Directing Innovation through Ownership Rights: Evidence on Private Returns vs. Social Spillovers

Panos Vorgias*

October 28, 2026

[Latest version]

Abstract

This paper studies how the allocation of intellectual-property rights over publicly funded inventions shapes the direction of private innovation and the diffusion of technological knowledge. I exploit the 1983 Presidential Memorandum on Government Patent Policy, which expanded contractors' ability to retain ownership of future federally funded inventions while leaving the underlying research funding mechanism largely unchanged. Using a matched difference-in-differences design at the firm–technology level, I find that stronger contractor ownership reduced cross-technology spillovers by approximately 25–30 percent, lowered the share of technologically similar follow-on innovation occurring outside the originating firm, and increased private patent market value by 30–40 percent. The findings document a tradeoff between private returns and social spillovers: stronger contractor ownership shifts innovation toward inventions with greater private value but narrower external technological influence. They identify ownership allocation as an important innovation-policy instrument for shaping this tradeoff, affecting not only firms' incentives to innovate but also the social reach of the innovations they choose to pursue.

Keywords: Intellectual property rights, Public R&D, Direction of innovation, Knowledge spillovers, Patent ownership

JEL Classification: O31, O34, O38, H54, D62.

*Boston University, panosv@bu.edu.

I am indebted to Martin Fiszbein, Kevin Lang, and Tarek Hassan, for exceptional guidance and support. I am also grateful to Benjamin Marx and to the participants of the Boston University Empirical Micro Seminar for helpful comments and suggestions. All remaining errors are my own.

1 Introduction

Public R&D policy faces a fundamental question: who should own inventions produced with government funding? Allowing private contractors to retain patent rights may strengthen incentives to commercialize federally funded inventions by increasing private appropriability. But transferring ownership from the government to the firm may also reduce the diffusion of knowledge generated by publicly financed research. Beginning with the Bayh–Dole Act of 1980, U.S. patent policy moved toward granting recipients of federal R&D funding stronger rights to retain ownership of the inventions they produced. Despite decades of debate over these reforms, we know surprisingly little about whether stronger ownership rights change the type of innovation firms choose to pursue. This paper combines newly assembled patent-ownership data with a previously unexplored federal patent-policy reform to provide causal evidence on this question.

I study the 1983 Presidential Memorandum on Government Patent Policy, which sharply expanded contractors' rights to retain title over federally funded inventions. Exploiting this policy change, I estimate how stronger contractor ownership affected the social and private returns to research of federally funded firms. I find that stronger ownership rights redirect innovation toward inventions with greater private returns but weaker external knowledge diffusion. This pattern emerges consistently across citation-based, text-based, and market-based measures of innovation.

Despite the large literature on spillovers, innovation, and growth, the role of patent ownership remains unresolved for at least two reasons. First, there are limited and often inconsistent data on both whether a patent resulted from federal support and who held title to that invention. Although patents are required to disclose government support, noncompliance remains substantial, and the absence of a meaningful monitoring mechanism has led to persistent under-reporting.[1] Even among patents that disclose federal funding, existing sources often report inconsistent ownership status: some list the contractor, others the agency, some both, and some neither. Second, identifying the causal effect of patent rights on innovation outcomes is inherently difficult because it requires a credible counterfactual for the path of innovation under an alternative ownership regime.

This paper addresses both challenges by combining newly assembled data on the ownership status of federally funded patents with a policy change that sharply expanded contractor title rights in 1983. I use administrative-data-based ownership records from Gross and Sampat (2025), which resolve inconsistencies in earlier ownership sources, and combine them with the patent-to-

[1] See, for example, Rai and Sampat (2012).

firm links in Dyevre (2023), which, to my knowledge, provide the most comprehensive publicly available source for linking patents to firms at the scale required for this analysis. The matched data allow me to build the counterfactual group while the 1983 Presidential Memorandum provides the source of plausibly exogenous variation, to study how contractor ownership affects the direction and quality of innovation.

Identifying the effect of contractor patent ownership is challenging for two reasons. First, firms receiving federal R&D support differ systematically from firms that do not, including in scale, research intensity, and innovative capacity. Second, under the pre-reform regime, agencies could retain title selectively, so ownership of federally funded inventions was not randomly assigned. I address the first concern by constructing a matched control group within narrow technology–industry cells based on pre-policy characteristics and outcome trends. I then estimate a difference-in-differences specification at the firm–technology level with firm–technology and technology-industry-year fixed effects, while allowing baseline observables to follow flexible time paths. Firm–technology fixed effects absorb permanent level differences between treated and matched control cells, while technology–industry–year fixed effects absorb annual shocks common to firms within the same technology–industry matching cell. The 1983 Presidential Memorandum is useful in addressing the second concern because it generated a rapid, cross-agency shift toward contractor title, reducing the role of agency-specific patent policies.

A remaining agency-related concern is that they may have retained title before 1983 for inventions with systematically different characteristics. However, the relevant margin for this study excludes inventions subject to persistent national-security or similar exemptions, which were less likely to be affected by the reform. Among the inventions whose ownership status plausibly changed, pre-reform agency retention may have been concentrated among patents with broader applicability or weaker private appropriability, that is, patents with relatively high ex ante spillover potential. Under that selection pattern, selective retention cannot readily explain the post-reform decline in spillovers and would instead tend to attenuate the estimated effect.

The link between contractor ownership and the direction of innovation is established by the prospective basis of patent assignment that the 1983 Reagan Memo brought. Prior to the reform, the decision on a patent's ownership status was taken from the funding agency on a case-by-case basis and only after an invention was made. Post 1983, contractors knew that any resulting inventions would be owned by them already at the project selection stage. Thus, a project's expected return was factored-in in their decisions on research. Critically, the institutional setting

also helps distinguish changes in knowledge diffusion from changes in the nature of innovation itself. The 1983 reform expanded contractors' ownership rights but did not alter the disclosure requirements of the patent system: patented inventions continued to become publicly available upon grant. Thus, the mechanism through which patented knowledge became publicly accessible remained unchanged.

My main outcome is cross-technology knowledge spillovers, measured as forward citations received over the subsequent ten years from inventions in other technology classes. Following the reform, these spillovers from treated firm–technology units decline by about 25 to 30 percent relative to comparable untreated units. Text-based measures of patent significance proposed by Kelly et al. (2021), which compare an invention's textual similarity to other patents, also decline. Moreover, the share of technologically similar follow-on innovation occurring outside the originating firm falls by about 16 to 18 percent. In contrast, the stock-market-based measure of patented inventions developed by Kogan et al. (2017) increases by up to 40 percent. The decline in spillovers is also supported by a synthetic difference-in-differences analysis following Arkhangelsky et al. (2021), which constructs a counterfactual using data-driven weights for control units. The estimates are similar in magnitude to those obtained from the baseline matched difference-in-differences specification.

The average treatment effect masks substantial heterogeneity across firms. Among firms with relatively limited pre-treatment R&D resources, contractor ownership reduces the production of inventions generating the largest knowledge spillovers. By contrast, firms with substantial research capabilities expand patenting primarily in the lower half of the spillover distribution. The evidence is consistent with two complementary mechanisms: a deterioration in the production of high-spillover inventions among resource-constrained firms and a dilution of innovation among large R&D performers through increased production of more privately appropriable research. The market-value results closely track these changes in patent production. Private value increases among low-spillover inventions in both R&D groups, while high-spillover inventions lose substantial value among low-R&D firms. The estimated decline in spillovers is robust to alternative measures of innovation quality, sample windows and approaches to constructing the counterfactual.

This paper makes three contributions. First, it provides causal evidence that the ownership status of publicly funded inventions shapes not only the commercialization of research, but also the direction of research itself. Second, it contributes to the literature on innovation policy and the assignment of intellectual property rights by identifying a tradeoff that has received less attention: stronger contractor control may narrow the private–social return wedge not only by increasing pri-

vate appropriation, but also by reducing an important component of the social return to publicly funded R&D, namely the diffusion of knowledge spillovers. Third, the results show that ownership rights reallocate innovative effort differently depending on firms' absorptive capacity: firms with stronger R&D capabilities are better able to convert ownership rights into privately valuable innovation, whereas firms with weaker capabilities reduce activity precisely in the segment of research that generates the largest external benefits.

Taken together, the results highlight a central tension in innovation policy. Knowledge spillovers constitute one of the principal channels through which innovation generates social returns and, consequently, one of the central justifications for public support of R&D in modern theories of innovation and endogenous growth. Institutions designed to strengthen incentives for commercialization may therefore weaken an important mechanism through which R&D contributes to long-run growth in two ways. First, by redirecting research away from inventions with the largest technological spillovers, they reduce both the social returns to innovation and the opportunities for other firms to benefit from externally generated knowledge. Second, by shifting innovative effort toward more narrowly applicable research, they may gradually narrow firms' effective knowledge base, limiting their ability to recognize, combine, and exploit external technological advances in the future.

Related Literature. This paper relates most directly to the literature studying the private and social returns to innovation. Beginning with Arrow (1962) and extending through Bloom et al. (2013) and Jones and Summers (2022), this literature emphasizes that spillovers from R&D create a wedge between private and social returns. While proposing different methods for measuring spillovers, these studies consistently find that social returns substantially exceed private returns.[2] This literature focuses primarily on corporate R&D and documents both positive technological spillovers and negative effects through product-market rivalry. The current paper extends this literature by identifying an additional institutional margin operating through publicly funded R&D: the allocation of ownership rights over the resulting inventions. I find that contractor ownership not only increases the private return to patented innovation but also reduces an important component of its social return by changing the direction of innovative effort.

The second literature to which this paper contributes estimates the social returns to publicly funded R&D. A large empirical literature focusing on agency-specific settings or case studies shows that government-funded research generates substantial downstream innovation, private R&D, and

[2]For example, Lucking et al. (2019) estimate a social-to-private return ratio of approximately 4-to-1, while Jones and Summers (2022) obtain a ratio of 5-to-1 under their conservative scenario.

commercialization.[3] Recent work exploiting cross-agency variation or aggregate government R&D reaches similar conclusions (Fieldhouse and Mertens, 2023; Dyevre, 2023; Fieldhouse and Mertens, 2026). This literature has primarily studied one margin of public R&D policy: how much research the government funds. I study a second margin: who owns the resulting inventions. By exploiting a federal policy change that altered ownership rights over publicly funded inventions across agencies, I show that the social returns to public R&D depend not only on the amount of public funding but also on the allocation of ownership rights over the resulting inventions.

A third related literature studies the changing direction of private-sector research. Budish et al. (2015) show that private research can be distorted away from long-horizon projects, while Arora et al. (2021) document a shift in corporate innovation away from research and toward development. My findings complement this evidence by showing that stronger ownership rights over publicly funded inventions redirect innovative activity toward outcomes with greater private appropriability but weaker knowledge diffusion. The paper therefore provides causal evidence on an institutional channel through which the balance between research-oriented and development-oriented innovation can shift.

Finally, the paper relates to the literature on patents as an intellectual property strategy and their effects on research investment. As emphasized by Williams (2017), credible causal evidence on how patent rights affect innovation remains limited, largely because it is difficult to observe a convincing counterfactual path of technological development.[4] My paper addresses this challenge by exploiting a sharp change in federal patent policy and combining it with a matched difference-in-differences design. This provides causal evidence that strengthening firms' ownership rights over publicly funded inventions affects not only the private value of innovation, but also its social

[3] Azoulay et al. (2019) study NIH funding and find that an additional $10 million generates three to four additional patents, roughly two of which are in different technological areas. Myers and Lanahan (2022) study the DOE and find approximately three additional patents for each patent directly funded. Moretti et al. (2025) estimate that a 10% increase in government-funded R&D raises privately funded R&D by 5–6%. Using U.S. Air Force SBIR awards, Howell et al. (2025) find that receiving an award increases subsequent venture-capital investment by 128%, patenting by 79%, and high-originality patenting by 194%.

[4] An important exception is Lerner (2009), who studies stronger patent property rights at the country level and finds no effect on R&D spending. A separate branch of the literature studies the effect of patent litigation on subsequent innovation; see, for example, Galasso and Schankerman (2018). From the management literature, recent work by Masclans et al. (2025) constructs a measure of the "commercial potential" of science using text-processing algorithms applied to scientific abstracts. For comparative studies of firms' intellectual property strategies, see, for example, Mezzanotti and Simcoe (2023).

value through knowledge spillovers. More broadly, the results suggest that patent policy should be evaluated not only in terms of whether it strengthens incentives to innovate, but also in terms of how it changes the composition of innovation between privately appropriable and socially diffusive research.

2 Background and Conceptual Framework

2.1 Historical and Institutional Background

The U.S. debate over patent rights in publicly funded research emerged from the postwar challenge of reconversion from war to peace. As World War II ended, policymakers faced two closely related questions: how to convert wartime scientific and industrial capacity into peacetime prosperity, and what role the federal government should play in supporting innovation. Wartime experience had demonstrated the value of large-scale public research investment, strengthening the case for a continuing federal role. Yet the emerging Cold War strengthened opposition to centralized control over science and a more egalitarian distribution of funding, increasingly associated with state planning rather than the market-oriented model policymakers sought to preserve. The debate was therefore not whether government should support research, but how that support should be organized and who should benefit from the resulting discoveries.

This tension is reflected in the *Bush–Kilgore debate*.[5] Bush emphasized scientific autonomy and competitive allocation of research funds, whereas Kilgore favored stronger public oversight and was more concerned about private appropriation of publicly financed discoveries.[6] Both positions reflected the same historical reality: wartime experience had made a permanent federal research role difficult to deny.

The resulting compromise departed fundamentally from prewar policy. Civilian research had largely been conducted in industrial laboratories and universities supported by industry and philanthropy, with federal involvement limited by concerns over scientific freedom. Postwar policymakers abandoned this reluctance toward federal support but did not centralize the research system itself.

[5] Vannevar Bush, director of the Office of Scientific Research and Development (OSRD), advocated a merit-based system insulated from day-to-day political influence in *Science, the Endless Frontier* (1945), widely regarded as the blueprint for the postwar U.S. innovation system. Senator Harley Kilgore favored public oversight and expressed concerns about the concentration of research funding among a few universities and large corporations.

[6] The NSF's official history identifies patent ownership, administrative control, and the balance between basic and applied research as the principal points of contention.

Instead, individual agencies retained substantial discretion over funding, contracting, and patent policy. This agency-by-agency approach shaped innovation policy for three decades and became central to the evolution of federal patent rights.

A central question followed: who should own federally funded inventions? The immediate post-war answer favored government ownership, reflecting concerns that exclusive contractor rights could amount to subsidies benefiting selected firms.[7] Over time, however, dissatisfaction grew with fragmented agency policies. By the late 1970s, policymakers increasingly argued that federally funded inventions were not being developed or commercialized efficiently where agencies retained title and licensed inventions nonexclusively. A 1998 Government Accountability Office review summarizes this contemporary diagnosis: before Bayh–Dole there was "no uniform federal policy," and by 1980 fewer than five percent of roughly 28,000 federally owned patents had been licensed, compared with substantially higher rates where contractors retained title.

The movement toward contractor ownership was gradual rather than abrupt. Kennedy's 1963 Statement of Government Patent Policy established the first government-wide ownership framework. Government ownership remained the general presumption, but agencies could grant contractors greater rights to facilitate practical application. President Nixon's 1971 revision extended this commercialization-oriented approach by giving agencies broader discretion to permit stronger contractor rights where these were expected to encourage utilization.[8]

Bayh–Dole marked the next major step. Enacted in 1980, it reversed the government-ownership presumption for federally funded inventions made by universities, nonprofit organizations, and small businesses. Its statement of policy is revealing: Congress presented patent rights as tools for promoting commercialization and collaboration with industry, but also employment and productivity growth.[9] The Act established the first government-wide statutory framework governing ownership of federally funded inventions. Rather than relying on agency-specific policies, it created common ownership rules for all research performers within its scope, substantially reducing institutional fragmentation.

[7] The *National Patent Planning Commission* (1945) advocated government ownership with broad commercial licensing. The Attorney General's 1947 Report to the President similarly advocated uniform federal policy and government ownership to *"avoid government favoritism toward particular firms"* (Eisenberg, 1996).

[8] Nixon's memorandum sought *"additional authority to permit contractors to obtain greater rights to inventions where necessary to achieve utilization,"* alongside improved evaluation of federal patent policy (August 23, 1971).

[9] 35 U.S.C. §200 calls for using patents "to promote the utilization" of federally supported inventions and "to promote the commercialization and public availability" of U.S. inventions, while preserving government rights against nonuse or unreasonable use.

Equally important, Bayh–Dole standardized prospective ownership arrangements. Ownership rights became part of the research contract, reducing reliance on post-invention negotiations. Contractors disclose subject inventions to the funding agency, elect whether to retain title within statutory deadlines, and pursue patent protection under common rules. From an economic perspective, this prospective allocation shapes incentives before research begins. Greater certainty over appropriation can influence research planning, project selection, organizational investment, and ultimately the inventions firms pursue. Patent policy therefore affects not only commercialization but also the direction of innovative activity.

Bayh–Dole, however, did not establish a uniform regime across all federally funded research. Its provisions applied to universities, nonprofit organizations, and small businesses, leaving large industrial contractors under earlier agency-specific rules. That distinction changed with President Reagan's February 18, 1983 Memorandum on Government Patent Policy, which extended its contractor-oriented philosophy to all federally funded performers.[10] Although contractor ownership had expanded over decades, the memorandum represented a distinct institutional break, extending a common regime to nearly all federally funded research.

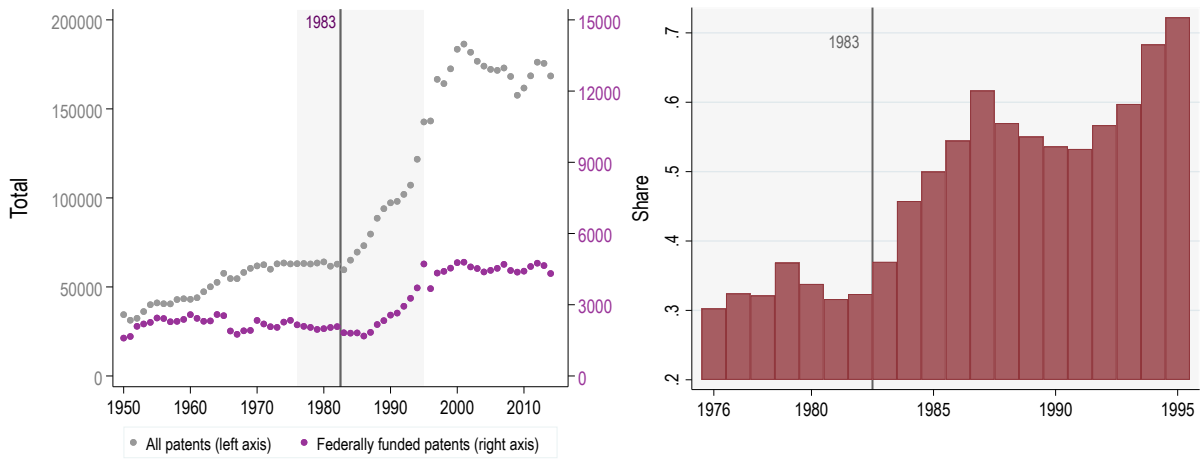
Figure 1 places the reform in historical perspective. Panel (a) shows the evolution of federally funded patents along total patenting. The shaded area corresponds to my analysis window, 1976-1995. Panel (b) reports the share among funded patents that is owned by contractors over 1976-1995. The figures illustrate the transition to contractor ownership and the central role of the 1983 policy change in this process. This expansion beyond the groups originally covered by Bayh–Dole provides the basis for the paper's empirical strategy.

Taken together, postwar U.S. innovation policy increasingly combined a commercialization mandate with stronger contractor property rights. The underlying premise was that allowing firms to appropriate a larger share of returns from federally funded inventions would encourage commercialization and crowd in private research investment, translating public research into productivity growth. Yet stronger rights may redirect effort away from projects with the largest external benefits, reducing knowledge diffusion and weakening spillovers, a principal engine of productivity growth. Whether contractor ownership promoted commercialization without reducing knowledge diffusion remained largely an empirical question. This paper examines that question.

[10] The 1984 Trademark Clarification Act subsequently codified this change by replacing the statutory reference to the 1971 Nixon policy.

Figure 1: Patent Funding & Contractor-ownership Timeline

(a) Federally-funded patents (total) (b) Contractor-owned funded patents (share)



Notes: Panel (a) shows the annual number of U.S. patents and federally funded patents over 1950–2019. The shaded area marks the analysis period (1976–1995). Panel (b) reports the share of federally funded patents owned by contractors over 1976–1995. The vertical line marks the 1983 Presidential Memorandum.

2.2 Conceptual Framework

Technological knowledge has long been understood as a public good because it is nonrival and only imperfectly excludable. Its benefits were formally expressed in Romer (1986), who argued that knowledge spillovers from investment can generate increasing returns at the aggregate level, helping launch the new growth literature. Modern endogenous growth theory still places technological progress at the center of long-run growth, with own-R&D and knowledge spillovers as its main engine (Romer, 1990; Akcigit et al., 2021).

The semi-public nature of knowledge creates a fundamental tension. From the perspective of social welfare, imperfect appropriability implies that the social return to innovation exceeds the private return, creating the well-known private-social wedge (Arrow, 1962). From the firm's perspective, however, outgoing spillovers is a negative externality benefiting rivals and reducing the profitability of innovation through business stealing (Aghion, 1992; Bloom et al., 2013; Lucking et al., 2019). Although these approaches differ in their interpretation of spillovers, both ultimately originate from the same underlying tension between knowledge as a semi-public good and firms' incentives to privately appropriate its returns.

The same tension is expressed in theories of firms' own research investment decisions. Incom-

ing spillovers complement own-R&D[11] and own-R&D strengthens firms' absorptive capacity of spillovers (Cohen et al., 1990), while outgoing spillovers weaken incentives to invest by facilitating competitors' innovation (Arora et al., 2021). Again, the contrasting role assigned to spillovers reflects the same underlying conflict between socially valuable knowledge diffusion and private appropriation.

The common feature of these theories is that they characterize the consequences of imperfect appropriability within a given institutional environment. As ownership in corporate-R&D is typically determined before research objectives are set, it affects the expected private return to alternative research trajectories at the stage of project selection already. If stronger ownership rights disproportionately increase the returns to projects whose benefits can be retained within the firm, they may shift effort toward those projects and away from research that generates more broadly usable knowledge. This paper examines this margin through the allocation of ownership rights over publicly funded inventions.

Publicly funded patents provide a natural setting for studying this institutional margin because a change in the degree of appropriation is clearly identifiable while the source of funding is held fixed. The 1983 Presidential Memorandum is particularly well suited to examining this mechanism because it changed the allocation of ownership rights prospectively. Firms could therefore anticipate greater control over successful inventions before undertaking research, allowing ownership expectations to influence project selection, commercialization incentives, and ultimately the type of innovation produced.

