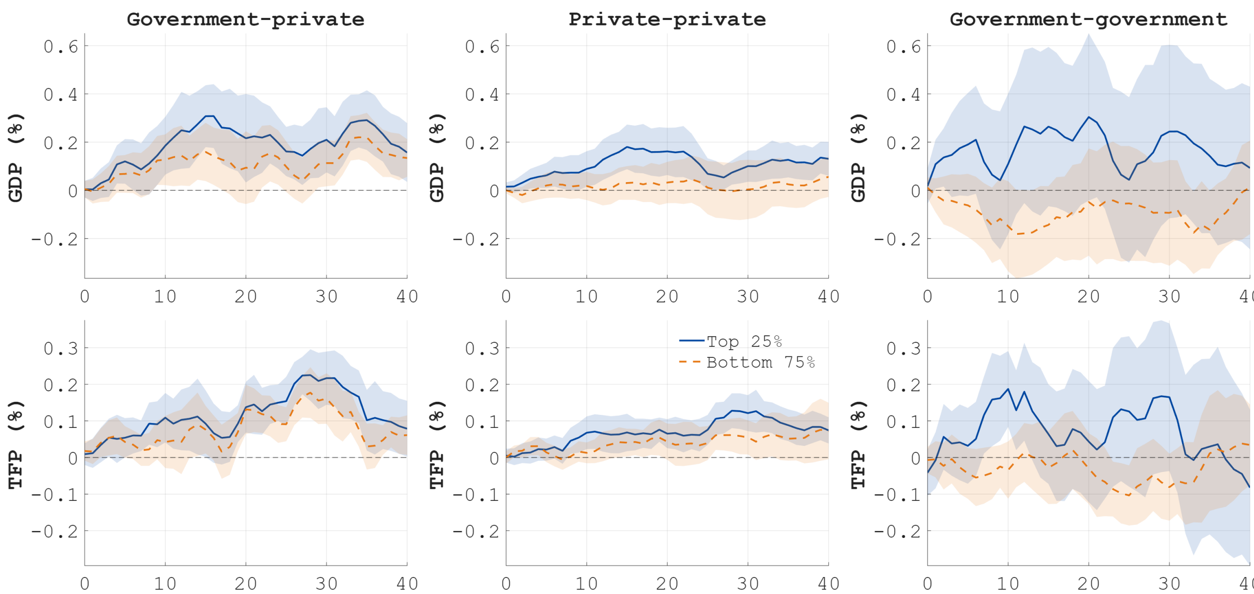

Note. This figure represents the dynamic effects of innovation shocks to the top quartile of patents ranked by the extended Kogan et al. (2017) patent value measure (described in the text) versus other patents in each category of patents (public-private, private-private, public-public; by column) on (log) real per-capita GDP and (log) utilization-adjusted TFP (by row). The estimation by local projections follows eq.(5). The size of the shock is normalized such as to increase total patents by 1% on impact. The set of controls includes 4 lags of the patent group shocked and real per-capita GDP, TFP, real per-capita investment, real stock prices, the T-bill, real per-capita R&D expenditure, and the number of patents in other groups. All variables except the T-bill are in logs. The solid blue (dashed orange) line represents the point estimate for top 25% (bottom 75%) patents by value, while the corresponding shaded areas report 90% confidence intervals computed from heteroskedasticity-robust standard errors. Sample: quarterly data 1950:Q1-2015:Q4.

G Definitions of industry and research field patent groups

This Appendix provides detailed information on the construction of industry and research fields by specifying the exact CPC codes that enter each category. We also show the composition of patents within agencies, universities, research institutes, and private-private by all research fields and the four research fields we focus on in the paper.

G.1 Industry groups

A patent is assigned to a sector if its primary CPC code in the USPTO Cooperative Patent Classification (CPC) Master Classification Files for US patent grants matches any of the criteria listed below.

Table G.1: Sector names and associated CPC codes

Sector Name	**CPC Codes**
Energy	B01D, C01B, C10L, C25B, F03B, F03D, G21, H01M, H02J, H02K, H02S, Y02C, Y02E, Y02P, Y02T
Aerospace	B64, F02K, F03H, F42B, F42E, F42F, G05D 1/xx
Health	A61B, A61F, A61J, A61K, A61L, A61M, A61N, A61P, A61Q, B01L, C07H, C07K, C12M, C12N, C12P, C12Q, G01N, G16H
Manufacturing	B08, B21-B30, B29C, B29D, B32B, B33B, B33Y, B65, B66, C21-C25, G05B 19/xx, H05K, Y02P
Education	*Based on assignee name keywords, not CPC codes*

The education sector is not constructed from CPC codes. It is based on assignee-name information and corresponds to patents assigned to universities. The classification is described in Appendix I.

