Innovation Policy and Inventors' Productivity: Evidence from Global AI Initiatives*

Leo Liu[1], Fariborz Moshirian[2], Vikram Nanda[3], Sheng Xu[2]

November 2025

[1]University of Technology Sydney, UTS Business School, Sydney, Australia
[2]University of New South Wales, UNSW Business School, Sydney, Australia
[3]University of Texas at Dallas, Naveen Jindal School of Business, Richardson, USA

Abstract

We construct novel datasets on AI innovation policies across 42 countries and global AI patents to examine the impact of these policies on the productivity of AI scientists. We find that scientists experience a *decrease* in productivity after their country initiates AI-supporting policies. We argue this decline could be driven by a shift in the nature of innovation being conducted; government support incentivizes scientists to pursue more novel and exploratory inventions. While these projects hold the potential to foster long-term growth, they are inherently characterized by longer development timelines and a greater risk of failure, resulting in a temporary decline in the average quantity and quality of innovations. This effect is potentially compounded by the government's inability to perfectly identify and fund the most promising radical innovations, leading to a misallocation of resources. We develop a general equilibrium framework that formalizes these dynamics. Our results highlight the transitional risks associated with government support for innovation, particularly as an economy navigates a new technological paradigm.

Keywords: Innovation, Patents, Artificial Intelligence, Public Policy, Scientists, Inventors

*Corresponding author: vikram.nanda@utdallas.edu (Nanda); Other authors: leo.liu@uts.edu.au (Liu), f.moshirian@unsw.edu.au (Moshirian), sheng.xu3@unsw.edu.au (Xu).

1 Introduction

Innovation is inherently risky and costly; consequently, it is often conducted at a socially suboptimal level as innovators tend to internalize these burdens. This creates significant scope for government intervention to correct market failures. Yet, we know relatively little about the effect of innovation-supporting policies that target select risky innovation projects. The current literature largely focuses on untargeted R&D tax incentives, which mitigate innovation underinvestment by lowering research costs but do not selectively support specific sectors or technologies, leaving project choices to the market (Howell, 2024).

In contrast, the quality of evidence on "mission-oriented" policies that focus on risky projects, such as the Apollo Project and DARPA, is low due to data and sample limitations (Bloom, Van Reenen, and Williams, 2019). However, these high-risk innovation projects play an increasingly important role due to the nature of modern technologies such as *artificial intelligence (AI)*. These technologies are often winner-takes-all which requires governments to pick potential "winners"—high-risk, high-reward projects that the private sector would otherwise underfund.

In addition, existing research on innovation is heavily concentrated in a few developed countries, including the United States, European nations, and Japan. This narrow geographical focus is a significant limitation, as it overlooks the growing importance of emerging economies such as China, which have become key contributors to the global innovation landscape.

To fill these gaps, we take advantage of the fact that there has been a surge in innovation policies specifically designed to foster Artificial Intelligence (AI). Since 2000, more than 200 AI-related government policies have been introduced worldwide to support AI development. These policies are mission-oriented by design: they target exploratory projects characterized by features such as high risk, high reward, and long time horizons with uncertain payoffs. For example, China's 2017 National New Generation AI Plan calls for "exploratory research where there is no consensus,"[1] and the U.S. AI R&D Strategic Plan emphasizes "long-term, high-risk, high-reward research"[2] — both aiming to foster transformative rather than incremental innovation. The U.S. National Robotics Initiative (NRI) in 2011 similarly prioritizes "research on co-robots that augment human capabilities" and supports "breakthroughs in autonomous decision making and perception," illustrating a focus on early stage, high risk innovation. Such policies are particularly important because

[1] See China's National New Generation AI Plan (2017) available at https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/

[2] See National Artificial Intelligence Research and Development Strategic Plan (2016) available at https://www.nitrd.gov/pubs/national_ai_rd_strategic_plan.pdf

AI is widely regarded as a transformative general purpose technology and a major driver of global economic growth in the twenty-first century. This paper evaluates the impact of these AI policies on innovation outcomes, focusing on scientists' productivity, as human capital is central to the innovation process (Lucas, 1988).

We construct two datasets enabling us to study this question. The first is AI policies from the OECD AI Policy Observatory,[3] which we refine to include only national-level initiatives.[4] Second, to measure innovation, particularly AI innovation, we utilize a global patent database from Google Patents, extending an established dataset of US AI patents (Giczy, Pairolero, and Toole, 2022) to a global context to assess AI development across countries.

Our primary finding is that the introduction of national AI support policies is associated with a *decline* in the innovation productivity of AI scientists. It is important to note that this is a relative decline, as estimated within a difference-in-differences (DID) framework. The absolute level of productivity in policy-adopting countries still increases, but at a slower rate than in comparable countries without such policies. Specifically, we document a reduction in the number of patents filed and a concurrent fall in their average quality, as measured by subsequent citations. This finding is robust across numerous specifications, and our validity checks show little evidence of confounding pre-existing trends. Furthermore, this pattern holds at the aggregate level: countries that implement AI policies subsequently experience a relative decline in their national AI patenting activity.

Our premise is that these results may be explained by the nature of the innovations being supported; AI projects are often explorative, "moon-shot" efforts and are inherently risky. For example, Cong, Lu, Shi, and Zhu (2024) finds AI follows a J-curved development path, as initial development is costly before firms can benefit from it. Brynjolfsson, Rock, and Syverson (2017) finds that initial AI development is subject to implementation lags — that it can be difficult to find use cases when an AI innovation first emerges. This is consistent with daily observation. The "Transformer" technology that backs today's large language models, for instance, first appeared in 2017. Despite being one of the most influential technologies of recent times, it did not achieve large-scale commercial success until ChatGPT was released in late 2022.

Another dynamic in AI is that it exhibits winner-takes-all properties (Babina, Fedyk, He, and Hodson, 2024). Training a large language model is very costly, requiring months

[3] Available at `https://oecd.ai`. Examples include the United States' *National Robotics Initiative (2011)*, *National AI R&D Strategic Plan (2016)*, and China's *New Generation AI Development Plan (2017)*, all embodying the mission-oriented design analyzed in this paper.

[4] The OECD database contains regional policies, some of which target specific groups or are unrelated to AI. We manually review each policy to ensure that it is (1) directly relevant to AI, (2) national in scope, and (3) supports innovation across all entities.

of training on Graphic Processing Unit (GPU) clusters, yet only the very top models (the "winners") can capture market share as they compete globally on single benchmarks. Pursuing AI therefore requires taking significant risks and paying high costs; success is otherwise unlikely. These dynamics suggest that supporting AI projects could potentially increase the risk profile of the innovation being conducted and, in turn, lead to a potential initial decline in innovation output.

To fix ideas, we develop a parsimonious general equilibrium model inspired by the endogenous growth literature (e.g. Romer, 1990; Aghion and Howitt, 1992; Acemoglu, Akcigit, Alp, Bloom, and Kerr, 2018). Our framework incorporates a key characteristic of modern AI policies: targeted support for high-impact, radical innovation, rather than general R&D tax relief.[5]

The model posits that while government policies can correct market failures related to radical innovation, they also introduce "transitional risks." These risks are twofold: First, by incentivizing a shift toward high-risk, high-reward projects, these policies increase the volatility of innovation outcomes, which can lower average quality in the short term as many novel attempts fail. Second, these risks are compounded by the government's imperfect ability to screen for and select the most promising projects.

Our empirical findings are broadly consistent with the predictions of this framework. We find strong evidence that after a country implements an AI policy, its scientists pursue more novel projects aimed at radical innovation. These novel projects are associated with longer development times and display significantly greater quality variance: on one hand, they are more likely to receive zero forward citations, suggesting a higher risk of complete failure; on the other, they are disproportionately likely to become high-impact "home run" innovations (e.g., top 1% by citations). Further supporting this, we find that AI scientists increasingly concentrate their efforts on a narrower set of technical domains, potentially indicating the deeper focus required for such breakthroughs. A relevant real-world example is Defense Advanced Research Projects Agency (DARPA)'s *Explainable AI (XAI)* program, launched in 2017 to advance interpretable machine-learning systems. The initiative funded a broad portfolio of university and industry teams developing methods for algorithmic transparency under significant uncertainty. Although many projects did not outperform standard deep-learning benchmarks, some produced influential conceptual frameworks and prototypes that advanced research on causal and interpretable AI(Gunning, Vorm, Wang, and Turek, 2021;

[5] In fact, very few of the AI policies are based on tax relief, in contrast with government support for general business R&D where tax is still the mainstream. See https://www.oecd.org/en/data/insights/statistical-releases/2025/04/rd-tax-incentives-continue-to-outpace-other-forms-of-government-support-for-rd-in-most-countries.html

Miller, 2019; Druce, Niehaus, Moody, Jensen, and Littman, 2021). This case suggests that mission-oriented programs yield uneven results, with only a subset of projects generating transformative advances.

These patterns also mirror the stated objectives of mission-oriented AI policies. For instance, China's *New Generation AI Development Plan* explicitly encourages "research with no consensus" and "cutting-edge, exploratory studies" in areas such as brain-inspired and quantum intelligence, while the *U.S. National AI R&D Strategic Plan* calls for "long-term, high-risk, high-reward research" that expands the frontiers of AI theory and methods. The NRI similarly emphasized "high-risk foundational research" and "transformative integration of robotics and intelligent systems." By design, these initiatives steer inventors toward riskier and more radical projects, inherently increasing the variance of innovation outcomes.

Our findings also support the model's assumption about imperfect government screening. While we find that governments tend to target projects with ex-ante indicators of novelty and often direct support toward high-impact inventors, this selection process is not perfect. Examining ex-post quality indicators suggests that government funding frequently fails to identify the highest-impact projects while supporting a handful of ventures that yield low-impact outcomes. This finding is consistent with the argument that government officials are typically not well-positioned to "pick winners," as they face the same or even greater information asymmetries than private sectors (Lerner, 2009; Howell, 2024).