This framework highlights three potential consequences of allocating ownership rights over publicly funded inventions.

Implication 1. (*The Trade-off*): Stronger ownership rights may increase private returns while reducing knowledge diffusion if they shift research toward projects whose benefits are more readily retained within the firm. Greater appropriability may therefore involve a trade-off between private appropriation and the diffusion of socially valuable knowledge. The empirical analysis examines whether greater contractor ownership is accompanied by higher private value and lower knowledge diffusion.

Implication 2. (*Ownership as a Policy Instrument*): The private–social wedge may itself be endogenous to institutional design. Although ownership rights seek to narrow this wedge by increasing private returns, they may also affect the external benefits generated by the resulting innovation. A

[11] several studies document strong complementarities with own-R&D. see, for example, Aghion and Jaravel (2015)

narrower wedge may therefore reflect higher private returns, lower social returns, or a combination of the two, with different welfare implications. Ownership should consequently be viewed as a policy instrument that can shape both research incentives and the diffusion of knowledge. An increase in private value alone does not establish whether the resulting innovation generates greater social benefits.

Implication 3. (*Dynamic Implications*): A further potential consequence follows from absorptive-capacity theory. If ownership affects the direction of innovative activity, it may also shape the knowledge base accumulated within firms and, consequently, their future ability to recognize and exploit external knowledge. The effects of ownership rights may therefore extend beyond the contemporaneous allocation of innovative effort to influence the dynamic evolution of firms' innovative capabilities. This dynamic mechanism provides a broader implication of the framework rather than an effect directly estimated in the empirical analysis.

3 Data and Descriptive Evidence

This paper combines complementary datasets to study how ownership of federally funded inventions affects firms' innovation strategies and the social value of their research. It links patent ownership, government funding, citations, textual characteristics, and firm financials to construct a firm–technology panel covering U.S. corporate innovation from 1950 to 2019. The empirical analysis focuses on 1976–1995, while the full 1950–2019 sample is retained for descriptive evidence on broader trends. Appendix A details the data sources and sample construction.

3.1 Data

The starting point is the patent-level ownership and government funding data assembled by Gross and Sampat (2025), the most comprehensive publicly available record of federally funded U.S. patents. Relative to earlier sources, their data substantially improve ownership classification using newly uncovered administrative records from the United States Patent and Trademark Office (USPTO) archives. These identify whether title remained with the government or was assigned to the contractor and, for a subset of patents, the funding agency. I combine these records with the patent-to-firm linkage developed by Dyevre (2023), which matches U.S. patents to publicly listed Compustat firms. This linkage allows me to identify firms exposed to federal research funding and combine patent-level outcomes with firm financial characteristics.

Citation data combine the historical series of Fleming et al. (2019) for patents filed before 1976 with USPTO records for patents filed from 1976 onward. Firm characteristics, including employment, sales, R&D expenditures, industry classification, and firm age, are obtained from Compustat. I also merge three complementary measures of innovation quality: text-based patent impact from Kelly et al. (2021), patent creativity from Kalyani (2024), and patent-level market value from Kogan et al. (2017).

The analysis follows patents from invention to the firms responsible for the underlying research. Throughout the paper, patents are assigned to their ultimate corporate owner at filing using the filing-date ownership version of the Dyevre (2023) linkage. This choice aligns innovation outcomes with the firms that made the original R&D investment and faced the policy incentives. The objective is to study research decisions when inventions are created, rather than how later mergers, acquisitions, or corporate reorganizations redistribute ownership.

The final matched dataset contains approximately two million patents assigned to Compustat firms among the 6.7 million patents filed in total over 1950–2019. Within the estimation window of 1976-1995, approximately 591,000 patents are matched to Compustat firms, of which roughly 248,000 enter the analysis sample after accounting for data availability on the control variables. Patents are aggregated to the firm–technology–year level, using three-digit Cooperative Patent Classification (CPC) classes.[12] This aggregation reduces idiosyncratic patent-level variation while closely approximating how funded research projects are organized and selected. Agencies fund projects within technological areas rather than individual patents, whose number and timing are unknown ex ante. Although projects are not observed, the firm–technology cell provides the closest observable counterpart to a technological research program within a firm. Appendix Table A1 summarizes the composition of the full patent database and the analysis sample.

The focus on publicly listed firms is standard in the innovation literature. These firms provide consistent longitudinal accounting data over several decades and account for a disproportionate share of U.S. innovation, making them well suited to studying how changes in intellectual-property rights affect corporate research.

The ownership classifications in Gross and Sampat (2025) form the basis of the estimation sample. Their records distinguish federally funded inventions whose title remained with the government from those assigned to contractors subject to a government license. This distinction is central because the 1983 reform expanded contractor ownership, changing intellectual-property rights while

[12] Technology classes correspond to the CPC section letter and two-digit class (for example, G02 for optics).

leaving the research funding mechanism largely unchanged.

A small fraction of funded patents have indeterminate ownership because the records do not permit unambiguous classification. Before 1983, these observations are very few and are excluded following Gross and Sampat (2025). After 1983, correspondence with the authors indicates that indeterminate patents assigned to business firms should, for practical purposes, be treated as contractor-owned. The baseline analysis therefore classifies them accordingly. Only 19 of these patents are assigned to non-firm organizations, and the results are robust to excluding them.

Because universities and nonprofits faced a different ownership regime following Bayh–Dole in 1980, I classify assignees using disambiguated PatentsView names. This distinguishes business firms from universities, government agencies, nonprofits, and mixed organizational forms. The baseline assigns patents by filing-date corporate ownership. As a robustness exercise, I exclude patents involving universities, nonprofits, government agencies, or mixed ownership. This verifies that the results are not driven by organizations already facing different intellectual-property incentives before 1983 or projects jointly directed by non-firm actors.

The estimating sample covers patents filed in 1976–1995, aggregated to the firm–technology–year level. The lower bound reflects more complete patent and assignee information with the start of PatentsView in 1976. It also avoids the structural changes of the early 1970s and provides a stable pre-reform period for defining treatment and pre-treatment characteristics. The upper bound precedes changes in patent eligibility and patenting behavior associated with software and business-method patents in the mid-to-late 1990s, which altered the composition of inventions. This window therefore maintains a relatively stable institutional and technological environment with consistent data coverage around the reform.

With the firm–technology panel constructed, I turn to the outcomes used to measure the external and private returns to innovation. The primary outcomes cover utility patents and measure spillovers through forward citations, the standard proxy for knowledge diffusion because they capture how subsequent inventions build on earlier discoveries. The baseline measures cross-technology citations, excluding those from the patent's own technology class. Aggregating these flows to the firm–technology–year level captures how research within a technological area generates knowledge that diffuses beyond its immediate field.

To verify that the effects are not specific to citations, I use the text-based measures of Kelly et al. (2021), which draw on a different source of variation: similarity across patent texts. Alongside their publicly available patent-level measures, I use the authors' proprietary patent-pair data with

similarity above 50%. Focusing on these highly similar pairs better aligns the measure with the paper's objective of tracing meaningful knowledge diffusion rather than broad textual overlap. They also reduce the mechanical tendency of aggregate similarity measures to classify prolific patentees' inventions as influential simply because they produce many subsequent patents. The pair-level data distinguish knowledge reuse within firms from diffusion to external innovators and underpin the external diffusion measures analyzed in the paper.

Finally, I measure the private value of innovation using the patent-level market valuation estimates of Kogan et al. (2017), which infer patent value from stock-market reactions to grants. Together, these complementary outcomes allow the analysis to distinguish whether stronger intellectual-property rights shift innovation toward inventions that are privately more valuable, socially more valuable, or both.

3.2 Descriptive Evidence

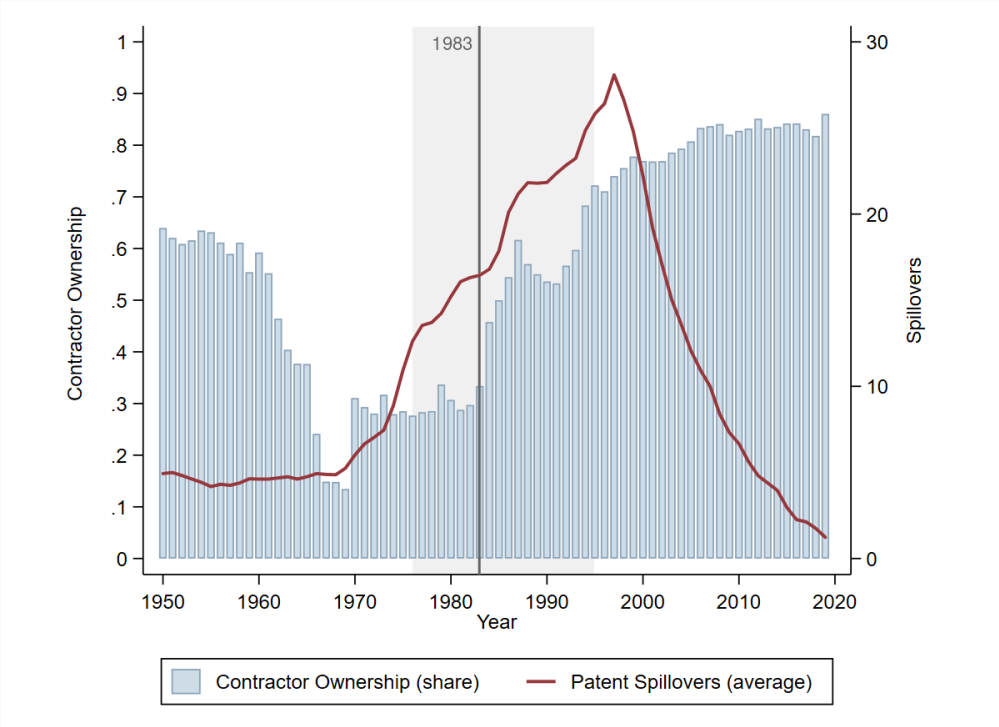
The descriptive evidence reveals a striking pattern: federally funded inventions generate substantial knowledge spillovers, but among these inventions, contractor ownership is associated with significantly narrower knowledge diffusion. This section documents this relationship and motivates the identification strategy developed in the following section.

Knowledge diffusion is measured using cross-technology citations. For each patent, I count forward citations received from outside its CPC technology class over the subsequent ten years.[13] Cross-class citations capture diffusion beyond an invention's immediate technological domain and provide a proxy for its broader external relevance.

Figure 2 provides a first indication that aggregate changes in contractor ownership and knowledge diffusion do not move together mechanically. It plots average spillovers among federally funded patents alongside the share under contractor ownership. Two patterns emerge. First, contractor ownership becomes substantially less prevalent in the late 1960s and 1970s before rising steadily after the early-1980s policy changes. Second, spillovers remain relatively stable through the 1950s and 1960s, increase sharply from the early 1970s through the late 1990s, and decline for recent cohorts. The strongest growth in diffusion occurs while contractor ownership remains relatively uncommon and changes little. This divergence highlights the difficulty of inferring ownership effects from aggregate time-series variation alone. Changes in diffusion may instead reflect the cumulative effects

[13] The ten-year window mitigates truncation concerns for recent patent cohorts, which have had less time to accumulate citations. Results are quantitatively and qualitatively unchanged using twenty- or thirty-year windows.

Figure 2: Knowledge Diffusion and Patent Ownership among Federally Funded Patents



of past research investments, evolving technological opportunities, or shifts in the composition of federally funded innovation.

To examine this relationship more directly, I estimate patent-level Poisson regressions relating spillovers to federal funding and contractor ownership under alternative controls and fixed effects:

$$\text{Spillovers}_i = \beta_0 + \beta_1 \text{Funded}_i + \beta_2 \text{ContractorOwnership}_i + \mathbf{X}_i'\gamma + \boldsymbol{\delta}_i + \varepsilon_i, \tag{1}$$

where i indexes patents, $\mathbf{X}_i$ contains firm- and patent-level controls, and $\boldsymbol{\delta}_i$ denotes the fixed effects. Across columns, these include filing-year, CPC technology-class, 2-digit SIC industry, CPC-by-filing-year, funding-agency, and agency-by-CPC-by-filing-year fixed effects. The increasingly demanding structures allow comparisons among patents exposed to similar technological and institutional environments. CPC-by-year effects compare patents within the same technology class and filing year, absorbing changes in citation practices and technological opportunities across fields over time. Agency-by-CPC-by-year effects further restrict comparisons to funded patents from the same agency.

Firm controls include labor productivity, measured as sales per employee, and R&D productivity, measured as R&D expenditures relative to patent output. Patent controls include total ten-year forward citations. Standard errors are clustered at the CPC-by-filing-year level to account for common shocks and correlated citation behavior within technological fields over time.

Table 1: Patent Ownership and Knowledge Diffusion, 1950–2019

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Government-funded	0.358***	0.367***	0.287***	0.344***			
	(0.044)	(0.050)	(0.041)	(0.047)			
Contractor-owned	-0.246***	-0.257***	-0.162***	-0.214***	-0.115***	-0.102**	-0.062
	(0.047)	(0.054)	(0.045)	(0.050)	(0.041)	(0.046)	(0.049)
Patent Controls	Y	Y	Y	Y	Y	Y	Y
Firm Controls	Y	Y	Y	Y	Y	Y	Y
Year FE	✓	✓			✓		
Technology FE	✓	✓			✓		
Industry FE		✓		✓			
Tech-Year FE			✓	✓		✓	
Agency FE					✓	✓	
Agency-Tech-Year FE							✓
SE	cpc-year	cpc-year	cpc-year	cpc-year	cpc-year	cpc-year	cpc-year
N	1,990,878	1,990,854	1,990,220	1,990,919	25,839	24,854	23,128

Notes: Patent-level Poisson pseudo-maximum likelihood (PPML) estimates. The outcome counts cross-CPC forward citations within ten years of filing. Firm controls include labor productivity (sales per employee) and R&D productivity (R&D expenditures per patent filed); patent controls include total ten-year forward citations. Standard errors are clustered by CPC–filing year. Year fixed effects refer to filing year. (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$)

Table 1 documents three stylized facts. First, federally funded patents generate substantially more cross-technology diffusion than non-funded patents. Across specifications, they receive approximately 32–43 percent more spillover citations, consistent with a large literature showing that publicly funded research generates substantial social returns and knowledge spillovers. Second, conditional on federal funding, contractor ownership is associated with significantly lower diffusion: approximately 17–23 percent fewer cross-technology citations in the first four specifications. This relationship remains remarkably stable as increasingly demanding industry and technology fixed effects are introduced, indicating that it is not driven simply by differences in citation intensity

across fields or industries.

Finally, restricting the sample to funded patents with identifiable agencies yields smaller, though generally still statistically significant, estimates. This attenuation primarily reflects sample composition rather than agency fixed effects. Re-estimating column (4) on the sample used in column (5) yields a contractor-ownership coefficient statistically indistinguishable from zero. Holding the sample fixed, introducing agency effects makes the relationship more negative. The smaller coefficients in columns (5)–(7) therefore reflect the distinct sample for which agency information is available, rather than agency differences explaining the negative association. Nevertheless, the negative relationship between contractor ownership and subsequent diffusion remains evident. Appendix Table C.1 further shows that the pattern is robust to excluding self-citations. The association becomes somewhat stronger in specifications with agency controls, bringing estimates closer to those in the broader sample.

These descriptive relationships should not be interpreted causally. Funding agencies may systematically assign different ownership arrangements across firms, technologies, or research projects, and these underlying differences may themselves affect subsequent diffusion. The following section therefore exploits the 1983 reform in government patent policy to isolate the causal effect of contractor ownership on knowledge diffusion.

4 Empirical Strategy

4.1 Identification

The objective is to estimate the causal effect of stronger contractor ownership rights over federally funded inventions on subsequent innovation outcomes. The central empirical challenge is that patent ownership is not randomly assigned. Firms performing federally funded research differ systematically from other firms, and federal agencies may direct funding toward projects with particular technological, scientific, or commercial characteristics. I address these concerns by exploiting the 1983 Reagan Memorandum on Government Patent Policy, which established a government-wide presumption that contractors should retain title to inventions arising from federally funded research.

Identification comes from a difference-in-differences design implemented at the firm–technology–year level, where a technology is defined by a CPC class. Federal research support is allocated to specific technological programs rather than uniformly across firms. Consequently, exposure to the ownership reform varies across technologies within the same firm, making the firm–technology cell

the appropriate unit of analysis.

Accordingly, a firm–technology cell is classified as treated if the firm received at least one federally funded patent in that technology before 1983. The comparison group consists of firm–technology cells belonging to firms that never receive federal funding during the sample period, while firm–technology cells first exposed to federal funding after 1983 are excluded from the estimation sample. Treatment status is therefore determined entirely by pre-reform exposure to federally funded research. Using never-treated controls ensures that the comparison group does not experience a contemporaneous change in the degree of appropriability during the analysis period.

To improve comparability between treated and untreated firm–technology cells, I construct the baseline sample using a two-step matching process combined with a common-support restriction. I first require exact matches within 2-digit SIC industry and CPC technology class and then match each treated firm–technology cell to the untreated cell with the nearest estimated propensity score, allowing controls to be matched with replacement. The propensity score is estimated using pre-treatment firm productivity, employment, firm age, and the pre-treatment level and trend of the spillover outcome.[14] Including the pre-treatment history of the outcome in the matching procedure helps ensure that treated cells are compared with controls that were similar not only in observable firm characteristics but also in their spillover dynamics prior to the reform.

Inspection of the propensity-score distributions motivates the strongest common support between treated and untreated cells is concentrated between propensity scores of 0.40 and 0.80. I therefore restrict the donor pool to untreated firm–technology cells within this interval before matching. This restriction limits reliance on controls with substantially different predicted probabilities of treatment while preserving the population of treated firm–technology cells. I retain this donor-pool restriction as the baseline to preserve broader coverage of treated firm–technology cells and provide a common reference sample for subsequent analyses. Comparability is assessed through pre-treatment covariate balance and outcome dynamics, while additional caliper restrictions are examined as robustness checks. Appendix A.2 provides additional details on the matching procedure and propensity-score overlap.[15] The baseline estimating equation is

[14] The use of pre-treatment spillover levels and trends is consistent with the emphasis in Imbens (2015) on incorporating lagged outcomes and other covariates expected to predict the outcome of interest.

[15] See Appendix Figure A1 for the propensity-score distributions of treated and matched control cells.

$$Y_{fkt} = \beta \left(Treated_{fk} \times Post_t \right) + \phi_{fk} + \lambda_{ks(f)t} + \mathbf{X}'_{fkt}\boldsymbol{\gamma} + \varepsilon_{fkt}, \tag{2}$$

where f, k, and t index firms, CPC technology classes, and filing years, respectively, and $s(f)$ denotes the two-digit SIC industry of firm f. $Treated_{fk}$ equals one for firm–technology cells exposed to federally funded research before 1983, while $Post_t$ equals one for filing years from 1983 onward. The specification includes firm–technology fixed effects, ϕ_{fk}, and CPC-by-SIC2-by-year fixed effects, $\lambda_{ks(f)t}$. The vector $\mathbf{X}_{fkt}$ contains flexible time paths constructed by interacting pre-treatment quartile indicators for firm size, R&D intensity, technological breadth, patenting activity, innovation quality, and firm age with filing-year fixed effects. These interactions allow firm–technology cells with different pre-treatment characteristics to follow different subsequent trajectories rather than imposing a common relationship between these characteristics and the outcome over time.

The coefficient of interest, β, captures the differential post-reform change in previously federally funded firm–technology cells relative to comparable never-funded cells, conditional on the fixed effects and controls. Standard errors are clustered at the firm level, allowing for arbitrary correlation in outcomes across technologies and years within firms.

The identifying assumption is that, absent the 1983 change in ownership rules, treated and comparison firm–technology cells would have followed parallel trends conditional on the included fixed effects and flexible control paths. Two threats to this assumption merit particular attention.

The first concern is selection into federally funded research. Firms receiving federal support differ systematically from never-funded firms in characteristics that may also influence subsequent innovation outcomes. The common-support restriction and matching procedure mitigate this concern by comparing treated cells with observationally more similar controls within the same industry and technology. Importantly, the design does not require matching to eliminate all pre-treatment differences in observable characteristics. Firm–technology fixed effects absorb persistent level differences across cells, while the flexible control paths allow remaining differences in pre-treatment characteristics to be associated with different subsequent outcome trajectories. Section 4.2 examines the resulting covariate balance and pre-treatment outcome dynamics directly.

The second concern is that the early 1980s coincided with broader changes in U.S. innovation policy, including the implementation of Bayh–Dole and a general strengthening of patent rights.

The design isolates differential exposure to a reform that specifically altered ownership rights over federally funded inventions. Matching restricts comparisons to firm–technology units operating within the same industry and technology, while CPC-by-SIC2-by-year fixed effects absorb shocks common to firms operating in the same technology and industry in a given year. Identification therefore comes from differences in the evolution of treated and control units within the same technology–industry–year cell, after accounting for persistent firm–technology differences and flexible time paths based on pre-treatment characteristics.

The timing of the reform further supports its interpretation as a discrete policy shock rather than the continuation of a widely anticipated trend. In June 1980, the Senate rejected by a vote of 60–34 an amendment that would have extended Bayh–Dole's ownership provisions to large federal contractors.[16] The subsequent expansion of contractor ownership through executive action therefore followed a policy path that Congress had rejected, reducing concerns that firms anticipated the reform well in advance. Consistent with this interpretation, Section 3.2 shows little evidence of anticipatory changes in federally funded patenting or contractor ownership.

4.2 Validation

The identifying assumption requires that, absent the 1983 reform, spillovers from treated and control firm–technology cells would have evolved similarly. I assess its credibility in three steps. I first present unmatched outcome means to document the underlying spillover dynamics before any adjustment. I then examine covariate balance in the matched sample and, finally, estimate an event-study specification that directly evaluates differential pre-treatment trends conditional on the baseline specification.