G.2 Research fields

In Figures G.1 and G.2, we report the innovation composition of federal agencies and institutional players, such as research institutes and universities, by research field. In Table G.2, we show how individual CPC codes map into the major research fields that we focus upon.

Figure G.1: Share of patents by all research fields in each agency and player

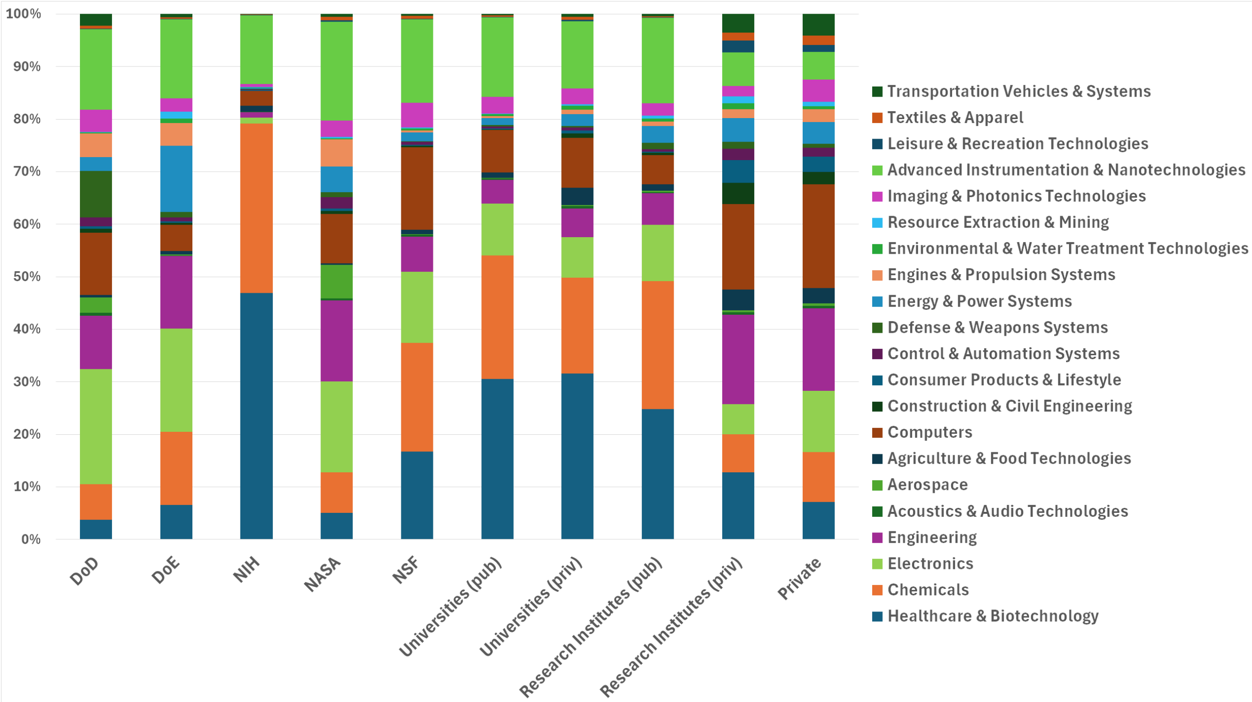

Note. This chart illustrates the proportions of patents in research fields for different groups. The first five bars (DoD, DoE, NIH, NASA, NSF) show the composition of patents for the five agencies we study, followed by public interest patents involving research institutes and universities. The final bar shows the proportions for all private-private patents. Sample: quarterly data 1950:Q1-2015:Q4.