Lastly, we study the welfare effects by looking at knowledge spillovers. Although the total amount of knowledge spillovers, measured by citations, appears to decline, we find that AI policies nonetheless promote greater international knowledge exchange. Specifically, a larger proportion of both backward and forward citations for a focal country's patents originate from researchers in other countries post-policy. This suggests that while the absolute volume of spillovers may decrease, policies enhance the relative share of cross-border knowledge exchange. Such findings align with policy objectives that emphasize global cooperation and standardization.[6]

Our study makes at least two contributions. First, we contribute to the innovation policy literature (see Bloom, Van Reenen, and Williams, 2019; Aghion, Bergeaud, and Van Reenen, 2023, for reviews) by offering novel evidence on modern "mission-oriented" policies. These policies are increasingly prevalent in high-tech sectors and around the globe but remain understudied. Our findings offer important policy implications, suggesting that

[6]For example, China's *New Generation AI Development Plan*, which calls for "actively participating in global research, development, and management of AI, and optimizing the allocation of innovative resources on a global scale"; the *U.S. National AI R&D Strategic Plan*, which stresses "developing shared public datasets and environments for AI training and testing"; and the NRI's focus on "interdisciplinary research partnerships and public dissemination of results."

policymakers should be mindful of potential short-term declines in innovation productivity, as the returns from such support may manifest in the longer term. While our conceptual framework provides a theoretical mechanism for this result, and our empirical analysis offers evidence consistent with it, we acknowledge that this may not be the sole explanation for the observed decline.

Second, we construct and introduce two new datasets valuable for future research. The first is a global AI patent dataset that extends the USPTO's AI patent classification system to multiple jurisdictions (Arora, Belenzon, Patacconi, and Suh, 2020). Given that AI development is a global phenomenon, this dataset is essential for international studies. The second is a refined dataset of national AI policies, built upon the OECD's AI Policy Observatory. While the original OECD catalog is a useful starting point, its varying scope limits its usefulness for empirical analysis. Our curation process addresses this by systematically filtering for national-level policies directly aimed at fostering AI innovation. Using these data, we provide stylized facts about global AI innovation and policies, opening avenues for future research.

2 Data

2.1 Constructing Global AI Patent Database

The first step in constructing our dataset was to develop a methodology to identify AI-related patent filings globally. We began by focusing on US AI patent filings using the Artificial Intelligence Patent Dataset (AIPD),[7] as provided by the USPTO (Giczy, Pairolero, and Toole, 2022).[8] We classify AI patents in accordance with the USPTO's guidelines.[9]

Next, we utilized US AI patents from the AIPD to identify international AI patent applications using semantic text vectors and a Support Vector Machine (SVM). We used

[7]The 2023 update of the AIPD identifies which of 15.4 million US patent documents (both patents and pre-grant publications, or PGPubs) published from 1976 through 2023 contain AI. These documents are separately categorized into eight AI component technologies from the AIPD: machine learning, vision, natural language processing, speech, evolutionary computation, AI hardware, knowledge processing, and planning and control.

[8]Note that this paper is written by Chief Economists at the USPTO.

[9]A patent application is deemed AI-related if its predictive AI score on AIPD was at least 0.5. AIPD classifies documents with a prediction score of 0.5 or higher as AI-related; documents scoring below this threshold are not considered AI-related. The distribution of these prediction scores is highly right-skewed. Consequently, increasing the classification threshold from 0.5 to 0.8 results in only a 5% reclassification of patents from AI to non-AI across the entire database. Our results are robust to this threshold up to a level of 0.9.

text vectors[10] provided by Google Patents to train an SVM using labeled AI patents from the USPTO. The trained SVM model was then able to identify AI patents in other patent offices.

Importantly, this classification is not US-specific because inventors around the globe file patents in the USPTO. In recent years, more than half of the inventors filing in the US were domiciled in other countries; therefore, the classifier is capable of detecting AI inventions globally for our purpose.

Furthermore, we categorized these patents into eight domains (machine learning, vision, natural language processing, speech, evolutionary computation, AI hardware, knowledge processing, planning & control) based on the highest prediction scores for these domains from AIPD. Finally, during the sample period from 2006 to 2019, we identified 1,568,818 AI patent filings with unique family IDs, including 455,461 filed by inventors from the United States.

2.1.1 Descriptive Statistics of AI Patent Applications

In analyzing our global AI patent applications, we provide a detailed breakdown at the country level in Table B3. The data show that the United States, Japan, China, and South Korea are among the most prolific, with approximately 455,461, 370,089, 302,758, and 152,437 applications, respectively. Furthermore, Figure 2 illustrates the time series variation in AI patent applications in 42 countries from 1990 to 2022. We find divergent trends between China and the United States: after 2018, China surpassed the United States to become the leading AI innovation by the number of AI patent applications, accounting for approximately 50% of global applications by 2019, compared to 20% for the United States. Historically, Japan led AI patent applications throughout much of the twentieth century, but experienced a decline from 1990 to 2022, with the United States surpassing it around 2002.

Furthermore, we examine the domain-specific distribution of AI patent applications, categorized into eight classifications by the AI Patent Database (AIPD), as shown in Table B3. The majority of patent filings from the United States (32.55%), China (30.70%), South Korea (27.50%) ,Germany (37.86%) and the United Kingdom (36.03%) fall under

[10] In the field of natural language processing, these vectors are known as text embeddings. Google provides 64-dimensional representations to capture the semantic meaning of a patent document.

"Planning and Control."[11] In contrast, Japan, Taiwan and Netherlands primarily focus on "Vision" technologies, which account for 34%, 35% and 30%of their respective applications.

2.2 Measuring Scientist-level AI Productivity

We use scientists (inventors) as the unit of analysis, following the innovation literature (e.g. Moretti, 2021), because they are the cornerstone of innovation.[12] Furthermore, this focus aligns with the nature of the policies studied, which typically support AI research and thus directly impact AI researchers. Having identified global AI patent applications, we use this dataset to measure the productivity of AI inventors. We define AI inventors as those individuals listed on at least one identified AI patent application. To measure productivity, we employ two metrics for each inventor i and year t.

The first metric is a simple count: the number of AI patent applications filed by inventor i in year t. Recognizing that patent value is heterogeneous, Our second metric is citation-weighted counts which is the sum of forward citations[13] received by inventor i's patents that were filed in year t. Citations are counted up to the end of 2024.

Patent citation data are subject to truncation issues which are a well-acknowledged challenge in innovation studies (Hall, Jaffe, and Trajtenberg, 2005). To mitigate this, our analysis period for patent filings extends only up to 2019. This cutoff provides a five-year window for patent applications filed by the end of 2019 to potentially be granted and to accumulate forward citations. This is generally considered adequate, as forward citations typically begin accruing relatively quickly after patent publication. Following Hall, Jaffe, and Trajtenberg (2005), we add year fixed effects to all our empirical tests to adjust for systematic differences in citation patterns across years.

Table 1 presents summary statistics for the key variables used in our analysis, which are constructed at the inventor-year level. It is important to note that the number of observations

[11] The AIPD classifies AI patents into eight categories: Planning & Control: processes for identifying, creating, and executing activities to achieve specified goals; Knowledge Processing: representing and deriving facts about the world for use in automated systems; AI Hardware: physical computing components designed for increased processing efficiency and speed; Computer Vision: extracting and interpreting information from images and videos; Machine Learning: computational models that learn from data; Natural Language Processing: understanding and processing written language data; Speech Recognition: techniques for interpreting spoken language, including responding to commands; Evolutionary Computation: computational routines inspired by natural evolution.

[12] It's no exception for AI innovation. Anecdotal evidence points to intense competition for top talent in the AI industry. In mid-2025, reports surfaced of Meta offering compensation packages reaching $100 million to poach researchers from competitors like OpenAI. This "cash-first" approach, often personally championed by CEO Mark Zuckerberg, has seen him directly negotiate with top-tier talent, reportedly tabling offers as high as $200 million to lure individuals away from rival labs.

[13] Forward citations refer to the number of times a focal patent is cited by subsequent patents.

reported here is smaller than that in our regression tables given we use "stacked" regressions, as detailed in the methodology section. Variable definitions are in Table B1.

2.3 National AI-Related Strategies and Policies Data

Our primary data source for analyzing national AI strategies and policies is the OECD AI Policy Database. This comprehensive dataset provides detailed information on AI-related policies across countries, including average annual budgets, policy types, and implementation periods. Policies are classified into four main categories: *governance*, *financial support*, *guidance & regulation*, and *AI enablers & other incentives*, each of which contains several subtypes.[14] Each category represents a different type of policy approach that shapes how governments support, regulate, and guide AI-related innovation efforts. Governance policies provide strategic direction and institutional oversight, aligning AI development with public priorities through coordination and state-led initiatives. Financial support policies act as direct subsidies for innovation inputs, including public R&D funding and capital investment. Guidance and regulation policies aim to manage externalities by establishing legal and ethical standards for AI development. AI enablers and other incentives focus on building long-term innovation capacity through investments in infrastructure, human capital, and collaborative ecosystems.

For instance, the governance category includes four subtypes: national AI strategies, coordination and monitoring bodies, public consultations, and public sector AI adoption. In our main analysis, we construct a *treatment-post* dummy variable for each policy and policy category, coded as 1 during the period of implementation and 0 otherwise. As part of our robustness checks, we also incorporate each policy's average annual budget. Due to the lack of granular financial allocation data, we assume an equal distribution of budget across subcategories within each policy type.

The original dataset contains over 1,000 unique policies from 69 countries, spanning 1969 to 2024. Since some policies are general innovation policies, rather than specifically AI-focused, we use ChatGPT GPT-4o (see Table B5 for exact prompt used) to analyze each

[14] Governance includes national AI strategies, coordination and monitoring bodies, public consultations with stakeholders or experts, and the use of AI in the public sector. Guidance and Regulation covers standards and certification for technology development and adoption, labor mobility regulations and incentives, regulatory oversight and ethical advisory bodies, and emerging AI-related regulations. Financial Support includes institutional funding for public research, project grants, business R&D grants, innovation-related loans and credits, centers of excellence, procurement programs for AI R&D and innovation, fellowships and scholarships, equity financing, and indirect financial support. AI Enablers and Other Incentives encompass innovation challenges, prizes and awards, knowledge transfer and advisory services, networking and collaborative platforms, AI computing and research infrastructure, data access and sharing, public awareness campaigns, civic participation activities, labor market policies, and AI skills and education programs.

policy description and identify those related to AI. Each description was evaluated ten times, and we retained only policies that were confirmed in at least seven out of ten evaluations as: (1) AI-related, (2) national in scope, and (3) applicable to all entity groups. Point (2) and (3) are to ensure that the selected policies align with our dependent variable, which includes all inventors.