4.2.1 Unmatched Outcome Dynamics

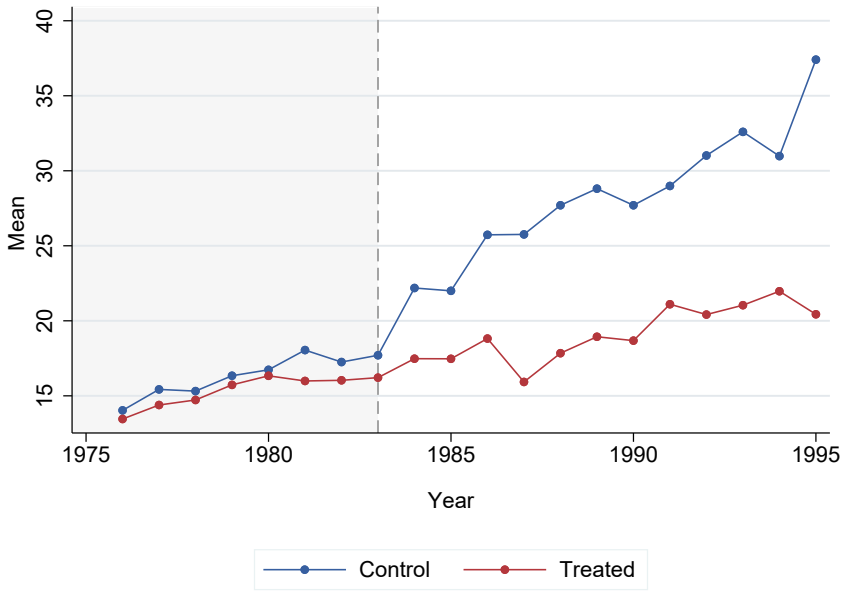
Figure 3 compares average spillovers among treated firm–technology cells and the full set of eligible, never-funded control cells. These are unconditional means and therefore incorporate the underlying differences in the composition, scale, and patenting activity of the two groups. Control cells exhibit higher average spillovers throughout the sample. At the firm–technology level, this ordering reflects

[16] See Letter from Sen. Gaylord Nelson, Chairman, Senate Select Committee on Small Business, to Rep. Robert W. Kastenmeier, Chairman, Subcommittee on Courts, Civil Liberties, and the Administration of Justice of the House Committee on the Judiciary (June 1980), reprinted in *House Hearings on President's Industrial Innovation Program*, pp. 877–878. I became aware of this source through Eisenberg (1996).

both baseline selection into federal funding and the different composition of treated and never-funded firms.[17] The level difference motivates the matching procedure and the inclusion of firm–technology fixed effects rather than constituting evidence about the effect of the reform.

The relevant feature for the difference-in-differences design is the evolution of the gap. Following the 1983 reform, spillovers from treated cells decline relative to those from never-funded cells. Because the figure is entirely unadjusted, it serves as a descriptive benchmark. The balance and event-study evidence below assess whether the post-reform divergence remains after constructing a more comparable control group and conditioning on the baseline specification.

Figure 3: Unmatched Spillover Dynamics by Treatment Status



Notes: Unconditional mean spillovers for treated and all eligible never-funded firm–technology cells over the estimation period. Spillovers count cross-technology forward citations over ten years. No matching weights, fixed effects, or flexible controls are applied. The vertical line marks the 1983 reform.

4.2.2 Covariate Balance

Table 2 reports covariate balance between treated and matched control firm–technology cells. Matching produces close balance in pre-treatment spillover levels and trends, the dimensions most directly related to the identifying variation in the outcome. Productivity is likewise well balanced. The remaining imbalance is primarily one of firm scale, reflected also in patenting activity. Patent

[17] The earlier descriptive analysis compares individual funded and non-funded patents, whereas Figure 3 gives equal weight to firm–technology cells. The two comparisons therefore use different units and weights and describe different margins of the data.

creativity, the table's measure of pre-treatment innovation quality, exhibits a considerably smaller difference.

Table 2: Covariate Balance

	Levels					Trends				
	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.
Spillovers	15.80 (24.56)	14.48 (20.28)	16.37 (23.77)	0.249	0.058	0.51 (7.35)	0.02 (6.06)	-0.08 (16.08)	0.441	0.073
# patents filed	10.77 (15.50)	3.26 (5.55)	1.89 (2.39)	$< 0.001^{***}$	0.645	-0.016 (1.58)	-0.407 (1.66)	-0.02 (0.61)	0.001^{***}	0.241
Patent creativity	1.15 (0.40)	1.02 (0.71)	1.09 (0.71)	0.650	0.220	0.005 (0.22)	-0.012 (0.28)	-0.02 (0.36)	0.067^{*}	0.069
Employees	110.16 (150.74)	22.23 (21.86)	23.32 (82.87)	$< 0.001^{***}$	0.816	0.32 (6.27)	0.22 (3.50)	0.27 (2.18)	0.648	0.020
TFP	114.44 (120.21)	116.48 (103.23)	106.80 (136.02)	0.769	-0.018	0.67 (8.71)	0.52 (5.61)	0.35 (10.29)	0.950	0.020
R&D intensity	3.72 (19.00)	4.69 (10.28)	6.30 (32.22)	0.078^{*}	-0.064	0.217 (2.12)	0.657 (4.31)	0.27 (3.16)	0.126	-0.130
Age (in 1983)	3.77 (3.22)	4.98 (5.80)	4.50 (5.53)	$< 0.001^{***}$	-0.257	—	—	—	—	—
# tech classes	22.29 (11.26)	5.23 (3.81)	4.66 (4.53)	$< 0.001^{***}$	2.031	-0.169 (1.15)	-0.032 (0.81)	0.12 (0.89)	$< 0.001^{***}$	-0.138
N	9,118	4,763	14,008			9,118	4,763	14,008		

Notes: The table reports pre-treatment means for treated firm–technology cells, matched control firm–technology cells, and all non-treated firm–technology cells. Standard deviations are reported in parentheses. Standardized differences compare treated and matched control observations. Patent creativity is measured using the creativity index developed by Kalyani (2024). Employees are reported in thousands, and R&D intensity is R&D expenditure, in millions of dollars, per CPC class. Firm age is measured as the number of years since the earliest year in which the firm can be identified in the data, defined as the minimum of its IPO year, first patent filing year, and first appearance in Compustat. The number of technology classes is the number of distinct CPC classes in which a firm filed patents during a given year. Covariates with standardized differences above 0.10 in either levels or pre-treatment trends enter the baseline specification through interactions of their pre-treatment quartile indicators with filing-year fixed effects.

Despite this smaller imbalance, patent creativity warrants separate attention because it allows the specification to account for pre-existing differences in technological novelty across firm–technology cells, a dimension of patent quality directly relevant to this analysis. Developed by Kalyani (2024), the measure identifies novel combinations of text bi-grams in patent documents

relative to the prior patent corpus. It thus captures a dimension of innovation distinct from citations and full-text similarity measures. Because it is constructed entirely from patent text and the prior corpus, without using post-filing citation information, it captures pre-treatment innovation quality without mechanically incorporating the forward-citation variation used to construct the spillover outcome.

Firm–technology fixed effects absorb persistent level differences across cells. The remaining concern is that firms with different pre-treatment characteristics might have followed different trajectories after 1983 even absent the reform. I therefore allow characteristics with standardized differences above 0.10 to generate flexible differential time paths through interactions of their pre-treatment quartile indicators with filing-year fixed effects.

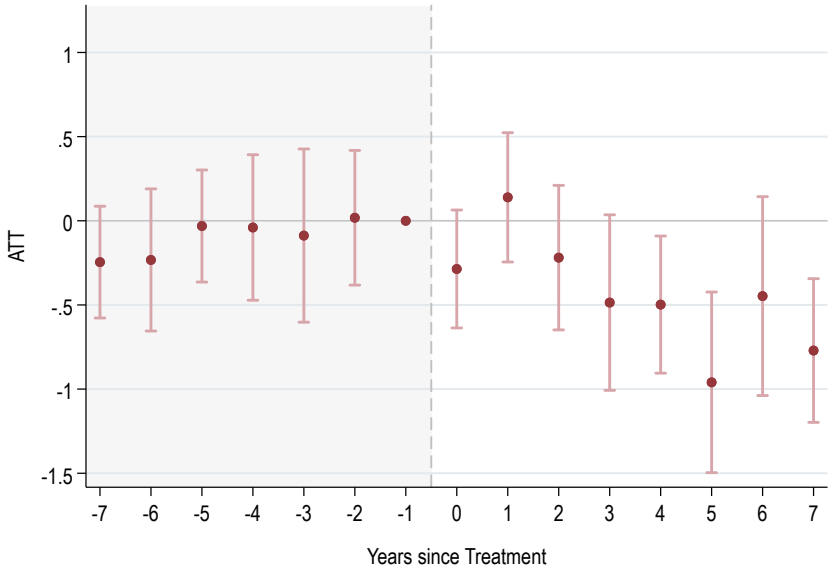
At the same time, flexible trend adjustments must be used carefully. As emphasized by Kahn-Lang and Lang (2020), such adjustments can affect difference-in-differences estimates and may introduce bias when the imposed paths fail to approximate the relevant counterfactual trajectories. I address this concern by introducing the flexible control paths sequentially in Table 3. The estimated effect is already negative and statistically significant at the five percent level in the fixed-effects-only specification, before any flexible covariate paths are added, and changes relatively little as the controls are introduced. The flexible trends therefore refine the comparison without determining the substantive conclusion. The fully controlled specification remains preferred because it allows the observable differences documented in Table 2 to generate different post-1983 trajectories.

4.2.3 Event-Study Evidence

I next estimate a dynamic event-study specification that replaces the post-treatment indicator with a full set of leads and lags relative to the 1983 reform. Unlike the unmatched comparison in Figure 3, the event study conditions on firm–technology, technology–year, and industry–year fixed effects together with the flexible time paths included in the preferred specification. It therefore provides the closest graphical analogue to the identifying assumption underlying the estimating equation.

Figure 4 reports the resulting estimates. The pre-treatment coefficients are centered around zero and display no systematic pattern, providing direct support for the parallel-trends assumption. Following 1983, the estimates become negative and remain below zero through the displayed horizon. The timing and persistence of this divergence reinforce the interpretation of the baseline difference-in-differences estimate as an effect of the reform.

Figure 4: Event-Study Estimates



Notes: Dynamic effects of the 1983 reform over 1976–1990 from the baseline Poisson specification with firm–technology, technology–year, and industry–year fixed effects and flexible time paths for imbalanced pre-treatment covariates. Bars show 95% confidence intervals clustered by firm. Event year −1 (1982) is omitted.

Together, the evidence supports the research design while making clear where adjustment remains necessary. The unmatched means document the raw post-reform divergence and the initial differences between treated and never-funded cells. Matching closely aligns their pre-treatment spillover dynamics, while the preferred specification accommodates the remaining observable differences through flexibly controlled time paths. The event-study estimates then show stable conditional dynamics before 1983 and a persistent relative decline beginning after the reform.

5 Main Results

This section presents the main evidence on the consequences of the 1983 ownership reform for innovation in firm–technology cells exposed to federal funding before the reform. I begin by examining its effect on cross-technology knowledge spillovers. I then investigate whether the decline reflects a broader reduction in technological significance using text-based measures of follow-on innovation from Kelly et al. (2021). Finally, I examine whether weaker external diffusion was accompanied by changes in the private market value of patents produced in treated firm–technology cells.

5.1 Effects on Cross-Technology Spillovers

The conceptual framework developed in the previous section predicts that stronger contractor ownership affects innovation through two complementary channels. First, it increases firms' ability to appropriate the returns from federally funded inventions. Second, by changing the relative profitability of different research projects, it shifts innovation toward inventions with greater private appropriability but potentially narrower technological scope. The empirical analysis begins by examining whether these changes are reflected in cross-technology knowledge spillovers.

Table 3 reports the baseline difference-in-differences estimates. All specifications include firm–technology fixed effects, technology-by-year fixed effects, and industry-by-year fixed effects. The columns sequentially introduce flexible time paths based on pre-treatment firm and innovation characteristics to allow treated and control observations with different initial conditions to follow different underlying trajectories over time.

Table 3: Effects on Cross-Technology Spillovers

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.307***	-0.281***	-0.230***	-0.272***	-0.261***	-0.314***	-0.317***
	(0.095)	(0.092)	(0.081)	(0.089)	(0.084)	(0.105)	(0.107)
Tech-Industry-Year FE	✓	✓	✓	✓	✓	✓	✓
Firm-Tech FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	13,879	13,879	13,871	13,871	13,480	12,990	12,990

Notes: The dependent variable is the number of forward citations received from patents outside the focal CPC technology class within ten years of the patent filing year. All specifications are estimated using Poisson pseudo-maximum likelihood (PPML). Control variables enter as interactions between pre-treatment quartile indicators and year fixed effects, allowing for flexible differential time paths across baseline covariate groups. Patent creativity is measured using the index developed by Kalyani (2024). Standard errors are clustered at the firm level. (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$)

The fixed-effects-only specification indicates a 19 percent decline in cross-technology spillovers, statistically significant at the five percent level.[18] The negative effect is therefore already present

[18] For PPML models, percentage effects are calculated as $\exp(\beta) - 1$.

before introducing flexible time paths for pre-treatment characteristics. Across the remaining specifications, the implied declines range from approximately 22 to 28 percent and are statistically significant at the one percent level. The larger estimate in column (2) primarily reflects the sample restriction due to data availability rather than the inclusion of the creativity controls: re-estimating the fixed-effects specification on the same sample yields a coefficient of -0.337, compared with -0.322 when the creativity controls are included. The sign, magnitude, and statistical precision of the result remains robust as the specification allows for differential time paths based on firm size, patenting activity, R&D expenditures, technological breadth, and firm age.

In the preferred specification, the 1983 reform reduces cross-technology spillovers in treated firm–technology cells by approximately 25 percent. This is a quantitatively meaningful decline given that cross-technology citations capture knowledge flows extending beyond the technological area in which an invention originated. The persistence of the negative estimate across specifications also shows that the result is not generated by the inclusion of flexible covariate adjustments. The event-study estimates presented in Figure 4 reinforce this interpretation. As discussed in Section 4.2, treated and control firm–technology cells exhibit parallel pre-reform trends, while the decline in cross-technology spillovers emerges only after the 1983 policy change.

Appendix Tables A3-A5 examine alternative definitions of external spillovers that sharpen the interpretation of the main result. First, excluding patents jointly assigned to non-firm organizations increases the estimated reduction to about 27 percent (Table A3). This restriction isolates firms' own research decisions and addresses the possibility that non-firm co-assignees follow different research agendas and, importantly, that university research incentives were already affected by the Bayh–Dole Act of 1980, potentially introducing pre-treatment exposure to strengthened ownership rights. Second, restricting spillovers to cross-technology citations originating from other firms strengthens the estimated effect to a 30 percent fall (Table A4), indicating that the decline is particularly pronounced for knowledge flows beyond the boundaries of the originating firm. Finally, imposing both restrictions simultaneously—limiting the sample to firm-only assignees and measuring only cross-firm cross-technology citations—produces a similarly negative and somewhat larger effect (Table A5). Together, these exercises show that the main result is robust to alternative definitions of external spillovers and, more importantly, that the decline in technological influence is concentrated in knowledge diffusing beyond the originating firm.

Although cross-technology citations provide a natural measure of knowledge spillovers, citation counts may also reflect changes in citation practices or in firms' propensity to build on their

own prior inventions. The next subsection therefore examines whether the decline in spillovers is also evident in text-based measures of technological significance that are independent of citation behavior.

5.2 Text-Based Evidence on Technological Significance

I next examine whether the decline in spillovers extends beyond citation-based measures by using the patent significance measure developed by Kelly et al. (2021). Rather than relying on patent citations, this approach measures patent influence through the textual similarity of patents and therefore provides an independent measure of whether inventions generate technologically related follow-on innovation. Formally, patent's j significance is defined as:

$$q_j^\tau = \frac{(Forward\ Similarity)_j^\tau}{(Backward\ Similarity)_j} \tag{3}$$

where τ denotes the post-filing horizon, in years, over which similarity is calculated. Following Kelly et al. (2021), the baseline analysis sets $\tau = 10$. Forward similarity measures the aggregate textual similarity between patent j and subsequent patents during this horizon, while backward similarity measures its similarity to prior patents.

The original Kelly measure has two features preventing it to be readily adopted as a measure of knowledge spillovers for the current study: First, it combines text similarity of future patents with that of past patents. A decline in patent significance could therefore reflect weaker technological diffusion, greater similarity to existing technologies, or both. This concern is directly linked to the effect of interest; if contractor ownership shifted firms' innovation to projects more fully appropriated, this is consistent with higher backward similarity as research may shift more to within-firm development and refinement. In this setup, Kelly measure can decrease post-treatment although knowledge spillovers have remained unchanged.

The second limitation for this study is that the Kelly measure aggregates textually similar inventions regardless of whether they are developed by the same firm or by external innovators. It therefore combines technological influence with development and refinement, which is more likely to be only relevant within-firm, making it less suited to identifying changes in the knowledge diffusion of innovation. To bring the measure closer to a knowledge spillovers proxy, I focus on the external

reach of new inventions, which is more likely to reflect actual technological influence.

To isolate this external component, I use micro-level data generously provided by the authors containing the underlying patent-pair similarity indices for patent pairs with text similarity exceeding 50 percent. These data allow me to identify the ownership of both patents in each pair and distinguish within-firm innovation from technological similarity extending to other firms.

The available patent-pair data have one remaining aspect that needs to be addressed. A measure based on the sum of similarity indices assigns greater weight to within-firm pairs, which can more easily reach very high similarity scores, than to cross-firm pairs, whose similarity is naturally lower. In a dataset that only the high-similarity part of the full distribution is available, this can systematically assign more weight to within-firm pairs.[19]

To address these limitations, I focus on forward patent similarity generated outside the originating firm and specifically, on its share of external forward similarity among total forward similarity. This measure captures the allocation of technological influence on follow-on innovation between the originating firm and external innovators. Using shares rather than levels also reduces the influence of the upward trend in patent activity observed during the time window of the analysis; increasing patent production at the firm-tech level would mechanically raise the number of forward patent pairs. Last, I replace the sum of similarity indices with the count of patent pairs in this highly selective dataset to assign equal weight to each highly similar patent pair. Formally,

$$External\ Diffusion\ Share_j^{\tau} = \frac{(External\ Forward\ Pairs)_j^{\tau}}{(Total\ Forward\ Pairs)_j^{\tau}} \tag{4}$$

where external forward pairs counts the high-similarity pairs between patent j and subsequent patents assigned to other firms over the same τ-year horizon.

The resulting measure captures the breadth of external diffusion by measuring how frequently technologically similar follow-on inventions arise outside the originating firm.

Before proceeding to the estimation table, in appendix table A6, I estimate my baseline using three versions of the Kelly measure, as a validation exercise. The publicly available measure produces small and statistically insignificant estimates. The re-constructed measure using the

[19] Kelly et al. (2021) report that the average similarity score in their data is 10.2 percent, while the 95th percentile is only 22.9 percent. Patent pairs with similarity exceeding 50 percent therefore represent an extreme upper tail of the similarity distribution.

high-similarity patent-pairs data yields much larger negative estimates, significant at five or ten-percent level, while excluding same-firm pairs provides the strongest estimates, both in magnitude and significance. Although the measure loses much of its strength in the more saturated technology-industry-year fixed effects and as additional controls are added, the pattern remains clear: moving from the public measure of patent significance to measures that progressively isolate the external diffusion component strengthens the estimate. Once exclude own-firm pairs the estimated decline in knowledge spillovers is comparable to the effect estimated using cross-citations; about 20-25 percent, albeit sensitive to the empirical specification.

Next, table 4 reports the estimated effect to my preferred outcome of knowledge spillovers; the share of external diffusion. The table shows a robust decline in the external diffusion share following the reform. Across specifications, the preferred estimates imply that the share of technologically similar follow-on innovation generated outside the originating firm falls by approximately 17–19 percent. The estimates remain stable as flexible controls are sequentially introduced. Figure 5 further shows that the decline emerges only after the 1983 reform with no evidence of differential pre-treatment trends.

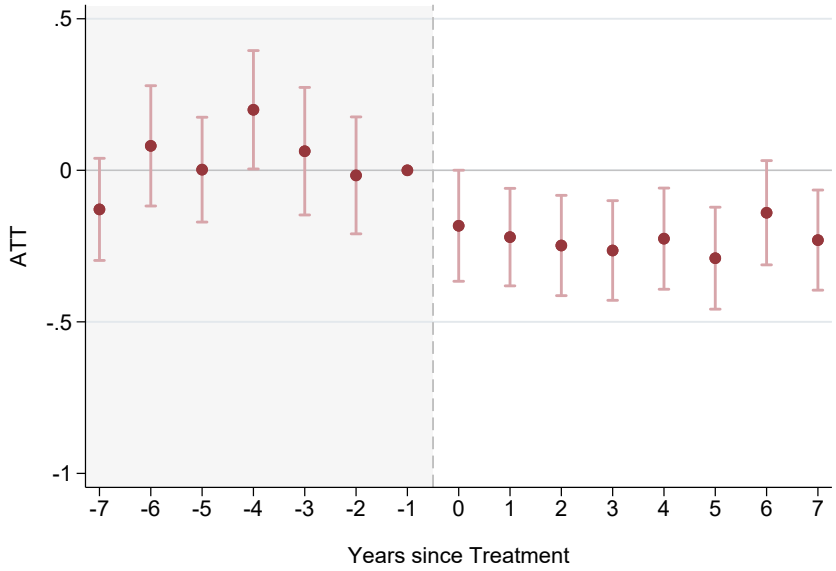
Table 4: Effects on the External Diffusion Share

	External Diffusion share						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.197***	-0.196***	-0.157***	-0.179***	-0.193***	-0.182***	-0.184***
	(0.048)	(0.048)	(0.043)	(0.042)	(0.045)	(0.043)	(0.042)
Tech-Industry-Year FE	✓	✓	✓	✓	✓	✓	✓
Firm-Tech FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	11,859	11,859	11,859	11,859	11,542	11,388	11,388

Notes: The dependent variable is the external diffusion share, defined as the share of high-similarity forward patent pairs assigned to other firms. All specifications are identical to those in the cross-citations estimates, table 3. Standard errors are clustered at the firm level. (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$)

The conclusions are robust to several alternative definitions of external diffusion. Appendix Tables A11–A15 first show that the results are virtually unchanged when the external diffusion share is constructed using similarity weights rather than counts. More importantly, I redefine external diffusion as technologically similar follow-on innovation occurring outside both the originating firm and the focal CPC technology class, mirroring the citation-based spillover measure used throughout the paper. The estimated effects remain quantitatively similar under this more restrictive definition. Finally, excluding patents jointly assigned with universities and other non-business organizations strengthens the results somewhat. Across all alternative definitions, the estimated effects remain negative, statistically significant at the one-percent level, and economically similar to the baseline.

Figure 5: Event-Study Estimates: External Diffusion Share



Notes: Same specification as the previous event-study figure, with the external diffusion share as the dependent variable. Vertical bars denote 95% confidence intervals. The year before the policy change is omitted (year -1).

The text-based evidence is consistent with the decline in knowledge spillovers documented using cross-technology citations. While the Kelly significance measures provide suggestive evidence of a decline in external technological influence, the external diffusion share offers a more direct measure of how follow-on innovation is allocated across firm boundaries. The results indicate that, following the reform, a smaller share of technologically similar inventions was generated outside the originating firm. This is consistent with stronger contractor ownership shifting the technological influence of federally funded inventions toward within-firm development and away from external innovators.