Figure G.2: Share of patents by main research fields in each agency and player

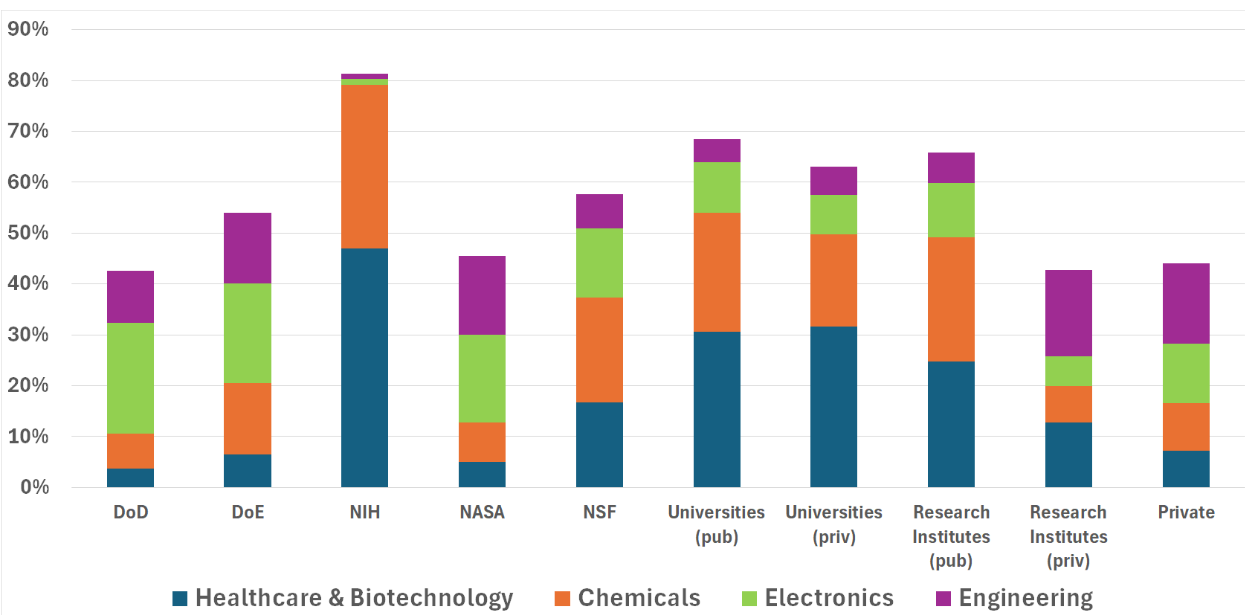

Note. This chart illustrates the proportions of patents in the four research fields studied in the paper, and detailed in Table G.2. The first five bars (DoD, DoE, NIH, NASA, NSF) show the composition of patents for the five agencies we study, followed by public interest patents involving research institutes and universities. The final bar shows the proportions for all private-private patents. Figure G.1 below shows the total composition of patents by research field. Sample: quarterly data 1950:Q1-2015:Q4.

Research field name	CPC codes
Chemicals	B01, C01, C07, C08, C09B, C09D, C09J, C09K
Electronics	H01, H03, H05, H10
Engineering	B03, B04, B05, B06, B07, B21, B22, B23, B24, B25, B26, B29, B30, B31, B32, B33, B41, B65, B66, B67, C03, C04, C21, C22, C23, C25, D21, F04, F15, F16, F26, F27
Healthcare & Biotechnology	A61B, A61C, A61D, A61F, A61G, A61H, A61J, A61K, A61L, A61M, A61N, A61P, C12M, C12N, C12P, C12Q

Table G.2: Research field names and associated CPC codes

Note. This chart shows the proportions of patents in the four research fields studied in the paper. A patent is assigned to a research field if its primary CPC code in the USPTO Cooperative Patent Classification (CPC) Master Classification Files for US patent grants matches any of the criteria listed below.

H The Composition of Basic R&D in the United States

National statistics on R&D funding and performance in the US are provided by the National Center for Science and Engineering Statistics of the NSF. The Center compiles the *National Patterns of R&D resources*, an annual statistical report that provides an integrated overview of the U.S. research and development landscape. Data contained in the National Patterns are based on the following annual surveys: *Business Enterprise Research and Development Survey*, *Annual Business Survey*, *Higher Education Research and Development Survey*, *Survey of Federal Funds for Research and Development*, *FFRDCs Research and Development Survey*, and *Survey of State Government Research and Development*.

R&D is typically divided into three categories: (i) basic research, aimed at acquiring new knowledge on the foundations of observable facts and phenomena without a specific application as outlet; (ii) applied research, directed primarily towards a specific solution or objective; (iii) experimental development, translating research findings into new or improved products, processes, or services, and implying prototyping, testing, and iterative improvements. Panel (a) of Figure H.1 reports the breakdown of annual R&D expenditures since 1953 by these three categories.