This process yielded a refined dataset of 188 AI-specific policies from 42 countries and regions, spanning 2011 to 2022. We further enhanced data quality by employing a team of research assistants to incorporate reliable government budget disclosures and reports, including information from social media sources. For example, many US AI policies lacked explicit budget information; therefore, we gathered data on annual federal non-defense AI R&D investments from the Networking and Information Technology Research and Development Program (NITRD). Since the FY2020 budget request, NITRD has reported these investments annually in response to Executive Order 13859, "Maintaining American Leadership in Artificial Intelligence."

2.3.1 Descriptive Study of AI Policies

Our analysis includes national AI-related policies from 42 countries, covering the period 2011 to 2023. Figure 1 displays a word cloud generated from the descriptions of the national AI policies examined, highlighting key terms such as "artificial intelligence,“ development", "research," "government," and "national". Further quantitative details on these policies are provided in Table B2. This table summarizes policy statistics, including average annual budgets categorized into areas like governance, guidance & regulation, financial support, and AI enablers & other incentives (alternatively termed non-monetary incentives). Notably, the United States and Saudi Arabia allocate the largest budgets, at 6,546 million and 4,936 million USD, respectively. Finally, to illustrate the qualitative focus of different national strategies, Table B6 summarizes policy descriptions for a set of major countries.

For the US, we identify the United States' National Robotics Initiative (NRI), launched in 2011, as the first major national-level AI policy in our analysis. According to the National Artificial Intelligence Research and Development Strategic Plan (2016), the NRI is explicitly highlighted as one of the earliest strategic R&D initiatives relevant to AI.[15] Although the plan does not explicitly label the NRI as the first AI-related policy, it documents the NRI's start date as earlier than other AI initiatives, suggesting its important role in shaping subsequent national AI policy frameworks. This pattern persists with our countries as well, we observe

[15] See National Artificial Intelligence Research and Development Strategic Plan (2016) available at https://www.nitrd.gov/pubs/national_ai_rd_strategic_plan.pdf. The relevant discussion is from page 6 to page 7.

countries continue to implement AI promoting policies after the first one.

Additional evidence further supports identifying the NRI as an initial policy shock. The National Artificial Intelligence Research and Development Strategic Plan (2016) documents a notable increase in AI-related publications and patents — particularly in deep learning — following the initiative's introduction,[16] implying a stimulative effect on AI innovation. Complementing this finding, we also document a significant increase in AI-specific research funding provided by the National Science Foundation (NSF) after the NRI's launch. Figure 5 shows that the proportion of NSF awards dedicated to AI research rose from approximately 7% in 2010 to around 10% in 2012, with a sharper increase immediately after the 2011 introduction of the NRI. This structural shift provides empirical support for treating the NRI as the initial national-level AI policy shock, marking the increased governmental commitment and public investment in AI research and development.

Figure 3 outlines the evolution of major U.S. AI-related policies. Following the NRI, the next pivotal policy was the National Artificial Intelligence Research and Development Strategic Plan (2016), which explicitly guided government investment toward underfunded research areas and aimed to strengthen the AI talent pipeline. Since 2016, numerous policies have emerged, including the Executive Order on Maintaining American Leadership in AI, reinforcing the government's strategic commitment to AI. The most recent example at the time of writing, although outside our sample period, is President Trump's Executive Order 'Removing Barriers to American Leadership in Artificial Intelligence', which creates a supportive environment for initiatives like Stargate by prioritizing deregulation and strengthening the leadership of American AI. Stargate is a $500 billion public-private partnership funded by a combination of government and private investment.

For emerging countries also major player in AI like China, its *National New Generation AI Plan*, launched in 2017, outlines ambitious targets for AI R&D and aims to establish China as a global leader in AI by 2030. This plan is not only advancing R&D but also accelerating industrialization and talent development. Meanwhile, Canada's *Pan-Canadian AI Strategy* of 2017, with a CAD 125 million investment, focuses on retaining and attracting top academic talent, improving research capacity and fostering commercialization. Additionally, Singapore's *AI Research Program* in 2018 launched high-quality research in fundamental AI technologies, encouraging national collaborations and nurturing local AI talents. Similarly, Japan's *High Performance Computing Infrastructure Project*, featuring the Fugaku supercomputer since 2012, supports extensive AI research capabilities.

The first country by policy budget, Saudi Arabia's *Saudi Data and Artificial Intelligence*

[16] See the relevant discussion available at https://www.nitrd.gov/pubs/national_ai_rd_strategic_plan.pdf from page 13 to page 14.

Authority (SDAIA), established in 2019, is the lead entity responsible for advancing the nation's data and AI agenda. Mandated with unlocking the value of data and AI, SDAIA aims to elevate Saudi Arabia into the elite league of data-driven economies. Its work spans innovation, infrastructure, and capability-building, while supporting the values-based *G20 AI Principles* and fostering a robust digital ecosystem. SDAIA oversees three key entities: the National Information Center, the National Data Management Office, and the *National Center for Artificial Intelligence (NCAI)*—the latter also launched in 2019 to orchestrate AI research, develop scalable solutions, advise the government on AI strategies, and promote AI education and public awareness. With its *National Strategy for Data and AI* launched in 2020, SDAIA seeks to position Saudi Arabia as a global leader in AI by 2030 through adoption, international collaboration, and human capital development.

One of the primary objectives of this paper is to evaluate the impact of these policies. However, it is important to recognize their substantial differences in scope, design, and implementation. Consequently, our analysis focuses on estimating the *Average Treatment Effects (ATE)* across the sample, offering a broad understanding of their overall influence. In later sections, we complement this aggregate analysis by examining the effects of individual countries' policies where feasible.

3 Empirical Method and Baseline Results

3.1 Empirical Model

A significant development in "Difference-in-Difference" methods since Bertrand, Duflo, and Mullainathan (2004) is that the two-way fixed effect (TWFE) approach is regarded as potentially biased when early-treated observations serve as a control for later-treated observations (Baker, Larcker, and Wang, 2022; Callaway and Sant'Anna, 2021; Sant'Anna and Zhao, 2020). Hence, we start our baseline analysis by using the recommended stacked "difference-in-differences" (DID) regression (see, e.g., Gormley and Matsa, 2011; Cengiz, Dube, Lindner, and Zipperer, 2019; Deshpande and Li, 2019; Baker, Larcker, and Wang, 2022):

$$E(Y_{ijct}) = exp(\beta_0 + \beta_1 \text{AI Policy}_{jct} + \phi_{ic} + \lambda_{tc} + \epsilon_{icjt}) \tag{3.1}$$

In this model, y represents one of several dependent variables of interest for inventor i in year t. The variable *AI Policy* is an indicator equal to one if inventors in country j are subject to a national AI-related policy implemented in or before year t. Inventor locations j are determined from the home addresses recorded on their patents. We define the post periods as all years from the implementation of a country's first national AI policy onwards;

this reflects the observation that countries typically maintain continuous support for AI and enact follow-on policies. Given that our dependent variables are count data, we estimate the model using Poisson regression (Cohn, Liu, and Wardlaw, 2022).

In the stacked difference-in-differences (DID) approach, each treated country is paired with control countries that are either "yet to treat" or "never treated". These control groups are unaffected by bias from the traditional two-way fixed effects (TWFE) estimator, as discussed above. The treated countries and their corresponding clean controls define a cohort. For an inventor in a treated country, this indicator switches from zero to one when the first national AI-related policy is implemented.

We include inventor $\times$ cohorts (ϕ_{ic}) fixed effects to control for any fixed differences between inventors and cohort $\times$ year (λ_{tc}) fixed effects to account for time trends at the cohort level. Standard errors are clustered at the country level, following Bertrand, Duflo, and Mullainathan (2004). This clustering is essential because our unit of analysis is at the inventor level with a large number of observations. Clustering at the country level ensures that statistical significance is not driven solely by the large N. Instead, test statistics are computed based on the number of countries. Finally, β_1 is the DID estimator, with fixed effects that absorb individual terms such as $Treat$ and $Post$.

We also estimate the structural approach as in Callaway and Sant'Anna (2021) to address the biases in TWFE, reported in Panel B of Table 2. We use stacked DID as our main specification due to its flexibility. The Callaway and Sant'Anna (2021) method requires inverse probability weighting using OLS; however, estimating with count data necessitates a count data model. Therefore, we estimate Callaway and Sant'Anna (2021) using the logarithm of the dependent variable which is count data. While the results are consistent, their interpretation requires caveats due to potential biases from the log transformation of the dependent variable (Cohn, Liu, and Wardlaw, 2022).

The key identifying assumption for difference-in-differences (DID) analyses is the parallel trends assumption, which requires that the treated and control groups exhibit similar trends in the dependent variable prior to the intervention. For all dependent variables analyzed, the trends—specifically, the interaction of the treatment indicator with time periods from t-3 to t+3—are reported in Table B8. Overall, no significant pre-treatment trends are found. The DID graph is also plotted with Callaway and Sant'Anna (2021) in Figure 4a, and again the results are consistent.

3.2 Baseline Results

Our findings indicate a decline in both the quantity and quality of innovation output following the implementation of national AI policies. Specifically, Panel A of Table 2 presents our DID estimates of the national AI policy's impact on inventor productivity and patent quality. A key finding is a negative effect on both the quantity and quality of inventor output. Specifically, Column 1 (Panel A) indicates the policy is associated with a 9.1% relative decrease in the number of AI patents filed by treated inventors post-implementation. It is important to interpret this within the DID framework. This result does not mean that absolute productivity fell in treated countries; rather, it indicates that productivity for inventors in treated countries increased less than it did for inventors in control countries after the policy implementation. In fact, inventors in both treated and control groups show an increasing trend in absolute productivity over time. As shown in Figure 4b, inventors in both groups exhibited an upward trend in absolute productivity over time after the policy implementation; however, the increase was smaller for inventors in treated countries relative to those in the control countries.