5.3 Private Market Value

The preceding results show that stronger contractor ownership reduced the external technological reach of funded firm–technologies. Whether this decline reflects a reduction in the overall value of innovation or a reallocation of returns from social to private value remains an open question. To examine this issue, I use the patent-level market-value measure developed by Kogan et al. (2017), aggregated to the firm–technology–year level. Because the Kogan measure is constructed from the stock-market response around the patent grant, I follow Kelly et al. (2021) in using patent publication year—which coincides with grant year during the sample period—for estimation.[20]

Table 5: Effects on Patent Market Value

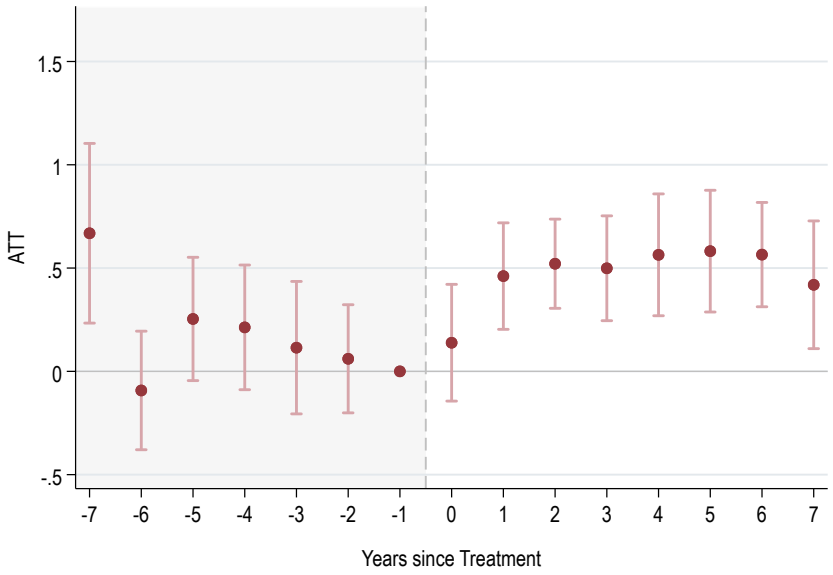
	Patent Market Value						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	0.242**	0.311***	0.320***	0.338**	0.381***	0.404***	0.376**
	(0.124)	(0.121)	(0.112)	(0.136)	(0.135)	(0.140)	(0.164)
Tech-Industry-Year FE	✓	✓	✓	✓	✓	✓	✓
Firm-Tech FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Market Capitalization			✓	✓	✓	✓	✓
Employees				✓	✓	✓	✓
Patents filed					✓	✓	✓
Stochastic Volatility						✓	✓
R&D							✓
N	11,001	11,001	10,819	10,819	10,730	10,617	10,399

Notes: Outcome variable is the patent market-value measure of Kogan et al. (2017), aggregated to the firm–technology–year level. Market values are matched by grant year. Stochastic volatility is calculated as the residual from regressions of excess stock returns on market returns and the Fama–French three factors using CRSP daily data. Market capitalization equals shares outstanding multiplied by the stock price at the end of the fiscal year. All specifications include firm–technology, technology×year, and SIC2×year fixed effects. Standard errors are clustered at the firm level. (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$)

[20] Throughout the sample period, U.S. patent applications were generally confidential until grant; the 18-month publication provisions introduced by the *American Inventors Protection Act* applied to applications filed on or after November 29, 2000. Grant therefore generally marked the first public disclosure of the patent document through the USPTO. Ownership had also been determined by grant under both the pre- and post-reform regimes, so the measure does not compare unresolved ownership before the reform with resolved ownership afterward.

Table 5 reports a positive and economically meaningful effect of the reform on patent market value. Once flexible time paths are introduced for pre-treatment patent creativity, market capitalization, firm size, citation intensity within the CPC class, stock-market volatility, and R&D, the estimated effect becomes positive, stable, and statistically significant. The preferred estimates imply an increase in private patent value of approximately 35–40 percent. The event-study estimates in Figure 6 show no clear evidence of differential pre-treatment trends with a statistically insignificant positive but downward pre-trend for the treated group. The post-treatment profile indicates that there is a persistent increase in market value only after the 1983 reform. If anything, the timing of the downward pre-trend offers an explanation of possible trend reversal consistent with a policy effect.

Figure 6: Event-Study Estimates: Patent Market Value



Notes: Same specification as the previous event-study figure, with patent market value as the dependent variable. Flexible time paths are additionally included for pre-treatment market capitalization and idiosyncratic stock-return volatility. Vertical bars denote 95% confidence intervals.

Taken together, the evidence paints a consistent picture. Stronger contractor ownership reduced the external technological significance and diffusion of federally funded inventions while increasing their private market value. This combination of results is consistent with the conceptual framework developed in Section 2.2: by increasing the returns to private appropriation, the reform shifted innovative effort toward inventions generating greater private value but narrower external technological relevance. The evidence therefore suggests that allocating ownership rights over publicly funded

research induces a trade-off, affecting not only the distribution of returns from innovation but also the direction of innovative effort itself.

5.4 Additional Analysis (Synthetic D-i-D)

The baseline analysis relies on a matching process along flexible controls to construct the counterfactual of the treatment group. To assess the stability of these findings I change the estimator and use the Synthetic Difference-in-Differences estimator of Arkhangelsky et al. (2021). Using 500 placebo iterations, this estimator constructs a weighted control group to reproduce the pre-treatment outcome dynamics of the treated units. Table 6 reports estimates for all outcomes discussed in baseline analysis using both the baseline matched donor pool and the full pool of eligible donor units.

Table 6: Synthetic Difference-in-Differences: All Outcomes

	Cross-Citations		External Diffusion		Patent Market Value	
	Matched	Full	Matched	Full	Matched	Full
Treatment effect	-0.233^{***}	-0.231^{***}	-0.054^{***}	-0.040^{***}	0.419^{***}	0.324^{***}
	(0.024)	(0.022)	(0.010)	(0.003)	(0.023)	(0.006)
Implied change (%)	−21.92	−21.88	−11.72	−8.65	70.47	50.95
Treated units	1,598	1,736	1,598	1,736	1,575	1,693
Donor units	2,895	14,008	1,862	14,008	1,866	11,611
N	89,860	314,880	69,200	314,880	68,820	266,080

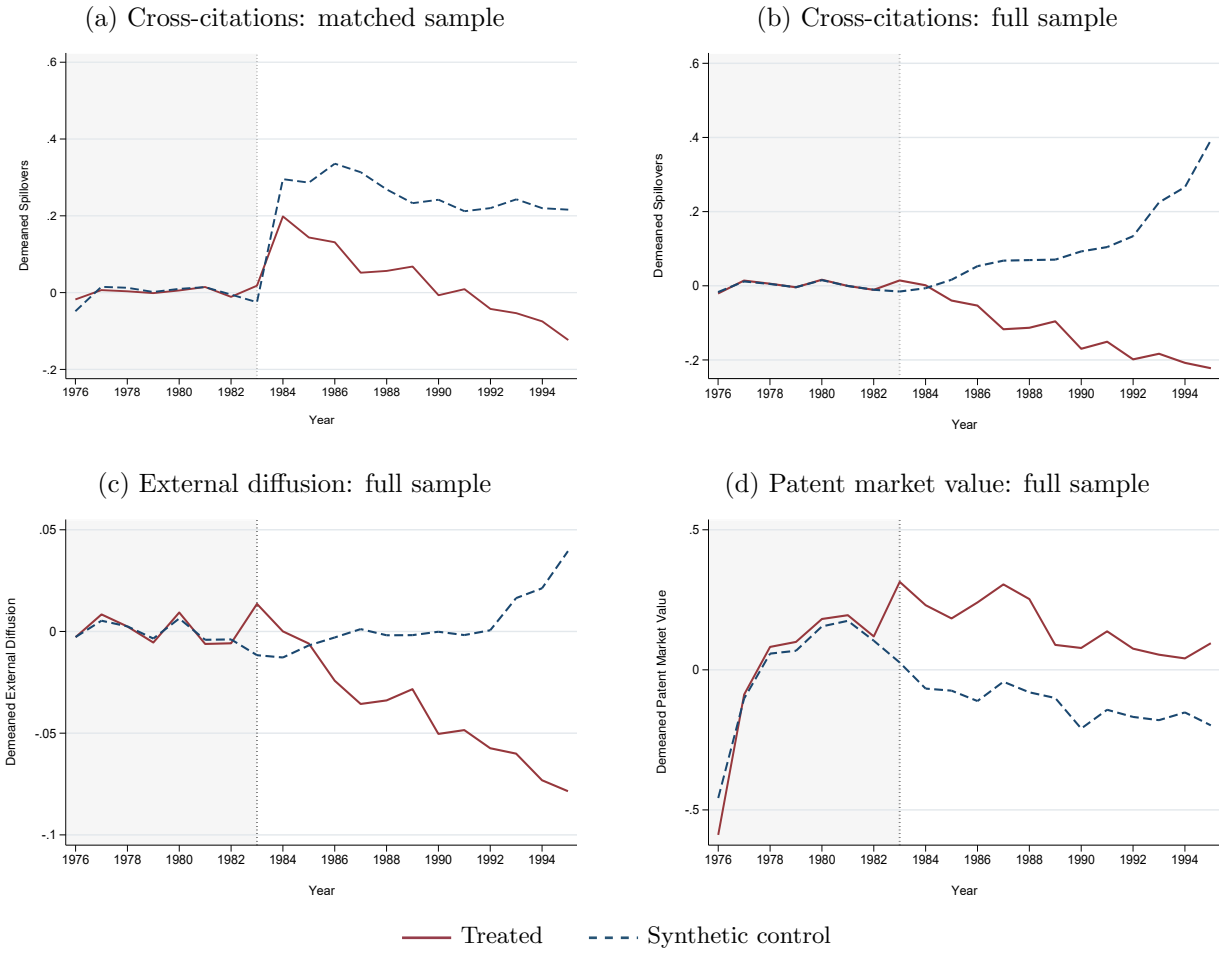
Notes: The table reports estimates from the Synthetic Difference-in-Differences (SDID) estimator of Arkhangelsky et al. (2021). Matched specifications use the baseline matched donor pool; full specifications use all eligible donor units. Outcomes are transformed using $\log(1+Y)$ and demeaned by their firm–technology pre-treatment (1976–1982) means. Years without observed outcomes are assigned zero to construct balanced panels. Implied percentage changes are calculated by transforming the estimated treatment effects back to the original outcome scale, using treated units' pre- and post-treatment means. Standard errors in parentheses ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

The SDID estimates of cross-citations imply declines of approximately 19 percent with the matched donor pool and 21 percent with the full donor pool. Both are statistically significant at the one-percent level. These estimates corroborate the economic magnitude of the approximately 25 percent baseline decline. Their similarity across donor pools also indicates that substantially expanding the comparison group does not materially change the SDID estimate. Donor-weight diagnostics show that weights are broadly distributed, with no single unit receiving disproportionate

weight (Appendix Table A59). Equivalently, the text-based external diffusion estimate implies a somewhat smaller but highly significant decline, relative to the 16 percent baseline decline. On the contrary, estimates on patent market value find a significantly stronger positive effect: 50-70 percent relative to 30-40 percent in the baseline.

Figure 7 shows the evolution of each of these outcomes: figures (a) and (b) refer to matched and full sample estimations for cross-citations, while figures (c) and (d) report full sample estimations for external diffusion and patent market value respectively.

Figure 7: Synthetic Difference-in-Differences Paths



Notes: The figure plots treated and synthetic-control paths from the SDID estimator of Arkhangelsky et al. (2021). Outcomes are transformed using $\log(1+Y)$ and demeaned by each firm–technology cell's 1976–1982 mean. Missing annual outcomes are assigned zero to construct balanced panels. The vertical dotted line marks the 1983 reform.

Across figures, treated and synthetic-control paths closely track one another before the reform and diverge afterward. The post-reform pattern is common across spillover measures, regardless of sample choice. Following the reform, knowledge spillovers decline relative to their synthetic counterfactual, while patent market value remains higher.

6 Heterogeneity

The previous section established that expanding contractor ownership reduced knowledge spillovers while increasing the private market value of federally funded inventions. These average effects are consistent with a reallocation of innovative effort: by increasing firms' ability to appropriate returns from federally funded research, stronger ownership rights may have increased the attractiveness of projects generating greater private value but narrower external technological influence.

This section examines this mechanism. I first study where in the spillover distribution the decline occurs and whether patenting shifts toward lower-spillover inventions. I then examine how patent production responds across firms with different pre-treatment research intensity and whether these changes are accompanied by changes in the private market value of patented inventions.

6.1 Reallocation Across the Spillover Distribution

The average decline in spillovers need not imply that contractor ownership reduced the external influence of all inventions uniformly. If stronger ownership rights changed the direction of research, the response may instead be concentrated among inventions generating broad technological spillovers and accompanied by a shift in patent production toward inventions with narrower external influence.

To examine these margins jointly, I classify patents into quartiles of the spillover distribution within CPC class and filing year. I retain the baseline matched sample for all quartile-specific outcomes in both panels so that both margins are examined within the same population. Table 7 reports two complementary exercises. Panel A estimates the effect of contractor ownership on mean spillovers within each quartile. Panel B examines how the share of patents belonging to each quartile changes following the reform. The first identifies where in the distribution the decline in spillovers occurs, while the second asks whether the composition of patent production shifts across the distribution.

Table 7: Spillovers Across the Distribution

		Spillovers				
Panel A: Means						
Top 25%	Treated × Post	-0.229**	-0.250**	-0.195**	-0.181**	-0.170*
		(0.094)	(0.098)	(0.085)	(0.087)	(0.087)
	N	9067	9067	9062	9044	8831
50–75%	Treated × Post	-0.037	-0.035	-0.012	-0.010	0.003
		(0.024)	(0.022)	(0.022)	(0.023)	(0.024)
	N	9235	9235	9232	9162	8868
25–50%	Treated × Post	0.001	0.005	-0.001	-0.004	-0.002
		(0.033)	(0.032)	(0.032)	(0.031)	(0.031)
	N	8901	8901	8898	8791	8550
Bottom 25%	Treated × Post	0.189***	0.169**	0.195***	0.210***	0.215***
		(0.066)	(0.072)	(0.072)	(0.073)	(0.079)
	N	5854	5854	5854	5758	5669
Panel B: Shares						
Top 25%	Treated × Post	-0.415***	-0.441***	-0.444***	-0.435***	-0.420***
		(0.136)	(0.126)	(0.138)	(0.137)	(0.119)
	N	11975	11975	11975	11880	11631
50–75%	Treated × Post	-0.171**	-0.177*	-0.117	-0.133	-0.193*
		(0.085)	(0.095)	(0.097)	(0.100)	(0.100)
	N	11676	11676	11676	11557	11352
25–50%	Treated × Post	0.075	0.050	-0.048	-0.066	-0.112
		(0.132)	(0.132)	(0.149)	(0.146)	(0.152)
	N	12012	12012	12012	11863	11594
Bottom 25%	Treated × Post	0.297**	0.280**	0.229*	0.177	0.327**
		(0.129)	(0.129)	(0.133)	(0.131)	(0.137)
	N	10381	10403	10403	10492	10320
CPC–SIC2–Year FE		✓	✓	✓	✓	✓
Firm–Tech FE		✓	✓	✓	✓	✓
Patent Creativity			✓	✓	✓	✓
Employees				✓	✓	✓
Patents filed					✓	✓
R&D Expenditure						✓

Notes: Patents are classified into quartiles of the spillover distribution within CPC class and filing year. Panel A reports treatment effects on mean spillovers within each quartile. Panel B reports treatment effects on the share of patents belonging to each quartile. Standard errors are clustered at the firm level ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

The two panels reveal a common distributional pattern. Panel A shows that the decline in spillovers is concentrated at the upper end of the distribution. In the preferred specification, mean spillovers among patents in the top quartile fall by approximately 16 percent, while the effects in the two middle quartiles are small and statistically insignificant. At the bottom of the distribution, the estimated effect is instead positive. The average decline documented in the previous section therefore masks substantial heterogeneity across inventions.

Panel B shows that this change in spillover intensity is accompanied by a shift in the composition of patent production. Following the reform, the share of patents produced by treated firm–technology cells that belong to the top spillover quartile declines by approximately 34 percent, relative to control cells. Since quartiles are defined within CPC class and filing year, this result indicates that treated cells shift away from producing patents at the upper end of their contemporaneous technology-specific spillover distribution. At the bottom quartile, the estimated effect is positive, implying an increase of approximately 39 percent in its share of patent production. The effects in the two middle quartiles are comparatively smaller and generally statistically insignificant.

Together the two panels suggest that stronger contractor ownership changed both the spillover intensity and the composition of innovation. Treated firm–technology cells produce patents with lower mean spillovers at the upper end of the distribution, while allocating a smaller share of their patent production to this category and a larger share to the bottom quartile. These results provide evidence of a shift in the direction of innovation toward inventions with narrower external technological influence. The next subsection examines which firms account for this shift.

6.2 Which Firms Drive the Reallocation?

Firms may respond differently to stronger ownership rights depending on their capacity to finance and undertake research. To investigate this dimension, I classify firm–technology cells into four groups defined by the upper and lower halves of the pre-treatment distributions of mean spillovers and firms' R&D expenditures. Firms with average pre-treatment R&D expenditures below $64 million constitute the low-R&D group, while those above this threshold constitute the high-R&D group. Matching is performed separately within each spillover–R&D cell, ensuring that treated and control units belong to the same part of both distributions. To expand the donor pool available within each cell, I relax the baseline restriction limiting control propensity scores to the 0.40–0.80

interval.[21] Additionally, I replace the baseline technology–industry–year fixed effects with separate technology–year and industry–year fixed effects, as the smaller samples resulting from within-cell matching limit the use of the more saturated specification.

Table 8 reports the resulting treatment effects on patent production. The first row examines patents in the upper half of the spillover distribution and the second patents in the lower half. The response differs sharply across firms and across the spillover distribution. Among low-R&D firms, contractor ownership substantially reduces production of patents in the upper half of the spillover distribution while increasing production in the lower half. The response therefore takes the form of a substitution away from high-spillover and toward low-spillover inventions rather than a uniform contraction in patenting.

Table 8: Patent Production by Spillovers and Pre-Treatment R&D

	Pre-treatment R&D		Heterogeneity
Spillovers	Low 50% (<$64M)	High 50% (>$64M)	p-value
Top 50%	-0.523***	0.355*	0.000
Bottom 50%	0.693***	1.058***	0.003

Notes: The table reports treatment effects from Poisson difference-in-differences specifications estimated separately within spillover–R&D cells. Patents are classified into the upper and lower halves of the spillover distribution within CPC class and filing year. Firms are classified according to their average pre-treatment R&D expenditures, with $64 million separating the low- and high-R&D groups. Matching is performed separately within each spillover–R&D cell. All specifications include firm–technology, technology–year, and two-digit SIC industry–year fixed effects. Standard errors are clustered at the firm level. The heterogeneity column reports the p-value from a test of equality between the low- and high-R&D treatment effects. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

High-R&D firms respond differently. The clearest response is a substantial increase in the production of low-spillover inventions. Patent production also rises in the high-spillover group, although that estimate is smaller and only marginally significant. The pattern therefore suggests that the expansion among high-R&D firms is concentrated primarily in inventions generating narrower external technological influence. The differences between the low- and high-R&D treatment effects are statistically significant in both parts of the spillover distribution.

These results provide a firm-level counterpart to the distributional evidence above. For low-R&D firms, stronger contractor ownership changes the direction of innovative effort through substitution

[21] In the robustness analysis of the main specification, relaxing the donor-pool restriction yields a smaller estimated decline in mean spillovers.

across types of inventions. For high-R&D firms, it operates primarily through differential expansion, with patenting growing much more strongly among low-spillover inventions. In both cases, the resulting composition of innovation shifts toward the lower-spillover portion of the distribution.

6.3 Private Returns and the Direction of Innovation

If this reallocation reflects the incentives created by stronger ownership rights, the changes in patent production should be related to changes in the private returns generated by the corresponding inventions. I examine this prediction using the patent market-value measure of Kogan et al. (2017). Table 9 applies the same spillover–R&D classification and within-cell matching procedure used for patent production, allowing the changes in private value to be compared directly with the reallocation documented in Table 8.

Table 9: Patent Market Value by Spillovers and Pre-Treatment R&D

	Pre-treatment R&D		Heterogeneity
Spillovers	Low 50% (<$64M)	High 50% (>$64M)	p-value
Top 50%	-1.982***	0.141	0.003
Bottom 50%	0.378***	0.333***	0.008

Notes: The dependent variable is the patent market-value measure developed by Kogan et al. (2017). Estimates are reported separately within spillover–R&D cells defined as in Table 8, with matching performed separately within each cell. All specifications include firm–technology, technology–year, and two-digit SIC industry–year fixed effects. Standard errors are clustered at the firm level. The heterogeneity column reports the p-value from a test of equality between the low- and high-R&D treatment effects ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

The market-value results closely track the reallocation in patent production. Among low-R&D firms, the private value of patents in the upper half of the spillover distribution declines sharply, while market value increases among patents in the lower half. This mirrors the shift in patent production away from high-spillover and toward low-spillover inventions documented above.

Among high-R&D firms, market value increases significantly for patents in the lower half of the spillover distribution, while the effect for high-spillover patents is small and insignificant. The strongest expansion in patenting among high-R&D firms occurs in the low-spillover group, where the reform also increases private patent value. This pattern is consistent with innovation shifting toward inventions combining narrower external technological influence with greater private value.

Taken together, the patent-production and market-value results connect the two sides of the

mechanism. Stronger contractor ownership shifts innovative activity toward the lower-spillover portion of the distribution, where private patent value increases. This suggests that changes in the relative returns to research projects played a role in the reallocation of innovation.

The reallocation pattern is also reflected in text-based measures of technological significance. Appendix Table A10 examines the distribution of external technological significance using patent-pair data from Kelly et al. (2021). Consistent with the citation-based evidence, the estimates point to a decline in the production of inventions generating highly significant external follow-on innovation. The decline extends across the upper three quartiles of the significance distribution, which may reflect the highly selective nature of the available sample.

The preceding results support a unified interpretation: the 1983 reform changed the direction of innovation rather than the diffusion process operating after inventions were produced. The reform did not alter the patent system's disclosure requirements: patented inventions continued to become publicly available upon grant. Its central change was instead to make contractor ownership the prospective default allocation of patent rights for federally funded inventions. Firms could therefore anticipate stronger ownership rights when choosing research projects.