R&D expenditures have four main funders: the US government (on average 95% federal and 5% state/local) through its departments and agencies; higher education institutions (universities and colleges, here called "universities" for brevity); "private businesses" (98% per year by US businesses and 2% by foreign companies); "Other non-profit organizations" including private foundations, research institutes not affiliated with government or higher education, charitable and philanthropic organizations that conduct or fund research. Panel (b) of Figure H.1 reports the breakdown of annual R&D expenditures since 1953 by source of funds, focusing on basic research only. The federal government has historically been the primary source of funding for this type of R&D, although its share has gradually declined over time as private sector contributions have increased.

The R&D actors discussed above not only fund but also conduct research activities using either their own resources or external funding. Panel (c) of Figure H.1 focuses specifically on federally funded basic research, showing how these funds are distributed across different types of performers. Historically, universities have been the dominant performers of federally funded basic research, with their share increasing steadily over time. Another major group of performers is that of Federally Funded Research and Development Centers (FFRDCs), referred to here as ”research institutes,” which are typically operated by universities, non-profit organizations, or industrial firms under contract with federal agencies. Collectively, universities and research institutes account for approximately 70% of total federal spending on basic research, on average.

Figure H.1: R&D expenditures in the United States

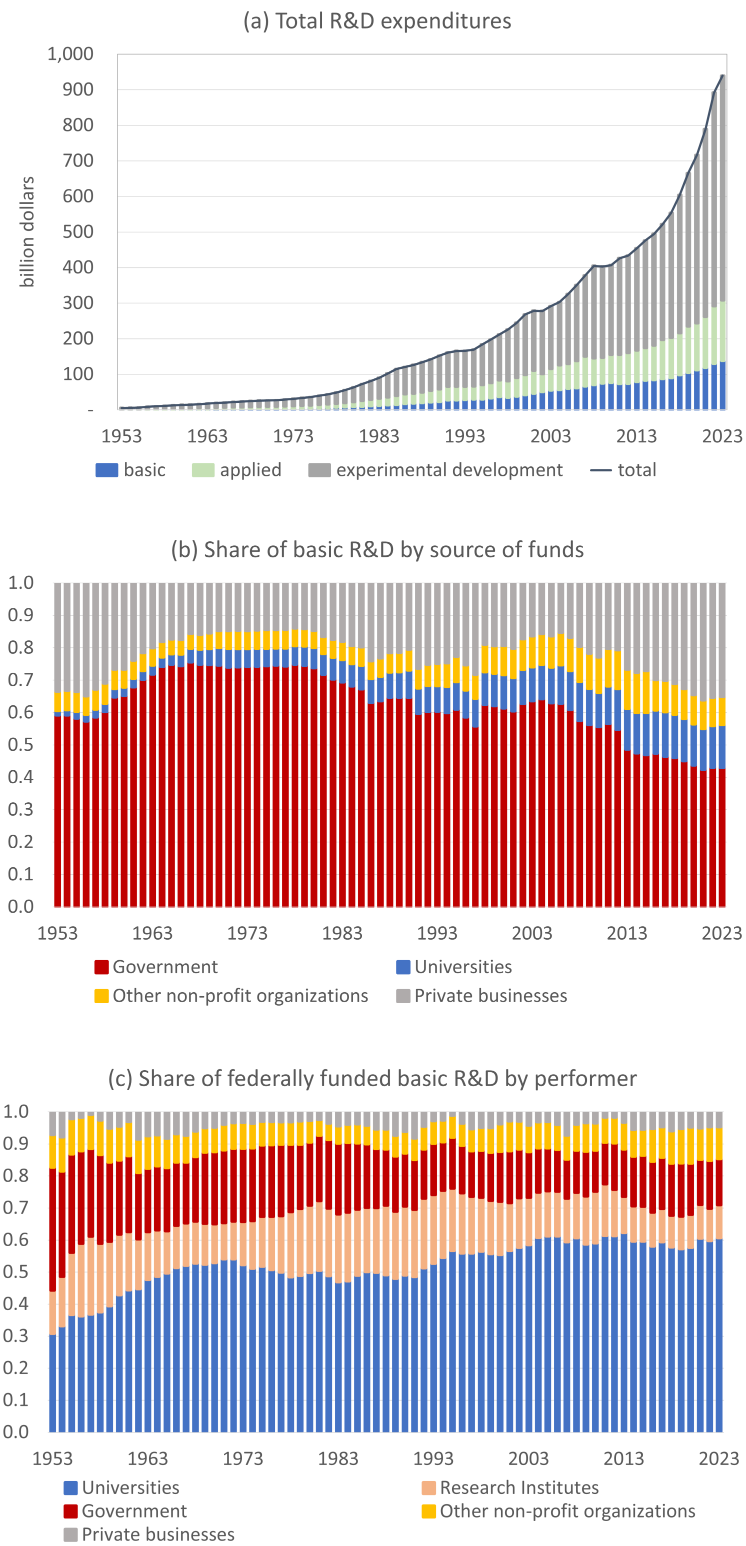

Source: National Patterns of R&D Resources *2022-2023, NSF. In Figure (b) and (c), Government includes both federal government and state and local governments, when the former accounts on average for 95% of the provided funds (b) and almost 100% of the expenditures (c). The share of basic research performed by the federal government is done with its departments and agencies.*

I Further Details on Key Institutional Players

This Appendix provides additional details on the construction and composition of the institutional-player groups used in Section 8. We focus on research entities, defined as universities and research institutes, and on start-ups among for-profit assignees. We first describe how we identify universities and research institutes and how we extend the classification back to 1950. We then report four descriptive tables. Table I.1 describes the largest research institutes and national-lab operators in the research-institute group. Table I.2 decomposes research-entity public-private patents into universities and research institutes by funding agency. Table I.3 reports, by agency, the share of public-private patents involving research entities and the share of start-ups among for-profit participants. Finally, Table I.4 describes the top start-up assignees among public-private innovators.