Our results indicate a decline in the quantity of innovation output. While one might hypothesize that this is driven by a shift toward fewer, higher-quality innovations, the evidence from Panel A of Table 2 contradicts this notion. In fact, our findings suggest that patent quality also declines. Specifically, Columns 2 and 3 show the policy is associated with decreases of 37.6% in citation-weighted count and 30% in average forward citations, respectively.

In Panel B of Table 2, we employ the method from Callaway and Sant'Anna (2021) to address potential biases in standard TWFE models. The results are qualitatively consistent with our primary findings. Across all columns, the introduction of AI policies is associated with a statistically significant decline in AI patent citation-weighted count, with estimates ranging from -18.2% to -38.8%, depending on the specification. Because the log transformation of the dependent variable in these models makes interpretation of the magnitude difficult, we rely on Panel A for economic magnitude, while Panel B serves as a robustness check.

Furthermore, we assess the robustness of our findings using alternative model specifications. First, in Table B9, we include the natural logarithm of GDP per capita as an additional control in the baseline Poisson regression to account for cross-country differences in economic development and underlying innovation capacity. Second, in Table B10, we replace the AI policy indicator with the natural logarithm of the policy budget as the main independent variable to better capture the intensity of government support. Finally, in Table B11, we estimate an ordinary least squares (OLS) regression using the AI policy indicator, where the dependent variables—AI patent metrics—are log-transformed to improve model fit and facilitate interpretation in percentage terms. Across all specifications, the results remain qualitatively consistent with our baseline findings.

Furthermore, we assess the robustness of our findings using alternative model specifications. First, in Table B9, we include the natural logarithm of GDP per capita as an additional control in the baseline Poisson regression to account for cross-country differences in economic development and innovation capacity. Second, in Table B10, we replace the AI policy indicator with the natural logarithm of the policy budget as the main independent variable to capture the intensity of government support. Finally, in Table B11, we estimate an ordinary least squares (OLS) regression using the AI policy indicator, where the dependent variables—AI patent metrics—are transformed using natural logarithms. Across all specifications, the results remain qualitatively consistent with our baseline findings.

3.2.1 Heterogeneous Effect of National AI-Related Policy on AI Innovation

We examine the effects of four distinct types of national AI policies, as defined by the OECD. First, *Governance* policies focus on overarching national strategies and institutional oversight. Second, *Financial Support* involves direct funding mechanisms, such as R&D grants and investment capital. Third, *Guidance and Regulation* establishes legal and ethical frameworks for AI development. Finally, *AI Enablers and Other Incentives* aim to build long-term capacity through investments in infrastructure, talent development, and collaborative platforms.

Our analysis indicates that the most of these policy types are associated with a decline in innovation productivity, as presented in Table B12. Policies focused on AI Enablers show the largest negative impact, potentially because their broad scope dilutes resources and their benefits take longer to materialize. Similarly, Financial Support and Guidance and Regulation are linked to significant declines; the former may encourage safer, incremental projects over radical breakthroughs, while the latter can deter novel research through burdensome compliance. The sole exception is Governance policies, which show no significant adverse effect on productivity.

3.3 Country-Level Analysis

A key question is whether these micro-level dynamics translate into a broader decline in national innovation output. To answer this, we examine the impact of national AI policies on country-level innovation aggregates. The Poisson estimates, presented in Panel A of Table 3, suggest a significant decline in both the quantity (patent counts) and quality (citation-weighted counts) of patents following policy adoption. Crucially, the pre-trend analysis in Panel B indicates that this decline is not driven by pre-existing downward trends, as there is no evidence of a negative trajectory prior to the implementation of the policy.

We then explore alternative explanations for this national-level decline beyond the observed drop in individual inventor productivity. One possibility is that the decline reflects a reduction in the total number of AI inventors. Although the policy's effect on inventor counts is negative, as shown in Panel A, Column 4, this relationship appears to be driven by selection rather than causation. Specifically, the Poisson estimate of –0.292 in Column 4 of Panel A implies a 25.32% reduction in the probability of having an additional AI inventor following policy adoption. Moreover, Column 4 of Panel B shows that countries already experiencing a downward trend in inventor counts were more likely to implement AI-related support policies. These findings suggest that the decline in inventor counts cannot be causally attributed to the policy itself.

A second potential alternative explanation is that national AI policies may induce an inflow of lower-impact inventors and crowd out high-impact inventors, thereby reducing aggregate innovation output. To examine this possibility, we analyze the composition of the inventor workforce based on pre-policy impact. Specifically, inventors are classified annually as high-impact if their AI patents rank in the top 10% in terms of citation-weighted count over the preceding five years. We compute the high-impact inventor share as the number of high-impact inventors divided by the total number of unique inventors filing AI patents in a given year; the remainder constitutes the low-impact inventor share. As reported in Table 3, we find no significant changes in the shares of high- or low-impact inventors following policy adoption, suggesting that a shift in inventor composition is unlikely to explain the observed decline in innovation output.

Overall, after investigating these alternative channels, our findings suggest that the decline in individual inventor productivity is a plausible contributing factor to the observed decline at the national level.

3.4 Placebo Test Using Non-AI Patents

Although the parallel trends assumption holds in our pre-treatment periods, this does not rule out the possibility that an unobserved confounding factor could simultaneously affect the treated countries precisely at the time of their treatment, biasing our results. To lend further confidence in the link between policy and innovation output, we conduct a placebo test using non-AI patents as the outcome variable. The underlying logic is that national

AI policies should primarily affect AI innovation, with smaller or little impact on other technological domains.

The results are presented in Table B13, which reports the policy effects on non-AI patent counts, citation-weighted counts, and average forward citations using both our stacked DID (Panel A) and the Callaway and Sant'Anna (2021) approach (Panel B). We find that the estimated policy effects on non-AI patents are negative but mostly statistically insignificant, especially in our baseline stacked DID specification.

While we observe some small, statistically significant negative effects in the more sensitive Callaway and Sant'Anna (2021) estimations, this is not entirely unexpected. AI innovations are often built upon non-AI technologies and can spur new applications in other fields (Babina et al., 2024). Therefore, a policy-induced reduction in AI innovation could plausibly create negative spillovers to non-AI domains. Nevertheless, the economic and statistical significance of these effects on non-AI patents is substantially smaller than the impact we document for AI patents.

3.5 Placebo Simulation: Randomized Policy Timing

To ensure that our results are not driven by random policy timing, we conduct placebo simulations using the same stacked difference-in-differences (DID) specification as in Table 2. In each iteration, we randomly assign pseudo-policy adoption years across countries and re-estimate the model. This procedure is repeated 100, 500, and 1,000 times, and we record the corresponding t-statistics for the placebo treatment effects. In Figure B1, the simulated distribution centers around zero, whereas the actual estimate lies far in the left tail. Only about 3% of the simulations produce a negative and statistically significant effect comparable to our baseline results, indicating that the observed findings are highly unlikely to have occurred by chance.

4 Conceptual Framework

Our results so far present an intriguing case that innovation policies that designed to promote innovation has in fact negative effect. To probe the reasons behind this negative effect, we develop a simple general equilibrium framework. Classic endogenous growth models (e.g., Romer, 1990) emphasize that technological progress stems from intentional R&D investments, suggesting a valuable role for policy guidance. Schumpeterian growth models (e.g., Aghion and Howitt, 1992) further highlight the role of creative destruction, where radical innovations displace obsolete technologies. Building on these insights, our model incorporates a key feature of modern innovation policy as discussed earlier: targeted support for radical innovations, as opposed to broad, untargeted R&D subsidies. We distinguish between two types of innovation: radical innovations, which offer the potential for significant gains but carry inherent risks, and incremental innovations, which provide safer, more modest returns.[17] Our model incorporates heterogeneous inventors, consumers, firms, and government interventions to assess how policy affects inventors' strategic choices and aggregate innovation outcomes.

The central insight from the model is that by subsidizing riskier projects with higher potential returns, governments can induce a strategic shift among inventors. Talented researchers, who might otherwise pursue safer projects, are incentivized to undertake more ambitious, "moonshot" endeavors. This shift toward riskier projects naturally leads to a higher rate of failure, explaining the observed decline in average productivity. However, it also increases the likelihood of major breakthroughs, creating a higher variance in innovation outcomes. The framework thus provides a rationale for why policies aimed at fostering long-term technological leadership might generate seemingly negative short-term outcomes. In an extension, we show how these outcomes can be compounded by the government's imperfect ability to screen for and select the most promising projects, which can lead to a misallocation of resources.

We proceed by laying out the economic environment, including the key agents and their objectives. We then derive the inventors' optimal innovation strategy and make sure the general equilibrium market-clearing conditions satisfy. Using this model, we analyze how government grants targeting radical R&D alter the equilibrium allocation of talent between incremental and radical innovation. We conduct comparative statics to explore these policy effects, drawing on insights from the literature on endogenous growth (Romer, 1990; Aghion and Howitt, 1992) and innovation incentives (e.g. Acemoglu et al., 2018). Finally, we extend the analysis to show how these policies address market failures and how their effectiveness interplays with the government's imperfect information about innovation quality. For convenience, we provide a table detailing the parameters in model and their definitions in Table B4.

[17] In fact, this way of modeling innovation is not new and has been heavily used in policy and innovation studies (e.g. Aghion, Bergeaud, and Van Reenen, 2023; Atkeson and Burstein, 2019).

4.1 Model Setup

4.1.1 Players

The model economy is populated by four types of agents: a continuum of heterogeneous inventors, perfectly competitive final goods firms, consumers, and a government. Time is discrete and, for the core of the analysis, we focus on a single period to cleanly illustrate the static choice problem faced by inventors.

Consumers: They represents the general population or households of the economy. This group constitutes the production workforce, supplying a total of L units of labor inelastically to firms. Their primary economic activity is to power the production of physical goods and services; they do not engage in innovation. A perfectly competitive final goods sector uses this labor to produce a single homogeneous good, Y_t, using a linear production function:

$$Y_t = A_t L_Y \tag{4.1}$$

where A_t represents the economy's aggregate stock of technology, and L_Y is the amount of production labor employed. Firms in this sector are price-takers and maximize profits, which leads to a simple determination of the equilibrium wage for production labor:

$$w_t = A_t \tag{4.2}$$

This wage, w_t, represents the baseline opportunity cost of labor in the economy and serves as a benchmark for the returns to innovation. This setup also ensures the goods and labour market clears so we can abstract away from market clearing conditions and focus on inventor choices.