The empirical patterns are consistent with this interpretation. Low-R&D firms substitute away from high-spillover inventions, whereas high-R&D firms expand patenting disproportionately among low-spillover inventions. Although operating through different margins, both responses redirect innovative effort toward inventions with narrower external technological influence, precisely where private value increases following the reform. The evidence therefore suggests that expanding contractor ownership influenced not only who captured the returns from federally funded research but also which inventions firms chose to pursue.

7 Robustness Tests

This section examines three potential concerns with the baseline estimates: comparability between treated and matched control units, sensitivity to the post-treatment window and pre-reform access to contractor title among firms with an established position. I first vary donor-pool restrictions and matching rules, and repeat selected checks over a shorter 1976–1990 window (in Appendix). I then exclude contractors that may have had greater access to title before the reform. The appendix reports complete regression results and additional robustness exercises, together with covariate balance tables and matching diagnostics for the alternative donor-pool and matching specifications.

Table 10 examines sensitivity to donor eligibility and the construction of matched controls. Baseline matching requires exact matches within SIC–CPC cells and restricts eligible controls to propensity scores in $(0.4, 0.8)$. First, I impose restrictions based directly on observable comparability. The "p1–p99 covariates" specification excludes potential controls whose pre-treatment characteristics fall outside the 1st–99th percentiles of the treated distribution across the baseline covariates, including pre-treatment spillover levels and trends. Second, I change the eligible donor pool towards the two extremes of the baseline: by removing the score-range restriction entirely and by restricting both treated and control units to their common support. These alternatives assess whether the decline depends on the baseline restriction on donor eligibility.

Third, I examine sensitivity to propensity-score distances and the exact-matching requirement. The "Baseline + Austin caliper" specification retains baseline matching and donor eligibility but imposes a caliper of 0.2 times the pooled standard deviation of logit propensity scores.[22] The "PSM only" specification removes SIC–CPC exact matching, applies the $(0.4, 0.8)$ restriction to both treated and control units, and imposes the same Austin caliper. This changes both the matching rule and the eligible population. Last, I fix the matched pairs of treatment and controls so that each firm-technology is matched to a single control firm-technology. This restriction benefits from comparing fixed pairs at the cost of a significant limitation of the controls available.

Altering the donor pool preserves the negative effect. Both removing the score-range restriction and restricting both groups to their common support yields a decline of approximately 18 percent. Excluding tail observations across observable covariates also yields an equivalent decline in the mean estimated effect. Adding an Austin caliper to the baseline matching produces an estimated decline of approximately 16 percent. Dropping the exact–matching within industry–technology cells requirement and relying only in propensity-score-matching with an Austin caliper produces a decline of approximately 34 percent. Caliper restrictions also materially change the composition of the treated sample. Under baseline exact matching, applying a caliper reduces mean employment among treated observations from approximately 115,000 to 45,000 employees (Table A41). These exercises assess sensitivity to propensity-score distances and treated sample composition.

The negative effect survives both the exclusion of distant propensity-score matches and the removal of the SIC–CPC exact-matching requirement. Finally, the estimated effect is robust also to matching treated firm-technologies to a single control.[23]

[22] See Rosenbaum and Rubin (1985) for propensity-score caliper matching and Austin (2011) for the caliper-width

Table 10: Spillovers: Alternative Control Matches

	Spillovers			Notes
	(1)	(2)	(3)	
Baseline	-0.307***	-0.281***	-0.317***	N: 12,990 (in full controls), donor
	(0.095)	(0.092)	(0.107)	p-scores $\in (0.4, 0.8)$
p1-p99 covariates	-0.224***	-0.188**	-0.195**	N: 10,817, in terms of all covariates
	(0.079)	(0.078)	(0.087)	
Unrestricted Overlap	-0.239***	-0.202***	-0.195**	N: 18,466 (full matched sample)
	(0.079)	(0.078)	(0.097)	
Symmetric overlap	-0.232***	-0.205***	-0.199***	N: 8,883, Baseline match
(0.4, 0.8)	(0.073)	(0.074)	(0.077)	both groups restricted
Baseline + Austin caliper	-0.190***	-0.166**	-0.172**	N: 7,923, Baseline match, donor restriction
	(0.071)	(0.074)	(0.078)	Austin: 0.2×(pooled SD of logit p-scores)
PSM only,	-0.238**	-0.188*	-0.362**	N: 12,920, Propensity-score match,
(0.4,0.8) overlap + Austin	(0.104)	(0.104)	(0.154)	symmetric overlap restriction, Austin
Fixed pre-83 matches	-0.203***	-0.165**	-0.224***	N: 12,185, baseline match otherwise
	(0.074)	(0.075)	(0.082)	
Tech-Industry-Year FE	✓	✓	✓	
Firm-Tech FE	✓	✓	✓	
Patent Creativity		✓	✓	
Full controls			✓	

Notes: Observations refer to full-control specifications. Full controls add filing-year effects by pre-treatment bins of: creativity, employment, patenting, R&D, technological breadth, and age. "p1–p99 covariates" additionally includes spillover levels and trends. Donor restrictions precede matching. Where indicated, the Austin caliper is 0.2 pooled SD of logit propensity scores (Austin, 2011). Standard errors are clustered by firm ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Appendix table A47 repeats selected robustness checks over 1976–1990. Full-control estimates remain stable across donor-pool definitions, implying spillover declines of 29–30 percent. Adding the Austin caliper under baseline or propensity-score-only matching yields similar, smaller declines of 17–18 percent. Both specifications retain the demanding fixed-effects structure with substantially fewer observations; the former is significant at 10 percent, while the latter is not statistically significant. Fixing matched pairs of treated and control units estimate a 25 percent drop in spillovers,

rule used here.

[23] Complete regression tables for each specification can be found in appendix: for alternative donor restriction ranges, tables (A18-A21), for restriction on covariates, table (A16)

significant at one percent level. The estimates thus remain consistently negative, with weaker statistical support under the caliper restrictions.

A separate concern arises from the institutional environment preceding the reform. Under pre-1983 federal patent policy, a contractor's established position in the relevant field was an explicit consideration in determining whether retaining title was appropriate. Contractors with a strong pre-existing market or technological position may therefore have had greater access to ownership rights before the reform. If they also experienced different subsequent spillover dynamics, their inclusion could confound the estimated effect.

Table 11 considers several proxies for an established position, constructed from pre-treatment information. I first exclude treated firm–technology cells associated with market-share leaders within their four-digit SIC industry, leaders or top-decile firms in patent stocks within the relevant CPC class, and leaders or top-decile firms in accumulated R&D.

Table 11: Spillovers: Exclude Contractors with Established Position

	Spillovers	Notes
Baseline	-0.317*** (0.107)	N: 12,990
Among all firm-technologies excluding:		
Market leader in 4-digit SIC	-0.186* (0.108)	N: 10,733
Leader in patent stock in CPC class	-0.239** (0.109)	N: 13,105
Top-10% in patent stock in CPC class	-0.282*** (0.108)	N: 7,900
Leader in R&D stock	-0.229*** (0.103)	N: 13,430
Top-10% in R&D stock	-0.132** (0.114)	N: 9,749
Among funded firm-technologies excluding:		
Leader in patent stock in CPC class	-0.286** (0.101)	N: 12,913
Top 2 in patent stock in CPC class	-0.262** (0.112)	N: 12,313
Leader in R&D stock	-0.219** (0.111)	N: 12,834
Top 2 in R&D stock	-0.280** (0.109)	N: 12,478

Notes: Estimates include firm–technology and technology–industry–year fixed effects, and full controls: flexible filing-year effects by pre-treatment bins of patent creativity, employment, patenting, R&D, technological breadth, and firm age. Eligible controls have propensity scores in $(0.4, 0.8)$. All established-position proxies use pre-treatment data. Standard errors are clustered by firm. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

These rankings may imperfectly capture agencies' assessments of federal contractors. A firm's position within the full population need not coincide with its position among firms performing federally funded research. I therefore also rank firms relative only to other federally funded firm–technology cells and exclude the leading or top-two firms by patent and R&D stocks. Throughout these exercises, I retain the baseline $(0.4, 0.8)$ donor-pool restriction before matching.

The estimates remain negative and economically substantial under every definition of an established position. They are statistically significant in all but one specification: excluding the single patent-stock leader within each CPC class reduces precision but leaves a substantial negative estimate. The pattern holds across product-market, patent-stock, and R&D definitions, and whether rankings use the full population or only federally funded units. The decline therefore persists after excluding contractors whose pre-reform position may have given them greater access to title rights.

The same conclusion holds over 1976–1990. Across the established-position exclusions, estimates remain negative, economically substantial, and statistically significant (Appendix Table A57). Removing self-citations also yields uniformly negative and generally larger estimates (Appendix Table A58), consistent with the main finding of a stronger decline in knowledge flows beyond the originating firm.

8 Conclusion

This paper studies how the prospective allocation of intellectual-property rights over publicly funded inventions shapes firms' research choices and the subsequent diffusion of technological knowledge. I exploit the 1983 Presidential Memorandum on Government Patent Policy, which prospectively expanded contractors' ability to retain ownership of future federally funded inventions. At the same time, the underlying research funding mechanism and the public disclosure requirements of the patent system remained largely unchanged. By prospectively increasing the expected private returns to successful innovation, the reform provides a natural experiment for studying how ownership rights affect the direction and social reach of innovative activity.

The evidence shows that expanding contractor ownership reduced the external technological influence of patented inventions. Following the reform, treated firm–technology units generated innovation associated with fewer cross-technology spillovers, lower rates of technologically similar follow-on innovation outside the originating firm, and higher private market value. These findings are robust across alternative measures of knowledge diffusion, alternative definitions of external

spillovers, and a synthetic difference-in-differences estimator. Because the reform left the public disclosure regime governing patented inventions unchanged, the evidence is more naturally interpreted as reflecting changes in the composition of innovation than changes in the diffusion of a given invention.

The paper also provides evidence on the mechanism underlying these changes. Rather than uniformly reducing innovative activity, stronger ownership rights altered the composition of research. Firms became relatively less likely to produce inventions in the upper tail of the spillover distribution and more likely to generate innovations with narrower external applicability. These compositional changes are concentrated among firms with high research intensity, consistent with appropriability changing the relative profitability of research projects rather than aggregate inventive effort.

These findings provide causal evidence on the fundamental tradeoff between stronger private appropriability and the diffusion of technological knowledge. The partially public-good nature of technological knowledge implies that stronger appropriability can simultaneously increase private incentives to innovate while reducing the knowledge available to future innovators. The empirical evidence supports this prediction: expanding contractor ownership increased the private returns to innovation but reduced their external technological influence. The evidence therefore suggests that stronger appropriability narrows Arrow's private–social wedge from both directions: by increasing the private returns to innovation while simultaneously reducing the external benefits generated by new knowledge.

From a policy perspective, the paper highlights ownership allocation as an important but often underappreciated dimension of innovation policy. Policies governing who owns publicly funded inventions do more than determine the distribution of returns after discoveries are made; they shape the kinds of discoveries firms choose to pursue in the first place. By changing the relative profitability of research projects with different private and social returns, ownership rules influence the balance between private appropriation and knowledge diffusion. The allocation of intellectual-property rights over publicly funded research should therefore be viewed not merely as a legal or administrative choice, but as a policy instrument that helps determine the social returns generated by public R&D investment.

Finally, the analysis points to important dynamic implications. By shifting innovative effort toward technologies with narrower applicability, stronger ownership rights alter the composition of the knowledge accumulated by firms, leading to a more specialized knowledge base. As firms increasingly build capabilities in narrower technological domains, the stock of broadly applicable

knowledge grows more specialized, limiting the range of ideas that future innovators can combine and extend. At the same time, generating fewer spillovers leaves firms with fewer opportunities to benefit from knowledge created elsewhere. To the extent that the returns to R&D are affected by the availability of external knowledge, as emphasized by the literature on absorptive capacity, these dynamics may also weaken incentives for future R&D investment and, over time, the development of firms' absorptive capacity. More broadly, these dynamic effects suggest that ownership policies governing publicly funded research may influence not only the contemporaneous allocation of private and social returns from innovation, but also the long-run accumulation and composition of the knowledge base that underlies technological progress and productivity growth.

References

Aghion, P, "Howitt,., 1992. A Model of Growth Through Creative Destruc'tion," *Econometrica*, 1992, *60* (2), 323.

Aghion, Philippe and Xavier Jaravel, "Knowledge spillovers, innovation and growth," *The Economic Journal*, 2015, *125* (583), 533–573.

Akcigit, Ufuk, Douglas Hanley, and Nicolas Serrano-Velarde, "Back to basics: Basic research spillovers, innovation policy, and growth," *The Review of Economic Studies*, 2021, *88* (1), 1–43.

Arkhangelsky, Dmitry, Susan Athey, David A Hirshberg, Guido W Imbens, and Stefan Wager, "Synthetic difference-in-differences," *American economic review*, 2021, *111* (12), 4088–4118.

Arora, Ashish, Sharon Belenzon, and Lia Sheer, "Knowledge spillovers and corporate investment in scientific research," *American Economic Review*, 2021, *111* (3), 871–898.

Arrow, Kenneth Joseph, "Economic welfare and the allocation of resources for invention," in "Readings in industrial economics: Volume two: Private enterprise and state intervention," Springer, 1962, pp. 219–236.

Austin, Peter C., "Optimal caliper widths for propensity-score matching when estimating differences in means and differences in proportions in observational studies," *Pharmaceutical Statistics*, 2011, *10* (2), 150–161.

Azoulay, Pierre, Joshua S Graff Zivin, Danielle Li, and Bhaven N Sampat, "Public R&D investments and private-sector patenting: evidence from NIH funding rules," *The Review of economic studies*, 2019, *86* (1), 117–152.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying technology spillovers and product market rivalry," *Econometrica*, 2013, *81* (4), 1347–1393.

Budish, Eric, Benjamin N Roin, and Heidi Williams, "Do firms underinvest in long-term research? Evidence from cancer clinical trials," *American Economic Review*, 2015, *105* (7), 2044–2085.

Cohen, Wesley M, Daniel A Levinthal et al., "Absorptive capacity: A new perspective on learning and innovation," *Administrative science quarterly*, 1990, *35* (1), 128–152.

Dyevre, Arnaud, "Public R&D spillovers and productivity growth," Technical Report, Working paper. Eaton, J., & Kortum, S.(1999). International technology ... 2023.

Eisenberg, Rebecca S, "Public Research and Private Development: Patents and Technology Transfer in Government–Sponsored Research," *Va. L. Rev.*, 1996, *82*, 1663.

Fieldhouse, Andrew J and Karel Mertens, *The returns to government R&D: Evidence from US appropriations shocks*, Federal Reserve Bank of Dallas, Research Department, 2023.

_ **and** _ , "The social returns to public R&D," *Entrepreneurship and Innovation Policy and the Economy*, 2026, *5* (1), 9–39.

Fleming, Lee, Hillary Greene, G Li, Matt Marx, and Dennis Yao, "Government-funded research increasingly fuels innovation," *Science*, 2019, *364* (6446), 1139–1141.

Galasso, Alberto and Mark Schankerman, "Patent rights, innovation, and firm exit," *The RAND Journal of Economics*, 2018, *49* (1), 64–86.

Gross, Daniel P and Bhaven N Sampat, "The Government Patent Register: A new resource for measuring US government-funded patenting," *Research Policy*, 2025, *54* (1), 105142.

Howell, Sabrina T, Jason Rathje, John Van Reenen, and Jun Wong, "Opening up military innovation: causal effects of reforms to US defense research," *Journal of Political Economy*, 2025, *133* (11), 3605–3651.

Imbens, Guido W., "Matching Methods in Practice: Three Examples," *Journal of Human Resources*, 2015, *50* (2), 373–419.

Jones, Benjamin F and Lawrence H Summers, "1 A Calculation of the Social Returns to Innovation," in "Innovation and public policy," University of Chicago Press, 2022, pp. 13–60.

Kahn-Lang, Ariella and Kevin Lang, "The promise and pitfalls of differences-in-differences: Reflections on 16 and pregnant and other applications," *Journal of Business & Economic Statistics*, 2020, *38* (3), 613–620.

Kalyani, Aakash, *The creativity decline: Evidence from US patents*, Federal Reserve Bank of St. Louis, Research Division, 2024.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, "Measuring Technological Innovation Over the Long Run," *American Economic Review: Insights*, 2021, *3* (3), 303–320.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological Innovation, Resource Allocation, and Growth," *Quarterly Journal of Economics*, 2017, *132* (2), 665–712.

Lerner, Josh, "The empirical impact of intellectual property rights on innovation: Puzzles and clues," *American Economic Review*, 2009, *99* (2), 343–348.

Lucking, Brian, Nicholas Bloom, and John Van Reenen, "Have R&D spillovers declined in the 21st century?," *Fiscal Studies*, 2019, *40* (4), 561–590.

Masclans, Roger, Sharique Hasan, and Wesley M Cohen, "Measuring the commercial potential of science," *Strategic Management Journal*, 2025, *46* (9), 2199–2236.

Mezzanotti, Filippo and Timothy Simcoe, "Innovation and appropriability: Revisiting the role of intellectual property," Technical Report, National Bureau of Economic Research 2023.

Moretti, Enrico, Claudia Steinwender, and John Van Reenen, "The intellectual spoils of war? Defense R&D, productivity, and international spillovers," *Review of Economics and Statistics*, 2025, *107* (1), 14–27.

Myers, Kyle R and Lauren Lanahan, "Estimating spillovers from publicly funded R&D: Evidence from the US Department of Energy," *American Economic Review*, 2022, *112* (7), 2393–2423.

Rai, Arti K and Bhaven N Sampat, "Accountability in patenting of federally funded research," *Nature biotechnology*, 2012, *30* (10), 953–956.

Romer, Paul M, "Increasing returns and long-run growth," *Journal of political economy*, 1986, *94* (5), 1002–1037.

_ , "Endogenous technological change," *Journal of political Economy*, 1990, *98* (5, Part 2), S71–S102.

Rosenbaum, Paul R. and Donald B. Rubin, "Constructing a Control Group Using Multivariate Matched Sampling Methods That Incorporate the Propensity Score," *The American Statistician*, 1985, *39* (1), 33–38.

Williams, Heidi L, "How do patents affect research investments?," *Annual Review of Economics*, 2017, *9*, 441–469.

Appendix

A Data Construction

A.1 Patent ownership data

The ownership and government funding information used throughout the paper comes from the patent-level dataset assembled by Gross and Sampat (2025). Their contribution substantially expands earlier compilations of federally funded patents by exploiting a previously unused administrative source preserved in the USPTO archives. These archival records identify whether title to a federally funded invention remained with the federal government or was assigned to the contractor while the government retained a non-exclusive license. For a subset of patents, the records also identify the funding agency.

The ownership distinction is the central institutional variable in this paper. Government-title patents correspond to inventions for which ownership remained with the sponsoring federal agency, whereas contractor-title ("license") patents correspond to inventions for which ownership was assigned to the performing organization subject to the government's retained rights. Throughout the paper, I adopt the ownership classifications provided in Gross and Sampat (2025), except for the treatment of post-1983 indeterminate observations discussed in Section A.4.

Although the full ownership dataset spans patents filed between 1950 and 2019, the empirical analysis exploits variation generated by the 1983 policy reform and therefore restricts estimation to patents filed between 1976 and 1995. Earlier and later years are nevertheless retained whenever necessary to construct citation- and text-based outcome variables.

A.2 Patent–firm matching

Patent ownership information is linked to publicly listed firms using the patent-to-Compustat linkage developed by Dyevre (2023). Their algorithm combines information from PatentsView, historical assignee names, Compustat, and corporate ownership records to assign patents to firms while accounting for mergers, acquisitions, name changes, and other corporate reorganizations.

For the purposes of this paper, I use the version of the linkage that assigns patents to the firm's ultimate owner at the time of filing rather than to subsequent owners following corporate

← Back to Data

restructuring. This choice aligns patent ownership with the firm that undertook the underlying research investment and therefore faced the incentives generated by changes in innovation policy.

Because identifying federally funded inventions is particularly important for the empirical analysis, I extend the Dyevre (2023) linkage by matching previously unmatched patents to Compustat firms. Candidate matches are generated using the disambiguated assignee names available in PatentsView together with large-language-model assistance and are subsequently verified manually before inclusion in the final dataset.

This extension adds 70,679 patent-to-firm matches over the full 1950–2019 sample, including 7,147 federally funded patents. Restricting attention to the 1976–1995 estimation period, the extension contributes 17,484 additional patent matches, including 2,931 federally funded patents. While these observations account for only about 7% of patents in the final estimation sample, they increase the number of matched federally funded patents by more than 50% relative to the original linkage.

A.3 Assignee classification

The ownership records in Gross and Sampat (2025) identify the legal ownership status of federally funded patents but do not classify assignees by organizational type. Because the institutional reforms studied in this paper affected different organizations at different times, I construct additional assignee classifications using the disambiguated assignee names available in PatentsView.

Each assignee is classified into one of four broad organizational categories: business firms, universities, government agencies, and individuals. In addition, I separately identify patents jointly assigned to firms and universities and patents jointly assigned to firms and government agencies. These mixed organizational forms are useful for assessing whether the estimated effects are influenced by inventions whose research direction may have been jointly determined by organizations operating under different institutional environments.

To distinguish nonprofit organizations from business firms, I additionally classify organizations according to their tax status under Sections 501(c)(3) and 501(c)(4) of the Internal Revenue Code. Candidate classifications are generated using large-language-model assistance applied to the PatentsView assignee names and are manually verified before incorporation into the dataset.

These classifications are used primarily for robustness exercises. The baseline analysis assigns patents according to filing-date corporate ownership, while alternative specifications progressively exclude patents involving universities, nonprofit organizations, government agencies, and mixed

organizational forms.

A.4 Treatment of indeterminate ownership observations

A small fraction of federally funded patents in Gross and Sampat (2025) are classified as indeterminate because the available archival records do not permit ownership status to be assigned unambiguously. These cases arise when the patent documents do not clearly identify whether title remained with the government or was assigned to the contractor, or when the archival material contains conflicting or incomplete ownership information.

For patents filed before the 1983 policy reform, I follow the ownership classifications in Gross and Sampat (2025) and exclude indeterminate observations from the estimation sample. These patents constitute only a very small share of pre-reform federally funded patents, and their exclusion has no meaningful effect on the composition of the sample.

The post-1983 period requires different treatment. Following correspondence with the authors of Gross and Sampat (2025), indeterminate patents assigned to business firms should, for practical purposes, be interpreted as contractor-owned under the post-reform institutional regime. The public release retains these observations as indeterminate in order to remain faithful to the original archival records rather than impose an additional interpretation.