I.1 Identifying Universities and Research Institutes

The analysis in Section 8 compares patents assigned to research entities with patents assigned to other organizations. Research entities include universities and research institutes. The construction combines assignee information from *PatentsView*, the start-up and venture-capital mapping in Ewens and Marx (2024), the Government Patent Register of Gross and Sampat (2025), and historical USPTO assignee records used to extend the classification back to 1950.

Universities. Universities are identified using the assignee names rather than on CPC codes. A patent is flagged as university-assigned if its disambiguated assignee organization name, converted to lowercase, contains one of the following strings:

> `caltech`, `carnegie mellon`, `cmu`, `college`, `columbia`, `cornell`, `cornell research foundation`, `drexel`, `georgia inst tech`, `georgia tech`, `harvard`, `institute of technology`, `jhu`, `johns hopkins`, `massachusetts inst tech`, `mit`, `new york univ`, `nyu`, `ohio state`, `penn state`, `pennsylvania state`, `princeton`, `regents of the university of california`, `rutgers`, `suny`, `tamu`, `texas a&m`, `texas a&m university`, `u-m`, `uc berkeley`, `uc davis`, `uc irvine`, `uc los angeles`, `uc merced`, `uc riverside`, `uc san diego`, `uc santa barbara`, `uc santa cruz`, `uiuc`, `umich`, `university`, `usc`, `yale`.

For consistency, we use this same university flag in the institutional-player analysis and in the historical extension to the pre-1976 sample.

Research institutes. Research institutes are identified using a residual procedure designed to avoid classifying ordinary private firms as research institutions. We start from the universe of private assignees in *PatentsView* and merge these assignees with the mapping in Ewens and Marx (2024). We then keep assignees that are not mapped in Ewens and Marx (2024). This residual group contains many non-standard organizations, but it may also include private firms that are simply not covered by the mapping. We therefore clean the residual group by removing remaining private firms using name information, legal-form indicators, and manual inspection of the largest assignees.

The private-firm screen uses strings such as:

> `company`, `co`, `corporation`, `corp`, `inc`, `incorporated`, `llc`, `ltd`, `limited`, `lp`, `plc`, `gmbh`, `ag`, `sa`, `nv`, `electronics`, `semiconductor`, `systems`, `technologies`, `pharmaceuticals`, `pharma`, `biotech`, `chemical`, `manufacturing`, `industries`.

These strings are used as screening devices, not as automatic deletion rules. We manually restore assignees whose legal names contain corporate suffixes but whose institutional role is that of a research institute, research hospital, FFRDC operator, or national-lab operator. Examples include:

> *Dana-Farber Cancer Institute, Inc.*, *Brookhaven Science Associates, LLC*, *UChicago Argonne, LLC*, *Los Alamos National Security, LLC*.