Inventors: They are a specialized, high-skill class of individuals completely separate from the production workforce: the inventors. This population, which can be thought of as the economy's dedicated R&D specialists, scientists, and engineers, is normalized to a unit mass of one for convenience. Their sole economic function is to create new knowledge assets (i.e., patents and ideas). The income they earn from this activity, whether the safe return ϵ or the radical reward V_R, represents the present market value of the future profits these assets are expected to generate. This income is thus a claim on future output, financed by the capital market today, allowing it to coexist with the production workers' claim on current output. This modeling choice, which separates the agents of innovation from the agents of production, is a standard feature in Schumpeterian growth models and allows for a focus on the allocation of specialized innovative talent (Akcigit and Kerr, 2018).

Each inventor i possesses a unique, innate talent level, p_i, which determines their personal probability of success on the radical path. This parameter reflects the widely accepted notion that innovative potential is not uniform, and that individual characteristics are primary drivers of innovation outcomes. We assume that p_i is drawn from a known, continuous, and differentiable cumulative distribution function (CDF) $F(p)$ with support on $[0, \overline{p}]$, where $\overline{p} < 1$ to rule out certain success for even the most talented inventor.

Firms: A sector of perfectly competitive firms produces a single homogeneous final good, Y_t. Production requires two inputs: labor, L_Y, and the economy's aggregate stock of technology, A_t. We assume a linear production function, $Y_t = A_t L_Y$. Firms act as price-takers in both the output and labor markets, maximizing profits. The role of firms in this setup is to link the labor and goods markets: they hire production labor from consumers and produce the final good, which is then consumed, ensuring both markets clear. In this capacity, firms serve to determine the equilibrium wage. Profit maximization under perfect competition ensures that the wage is equal to the marginal product of labor, yielding Equation 4.2 as discussed earlier.

Government: The government's objective is to promote technological advancement for productivity gains, particularly for radical innovation where market failures are presumed to exist. The government offers a subsidy, G, for radical innovation, financed by lump-sum taxes on consumers. This allows us to isolate the incentive effect of the policy on the direction of innovation. The grant acts as a form of insurance, reducing the private cost of failure and encouraging risk-taking.

4.1.2 Additive Technological Progress

While the model is static, the equilibrium allocation of talent has direct consequences for the expected growth of the economy. We can formalize the law of motion for technology by assuming an additive process where incremental and radical innovations contribute differently to progress. Let δ be the small, certain productivity gain from an incremental success, and let γ be the large productivity gain from a radical success, with $\gamma \gg \delta$. The expected level of technology in the next period, A_{t+1}, is then given by:

$$E[A_{t+1}] = A_t + (1 - \psi^*) \cdot \delta + \psi^* \cdot E[p_i | p_i \geq p^*] \cdot \gamma \tag{4.3}$$

This equation makes the government's long-term objective explicit. By increasing the subsidy G, the government raises ψ^*, the faction of inventors who pursue the radical path (as defined in the next section in Equation 4.7)., knowingly reducing the certain, short-term gains from incremental innovation ($(1 - \psi^*)\delta$) in order to increase the chances of a high-impact, radical breakthrough ($\psi^* \cdot E[p_i | p_i \geq p^*] \cdot \gamma$) that drives long-run growth.

4.2 The Inventor's Optimization Problem and Implication

As discussed earlier, when peruse innovation projects, the inventor face two choices:

Incremental Innovation: By pursuing an incremental innovation, an inventor is guaranteed a small, certain payoff, we posit that the payoff is at least as large as the production wage: $\epsilon \geq w_t$. This implies that inventors are a specialized group whose opportunity cost is engaging in other R&D, not production work. For simplicity, we set $\epsilon = w_t$ without loss of generality as setting any $\epsilon \geq w_t$ would not alter model's key results.

Radical Innovation: An inventor i choosing the radical path incurs a uniform private cost c to undertake a high-risk project. The project succeeds with the inventor's idiosyncratic probability p_i, generating a large private value V_R (e.g., the capitalized value of monopoly profits from a patent). If the project fails, which occurs with probability $(1 - p_i)$, the return is zero.

4.2.1 The Indifference Condition

A risk-neutral inventor i will choose the radical innovation path if and only if the expected payoff from doing so, inclusive of the government grant, exceeds the payoff from their safe outside option. The decision rule is to choose radical innovation if:

$$p_i \cdot V_R - c + G \geq w_t \tag{4.4}$$

Given the continuous distribution of talent p_i, there will exist a marginal inventor, denoted by a cutoff probability p^*, who is exactly indifferent between the two paths. For this inventor, the expected utility from each path is equal:

$$p^* \cdot V_R - c + G = w_t \tag{4.5}$$

This equation represents the central trade-off in the model. It crisply captures how an inventor's ability (p^*), the rewards to innovation (V_R, w_t), the costs (c), and government policy (G) interact to determine the innovation choice. We can solve this equation directly for the equilibrium cutoff probability:

$$p^* = \frac{w_t + c - G}{V_R} \tag{4.6}$$

All inventors with a higher intrinsic success probability ($p_i \geq p^*$) will strictly prefer the risky but potentially lucrative radical path. Conversely, all those with a lower success probability ($p_i < p^*$) will opt for the safety of incremental innovation.

4.2.2 Equilibrium Allocation of Talent

The equilibrium fraction of the inventor population that chooses to pursue radical innovation, which we denote by ψ , is the measure of all inventors whose success probability is at or above the cutoff p^*. This is given by:

$$\psi^* = \int_{p^*}^{\overline{p}} f(p)dp = 1 - F(p^*) \tag{4.7}$$

where $F(\cdot)$ is the cumulative distribution function of success probabilities. The remaining fraction of inventors, $1 - \psi^*$, engages in incremental innovation.

Since we separate the inventor population from the production labor force, the allocation of innovative talent is determined within the community of inventors, as described above. The production labor market clears separately, with the full labor supply L being employed in the final goods sector ($L_Y = L$), determining the wage $w_t = A_t$.

4.2.3 Comparative Statics

An increase in the grant G lowers the required success probability p^*:

$$\frac{\partial p^*}{\partial G} = -\frac{1}{V_R} < 0 \tag{4.8}$$

As the grant increases, the required success probability p^* for a radical project to be privately viable falls. This induces a marginal group of inventors—those with abilities p_i just below the old cutoff but above the new one—to switch from the safe incremental path to the risky radical path. Consequently, the fraction of inventors pursuing radical innovation, $\psi^* = 1 - F(p^*)$, unambiguously increases with the size of the grant. The policy successfully directs the economy's innovative efforts toward more ambitious projects.

4.2.4 Impact on Innovation Output

Quantity: The total expected number of successful innovations (e.g. patents) in a period is the sum of certain incremental successes and probabilistic radical successes:

$$E[Quantity] = (1 - \psi^*) \cdot 1 + \psi^* \cdot E[p_i | p_i \geq p^*] \tag{4.9}$$

The first term represents the certain successes from the $1-\psi^*$ inventors pursuing incremental innovation. The second term represents the expected number of successes from the ψ^* inventors pursuing radical innovation. When policy G increases, ψ^* rises. This reallocates inventors from the first term (where each contributes one certain patent) to the second term (where each contributes $p_i < 1$ expected patents). The inventors who are induced to switch are, by definition, those with lower ability, having success probabilities p_i close to the new, lower cutoff p^*. If the average success probability of these marginal inventors is sufficiently low, the loss of certain incremental patents may not be fully offset by the gain in expected radical patents. This leads to a potential decline in the total expected number of innovations, providing a direct theoretical explanation for our empirical finding of a decrease in patent quantity post-policy.

Average Quality: By lowering the barrier to entry for ambitious projects, the policy attracts a new group of marginal inventors whose projects have a higher probability of failure. This influx of high-risk attempts could mechanically drag down the average innovation quality for the entire economy.

To formalize this, let an incremental patent have a certain, baseline quality of 1. A successful radical patent has a much higher quality of $R > 1$ (one can think of this as some measure of radicalness, like forward citations), while a failed radical project has a quality of 0. The impact on average quality can be understood by examining the choice of the marginal inventor. This inventor essentially gives up an incremental project with a certain quality of 1 for a radical attempt with an expected quality of $p^* \cdot R$. Consequently, the quality-weighted output can decrease if this expected gain from the radical attempt is less than the certain loss from the foregone incremental project, a condition met if $p^*R < 1$.

This results propagate to the economy level. The quality-weighted output of patents in the economy is:

$$E[\textit{Quality-weighted Output}] = (1-\psi^*) \cdot 1 + \psi^* \cdot E[p_i | p_i \geq p^*] \times R \tag{4.10}$$

The government grant G lowers the talent cutoff p^*, causing inventors at this new, lower margin to switch from the incremental to the radical path. When an inventor with talent p^* switches, the economy loses one certain patent of quality 1 and gains a radical attempt that yields an expected quality of $p^* \cdot R$. If this expected gain is less than the certain loss, the switch is quality-reducing in aggregate terms. This gives us that a decline in total quality occurs if $p^* \cdot R < 1$. This means that if the marginal inventor's probability of success (p^*) is low enough that even when multiplied by the high reward of a radical breakthrough (R), it does not equal the quality of a certain incremental success, then their decision to switch lowers the total quality output of the economy. This influx of lower-talent inventors into the radical R&D pool mechanically lowers the average success rate across all radical projects. This "dilution" of the talent pool, with a higher share of projects that ultimately fail and contribute zero quality.