Accordingly, the baseline analysis classifies post-1983 indeterminate patents assigned to firms as contractor-owned ("license" in the Gross–Sampat coding). This reassignment affects 1,149 federally funded patents in the estimation sample. This treatment is both institutionally and empirically motivated. Institutionally, it follows the interpretation recommended by the authors of Gross and Sampat (2025), who note that these observations remain classified as indeterminate in the public release solely to preserve fidelity to the archival records rather than because they correspond to a distinct ownership regime. Empirically, inspection of the disambiguated assignee names reveals that only 19 of these patents involve non-firm co-assignees, all of which are joint assignments between firms and universities. Excluding these 19 observations leaves the estimates virtually unchanged, indicating that the baseline results are not driven by mixed-assignee cases.

A.5 Citation construction

Patent citations are assembled from two complementary sources. For patents filed before 1976, I use the historical citation data constructed by **?**, which recover citation links for the early U.S. patent system. For patents filed from 1976 onward, citation information is obtained directly from

the USPTO through PatentsView.

Combining the two datasets yields a harmonized citation panel spanning patents filed between 1950 and 2019. The extended time horizon is important because forward citations accumulate gradually over time and because patents filed near the beginning or end of the estimation window require observations outside 1976–1995 to measure subsequent knowledge flows consistently.

The baseline spillover measure counts forward citations originating from patents assigned to different CPC technology classes than the cited patent. Unless otherwise noted, citations are accumulated over a ten-year forward window following the filing year. Alternative citation windows and alternative spillover definitions are examined as robustness exercises.

A.6 Innovation outcome measures

The paper employs four complementary classes of innovation outcomes that capture different dimensions of research performance.

Citation-based measures. The primary outcomes measure knowledge diffusion using forward patent citations. At the patent level, spillovers are measured as forward citations originating from patents outside the focal patent's own CPC technology class. At the firm–technology–year level, these citations are aggregated across all patents filed within the same firm and technology cell. Additional outcomes exclude self-citations to isolate external knowledge diffusion.

Text-based measures. To complement citation-based outcomes, I use the patent-level textual similarity measures developed by Kelly et al. (2021). Their publicly available dataset provides measures of patent impact and novelty derived from similarities between patent texts. I additionally use the authors' patent-pair similarity data for patent pairs with textual similarity exceeding 50%, which permit construction of measures distinguishing within-firm knowledge reuse from diffusion to external innovators. These pair-level measures form the basis of the external diffusion outcomes reported in the paper.

Patent creativity. I also employ the patent creativity index developed by Kalyani (2024), which measures novelty through new combinations of text bi-grams appearing in patent documents. Relative to citation-based measures and the full-text similarity measures of Kelly et al. (2021), this index captures a different dimension of inventive activity and is available for nearly the entire estimation sample. In the context of this paper, the measure is best interpreted as reflecting a combination

of technological novelty and the codification of knowledge into patentable form. Because both processes were already evolving prior to the 1983 reform, the empirical specifications control flexibly for pre-treatment differences in creativity across firm–technology cells. The creativity measure is therefore not interpreted as a standalone estimate of treatment effects but rather as an independent source of variation for evaluating whether the broader pattern of results is consistent across alternative measures of innovation.

Market value. Finally, I measure the private value of innovation using the patent valuation estimates of Kogan et al. (2017). These estimates infer the value of individual patents from abnormal stock-market reactions around patent grants. Although coverage is limited to publicly traded firms with available financial-market data, the market-value measure provides a useful complement to citation- and text-based outcomes by directly capturing the private economic returns to patented inventions.

A.7 Construction of the estimation sample

The empirical analysis is conducted at the firm–technology–year level, where technologies are defined using three-digit Cooperative Patent Classification (CPC) classes. Patents are assigned to firms according to filing-date ownership and then aggregated within firm, technology, and filing year.

Treatment status is determined entirely using the pre-policy period. A firm–technology cell is classified as treated if it receives at least one federally funded patent before the 1983 policy reform. Control observations consist of firm–technology cells that never receive federally funded patents during the sample period. Firm–technology cells first receiving federal funding after 1983 are excluded from the estimation sample in order to maintain a clean comparison between pre-existing recipients and never-treated firms.

The baseline estimation window spans 1976–1995. This period provides sufficiently long pre- and post-reform horizons while avoiding major institutional changes to the U.S. patent system during the late 1990s that could otherwise confound the interpretation of the policy effect.

B Figures

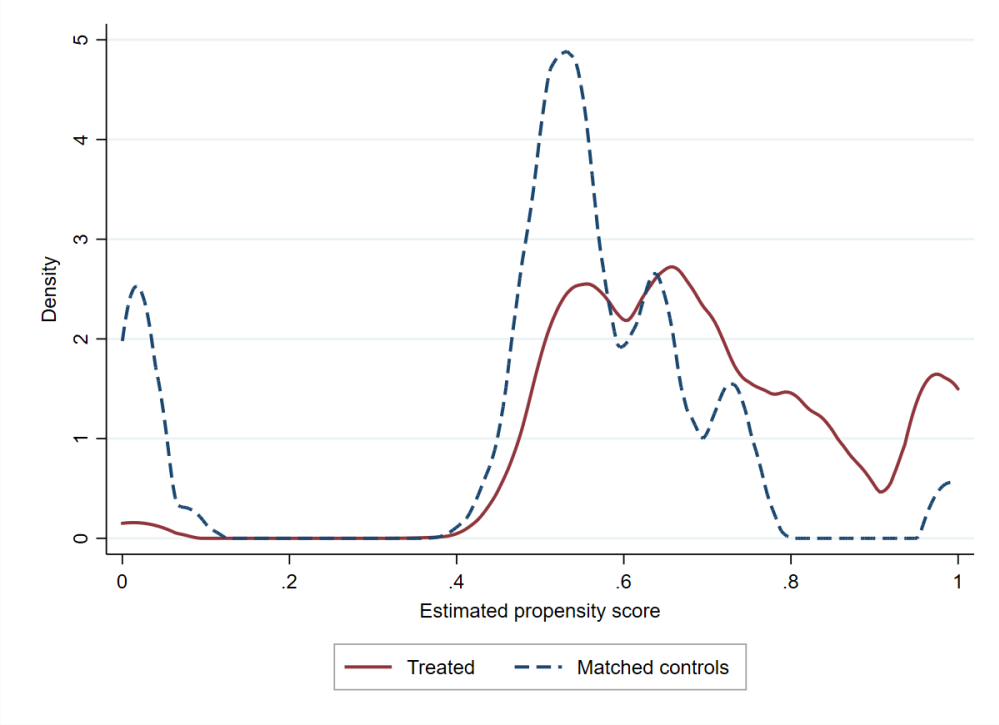


Figure A1: Propensity-Score Overlap Before Matching

Notes: The figure plots the estimated propensity-score distributions for treated and untreated firm–technology cells prior to matching.

← Back to Identification

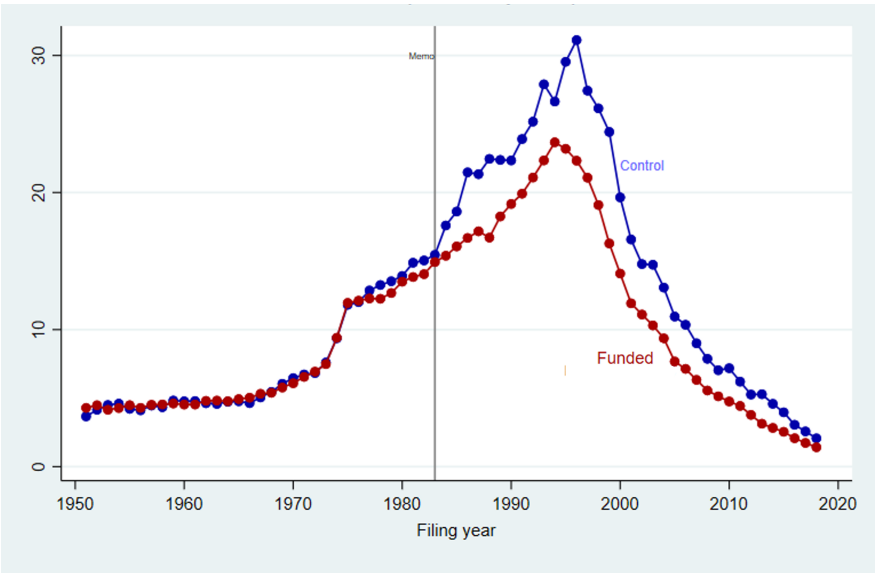


Figure A2: Spillovers by Matched Group, 1950–2019

Notes: The figure plots average cross-technology forward citations for treated and matched control firm-technology cells by filing year. The vertical line marks the 1983 reform.

← Back to Validation

C Tables

C.1 Descriptive Evidence

Table A1: Sample Composition

Sample	Observations	Firms	Firm–techs	CPC	SIC2	Patents
Panel A: Full Patent Database						
1950–2019	457,413	15,646	88,129	154	71	6,697,055
1976–1995	125,121	6,074	34,980	128	66	1,639,666
Panel B: Analysis Sample, 1976–1995						
All	48,205	4,368	15,744	128	66	247,158
Treated	17,040	226	1,736	104	27	185,864
Control	31,165	4,142	14,008	128	66	61,294

Notes: Observations refer to firm–technology–year records. Firms, firm–technology cells, CPC classes, and two-digit SIC industries are counted uniquely. Patents refer to the number of patents represented in each sample.

Table A2: Spillovers & Patent Ownership , 1950-2019, (excl. self-citations)

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Government-funded	0.302***	0.268***	0.269***	0.250***			
	(0.032)	(0.033)	(0.030)	(0.031)			
Contractor-owned	-0.227***	-0.183***	-0.166***	-0.128***	-0.162***	-0.132**	-0.095*
	(0.038)	(0.038)	(0.035)	(0.037)	(0.046)	(0.052)	(0.055)
Patent Controls	Y	Y	Y	Y	Y	Y	Y
Firm Controls	Y	Y	Y	Y	Y	Y	Y
Year FE	✓	✓			✓		
CPC FE	✓	✓			✓		
Industry FE		✓		✓			
CPC-Year FE			✓	✓		✓	
Agency FE					✓	✓	
Agency-CPC-Year FE							✓
SE	cpc-year	cpc-year	cpc-year	cpc-year	cpc-year	cpc-year	cpc-year
N	2,002,341	2,002,317	2,001,677	2,001,653	25,426	24,415	22,730

Notes: Spillovers defined as cross-CPC citations over next ten years excluding self-citations and assumed to be Poisson distributed. Controls: citations over next 10 years, TFP as sales/employees, R&D productivity as R&D expenditure over patents-per-year. Standard errors clustered at the year-cpc_class level. Year FE refer to the filing year. (* p <0.10, ** p <0.05, *** p <0.01)

← Back to Descriptive Section

C.2 Main Results

Table A3: Effects on Cross-Technology Spillovers (firms only), 1976-1995

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.071	-0.265***	-0.230***	-0.215**	-0.247***	-0.317***	-0.314***
	(0.110)	(0.086)	(0.088)	(0.090)	(0.087)	(0.108)	(0.108)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	24,783	19,198	19,181	19,181	18,448	17,838	17,838

Notes: Spillovers are defined as cross-CPC forward citations over the next ten years and assumed to be Poisson distributed. Patents with universities, state agencies or NGOs as co-assignees are excluded. All controls refer to pre-treatment quartile bins allowing for flexible time paths post-1983. Baseline FE include: firm-tech, cpc×year, sic2×year. Standard errors are clustered at the firm level (* p <0.10, ** p <0.05, *** p <0.01)

Table A4: Effects on Cross-Technology Spillovers (excl. self-citations), 1976-1995

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.078	-0.182**	-0.227**	-0.207**	-0.271***	-0.345***	-0.366***
	(0.084)	(0.077)	(0.091)	(0.096)	(0.089)	(0.108)	(0.107)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	26,224	20,032	19,991	19,991	19,252	18,817	18,817

Notes: Spillovers are defined as cross-CPC forward citations over the next ten years and assumed to be Poisson distributed. Identified same-firm citations are excluded. All controls refer to pre-treatment quartile bins allowing for flexible time paths post-1983. Baseline FE include: firm-tech, cpc×year, sic2×year. Standard errors are clustered at the firm level (* p <0.10, ** p <0.05, *** p <0.01)

← Back to Main Results

Table A5: Effects on Cross-Technology Spillovers (excl. self-citations, firms only), 1976-1995

	Patent Similarity Index						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.104	-0.319***	-0.317***	-0.293***	-0.317***	-0.372***	-0.372***
	(0.114)	(0.083)	(0.085)	(0.091)	(0.091)	(0.112)	(0.113)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	25,059	19,171	19,124	19,124	18,433	17,763	17,763

Notes: Spillovers are defined as cross-CPC forward citations over the next ten years and assumed to be Poisson distributed. Identified same-firm citations are excluded. Patents with universities, state agencies or NGOs as co-assignees are excluded. All controls refer to pre-treatment quartile bins allowing for flexible time paths post-1983. Baseline FE include: firm-tech, cpc×year, sic2×year. Standard errors are clustered at the firm level (* p <0.10, ** p <0.05, *** p <0.01)

← Back to Main Results

C.3 Text-Based Evidence on Technological Significance

Table A6: KPST (2021) Patent Significance Measure (alternative versions)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Kelly public							
Treated × Post	0.011	-0.008	-0.006	-0.005	-0.011	-0.004	-0.006
	(0.010)	(0.008)	(0.009)	(0.009)	(0.011)	(0.014)	(0.014)
N	25,262	19,363	19,323	19,323	18,584	17,910	17,910
Panel B: Kelly high similarity analog							
Treated × Post	-0.096**	-0.132**	-0.110**	-0.108**	-0.096*	-0.083	-0.086
	(0.046)	(0.052)	(0.052)	(0.049)	(0.051)	(0.053)	(0.053)
N	24,634	18,727	18,690	18,435	17,782	17,364	17,364
Panel C: Kelly high similarity excl-self pairs							
Treated × Post	-0.320***	-0.355***	-0.271**	-0.257**	-0.243*	-0.208	-0.219*
	(0.108)	(0.115)	(0.137)	(0.130)	(0.133)	(0.128)	(0.128)
N	25,266	19,288	19,248	19,248	18,532	17,872	17,872
Panel D: Kelly high similarity excl-self pairs (joint cell–year FE)							
Treated × Post	-0.140***	-0.158***	-0.112*	-0.126**	-0.096	-0.072	-0.072
	(0.050)	(0.057)	(0.067)	(0.064)	(0.069)	(0.072)	(0.073)
N	17,876	15,944	15,933	15,933	15,382	14,878	14,878
Panels A-C: Tech-Year FE	✓	✓	✓	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓	✓	✓	✓
Panel D: Tech-Industry-Year FE	✓	✓	✓	✓	✓	✓	✓
(All panels)							
Firm–Technology FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓

Notes: Entries report treatment coefficients, with firm-clustered standard errors in parentheses. All regressions include firm–technology fixed effects. Panels A–C additionally include separate CPC–year and SIC2–year fixed effects; Panel D instead includes joint CPC–SIC2–year fixed effects ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A7: KPST (2021) Patent Significance Measure

	Patent Significance						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	0.011	-0.008	-0.006	-0.005	-0.011	-0.004	-0.006
	(0.010)	(0.008)	(0.009)	(0.009)	(0.011)	(0.014)	(0.014)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	25,262	19,363	19,323	19,323	18,584	17,910	17,910

Notes: Patent significance is measured as text similarity of a patent with future patents according to Kelly et al. (2021). This table refers to the publicly available version of the measure which is calculated as the sum of similarity indexes of all future patents, both within and across firms. Thus, firms that file more patents will mechanically score higher patent significance as their patents will naturally have high text similarity. Baseline FE include: firm-tech, cpc×year, sic2×year. Standard errors are clustered at the firm level (* p <0.10, ** p <0.05, *** p <0.01)

Table A8: Effects on Text-Based Technological Significance

	Patent Significance				
	(1)	(2)	(3)	(4)	(5)
Treated x Post	-0.096**	-0.132**	-0.110**	-0.108**	-0.096*
	(0.046)	(0.052)	(0.052)	(0.049)	(0.051)
Baseline FE	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓
Employees			✓	✓	✓
Patents filed				✓	✓
R&D					✓
N	24,634	18,727	18,690	18,435	17,782

Notes: The dependent variable is the reconstructed KPST measure based on patent pairs with text similarity exceeding 50 percent, following Kelly et al. (2021). All similarity pairs, including those assigned to the same firm, are included. Baseline FE include: firm-tech, cpc×year, sic2×year. Standard errors are clustered at the firm level (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$)

← Back to Tech Significance

Table A9: Effects on External Text-Based Patent Significance (counts)

	Patent Significance						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.334***	-0.363***	-0.268*	-0.261*	-0.264*	-0.214	-0.217
	(0.123)	(0.126)	(0.149)	(0.140)	(0.141)	(0.136)	(0.137)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	25,327	19,270	19,227	19,227	18,492	17,847	17,847

Notes: Re-constructed KPST measure calculating forward and backward similarity by summing over counts of pairs with similarity above 50%. Identified same-owner pairs are excluded. (* p <0.10, ** p <0.05, *** p <0.01)

← Back to Tech Significance

Table A10: Technological Significance Across the Distribution

		Patent Significance (counts, excl. self-pairs)				
A: 0-25%	Treated x Post	0.274	0.408**	0.458*	0.511**	0.590
		(0.185)	(0.208)	(0.252)	(0.244)	(.)
	N	7788	7112	7112	7028	6946
B: 25-50%	Treated x Post	-0.091	-0.073	-0.254*	-0.342**	-0.394***
		(0.110)	(0.125)	(0.145)	(0.152)	(0.143)
	N	17397	15853	15836	15625	15163
C: 50-75%	Treated x Post	-0.312***	-0.360***	-0.225*	-0.177	-0.131
		(0.117)	(0.128)	(0.117)	(0.116)	(0.119)
	N	17986	16405	16367	16195	15816
D: 75-100%	Treated x Post	-0.228*	-0.333***	-0.218	-0.214	-0.114
		(0.116)	(0.127)	(0.145)	(0.138)	(0.142)
	N	17292	15690	15672	15556	15175
Baseline FE		✓	✓	✓	✓	✓
Patent Creativity			✓	✓	✓	✓
Employees				✓	✓	✓
Citations, within cpc					✓	✓
R&D Expenditure						✓

Notes: The dependent variable is the reconstructed patent significance measure, defined as the count of high-similarity patent pairs excluding pairs assigned to the same firm. Baseline FE refers to firm–technology, technology–year, and two-digit SIC industry–year fixed effects. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

← Back to Tech Significance

C.4 External Diffusion Share

Table A11: Effects on the External Diffusion Share (weights)

	Externalization Share				
	(1)	(2)	(3)	(4)	(5)
Treated x Post	-0.105**	-0.152***	-0.153***	-0.149***	-0.166***
	(0.041)	(0.045)	(0.045)	(0.044)	(0.043)
Baseline FE	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓
Employees			✓	✓	✓
Patents filed				✓	✓
R&D					✓
N	24,081	18,471	18,433	18,433	17,903

Notes: External diffusion is the share of a patent's similarity to future external patents in total future-patent similarity, reconstructed from the authors' high-similarity dataset. Following Kelly et al. (2021), similarity sums indexes across all future patents, within and across firms; firms with more patents therefore mechanically have higher significance. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A12: Effects on the External & Cross-CPC Diffusion Share (counts)

	Externalization Share				
	(1)	(2)	(3)	(4)	(5)
Treated x Post	-0.110***	-0.154***	-0.143***	-0.133***	-0.142***
	(0.038)	(0.042)	(0.050)	(0.047)	(0.047)
Baseline FE	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓
Employees			✓	✓	✓
Patents filed				✓	✓
R&D					✓
N	23,980	18,199	18,153	18,153	17,520

Notes: External diffusion is a patent's share of text similarity with future external patents in total future-patent similarity, reconstructed from the authors' high-similarity dataset. Following Kelly et al. (2021), similarity is summed across all future patents, within and across firms, so firms with more patents mechanically score higher significance. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

← Back to Externalization Share

Table A13: Effects on the External & Cross-CPC Diffusion Share (weights)

	Externalization Share				
	(1)	(2)	(3)	(4)	(5)
Treated x Post	-0.122***	-0.166***	-0.148***	-0.139***	-0.144***
	(0.039)	(0.043)	(0.048)	(0.046)	(0.046)
Baseline FE	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓
Employees			✓	✓	✓
Patents filed				✓	✓
R&D					✓
N	23,876	18,183	18,148	18,148	17,531

Notes: External diffusion is the share of a patent's future text similarity attributable to external patents, reconstructed from the authors' high-similarity dataset. Following Kelly et al. (2021), total similarity sums similarity indexes across all future patents, within and across firms; firms with more patents therefore mechanically score higher significance. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A14: Effects on the External & Cross-CPC Diffusion Share (firms only, counts)

	Externalization Share				
	(1)	(2)	(3)	(4)	(5)
Treated x Post	-0.136***	-0.198***	-0.181***	-0.168***	-0.180***
	(0.040)	(0.044)	(0.050)	(0.048)	(0.050)
Baseline FE	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓
Employees			✓	✓	✓
Patents filed				✓	✓
R&D					✓
N	23,884	18,185	18,172	18,172	17,548

Notes: External diffusion is the share of a patent's total future text similarity attributable to external patents, reconstructed from the authors' high-similarity dataset. Following Kelly et al. (2021), total similarity sums similarity indexes across future patents within and across firms; greater patenting therefore mechanically raises measured significance. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

← Back to Externalization Share

Table A15: Effects on the External & Cross-CPC Diffusion Share (firms only, weights)

	Externalization Share				
	(1)	(2)	(3)	(4)	(5)
Treated x Post	-0.141***	-0.198***	-0.191***	-0.184***	-0.194***
	(0.040)	(0.045)	(0.052)	(0.050)	(0.052)
Baseline FE	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓
Employees			✓	✓	✓
Patents filed				✓	✓
R&D					✓
N	23,857	18,191	18,178	18,178	17,545

Notes: External diffusion is measured as the share of text similarity of a patent with external future patents over total similarity of future patents. This table refers to the re-constructed version of the measure from the high-similarity dataset provided by the authors. Text similarity is calculated following Kelly et al. (2021) as the sum of similarity indexes of all future patents, both within and across firms. Thus, firms that file more patents will mechanically score higher patent significance as their patents will naturally have high text similarity. (* p <0.10, ** p <0.05, *** p <0.01)

← Back to Externalization Share

C.5 Alternative Donor Pool Matched Controls

Table A16: Spillovers Excluding Tail Control Matches

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.245***	-0.230***	-0.191**	-0.198**	-0.245***	-0.187**	-0.208**
	(0.075)	(0.071)	(0.079)	(0.080)	(0.083)	(0.089)	(0.090)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	20,835	15,579	15,570	15,570	15,545	15,539	15,539

Notes: Analysis period: 1976-1995. The donor pool is restricted prior to matching to control firm–technology cells lying within the 1st–99th percentiles of the treated distribution for pre-treatment firm size (employees), technological breadth (number of CPC classes), patenting activity, R&D expenditures, patent creativity, spillovers, and spillover trends. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A17: Spillovers Using Comparable-Sized Control Firms