From the cleaned residual set, we retain the assignees that appear among government-private patents. These retained assignees form the research-institute list. We then apply this same research-institute assignee list to private-private patents as well. Thus, a privately funded patent is classified as a research-institute patent if its assignee belongs to the list identified from the cleaned residual set of government-private assignees.

The following strings are used only to help locate candidate residual assignees and historical variants during manual review:

> `research institute`, `institute for`, `institute of`, `research center`, `research centre`, `laboratory`, `laboratories`, `national laboratory`, `national laboratories`, `medical research`, `cancer research`, `biological studies`, `battelle`, `argonne`, `brookhaven`, `cold spring harbor`, `dana farber`, `scripps`, `salk`, `sloan kettering`, `whitehead`, `wistar`, `los alamos`, `sandia`, `lawrence berkeley`, `lawrence livermore`, `ut battelle`, `uchicago argonne`.

The final research-institute list retains independent research institutes, research hospitals, research foundations, FFRDCs, and national-lab operators. It excludes ordinary private firms, corporate laboratories, and assignees whose primary role is commercial production rather than scientific research.

Patent-level classification. The classification is first performed at the assignee level and then mapped to patents. A patent is classified as a university patent if at least one assignee is identified as a university using the Appendix G classification. A patent is classified as a research-institute patent if at least one assignee belongs to the research-institute list described above. A patent is classified as a research-entity patent if it is assigned to either a university or a research institute. When a patent has both university and research-institute assignees, it is included in the combined research-entity group.

Extension to 1950. The modern assignee information is based on *PatentsView*. To extend the classification back to 1950, we use historical USPTO assignee records, retrieved through Google Patents Public Data on BigQuery. The BigQuery search is used to recover historical assignee names for U.S. patents filed between 1950 and 1975. It is not used as the final classifier.

For universities, we search the historical assignee names using the university strings listed in Appendix G and manually checked variants of large university names. For research institutes, we search historical assignee names using the post-1976 research-institute assignee list, together with manually verified historical name variants. In both cases, historical assignee names are standardized in the same way as modern assignee names before matching.

This procedure gives a harmonized research-entity indicator for the full 1950–2015 sample. The post-1976 component is based on *PatentsView*; the pre-1976 component is based on historical USPTO assignee records matched to the university and research-institute lists.

Validation. We validate the classification in three steps. First, we inspect the largest assignees classified as universities and research institutes to verify that they correspond to higher-education institutions, university research foundations, independent research institutes, research hospitals, FFRDCs, or national-lab operators. Second, we inspect the largest assignees excluded from the residual research-institute list to ensure that ordinary private firms are not retained simply because they are missing from the Ewens and Marx (2024) mapping. Third, we verify that the resulting research-entity patent series does not display mechanical discontinuities around 1976, when the main assignee source switches from historical USPTO records to *PatentsView*.

I.2 Start-ups in Public-Private Patenting

For start-ups and venture-capital-backed firms, we use the classification in Ewens and Marx (2024). A patent is classified as a start-up patent if its filing date is within three years of the assignee's founding year. A patent is classified as VC-backed if the assignee is recorded as having received venture-capital financing at any point in its life. Start-up status and VC backing are not mutually exclusive. In Section 8.2, "other firms" denotes non-start-ups in the start-up comparison and non-VC-backed firms in the VC comparison.