4.2.5 Impact on Expected Productivity and Long-Term Growth

While the policy may lower the average observed quality of innovations by encouraging riskier projects, its core purpose is to increase the total expected productivity growth of the economy. The total expected change in the technology stock, recall from Equation 4.3, is the sum of expected gains from both innovation types:

$$E[A_{t+1}] - A_t = (1 - \phi^*) \cdot \epsilon + \phi^* \cdot E[p_i | p_i \geq p^*] \cdot \gamma$$

Because the productivity gain from a radical innovation (γ) is substantially larger than that from an incremental one (ϵ), the expected gain from a radical attempt ($p_i \cdot \gamma$) can far exceed the certain gain from an incremental project (ϵ), even for projects with a modest probability of success. By inducing more inventors to pursue radical innovation (increasing ϕ^*), the policy shifts the economy's R&D portfolio towards activities with higher expected social returns. This increases the overall expected productivity growth, providing the justification for government implementing innovation policy.

4.3 Model Extensions 1: Knowledge Spillover

A central feature of the economics of innovation is that the production of knowledge is subject to a classic market failure, leading to a socially suboptimal level of investment in research and development (R&D). This failure arises because knowledge constitutes a quasi-public good, generating significant positive externalities, or "knowledge spillovers," that are not appropriated by the innovating agent.

New knowledge created through innovation exhibits two primary characteristics of a public good: First, the use of a piece of knowledge by one agent does not preclude its simultaneous use by another. An engineering principle or a line of code can be utilized by limitless firms without being depleted. It is difficult and costly to prevent third parties from benefiting from the discovery Bloom, Van Reenen, and Williams (2019). Despite legal frameworks like the patent system, knowledge tends to disseminate. The disclosure required by a patent application itself diffuses information, while other avenues like reverse-engineering, employee mobility, and academic publications ensure that the foundational ideas behind an innovation eventually become part of the public knowledge.

The consequence of these characteristics is a divergence between the private returns captured by the inventor and the total social returns generated by an innovation. We can formalize this distinction as follows.

Let V_R represent the private value of a successful radical innovation, as perceived by the inventor. This value is the capitalized stream of monopoly rents, licensing fees, or other private benefits that the inventor can legally secure.

The total social value of the innovation, denoted V_S, includes not only the inventor's private returns but also the monetized value of the knowledge spillovers, $V_{Spillover}$. These spillovers include the consumer surplus from lower prices or higher quality products, the value of follow-on innovations by competing or complementary firms, and the contribution of the new knowledge to the public intellectual commons. Thus, the social value can be expressed as:

$$V_S = V_R + V_{Spillover} \tag{4.11}$$

By definition, for any innovation with positive externalities, $V_{Spillover} > 0$, and therefore $V_S > V_R$. Empirical economic literature consistently finds that the social rate of return to R&D is substantially higher than the private rate of return, confirming the significance of this wedge.

A rational, risk-neutral inventor will only choose to pursue a radical innovation project if the expected private payoff exceeds the opportunity cost. Using the framework from Section 3.4, this private investment criterion for the marginal inventor is given by:

$$p^* \cdot V_R - c \geq w_t - G \tag{4.12}$$

This equation defines the private optimum. The inventor invests only if their expected, appropriable return covers the net costs.

However, a social planner aiming to maximize economic welfare would want the inventors to make the investment decision based on the full social return. The socially optimal investment criterion is:

$$p^* \cdot V_S - c \geq w_t \tag{4.13}$$

Substituting the definition of V_S, we get:

$$p^* \cdot (V_R + V_{Spillover}) - c \geq w_t \tag{4.14}$$

A comparison of the private and social investment criteria gives the source of the market failure. There exists a set of innovation projects for which the private calculus leads to rejection, while the social calculus would demand investment. Specifically, a project will be inefficiently foregone if the following condition holds:

$$p^* \cdot V_R - c < \epsilon \ \text{ and } \ p^* \cdot (V_R + V_{Spillover}) - c \geq w_t \tag{4.15}$$

This inequality describes valuable innovations that are not pursued because the inventor cannot capture a sufficient fraction of the total social value ($V_{Spillover}$) they create. Each inventor, acting rationally based on private incentives, makes a decision that is individually optimal but collectively results in an aggregate level of R&D that is below the social optimum. This provides a clear theoretical justification for public policies, such as the research subsidy G in this model, designed to close the gap between the private and social returns to innovation and encourage investment toward a more efficient level.

In an idealized scenario where the government has perfect information, it can set an optimal grant, G^*, that perfectly aligns private incentives with the social optimum. The goal is to encourage every project that is socially valuable but not privately profitable. This is achieved by setting the grant equal to the expected value of the knowledge spillover for the marginal inventor that society wishes to incentivize. If p^* is the talent level of this marginal inventor, the optimal grant would be:

$$G^* = p^* \cdot (V_S - V_R) = p^* \cdot V_{Spillover} \tag{4.16}$$

Such a grant precisely internalizes the externality for the marginal project, ensuring all socially desirable innovations are undertaken without wasting resources on projects that are already privately viable or those that are not socially valuable at all. The central takeaway is that a positive grant ($G^* > 0$) is required, which justifies government subsidies for innovation. Thus, our framework shows how private markets can systematically underinvest in radical innovation and how targeted subsidies can, in principle, correct this distortion to enhance aggregate welfare, a topic widely discussed in the endogenous growth literature (Romer, 1990; Atkeson and Burstein, 2019; Acemoglu and Restrepo, 2018).

4.4 Model Extension 2: Imperfect Government Screening and the Allocation of Talents

For modern AI policies, governments typically do not fund all applicants; they attempt to screen them. We now explicitly model this process, assuming a sequential framework where inventors apply for a grant, receive a decision, and then make their final project choice.

4.4.1 The Screening Mechanism

We assume the government designates projects as "high-quality" if the inventor's talent is above a threshold, $p_i \geq p_H$. However, its screening is imperfect and subject to two types of errors:

Type I Error (False Negative): A high-quality inventor ($p_i \geq p_H$) is denied a grant. This occurs with probability α.

Type II Error (False Positive): A low-quality inventor ($p_i < p_H$) is mistakenly awarded a grant. This occurs with probability β.

The parameters α and β measure the government's screening fallibility; lower values indicate greater accuracy.

4.4.2 Equilibrium Composition of the Radical Innovation Pool

The final pool of inventors who pursue radical innovation is composed of distinct groups based on their talent and grant status. The decision to proceed depends on two talent cutoffs: one for inventors who receive the grant (p_G^*) and a higher one for those who do not (p_N^*). (1) Cutoff with Grant (p_G^*) and (2) Cutoff with No Grant (p_N^*).

The total fraction of inventors pursuing radical innovation, $\psi_{Radical}$, is the sum of three groups: (1) High-Quality, Funded Inventors (S_1: High-quality inventors ($p_i \geq p_H$) who receive the grant (with probability $1-\alpha$) and choose the radical path. Assuming $p_H > p_G^*$, all will proceed. The size of this group is:

$$S_1 = (1-\alpha)\int_{p_H}^{\overline{p}} f(p)dp \tag{4.17}$$

(2) High-Quality, Unfunded Inventors (S_2) High-quality inventors ($p_i \geq p_H$) who are denied the grant (with probability α) but whose talent is so high ($p_i \geq p_N^*$) that they proceed anyway. The size of this group is:

$$S_2 = \alpha \int_{p_N^*}^{\overline{p}} f(p)dp \tag{4.18}$$

(3) Low-Quality, Funded Inventors (S_3): Low-quality inventors ($p_i < p_H$) who mistakenly receive the grant (with probability β) and whose talent is sufficient to proceed with the subsidy ($p_i \geq p_G^*$). The size of this group is:

$$S_3 = \beta \int_{p_G^*}^{p_H} f(p)dp \tag{4.19}$$

The total pool is $\psi_{Radical} = S_1 + S_2 + S_3$.

An increase in α (more false negatives) shrinks the pool of funded high-quality inventors (S_1). This leads to a net loss of high-quality radical projects, as talented inventors in the critical range ($p_H \leq p_i < p_N^*$) who are denied funding will rationally choose the incremental path, leading to *misallocation of talent*.

An increase in β (more false positives) expands the pool of funded low-quality inventors (S_3). This injects projects with a low probability of success into the R&D portfolio, leading to *misallocation of financial resources*.

Notice both larger α and β would thus could contribute to decline in average patent quality from inventors, as it increase the share of low-quality inventors relative to high-quality inventors.

5 Mechanisms Tests

5.1 Potential Channel 1: Supporting Radical Innovation Attempts

As established in our conceptual framework, the observed decline in inventor productivity may be attributed to national policies that lower the costs for scientists to pursue radical, high-impact but risky innovations. To investigate this channel, we first characterize these radical attempts. We argue that a necessary precursor to radical impact is novelty, defined as a significant departure from established knowledge. This concept is not only consistent with the literature (e.g. Kelly, Papanikolaou, Seru, and Taddy, 2021) but is also a fundamental tenet of innovation process.[18] It is important to emphasize that ex-ante novelty does not guarantee ex-post impact. Many novel projects are "trial-and-error" efforts that ultimately fail, and this high probability of failure is the defining characteristic of radical innovation's risk profile.

We first test this hypothesis by examining backward citations, a conventional proxy for an invention's reliance on prior art. The evidence suggests a reduced reliance on prior knowledge among treated inventors following the introduction of national AI policies. Specifically, the Poisson estimates in column 1 of Table 4 suggest a reduction of 10.6% in the probability that an inventor's AI patent filing will have a backward citation in the same year for treated inventors after policy implementation.

While intuitive, citation counts can be a coarse measure of knowledge dependence. To analyze this mechanism more rigorously, we employ a text-based approach that leverages recent advances in natural language processing (Kelly et al., 2021). The idea is that novel

[18] See, e.g., the novelty requirement under 35 U.S.C. § 102, as explained by the USPTO: https://www.uspto.gov/sites/default/files/about/offices/ous/Cooper_Union_20130610.pdf

inventions should be textually distinct from prior art. Our method uses semantic similarity, which captures conceptual closeness beyond simple word or phrase matching, to identify these connections. To construct our novelty metric, we first count the number of prior-year patents with a very high semantic similarity score (cosine similarity > 0.9) to the focal patent. We refer to this measure as the count of competing AI patents. We then take the negative of this count, such that a higher value for our metric indicates greater novelty.