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	0.107	-0.195	-0.246*	-0.262*	-0.289**	-0.403**	-0.411**
	(0.118)	(0.121)	(0.132)	(0.147)	(0.128)	(0.192)	(0.192)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	21,968	13,513	13,113	13,113	12,969	12,802	12,802

Notes: Analysis period: 1976-1995. The donor pool is restricted prior to matching to control firm–technology cells whose pre-treatment employment is at least as large as the 25th percentile of treated firms. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

← Back to Robustness

Table A18: Covariate Balance (full matched sample)

	Levels					**Trends**				
	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.
Spillovers	15.72	15.16	16.37	0.243	0.024	0.50	0.21	-0.08	0.043**	0.042
	(24.40)	(22.77)	(23.77)			(7.05)	(6.99)	(16.08)		
# patents filed	9.89	3.07	1.89	< 0.001***	0.598	-0.001	-0.355	-0.02	< 0.001***	0.203
	(15.26)	(5.28)	(2.39)			(1.89)	(1.59)	(0.61)		
Patent creativity	1.16	1.07	1.09	0.324	0.157	0.006	0.0003	-0.02	0.129	0.022
	(0.41)	(0.73)	(0.71)			(0.23)	(0.30)	(0.36)		
Employees	127.40	44.08	23.32	< 0.001***	0.545	0.31	0.23	0.27	0.347	0.013
	(170.11)	(133.38)	(82.87)			(6.81)	(4.44)	(2.18)		
TFP	121.03	118.73	106.80	0.729	0.018	0.84	0.55	0.35	0.303	0.035
	(135.46)	(114.97)	(136.02)			(9.58)	(6.90)	(10.29)		
R&D intensity	3.71	4.57	6.30	< 0.001***	-0.067	-0.06	0.33	0.27	< 0.001***	-0.110
	(16.59)	(7.13)	(32.22)			(3.97)	(2.87)	(3.16)		
Age (in 1983)	3.64	4.59	4.50	< 0.001***	-0.227	—	—	—	—	—
	(2.41)	(5.46)	(5.53)							
# tech classes	24.33	5.30	4.66	< 0.001***	1.995	-0.10	-0.025	0.12	< 0.001***	-0.070
	(13.21)	(3.76)	(4.53)			(1.34)	(0.84)	(0.89)		
N	15,706	10,654	14,008			15,706	10,654	14,008		

Notes: Analysis period: 1976-1995. Baseline 2-step matching: exact within cpc-by-sic2 cell, Nearest-neighbor propensity score matching within cell. No overlap restriction is applied. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A19: Full matched sample

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.107	-0.287***	-0.284***	-0.265***	-0.260***	-0.240**	-0.250***
	(0.107)	(0.085)	(0.090)	(0.093)	(0.088)	(0.095)	(0.088)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
Agegroup						✓	✓
# of CPCs							✓
N	24,911	19,263	19,222	19,222	18,466	18,466	18,466

Notes: Analysis period: 1976-1995. Matching and estimation follow the baseline specification. No overlap restriction is applied. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A20: Restrict Donor Pool to 0.1-0.9 Propensity scores

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.146*	-0.251***	-0.233***	-0.204**	-0.219**	-0.255**	-0.249**
	(0.088)	(0.077)	(0.089)	(0.096)	(0.098)	(0.109)	(0.111)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	21,631	15,586	15,578	15,578	15,186	15,109	15,109

Notes: Analysis period: 1976-1995. The donor pool is restricted prior to matching to controls with propensity score matching within the $(0.1, 0.9)$ range of the distribution. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A21: Restrict Donor Pool to 0.3-0.9 Propensity scores

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.172**	-0.295***	-0.299***	-0.263***	-0.286***	-0.338***	-0.334***
	(0.085)	(0.074)	(0.084)	(0.092)	(0.091)	(0.107)	(0.109)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	21,604	15,567	15,559	15,559	15,166	15,092	15,092

Notes: Analysis period: 1976-1995. The donor pool is restricted prior to matching to controls with propensity score matching within the $(0.3, 0.9)$ range of the distribution. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

← Back to Robustness

Table A22: Covariate Balance (Baseline with Austin caliper)

	Levels					**Trends**				
	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.
Spillovers	15.30	13.00	16.37	0.047**	0.102	0.45	0.026	-0.08	0.279	0.072
	(28.01)	(14.99)	(23.77)			(6.18)	(5.54)	(16.08)		
# patents filed	8.24	3.28	1.89	< 0.001***	0.492	-0.049	-0.425	-0.02	< 0.001***	0.233
	(13.11)	(5.60)	(2.39)			(1.51)	(1.72)	(0.61)		
Patent creativity	1.11	1.05	1.09	0.508	0.107	0.008	0.015	-0.02	0.994	-0.033
	(0.42)	(0.70)	(0.71)			(0.23)	(0.21)	(0.36)		
Employees	42.57	22.19	23.32	< 0.001***	0.769	0.46	0.36	0.27	0.450	0.030
	(31.60)	(20.19)	(82.87)			(2.36)	(3.95)	(2.18)		
TFP	98.54	126.02	106.80	< 0.001***	-0.262	1.01	0.81	0.35	0.021**	0.034
	(84.11)	(121.82)	(136.02)			(5.32)	(6.55)	(10.29)		
R&D intensity	2.08	4.74	6.30	< 0.001***	-0.301	0.21	0.78	0.27	< 0.001***	-0.156
	(6.68)	(10.51)	(32.22)			(2.42)	(4.62)	(3.16)		
Age (in 1983)	3.92	4.80	4.50	< 0.001***	-0.202	—	—	—	—	—
	(3.58)	(4.99)	(5.53)							
# tech classes	16.96	5.59	4.66	< 0.001***	1.750	-0.14	-0.07	0.12	0.016**	-0.076
	(8.32)	(3.91)	(4.53)			(1.05)	(0.83)	(0.89)		
N	7,904	6,051	14,008			7,904	6,051	14,008		

Notes: Analysis period: 1976-1995. Baseline matching; donor p-scores in $(0.4, 0.8)$. Austin caliper: 0.2 times the pooled SD of logit p-scores. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A23: Baseline with Austin caliper

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.153**	-0.230***	-0.208***	-0.217**	-0.146*	-0.208**	-0.208**
	(0.077)	(0.070)	(0.077)	(0.088)	(0.080)	(0.082)	(0.082)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	12,890	8,662	8,662	8,662	8,443	8,307	8,307

Notes: Analysis period: 1976-1995. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A24: Propensity-Score Matching Diagnostics (Austin caliper)

	Statistic	Value
Panel A. Austin caliper		
SD of logit propensity score, treated	σ_T	2.994
SD of logit propensity score, controls	σ_C	3.119
Pooled SD of logit propensity score	σ_P	3.089
Austin caliper ($0.2 \times \sigma_P$)	Caliper	0.618
Panel B. Caliper diagnostic		
Candidate matches	N	5,218
Within caliper	N	2,991
Outside caliper	N	2,227
Panel C. Propensity-score distance among accepted matches		
Matched treated observations	N	2,991
Mean absolute propensity-score distance	Mean	0.043
Standard deviation	SD	0.038
Median	p50	0.034
75th percentile	p75	0.070
90th percentile	p90	0.105
95th percentile	p95	0.124
99th percentile	p99	0.143
Maximum	Max	0.149
Panel D. Accepted matches by treated overlap status		
Treated observations inside overlap region	N	2,810
Mean propensity-score distance	Mean	0.044
Standard deviation	SD	0.040
Median	p50	0.032
95th percentile	p95	0.124
Maximum	Max	0.149
Treated observations below overlap region	N	3
Mean propensity-score distance	Mean	0.036
Standard deviation	SD	0.016
Median	p50	0.027
95th percentile	p95	0.055
Maximum	Max	0.055
Treated observations above overlap region	N	178
Mean propensity-score distance	Mean	0.038
Standard deviation	SD	0.009
Median	p50	0.038
95th percentile	p95	0.057
Maximum	Max	0.058

Notes: Analysis period: 1976-1995. The table reports diagnostics for propensity-score nearest-neighbor matching. The matching caliper follows Austin's rule and is set equal to 0.2 standard deviations of the logit of the propensity score, using the pooled standard deviation across treated and control observations.

Table A25: Covariate Balance (0.4-0.8 symmetric overlap restriction)

	Levels					**Trends**				
	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.
Spillovers	15.26	13.83	16.37	0.271	0.062	0.52	0.05	-0.08	0.237	0.080
	(26.11)	(19.59)	(23.77)			(6.07)	(5.66)	(16.08)		
# patents filed	7.73	3.33	1.89	$< 0.001^{***}$	0.511	-0.124	-0.434	-0.02	0.005^{***}	0.192
	(10.78)	(5.65)	(2.39)			(1.46)	(1.75)	(0.61)		
Patent creativity	1.13	1.03	1.09	0.720	0.172	0.006	0.001	-0.02	0.198	0.022
	(0.44)	(0.70)	(0.71)			(0.22)	(0.22)	(0.36)		
Employees	44.83	20.21	23.32	$< 0.001^{***}$	1.022	0.52	0.30	0.27	0.154	0.069
	(27.86)	(19.62)	(82.87)			(2.50)	(3.58)	(2.18)		
TFP	108.70	124.74	106.80	0.056^{*}	-0.145	1.35	0.65	0.35	$< 0.001^{***}$	0.112
	(105.48)	(115.73)	(136.02)			(6.32)	(6.20)	(10.29)		
R&D intensity	2.23	4.28	6.30	0.001^{***}	-0.255	0.21	0.58	0.27	0.004^{***}	-0.117
	(6.61)	(9.50)	(32.22)			(2.23)	(3.89)	(3.16)		
Age (in 1983)	3.86	4.69	4.50	$< 0.001^{***}$	-0.197	—	—	—	—	—
	(3.40)	(4.93)	(5.53)							
# tech classes	17.95	5.15	4.66	$< 0.001^{***}$	1.974	-0.21	-0.01	0.12	$< 0.001^{***}$	-0.203
	(8.42)	(3.65)	(4.53)			(1.16)	(0.79)	(0.89)		
N	9,164	6,697	14,008			9,164	6,697	14,008		

Notes: Analysis period: 1976-1995. Baseline matching; donor and treated p-scores in $(0.4, 0.8)$. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A26: 0.4-0.8 symmetric overlap restriction

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.143*	-0.211***	-0.151**	-0.114	-0.108	-0.174**	-0.174**
	(0.079)	(0.071)	(0.070)	(0.074)	(0.073)	(0.077)	(0.077)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	14,799	10,784	10,784	10,784	10,554	10,259	10,259

Notes: Analysis period: 1976-1995. Baseline matching; donor and treated p-scores in $(0.4, 0.8)$. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A27: Covariate Balance (0.4-0.8 overlap restriction and 0.2 caliper on p-scores)

	Levels					**Trends**				
	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.
Spillovers	15.59	12.83	16.37	0.009***	0.124	0.57	-0.04	-0.08	0.109	0.088
	(27.76)	(15.05)	(23.77)			(8.11)	(5.30)	(16.08)		
# patents filed	8.41	3.55	1.89	< 0.001***	0.481	-0.058	-0.513	-0.02	< 0.001***	0.272
	(12.97)	(5.92)	(2.39)			(1.48)	(1.85)	(0.61)		
Patent creativity	1.11	1.03	1.09	0.373	0.126	0.006	0.017	-0.02	0.785	-0.051
	(0.43)	(0.71)	(0.71)			(0.23)	(0.19)	(0.36)		
Employees	45.86	22.35	23.32	< 0.001***	0.820	0.61	0.38	0.27	0.423	0.064
	(34.75)	(20.91)	(82.87)			(3.26)	(3.83)	(2.18)		
TFP	101.98	127.26	106.80	< 0.001***	-0.243	0.99	0.85	0.35	0.031**	0.024
	(83.79)	(120.97)	(136.02)			(5.19)	(6.44)	(10.29)		
R&D intensity	2.14	4.93	6.30	< 0.001***	-0.321	0.21	0.86	0.27	< 0.001***	-0.176
	(6.42)	(10.53)	(32.22)			(2.31)	(4.70)	(3.16)		
Age (in 1983)	3.88	4.96	4.50	< 0.001***	-0.228	—	—	—	—	—
	(3.48)	(5.69)	(5.53)							
# tech classes	17.84	5.47	4.66	< 0.001***	1.798	-0.14	-0.09	0.12	0.043**	-0.056
	(8.88)	(3.96)	(4.53)			(1.08)	(0.83)	(0.89)		
N	8,957	6,417	14,008			8,957	6,417	14,008		

Notes: Analysis period: 1976-1995. Baseline matching; donor p-scores in $(0.4, 0.8)$. 0.20 caliper in raw propensity scores. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A28: 0.4-0.8 overlap restriction and 0.2 caliper on p-scores

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.208**	-0.293***	-0.285***	-0.256***	-0.218***	-0.293***	-0.293***
	(0.090)	(0.077)	(0.084)	(0.091)	(0.083)	(0.094)	(0.094)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	14,315	9,641	9,641	9,641	9,383	9,210	9,210

Notes: Analysis period: 1976-1995. Baseline matching; donor p-scores in $(0.4, 0.8)$. 0.20 caliper in raw propensity scores. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A29: Propensity-Score Matching Diagnostics (0.20 caliper in levels)

	N	Percent
Panel A. Caliper acceptance		
Within caliper	3,316	63.55
Outside caliper	1,902	36.45
Total	5,218	100.00
Panel B. Retained matches by treated propensity score		
Treated propensity score < 0.40	8	0.24
Treated propensity score 0.40–0.80	3,050	91.98
Treated propensity score > 0.80	258	7.78
Total	3,316	100.00

	Statistic	Propensity-score distance
Panel C. Distance between retained matched pairs		
Mean		0.055
Standard deviation		0.051
Median		0.037
75th percentile		0.083
90th percentile		0.138
95th percentile		0.159
99th percentile		0.192
Maximum		0.195

Panel D. Distance by treated propensity-score region

	N	Mean	[Median; 95th pct.]
$p < 0.40$	8	0.124	[0.171; 0.185]
$0.40 \leq p \leq 0.80$	3,050	0.054	[0.036; 0.158]
$p > 0.80$	258	0.070	[0.041; 0.177]

Notes: Analysis period: 1976-1995. The table reports diagnostics for propensity-score nearest-neighbor matching with a caliper. Propensity-score distance is the absolute difference in propensity scores between a treated observation and its matched control. Panel B reports the location of retained treated observations relative to the 0.40–0.80 propensity-score region. Panels C and D report the realized propensity-score distance between retained matched pairs.

← Back to Robustness

Table A30: 0.4-0.8 overlap restriction and 0.2 caliper on p-scores – Firms only

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.143*	-0.218***	-0.201**	-0.187**	-0.155*	-0.204**	-0.204**
	(0.085)	(0.072)	(0.082)	(0.091)	(0.083)	(0.091)	(0.091)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	14,231	9,642	9,642	9,642	9,389	9,216	9,216

Notes: Analysis period: 1976-1995. Baseline matching; donor p-scores in $(0.4, 0.8)$. 0.20 caliper in raw propensity scores. Patents with non-firm actors as co-assignees are excluded (ie Universities, NGOs, Agencies). Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A31: 0.4-0.8 overlap restriction and 0.2 caliper on p-scores – Self-citations excluded

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.228***	-0.314***	-0.308***	-0.275***	-0.242***	-0.277***	-0.277***
	(0.088)	(0.076)	(0.085)	(0.089)	(0.080)	(0.090)	(0.090)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	14,224	9,639	9,639	9,639	9,384	9,211	9,211

Notes: Analysis period: 1976-1995. Baseline matching; donor p-scores in $(0.4, 0.8)$. 0.20 caliper in raw propensity scores. Identified self-citations are excluded. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

← Back to Robustness

Table A32: 0.4-0.8 overlap restriction and 0.2 caliper on p-scores – Firms only, no self-citations

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.174**	-0.269***	-0.254***	-0.226***	-0.190**	-0.244***	-0.244***
	(0.081)	(0.072)	(0.080)	(0.085)	(0.080)	(0.088)	(0.088)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	14,220	9,637	9,637	9,637	9,383	9,210	9,210

Notes: Analysis period: 1976-1995. Baseline matching; donor p-scores in $(0.4, 0.8)$. 0.20 caliper in raw propensity scores. Patents with non-firm actors as co-assignees are excluded. Identified citations from own-firm are excluded. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

← Back to Robustness

Table A33: Covariate Balance (No exact matching, 0.4-0.8 overlap restriction and Austin caliper)

	Levels					**Trends**				
	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.
Spillovers	15.50	12.70	16.37	0.282	0.165	0.30	-0.26	-0.08	0.005***	0.093
	(20.77)	(11.89)	(23.77)			(7.15)	(4.67)	(16.08)		
# patents filed	9.06	1.98	1.89	< 0.001***	0.673	-0.086	-0.092	-0.02	0.182	0.006
	(14.62)	(2.70)	(2.39)			(1.57)	(0.60)	(0.61)		
Patent creativity	1.16	1.09	1.09	0.156	0.111	0.006	0.014	-0.02	0.311	-0.014
	(0.42)	(0.71)	(0.71)			(0.24)	(0.43)	(0.36)		
Employees	111.07	31.53	23.32	< 0.001***	0.969	0.92	0.15	0.27	0.003***	0.180
	(113.97)	(22.43)	(82.87)			(5.41)	(2.52)	(2.18)		
TFP	118.35	108.68	106.80	< 0.001***	0.085	0.96	0.82	0.35	< 0.001***	0.020
	(132.25)	(90.69)	(136.02)			(7.53)	(5.37)	(10.29)		
R&D intensity	2.94	5.89	6.30	0.017**	-0.223	0.16	0.25	0.27	0.079*	-0.057
	(15.16)	(11.04)	(32.22)			(1.70)	(1.49)	(3.16)		
Age (in 1983)	3.68	4.46	4.50	< 0.001***	-0.196	—	—	—	—	—
	(2.91)	(4.75)	(5.53)							
# tech classes	26.42	5.10	4.66	< 0.001***	1.895	-0.083	0.240	0.12	< 0.001***	-0.286
	(15.53)	(3.45)	(4.53)			(1.22)	(1.03)	(0.89)		
N	15,550	10,197	14,008			15,550	10,197	14,008		

Notes: Analysis period: 1976-1995. Nearest neighbor propensity score matching only; no cpc-sic2 matching. Donor p-scores in $(0.4, 0.8)$. Austin caliper: 0.2 times the pooled SD of logit p-scores. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A34: No exact matching, 0.4-0.8 overlap and Austin caliper

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.204	-0.369***	-0.338***	-0.325**	-0.354***	-0.549***	-0.549***
	(0.174)	(0.130)	(0.130)	(0.131)	(0.130)	(0.182)	(0.182)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	24,540	21,161	21,161	21,161	16,914	16,785	16,785

Notes: Analysis period: 1976-1995. Nearest neighbor propensity score matching; no cpc-sic2 matching. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A35: Propensity-Score Matching Diagnostics (Austin caliper)

	N	Percent
Panel A. Austin-caliper acceptance		
Within caliper	7,053	93.52
Outside caliper	489	6.48
Total	7,542	100.00
Panel B. Retained matches by treated propensity score		
Treated propensity score < 0.40	12	0.17
Treated propensity score 0.40–0.80	4,549	64.50
Treated propensity score > 0.80	2,492	35.33
Total	7,053	100.00

	Statistic	Propensity-score distance
Panel C. Distance between retained matched pairs		
Mean		0.044
Standard deviation		0.070
Median		0.0003
75th percentile		0.086
90th percentile		0.194
95th percentile		0.198
99th percentile		0.200
Maximum		0.200

Panel D. Distance by treated propensity-score region

	N	Mean	[Median; 95th pct.]
$p < 0.40$	12	0.068	[0.027; 0.145]
$0.40 \leq p \leq 0.80$	4,549	0.0002	[0.0001; 0.001]
$p > 0.80$	2,492	0.123	[0.119; 0.199]

Notes: Analysis period: 1976-1995. The table reports diagnostics for propensity-score nearest-neighbor matching with an Austin caliper, defined as 0.2 times the pooled standard deviation of the logit of the propensity score. Propensity-score distance is the absolute difference in propensity scores between a treated observation and its matched control. Panel B reports the location of retained treated observations relative to the baseline's 0.40–0.80 propensity-score region. Panels C and D report the realized propensity-score distance between retained matched pairs.

← Back to Robustness

Table A36: Covariate Balance (No exact matching, 0.4-0.8 symmetric overlap and Austin caliper)

	Levels					**Trends**				
	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.
Spillovers	15.00	12.46	16.37	0.272	0.135	0.38	-0.16	-0.08	0.003***	0.097
	(23.78)	(11.79)	(23.77)			(6.75)	(4.74)	(16.08)		
# patents filed	7.12	1.99	1.89	< 0.001***	0.689	-0.098	-0.090	-0.02	0.534	-0.007
	(10.18)	(2.71)	(2.39)			(1.37)	(0.61)	(0.61)		
Patent creativity	1.13	1.09	1.09	0.056*	0.074	0.004	0.002	-0.02	0.481	0.004
	(0.45)	(0.71)	(0.71)			(0.26)	(0.43)	(0.36)		
Employees	46.29	31.68	23.32	0.985	0.580	0.63	0.15	0.27	0.395	0.177
	(27.70)	(22.45)	(82.87)			(2.83)	(2.52)	(2.18)		
TFP	118.16	110.11	106.80	0.223	0.072	1.81	0.90	0.35	< 0.001***	0.128
	(126.74)	(93.52)	(136.02)			(8.24)	(5.57)	(10.29)		
R&D intensity	2.17	5.97	6.30	0.005***	-0.431	0.20	0.24	0.27	0.141	-0.022
	(5.66)	(11.11)	(32.22)			(2.07)	(1.50)	(3.16)		
Age (in 1983)	3.80	4.57	4.50	< 0.001***	-0.188	—	—	—	—	—
	(3.20)	(4.89)	(5.53)							
# tech classes	18.83	5.11	4.66	< 0.001***	2.110	-0.097	0.225	0.12	< 0.001***	-0.279
	(26.96)	(3.46)	(4.53)			(1.26)	(1.04)	(0.89)		
N	10,202	7,673	14,008			10,202	7,673	14,008		

Notes: Analysis period: 1976-1995. Nearest neighbor propensity score matching only. Donor and treated p-scores in $(0.4, 0.8)$. Austin caliper: 0.2 times the pooled SD of logit p-scores. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A37: No exact matching, 0.4-0.8 symmetric overlap restriction and Austin caliper

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.215*	-0.318***	-0.288**	-0.248**	-0.286***	-0.415***	-0.415***
	(0.126)	(0.120)	(0.123)	(0.123)	(0.109)	(0.142)	(0.142)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	16,641	13,081	13,081	13,081	11,354	11,226	11,226

Notes: Analysis period: 1976-1995. Nearest neighbor propensity score matching; no cpc-sic2 matching. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A38: Propensity-Score Matching Diagnostics (Austin caliper)

	Statistic
Panel A. Austin caliper	
SD logit propensity score, treated	2.995
SD logit propensity score, controls	3.119
Pooled SD logit propensity score	3.089
Austin caliper (0.2× pooled SD)	0.618
Panel B. Propensity-score distance	
Number of matched pairs	4,549
Mean	0.00024
Standard deviation	0.00032
Median	0.00012
75th percentile	0.00031
90th percentile	0.00064
95th percentile	0.00095
99th percentile	0.00152
Maximum	0.00210

Notes: Analysis period: 1976-1995. The table reports diagnostics for propensity-score nearest-neighbor matching with a caliper. Propensity-score distance is the absolute difference in propensity scores between a treated observation and its matched control. Panel B reports the location of retained treated observations relative to the 0.40–0.80 propensity-score region. Panels C and D report the realized propensity-score distance between retained matched pairs.