Table I.1: Top U.S. Research Institutes and National-Lab Operators

Institute / Entity	**Count**	**Description**	**FFRDC Status**
Battelle Memorial Institute	829	Major United States nonprofit science and technology research institute; manages or co-manages several national laboratories for the United States Department of Energy (DOE) and the Department of Homeland Security (DHS); deep ties to government and academia.	Yes
The General Hospital Corporation	795	Legal entity for Massachusetts General Hospital, the oldest and largest teaching hospital of Harvard Medical School; a world leader in biomedical research and clinical care.	No
The Scripps Research Institute	642	Leading nonprofit American medical research facility focusing on biomedical research; headquartered in La Jolla, California, and Jupiter, Florida; historically affiliated with universities and hospitals.	No
Los Alamos National Security, LLC	453	Consortium including the University of California, Bechtel National, BWXT Government Group, and URS Energy and Construction that managed Los Alamos National Laboratory for the United States Department of Energy; the laboratory is a premier United States nuclear research facility.	Yes
UChicago Argonne, LLC	442	Limited liability company formed by the University of Chicago to manage Argonne National Laboratory for the United States Department of Energy; closely tied to University of Chicago and federal research.	Yes
The Brigham and Women's Hospital	411	Major teaching hospital of Harvard Medical School in Boston, Massachusetts; recognized for biomedical research and clinical care.	No
Dana-Farber Cancer Institute, Inc.	347	Major cancer treatment and research center in Boston, Massachusetts; principal teaching affiliate of Harvard Medical School; member of the Dana-Farber/Harvard Cancer Center consortium.	No
Salk Institute for Biological Studies	319	Nonprofit scientific research institute in La Jolla, California, founded by Jonas Salk; world-renowned for biomedical research, especially in neuroscience and genetics; independent but collaborates with universities such as the University of California, San Diego and with the National Institutes of Health.	No
Sloan-Kettering Institute for Cancer Research	276	Biomedical research division of Memorial Sloan Kettering Cancer Center in New York City; world leader in cancer research and treatment; affiliated with multiple universities including Cornell University, Rockefeller University, and Weill Cornell Medical College.	No
Brookhaven Science Associates, LLC	275	Limited liability company formed by Battelle Memorial Institute and Stony Brook University to manage Brookhaven National Laboratory for the United States Department of Energy; strong ties to the federal government and academia, including collaborations with other universities and private sector partners.	Yes

Notes: This table reports the largest assignees in the research-institute group. Counts refer to patents assigned to each institute or national-lab operator in the sample used for the institutional-player analysis. The final column indicates whether the entity is, or operates, a Federally Funded Research and Development Center (FFRDC) or national-lab-related organization.

Table I.2: Breakdown of Research-Entity Public-Private Patents

	Universities		Research Institutes	
Agency	Frequency	Percent	Frequency	Percent
DOD	6,749	14.2%	996	9.1%
DOE	5,012	10.6%	3,623	33.0%
NIH	21,978	46.3%	5,231	47.6%
NASA	1,579	3.3%	163	1.5%
NSF	9,936	20.9%	422	3.8%

Notes: This table reports the distribution of public-private patents assigned to research entities across funding agencies. Research entities include universities and research institutes. Frequencies report the number of public-private patents assigned to each assignee type and funded by each agency. Percentages are computed within each assignee type, so the university and research-institute columns need not sum to 100 because the table reports only the main federal agencies.

Table I.3: Research-Entity and Startup Shares in Public-Private Patents by Agency

Agency	Research-entity share	Startup share
DoD	38.9%	3.2%
NASA	54.0%	4.2%
DoE	53.2%	3.6%
NIH	78.2%	5.9%
NSF	84.8%	6.8%

Notes: The first column shows the share of public-private patents involving research entities, defined as universities and research institutes, by funding agency. The second column reports the share of start-ups among for-profit participants in public-private patents. Start-up status is based on Ewens and Marx (2024).

Table I.4: Top Startup Innovators in Public-Private Patenting

Company	Patents	Description
Nanosphere, Inc.	21	Biotechnology company specializing in nanoparticle-based molecular diagnostics; developed the Verigene platform for multiplex genetic and infectious disease testing; acquired by Luminex in 2016.
Superior MicroPowders LLC	18	Advanced materials company based in Albuquerque, New Mexico; developed fine powders and inks (e.g., for fuel cells, batteries, and catalysts) via spray-based processes; acquired by Cabot Corporation in 2003.
Pacific Biosciences of California, Inc.	16	Developer of single-molecule, real-time (SMRT) DNA sequencing systems enabling long-read genomics applications in human, plant, and microbial biology.
ARCH Development Corporation	15	University of Chicago–affiliated technology commercialization and incubation arm established in 1986; incubated and managed startups to commercialize research from the university and Argonne National Laboratory.
Molecular Optoelectronics Corp. (MOEC)	10	Research and development firm founded in 1993 in Watervliet, New York; worked on optoelectronic materials and thin-film devices for display, sensing, and photonic applications.

Notes: This table reports the top start-up assignees among public-private patents, based on the start-up classification in Ewens and Marx (2024). The count column reports the number of public-private patents assigned to each start-up in the sample used for the start-up analysis.