Using this measure, we again find a shift among treated inventors toward filing patents that are more distinct from prior art. In Table 4, we regress count of competing AI patents on indicator of AI policy implementation (we flip the sign of the coefficients to gauge the effect of novelty). The estimates in Table 4 (Panel A, Column 3) indicate an 18.8% reduction (increase) in the number of competing patents (innovation novelty) from previous year for treated inventors following policy implementation. As a robustness check, we construct an alternative measure of competing patent counts by including similar patents from all prior years as well as the current year, as shown in Table B16. This broader window helps account for the possibility that inventors may draw on older or contemporaneous knowledge, not just the most recent year's inventions. This provides evidence for an increase in the novelty of AI patents. It is important to reiterate, however, that we define novelty as distinctiveness from existing inventions, not as a direct measure of ex-post technological impact.

5.1.1 Patent Novelty: High Risk, High Reward?

Are these novel inventions indeed riskier, while also possessing the potential for greater impact? To answer this, we investigate the characteristics of novel inventions, focusing on our text-based measure of novelty as it is methodologically more advanced.

We conduct a patent-level analysis to assess the risk–return trade-off associated with novel inventions. To do so, we define three key metrics. We capture the risk of failure by identifying "low-impact" patents as those receiving zero forward citations. We proxy for development cost using the patent's development duration — measured as the number of days between the filing date of the focal AI patent and the most recent filing date among its backward citations.[19] On the return side, we measure patent quality using citation-weighted counts. Over our 2006–2019 sample period, 54% of AI patents are classified as low-impact, with an average development time of 971 days.

Panel A of Table 5 Columns 1 and 3 show that more novel patents are significantly more likely to be low-impact (i.e., failures) and are associated with longer development times,

[19]The intuition for this proxy is analogous to the academic writing process, where most references are gathered during the initial stages of development. Consequently, the time elapsed between the dates of those references and the final submission date serves as a good proxy for how long the project was actively developed before reaching its final, ready state.

respectively. However, as shown in Column 2, these same novel patents are also more likely to be of higher quality, as measured by their citation-weighted counts. Together, these findings provide strong evidence of the risk–return trade-off inherent in the pursuit of novel innovation.

While novel patents have a greater risk of failure and require more time to develop, they are also more likely to be high-quality. To further examine this trade-off, we examine whether novel patents are more likely to become radical (i.e., extremely high-impact) innovations. We construct a series of binary indicators for high-impact patents based on increasingly stringent citation-weighted count thresholds (e.g., top 10%, top 1%, top 0.5%, and top 0.05%). Panel B of Table 5 presents regression results using these indicators as dependent variables. Across all thresholds, we find a consistent and statistically significant positive relationship between novelty and the likelihood of a patent being classified as high-impact. Specifically, more novel patents are 0.277 percentage points more likely to rank among the top 10% of high-impact AI patents and 0.006 percentage points more likely to fall within the top 0.05%.

To complement this regression-based evidence, we further investigate the risk profile of novel patents by analyzing the variance in forward citations across different levels of novelty. As shown in Table 6, Panel A groups patents by the number of competing filings in the prior year and reports the variance in forward citations within each group. We find that the most novel patents (Group 1, with zero competing patents) exhibit the highest variance in forward citations (2232.45), while less novel patents display considerably lower variance. We find that the most novel patents (Group 1, with zero competing filings) exhibit the highest variance in forward citations (2232.45), while less novel patents display substantially lower variance. For example, the variance for Group 1 is approximately seven times greater than that for Group 5 (which includes patents with more than three competing prior-year filings). This suggests that more novel patents are associated with greater uncertainty in their future impact, highlighting the inherent risk of pursuing novel innovation. Panel B confirms these differences statistically through F-tests, all of which indicate that Group 1 has significantly higher variance than other groups, which is statistically significant at 1% level.

Taking together, these findings suggest that while novel AI patents are more likely to fail — often become low-impact patent without forward citations — and require longer development periods, they are also disproportionately more likely to become radical innovations. This finding supports the view that government support can incentivize inventors to pursue novel and exploratory projects with high risk of failure, some of which yield high-impact outcomes. Accordingly, such policies may be effective in fostering breakthrough innovations, despite being accompanied by greater uncertainty.

5.1.2 High-Impact Inventors, Novelty, and Access to Government Support

The theoretical framework in Section 4.2 predicts that inventors respond to policy incentives favoring radical over incremental innovation. In particular, high-impact inventors face a trade-off between the certainty of incremental projects and the higher risk associated with radical innovation. Targeted government support, such as grants, can reduce this risk and thereby encourage these inventors to pursue more ambitious, novel projects. This theoretical prediction motivates key empirical questions: Do these policies incentivize high-impact inventors to pursue novel projects aimed at achieving radical breakthroughs? And are these inventors more likely to receive government funding that mitigates the risks involved in such endeavors?

To explore these questions, we begin by examining whether the observed effects are primarily driven by high-impact inventors. We classify inventors as high-impact if they rank annually in the top 10% of AI patents by citation-weighted counts over the preceding five years within their respective countries. Consistent with the theoretical model, our results show that high-impact inventors play a primary role in driving this shift, particularly by reducing their reliance on existing knowledge and engaging in more novel innovation.

Specifically, the Poisson estimates in Columns 1 and 2 of Table 7 show that, relative to other inventors, high-impact inventors have significantly more backward citations and higher average backward citations, reflecting stronger engagement with existing technological knowledge. However, the negative and statistically significant interaction terms between AI policy and high-impact inventor status indicate that, following the implementation of national AI policies, these inventors reduce their backward citation by 18.8% and average backward citations by 8.0%. The evidence align closely with the model's prediction, suggesting a reduced reliance on prior knowledge among treated high-impact inventors following the introduction of national AI policies.

To examine this shift more rigorously, we leverage our text-based novelty metric to assess whether high-impact inventors are more likely to undertake novel, high-risk projects. The Poisson estimate in Column 3 of Table 7 shows a 15% reduction in the likelihood that a high-impact inventor's AI patent has a textually similar filing from the prior year after policy implementation. This provides further evidence that national AI-related policies particularly encourage high-impact inventors to pursue more novel and potentially radical innovations. These empirical finding align with the model's prediction that government support encourages high-impact inventors to pursue novel innovation.

Furthermore, our baseline results show a decline in both the quantity and quality of innovation output following the implementation of national AI policies. To assess whether this decline is also driven by high-impact inventors, we examine the results in Table 8, which show that high-impact inventors also play a primary role in driving the observed reduction in AI innovation output.

Although our findings indicate that the observed effects are primarily driven by high-impact inventors, a key question remains: Is this because these inventors are more likely to receive government support, thereby lowering the effective risk threshold and encouraging them to pursue more radical innovation? To explore this possibility, we examine whether high-impact inventors are more likely to receive government funding for their innovation efforts. We utilize the USPTO government interest dataset from PatentView to identify AI-related patents that acknowledge government support and link them to the corresponding inventors. Our analysis reveals that high-impact inventors—those ranking in the top 10% and top 5% of citation-weighted patent counts—are significantly more likely to receive government funding. However, this relationship is not statistically significant for inventors in the top 1% category.

Specifically, Columns 1 and 2 of Table B15 indicate that being a top 10% inventor is associated with a 0.7 to 0.8 percentage point increase in the likelihood of receiving government funding, with both estimates statistically significant at the 1% level. Since the first U.S. national AI-related policy—the National Robotics Initiative—was introduced in 2011, we include an interaction term with a post-2011 indicator to assess potential changes over time. The interaction term is not statistically significant, suggesting that the relationship between high-impact status and government funding remained stable before and after 2011.

We also examine alternative thresholds for defining high-impact inventors. Columns 3 and 4 focus on inventors in the top 5% of citation-weighted patent counts. The coefficients are positive and statistically significant, and the interaction term with the post-2011 period is also significant at the 10% level. This suggests that the likelihood of receiving government support for top 5% inventors increased following the introduction of AI-related national policies.

In contrast, Columns 5 and 6 show no statistically significant relationship between being a top 1% inventor and receiving government funding, either before or after 2011. This may reflect greater heterogeneity or selection dynamics at the very top of the distribution, where the connection between citation-based impact and funding allocation may be weaker or more idiosyncratic.

In summary, these results provide empirical support for the model's predictions. High-impact inventors play a primary role in pursuing riskier and more exploratory projects—those that rely less on established knowledge and exhibit greater novelty or "venturing" behavior. Moreover, high-impact inventors, particularly those in the top 10% and top 5%, are significantly more likely to receive government funding to support their innovation efforts.

5.2 Potential Channel 2: Government Failure to Target Radical Innovation

In the conceptual framework in section 4.4, we predict that the inability of government to perfectly screen for promising R&D project would decrease patent output and quaility as this can induce low-quality inventors into the radical innovation pool Type II errors (funding low-potential projects) and leave high-quality projects unfunded Type I errors.

As we do not grant-level data for other countries, we focus on US, where we know whether a patent has related government grant recorded in the patent documents (government interests) and therefore can evaluate whether subsidized patent is of high quality and shed some lights on government screening ability. We first investigate whether government support is effectively target toward high-impact and novel innovation in Table 9. Panel A shows that, on average, government-funded AI patents are significantly more likely to rank within the top 10% and top 5% of forward citations and are less likely to be classified as low-impact (i.e., zero citations). They also exhibit greater novelty, suggesting that government funding tends to support more original research. However, these patents are no more likely to be in the top 1% of most-cited patents, indicating a limitation in identifying truly breakthrough innovations ex ante—consistent with the model's prediction of screening constraints.

Panel B introduces an interaction term with a post-2011 indicator to capture the period following the introduction of the U.S. National Robotics Initiative (NRI)—the first major national AI policy. The results reveal a significant decline in targeting precision after 2011: government-funded patents become less likely to fall into the top 10% and 5% of most-cited patents, and they receive fewer forward citations overall. Notably, the association between government support and novelty remains stable, suggesting continued support for original ideas but diminished effectiveness in targeting high-impact outcomes.

In conclusion, these findings support the theoretical prediction that imperfect screening can hinder the optimal allocation of public R&D funding. Initially, government support aligned reasonably well with the objective of fostering radical innovation by favoring patents with higher novelty and citation impact. However, its overall effectiveness declined following NRI in 2011. This decline indicates a potential misalignment in public R&D funding mechanisms: although novel or exploratory patents continue to receive support, such funding is no longer consistently associated with high-impact outcomes. These results point to a breakdown in the government's ability to reliably identify and support radical AI innovations over time.