← Back to Robustness

Table A39: Covariate Balance (No exact matching, 0.4-0.8 symmetric overlap and Austin caliper)

	Levels					**Trends**				
	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.
Spillovers	15.21	13.46	16.37	0.644	0.092	0.43	0.72	-0.08	0.414	-0.046
	(23.55)	(12.87)	(23.77)			(6.73)	(6.01)	(16.08)		
# patents filed	7.39	1.82	1.89	$< 0.001^{***}$	0.702	-0.137	0.002	-0.02	0.004^{***}	-0.107
	(11.03)	(2.11)	(2.39)			(1.54)	(1.02)	(0.61)		
Patent creativity	1.14	1.09	1.09	0.500	0.068	0.003	0.019	-0.02	0.740	-0.046
	(0.44)	(0.72)	(0.71)			(0.26)	(0.40)	(0.36)		
Employees	48.21	33.60	23.32	0.835	0.553	0.64	1.48	0.27	0.284	-0.140
	(29.44)	(22.97)	(82.87)			(2.82)	(7.97)	(2.18)		
TFP	117.32	121.07	106.80	0.014^{**}	-0.033	1.74	2.38	0.35	0.026^{**}	-0.074
	(124.97)	(105.30)	(136.02)			(8.13)	(9.11)	(10.29)		
R&D intensity	2.15	6.11	6.30	$< 0.001^{***}$	-0.427	0.20	0.50	0.27	0.881	-0.109
	(5.57)	(11.91)	(32.22)			(2.04)	(3.30)	(3.16)		
Age (in 1983)	3.79	4.62	4.50	$< 0.001^{***}$	-0.187	—	—	—	—	—
	(3.18)	(5.45)	(5.53)							
# tech classes	19.16	5.47	4.66	$< 0.001^{***}$	2.046	-0.128	0.147	0.12	$< 0.001^{***}$	-0.227
	(8.68)	(3.76)	(4.53)			(1.27)	(1.15)	(0.89)		
N	8,372	6,788	14,008			8,372	6,788	14,008		

Notes: Analysis period: 1976-1990. Nearest neighbor propensity score matching only. Donor and treated p-scores in $(0.4, 0.8)$. Austin caliper: 0.2 times the pooled SD of logit p-scores. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A40: No exact matching, 0.4-0.8 symmetric overlap restriction and Austin caliper

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.023	-0.109	-0.068	-0.058	-0.112	-0.203	-0.203
	(0.125)	(0.125)	(0.143)	(0.150)	(0.158)	(0.154)	(0.154)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	14,063	11,486	11,486	11,486	9,838	9,730	9,730

Notes: Analysis period: 1976-1990. Nearest neighbor propensity score matching; no cpc-sic2 matching. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A41: Propensity-Score Matching Diagnostics (Austin caliper)

	Statistic
Panel A. Austin caliper	
SD logit propensity score, treated	2.842
SD logit propensity score, controls	3.051
Pooled SD logit propensity score	3.002
Austin caliper (0.2× pooled SD)	0.600
Panel B. Propensity-score distance	
Number of matched pairs	4,678
Mean	0.00037
Standard deviation	0.00063
Median	0.00014
75th percentile	0.00044
90th percentile	0.00085
95th percentile	0.00144
99th percentile	0.00372
Maximum	0.00438

Notes: Analysis period: 1976-1990. The table reports diagnostics for propensity-score nearest-neighbor matching with a caliper. Propensity-score distance is the absolute difference in propensity scores between a treated observation and its matched control. Panel B reports the location of retained treated observations relative to the 0.40–0.80 propensity-score region. Panels C and D report the realized propensity-score distance between retained matched pairs.

← Back to Robustness

Table A42: Covariate Balance (Baseline with Austin caliper, 1976-1990)

	Levels					**Trends**				
	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.	Treated Firms	Matched Controls	All Controls	p-value	Std. Diff.
Spillovers	15.53	13.63	16.37	0.080*	0.084	0.58	0.05	-0.08	0.134	0.077
	(27.58)	(16.69)	(23.77)			(8.06)	(5.35)	(16.08)		
# patents filed	8.41	3.55	1.89	< 0.001***	0.485	-0.054	-0.487	-0.02	< 0.001***	0.255
	(12.93)	(5.89)	(2.39)			(1.48)	(1.90)	(0.61)		
Patent creativity	1.11	1.05	1.09	0.424	0.109	0.006	0.011	-0.02	0.619	-0.022
	(0.43)	(0.72)	(0.71)			(0.23)	(0.20)	(0.36)		
Employees	45.99	22.20	23.32	< 0.001***	0.835	0.62	0.41	0.27	0.430	0.059
	(34.56)	(20.76)	(82.87)			(3.27)	(3.78)	(2.18)		
TFP	102.07	126.73	106.80	< 0.001***	-0.239	0.96	0.64	0.35	0.017**	0.051
	(83.27)	(120.19)	(136.02)			(5.17)	(6.84)	(10.29)		
R&D intensity	2.14	5.09	6.30	< 0.001***	-0.329	0.21	0.83	0.27	< 0.001***	-0.169
	(6.38)	(10.05)	(32.22)			(2.29)	(4.65)	(3.16)		
Age (in 1983)	3.88	4.78	4.50	< 0.001***	-0.193	—	—	—	—	—
	(3.48)	(5.62)	(5.53)							
# tech classes	17.89	5.50	4.66	< 0.001***	1.808	-0.146	0.060	0.12	< 0.001***	-0.087
	(8.87)	(3.91)	(4.53)			(1.11)	(0.83)	(0.89)		
N	6,930	5,161	14,008			6,930	5,161	14,008		

Notes: Analysis period: 1976-1990. Nearest neighbor propensity score matching only. Donor p-scores in $(0.4, 0.8)$. Austin caliper: 0.2 times the pooled SD of logit p-scores ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A43: Baseline with Austin caliper, 1976-1990

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.069	-0.136*	-0.142*	-0.147	-0.131	-0.183*	-0.183*
	(0.076)	(0.078)	(0.086)	(0.095)	(0.086)	(0.096)	(0.096)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	11,248	8,320	8,320	8,320	8,079	7,913	7,913

Notes: Analysis period: 1976-1990. Baseline matching and Austin caliper. Standard errors are clustered at the firm level. ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A44: Propensity-Score Matching Diagnostics (Austin caliper)

	Statistic	Value
Panel A. Austin caliper		
SD of logit propensity score, treated	σ_T	2.842
SD of logit propensity score, controls	σ_C	3.051
Pooled SD of logit propensity score	σ_P	3.002
Austin caliper ($0.2 \times \sigma_P$)	Caliper	0.600
Panel B. Caliper diagnostic		
Candidate matches	N	5,243
Within caliper	N	3,361
Outside caliper	N	1,882
Panel C. Propensity-score distance among accepted matches		
Matched treated observations	N	3,361
Mean absolute propensity-score distance	Mean	0.055
Standard deviation	SD	0.052
Median	p50	0.037
75th percentile	p75	0.083
90th percentile	p90	0.136
95th percentile	p95	0.167
99th percentile	p99	0.193
Maximum	Max	0.200
Panel D. Accepted matches by treated overlap status		
Treated observations inside overlap region	N	3,111
Mean propensity-score distance	Mean	0.053
Standard deviation	SD	0.052
Median	p50	0.034
95th percentile	p95	0.166
Maximum	Max	0.200
Treated observations below overlap region	N	8
Mean propensity-score distance	Mean	0.096
Standard deviation	SD	0.056
Median	p50	0.133
95th percentile	p95	0.142
Maximum	Max	0.142
Treated observations above overlap region	N	242
Mean propensity-score distance	Mean	0.074
Standard deviation	SD	0.051
Median	p50	0.042
95th percentile	p95	0.179
Maximum	Max	0.194

Notes: Analysis period: 1976–1990. The table reports diagnostics for nearest-neighbor propensity-score matching with a caliper. Distance is the absolute propensity-score difference within matched pairs. Panel B reports retained treated observations relative to the 0.40–0.80 region; Panels C–D report distances among retained matches.

Table A45: Fixed pre-1983 Matched Pairs, 1976-1995

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.090	-0.208***	-0.174**	-0.158**	-0.168**	-0.230***	-0.225**
	(0.095)	(0.073)	(0.078)	(0.078)	(0.075)	(0.086)	(0.088)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
	15,535	14,393	14,393	14,393	14,081	13,927	13,927

Notes: Analysis period: 1976–1990. The table reports diagnostics for nearest-neighbor propensity-score matching with a caliper. Distance is the absolute propensity-score difference within matched pairs. Panel B reports retained treated observations relative to the 0.40–0.80 region; Panels C–D report distances among retained matches.

Table A46: Fixed pre-1983 Matched Pairs, 1976-1990

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.059	-0.169**	-0.195**	-0.197**	-0.221***	-0.299***	-0.293***
	(0.084)	(0.076)	(0.079)	(0.077)	(0.078)	(0.097)	(0.099)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	12,507	11,527	11,527	11,527	11,289	11,194	11,194

Notes: Analysis period: 1976–1990. The table reports diagnostics for nearest-neighbor propensity-score matching with a caliper. Distance is the absolute propensity-score difference within matched pairs. Panel B reports retained treated observations relative to the 0.40–0.80 region; Panels C–D report distances among retained matches.

Table A47: Spillovers: Alternative Control Matches (1976-1990)

	Spillovers			Notes
	(1)	(2)	(3)	
Baseline (1976-1995)	-0.209**	-0.322***	-0.293***	N: 15,070 (in full controls), donor
	(0.090)	(0.075)	(0.108)	p-scores $\in (0.4, 0.8)$
Panel B: 1976 - 1990				
Baseline	-0.162**	-0.262***	-0.361***	N: 12,556, donor
	(0.093)	(0.085)	(0.114)	p-scores within $(0.4, 0.8)$
Unrestricted Overlap	-0.158*	-0.251***	-0.339***	N: 15,446 (in full-controls)
	(0.081)	(0.076)	(0.115)	
Baseline + Austin caliper	-0.069	-0.136*	-0.183*	N: 7,913, Baseline match, donor restriction
	(0.076)	(0.078)	(0.096)	Austin: 0.2×(pooled SD of logit p-scores)
PSM only,	-0.023	-0.109	-0.203	N: 9,730, Propensity-score match,
(0.4,0.8) overlap + Austin	(0.125)	(0.125)	(0.154)	symmetric overlap restriction, Austin
Fixed pre-83 matches	-0.059	-0.169**	-0.293***	N: 11,194, restricting overlap to
	(0.084)	(0.076)	(0.099)	p-scores within 30th–90th percentiles
Baseline FE	✓	✓	✓	
Patent Creativity		✓	✓	
Full controls			✓	

Notes: Counts refer to full-control specifications. Baseline fixed effects are industry–year, technology–year, and firm–technology; full controls add filing-year effects by pre-treatment bins of creativity, employment, patenting, R&D, technological breadth, and age. "p1–p99 covariates" additionally includes spillover levels and trends. Donor restrictions precede matching. Where indicated, the Austin caliper is 0.2 pooled SD of logit propensity scores (Austin, 2011). Standard errors are clustered by firm ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

C.6 Excluding Contractors with Established Position

Table A48: Exclude Market Leader (by 4-digit SIC)

	Spillovers							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treated x Post	-0.214**	-0.214**	-0.269***	-0.184**	-0.233**	-0.184**	-0.186*	-0.186*
	(0.089)	(0.089)	(0.081)	(0.082)	(0.097)	(0.090)	(0.105)	(0.108)
Tech-Industry-Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Firm-Tech FE	✓	✓	✓	✓	✓	✓	✓	✓
Patent Creativity			✓	✓	✓	✓	✓	✓
Employees				✓	✓	✓	✓	✓
Patents filed					✓	✓	✓	✓
R&D						✓	✓	✓
# of CPCs							✓	✓
Agegroup								✓
N	15,876	15,876	11,122	11,114	11,114	10,789	10,733	10,733

Notes: Analysis period: 1976–1995. Firms funded while holding the highest pre-treatment (1976–1982) market share are excluded. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered by firm ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

Table A49: Exclude Technology Leader (by Patent Stock in CPC class)

	Spillovers							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treated x Post	-0.160**	-0.160**	-0.244***	-0.236***	-0.279***	-0.258***	-0.240**	-0.239**
	(0.080)	(0.080)	(0.081)	(0.080)	(0.097)	(0.094)	(0.106)	(0.109)
Tech-Industry-Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Firm-Tech FE	✓	✓	✓	✓	✓	✓	✓	✓
Patent Creativity			✓	✓	✓	✓	✓	✓
Employees				✓	✓	✓	✓	✓
Patents filed					✓	✓	✓	✓
R&D						✓	✓	✓
# of CPCs							✓	✓
Agegroup								✓
N	18,163	18,163	13,605	13,605	13,605	13,183	13,105	13,105

Notes: Analysis period: 1976–1995. Firms funded while holding the largest pre-treatment (1976–1982) patent stock are excluded; patent stock is defined as patents filed over the preceding 10 years. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered by firm ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$)

← Back to Robustness

Table A50: Exclude Technology Top-10% (by Patent Stock in CPC class)

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.402***	-0.411***	-0.228***	-0.265***	-0.237***	-0.283***	-0.282***
	(0.089)	(0.085)	(0.083)	(0.088)	(0.087)	(0.107)	(0.108)
Tech-Industry-Year FE	✓	✓	✓	✓	✓	✓	✓
Firm-Tech FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	8,199	8,199	8,194	8,194	7,960	7,900	7,900

Notes: Analysis period: 1976-1995. Firms funded while on top-10% of stock of patents in given cpc class during the pre-treatment period, 1976-1982, are excluded. Stock of patents is measured as the total of patents filled over past 10 years for any given year. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A51: Exclude Technology Leaders among Funded (by Patent Stock in CPC class)

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.322***	-0.328***	-0.275***	-0.336***	-0.300***	-0.279***	-0.286***
	(0.074)	(0.072)	(0.079)	(0.094)	(0.086)	(0.098)	(0.101)
Tech-Industry-Year FE	✓	✓	✓	✓	✓	✓	✓
Firm-Tech FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	13,385	13,385	13,380	13,380	12,995	12,913	12,913

Notes: Analysis period: 1976-1995. Firms funded while they had the highest stock of patents in given cpc class among funded units during the pre-treatment period, 1976-1982, are excluded. Stock of patents is measured as the total of patents filled over past 10 years for any given year. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

← Back to Robustness

Table A52: Exclude Top-2 Technology Leaders among Funded (by Patent Stock in CPC class)

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	0.102	-0.128	-0.212*	-0.229*	-0.261**	-0.427**	-0.462**
	(0.103)	(0.109)	(0.121)	(0.139)	(0.133)	(0.200)	(0.212)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	21,759	13,585	13,147	13,147	12,982	12,813	12,813

Notes: Analysis period: 1976-1995. Firms funded while on top-2 of patents in given cpc class among funded units during the pre-treatment period, 1976-1982, are excluded. Stock of patents is measured as the total of patents filled over past 10 years for any given year. Controls are resticted to be at least as large as bottom-25th percentile of treated firms in pre-treatment period. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A53: Spillovers: Exclude Technology Leaders (by R&D Stock)

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.236***	-0.222**	-0.189**	-0.236**	-0.224**	-0.226**	-0.229**
	(0.089)	(0.087)	(0.082)	(0.100)	(0.097)	(0.100)	(0.103)
Tech-Industry-Year FE	✓	✓	✓	✓	✓	✓	✓
Firm-Tech FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	13,910	13,910	13,906	13,906	13,517	13,430	13,430

Notes: Analysis period: 1976-1995. Firms funded while they had the highest R&D stock in given cpc class the pre-treatment period, 1976-1982, are excluded. Stock of R&D is measured as the total of R&D expenditures over past 10 years for any given year. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

← Back to Robustness

Table A54: Spillovers: Exclude Technology Top-10% (by R&D Stock)

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	-0.245***	-0.246***	-0.118	-0.146	-0.123	-0.133	-0.132
	(0.088)	(0.083)	(0.079)	(0.091)	(0.091)	(0.110)	(0.114)
Tech-Industry-Year FE	✓	✓	✓	✓	✓	✓	✓
Firm-Tech FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	10,110	10,110	10,107	10,107	9,808	9,749	9,749

Notes: Analysis period: 1976-1995. Firms funded while on top-10% of stock of R&D in given cpc class during the pre-treatment period, 1976-1982, are excluded. Stock of R&D is measured as the total of R&D expenditures over past 10 years for any given year. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A55: Spillovers: Exclude Technology Leaders among Funded (by R&D Stock)

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	0.031	-0.160	-0.252**	-0.268**	-0.276**	-0.413**	-0.419**
	(0.116)	(0.103)	(0.115)	(0.133)	(0.130)	(0.184)	(0.200)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	21,957	13,560	13,146	13,146	12,975	12,834	12,834

Notes: Analysis period: 1976-1995. Firms funded while they had the highest stock of R&D among funded units in given cpc class during the pre-treatment period, 1976-1982, are excluded. Stock of R&D is measured as the total of R&D expenditure over past 10 years for any given year. Controls are resticted to be at least as large as bottom-25th percentile of treated firms in pre-treatment period. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

← Back to Robustness

Table A56: Spillovers: Exclude Top-2 Technology Leaders among Funded (by R&D Stock)

	Spillovers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated x Post	0.008	-0.191*	-0.346***	-0.360***	-0.401***	-0.492***	-0.480**
	(0.106)	(0.110)	(0.117)	(0.137)	(0.132)	(0.189)	(0.196)
Baseline FE	✓	✓	✓	✓	✓	✓	✓
Patent Creativity		✓	✓	✓	✓	✓	✓
Employees			✓	✓	✓	✓	✓
Patents filed				✓	✓	✓	✓
R&D					✓	✓	✓
# of CPCs						✓	✓
Agegroup							✓
N	21,507	13,133	12,763	12,763	12,614	12,478	12,478

Notes: Analysis period: 1976-1995. Firms funded while on top-2 of R&D stock among funded units in given cpc class during the pre-treatment period, 1976-1982, are excluded. Stock of R&D is measured as the total of R&D expenditure over past 10 years for any given year. Controls are resticted to be at least as large as bottom-25th percentile of treated firms in pre-treatment period. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A57: Spillovers: Exclude Contractors with Established Position (1976-1990)

	Spillovers	Notes
Baseline	-0.339*** (0.115)	N: 15,466
Among all firm-technologies excluding:		
Market leader in 4-digit SIC	-0.353*** (0.119)	N: 15,421
Leader in patent stock in CPC class	-0.435*** (0.114)	N: 12,249
Top-10% in patent stock in CPC class	-0.305** (0.120)	N: 8,178
Leader in R&D stock	-0.313*** (0.113)	N: 12,578
Top-10% in R&D stock	-0.328*** (0.120)	N: 9,572
Among funded firm-technologies excluding:		
Leader in patent stock in CPC class	-0.288*** (0.108)	N: 12,088
Top 2 in patent stock in CPC class	-0.275** (0.110)	N: 12,568
Leader in R&D stock	-0.333*** (0.122)	N: 12,514
Top 2 in R&D stock	-0.432*** (0.122)	N: 12,140

Notes: The donor pool is restricted prior to matching to controls at least as large as the 25th percentile of the treated pre-treatment size distribution. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A58: Spillovers (excl. self-citations): Exclude Contractors with Established Position

	Spillovers	Notes
Baseline	-0.293*** (0.108)	N: 15,070
Among all firm-technologies excluding:		
Market leader in 4-digit SIC	-0.571*** (0.211)	N: 10,471
Leader in patent stock in CPC class	-0.435** (0.198)	N: 12,419
Top-10% in patent stock in CPC class	-0.290** (0.142)	N: 8,389
Leader in R&D stock	-0.461** (0.202)	N: 12,858
Top-10% in R&D stock	-0.361** (0.151)	N: 9,798
Among funded firm-technologies excluding:		
Leader in patent stock in CPC class	-0.373* (0.202)	N: 12,350
Top 2 in patent stock in CPC class	-0.500*** (0.191)	N: 12,845
Leader in R&D stock	-0.428** (0.200)	N: 12,841
Top 2 in R&D stock	-0.436** (0.200)	N: 12,460

Notes: The donor pool is restricted prior to matching to controls with propensity score matching within the 40th-80th percentiles of the distribution. Matching and estimation otherwise follow the baseline specification. Standard errors are clustered at the firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

← Back to Robustness

C.7 Synthetic Difference-in-Differences

Table A59: Synthetic D-i-D: Donor Weight Diagnostics

Treated units	Donor units	Positive donors	Mean weight	SD	Min	Median (p50)	p95	p99	Max
A. Matched Donor Pool									
1,598	3,003	3,003	0.000333	0.0000207	0.0003081	0.0003342	0.0003389	0.0004175	0.0006142
B. Full Sample									
1,736	14,008	14,008	0.0000714	0.00000764	0.0000608	0.0000709	0.0000776	0.0001025	0.000246

Notes: The table summarizes donor weights assigned by the Synthetic Difference-in-Differences (SDID) estimator on the matched donor pool. Donor units correspond to unique matched control firm–technology cells after deduplication. All donor units receive positive weight. The narrow distribution of donor weights indicates that the synthetic comparison group is constructed from a broad set of donor units rather than being driven by a small number of influential controls.

Table A60: Synthetic D-i-D: Spillovers (excl. indeterminates)

	Spillovers
ATT	-0.230^{***}
	(0.0215)
Observations	314680

Notes: The table reports estimates from the Synthetic Difference-in-Differences (SDID) estimator of Arkhangelsky et al. (2021). Estimation uses the full sample of eligible donor units. Inference is based on 500 placebo replications. Standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

← Back to Robustness