However, we restrain ourself to draw conclusion on whether this dynamics can explain our baseline results, that cross countries, on average, experience decline in output following initial AI fostering policies.

5.3 Potential Channel 3: Fewer Disclosures of AI Patent Filings

Our first proposed channel suggests that the decline in inventors' productivity is likely driven by government support that incentivizes scientists to pursue more novel and exploratory inventions. However, as shown in Table B8, the baseline effect begins at t=0, which could reflect reduced disclosure of AI patent filings rather than the pursuit of novel inventions, as inventors may not have had sufficient time to develop entirely new innovations capable of producing such an immediate effect following policy implementation. Nonetheless, given that approximately 30% of AI patents are developed and filed within one year, the immediate effect at t=0 could still be consistent with this channel if national AI policies spur novel innovations that were already near completion, thereby resulting in rapid post-policy filings.

Although our first channel remains plausible, reduced disclosure of AI patent filings could also be an alternative explanation for our baseline results. The introduction of national AI policies could incentivize firms and governments to delay or withhold certain patent disclosures, particularly for sensitive technologies. Such strategic withholding would reduce the number of patents recorded in public databases without necessarily diminishing the actual level of innovation activity.

To examine the disclosure mechanism, we analyze AI–military-related patents to assess whether the introduction of national AI policies is associated with fewer publicly disclosed AI patent filings. Military-oriented AI innovations may not be fully disclosed due to national security considerations, potentially reducing the observable number of AI–military patents and lowering measured patent quality.[20] Our empirical analysis does not support this concern. Using inventor-level data, we identify AI–military-related patents based on CPC codes related to weapons, defense systems, and military applications.[21] As reported in

[20] See, e.g., Summary of the 2018 Department of Defense Artificial Intelligence Strategy: https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF

[21] AI–military-related patents are identified using CPC codes covering: F41–F42 (Weapons & Ammunition), including firearms (F41A–C, F41F–G, F41H–J), ammunition, missiles, fuzes, and blasting technologies (F42B–D); B (Military Transport & Vehicles), including warships and submarines (B63G), military aircraft (B64C–D), and specialized military vehicles (B60P); G (Sensors & Control), including radar, navigation, target tracking (G01S), vehicle control (G05D), data processing (G06F), alarm systems (G08B), and simulators (G09B); and H (Military Communication), including antennas (H01Q), secure communication, electronic warfare, cybersecurity, and military networks (H04B, H04L, H04W).

Table 10, we find no significant change in either the number or quality of AI–military-related patent filings following the introduction of national AI policies.

6 Additional Findings

6.1 The Pursuit of Novelty via Domain Specialization

Since our results suggest that the observed effects are driven by a significant shift among treated inventors toward filing novel patents that are more distinct from prior art, we explore whether achieving such novelty involves a more focused and in-depth investigation within specific innovation areas. As supporting evidence for this channel, we examine whether inventors concentrate on a narrower range of technological domains following policy implementation.

To test this, we construct two measures of domain specialization: (1) the number of AI patent domains per inventor and (2) the Herfindahl-Hirschman Index (HHI) based on AI classifications from the AI Patent Database (AIPD). As reported in Table B17, column 1 shows a statistically significant decline in the number of AI patent domains per inventor post-policy, with a coefficient of -0.058. Column 2 shows a corresponding and significant increase in the HHI, indicating greater concentration.

In conclusion, these results suggest that treated inventors are more likely to narrow their technological focus, consistent with the pursuit of novel and exploratory innovation within specialized domains.

6.2 Knowledge Spillover and International Collaboration

Standard economic theory posits that, in the absence of market failures, investment decisions are best left to private firms. The primary justification for government intervention lies in the presence of knowledge spillovers. As Bloom, Van Reenen, and Williams (2019) notes, "ideas are promiscuous." Even under a well-designed intellectual property regime, innovators often fail to capture the full benefits of their inventions because spillovers enable others to benefit without bearing the full R&D costs. Therefore, it is important to examine how these policies influence knowledge diffusion.

Our baseline estimates suggest that knowledge spillovers—proxied by citations—have declined. However, spillovers may also occur through international collaboration, such as cross-country citations. To investigate this possibility, we examine changes in the share of cross-country backward and forward citations following the implementation of national

AI-related policies, as shown in Table 11. The OLS estimates in columns 1 and 2 indicate increases of 4.8% and 5.3% in the share of cross-country backward and forward citations, respectively. These results suggest a greater reliance on international knowledge sources and increased global reach of domestic innovation.

7 Conclusion

Using two novel datasets, we describe the global landscape of national promotion AI policies and AI innovation. Employing a stacked Difference-in-Differences framework, we observe that, at the inventor level, AI scientists produce fewer patents of lower average quality following the introduction of AI-support policies, despite an increase in the novelty of their work. This pattern aligns with the high-risk, high-reward dynamics of AI, where policies incentivize exploratory efforts with uncertain outcomes. Similar patterns emerge at the country level, indicating a consistent decline in both the quantity and quality of AI patents after policy implementation.

These findings highlight a important trade-off in the design of mission-oriented innovation policies. While such policies can address underinvestment in high-risk and high-reward projects , they may also lead to transitional inefficiencies as innovative activity shifts from incremental to exploratory projects. The short-term decline in inventor productivity should therefore not be interpreted as policy failure but as a natural by-product of structural experimentation. Over time, the reallocation of inventive effort toward more novel projects can produce a wider variance of outcomes—many unsuccessful attempts accompanied by a few transformative breakthroughs—that expand long-run technological potential.

Our results also suggest the limits of government screening in identifying high-impact projects exante. Although policymakers tend to select inventors and proposals that appear promising ex ante, the realized outcomes exhibit substantial heterogeneity. This pattern is consistent with imperfect information and selection frictions, suggesting that even well-intentioned targeting cannot fully eliminate uncertainty in radical innovation. Moreover, the evidence that AI policies enhance international knowledge exchange indicates that these programs may facilitate the global diffusion of ideas, consistent with the stated goals of many national AI strategies. Overall, Our findings inform policies and highlight the importance of balancing short-term productivity trade-offs with potential long-term gains when designing AI policies. Policymakers should recognize that nurturing innovation in rapidly evolving areas like AI may require tolerating initial inefficiencies and elevated failure rates.

Future research could extend our study in several directions. First, examining the long-run dynamics of inventor productivity and patent impact would clarify whether short-term declines give way to sustained innovation gains. Second, investigating heterogeneity across policy instruments—such as grants, research programs, or governance initiatives—could shed light on which forms of intervention most effectively balance exploration with productivity. Third, linking inventor-level changes to firm-level outcomes and technology diffusion could deepen our understanding of how mission-oriented policies reshape the broader innovation ecosystem. As countries increasingly compete for leadership in AI, understanding these policy trade-offs is essential for fostering innovation that is not only faster, but also more transformative and socially valuable.

REFERENCES

Acemoglu, D., U. Akcigit, H. Alp, N. Bloom, and W. Kerr. 2018. Innovation, reallocation, and growth. *American Economic Review* 108:3450–91–. doi:10.1257/aer.20130470.

Acemoglu, D., and P. Restrepo. 2018. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review* 108:1488–1542–.

Aghion, P., A. Bergeaud, and J. Van Reenen. 2023. The impact of regulation on innovation. *American Economic Review* 113:2894–2936–.

Aghion, P., and P. Howitt. 1992. A model of growth through creative destruction. *Econometrica* 60:323–51.

Akcigit, U., and W. R. Kerr. 2018. Growth through heterogeneous innovations. *Journal of Political Economy* 126:1374–443.

Arora, A., S. Belenzon, A. Patacconi, and J. Suh. 2020. The changing structure of american innovation: Some cautionary remarks for economic growth. *Innovation Policy and the Economy* 20:39–93.

Atkeson, A., and A. Burstein. 2019. Aggregate implications of innovation policy. *Journal of Political Economy* 127:2625–83. doi:10.1086/701414.

Babina, T., A. Fedyk, A. He, and J. Hodson. 2024. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151:103745–.

Baker, A. C., D. F. Larcker, and C. C. Wang. 2022. How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics* 144:370–95. ISSN 0304-405X.

Bertrand, M., E. Duflo, and S. Mullainathan. 2004. How much should we trust differences-in-differences estimates? ∗. *The Quarterly Journal of Economics* 119:249–75. ISSN 0033-5533. doi: 10.1162/003355304772839588.

Bloom, N., J. Van Reenen, and H. Williams. 2019. A toolkit of policies to promote innovation. *Journal of Economic Perspectives* 33:163–84–. doi:10.1257/jep.33.3.163.

Brynjolfsson, E., D. Rock, and C. Syverson. 2017. Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics. NBER Working Paper 24001, National Bureau of Economic Research. doi:10.3386/w24001.

Callaway, B., and P. H. Sant'Anna. 2021. Difference-in-differences with multiple time periods. *Journal of Econometrics* 225:200–30.

Callaway, B., and P. H. Sant'Anna. 2021. Difference-in-differences with multiple time periods. *Journal of econometrics* 225:200–30.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer. 2019. The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics* 134:1405–54.

Cohn, J. B., Z. Liu, and M. I. Wardlaw. 2022. Count (and count-like) data in finance. *Journal of Financial Economics* 146:529–51. ISSN 0304-405X.

Cong, L. W., Y. Lu, H. Shi, and W. Zhu. 2024. Automation-induced innovation shift. *Available at SSRN 5049949* .

Deshpande, M., and Y. Li. 2019. Who is screened out? application costs and the targeting of disability programs. *American Economic Journal: Economic Policy* 11:213–48.

Druce, J., J. Niehaus, V. Moody, D. Jensen, and M. L. Littman. 2021. Brittle ai, causal confusion, and bad mental models: challenges and successes in the xai program. *arXiv preprint arXiv:2106.05506* .

Giczy, A. V., N. A. Pairolero, and A. A. Toole. 2022. Identifying artificial intelligence (AI) invention: a novel AI patent dataset. *Journal of Technology Transfer* 47:476–505.

Gormley, T. A., and D. A. Matsa. 2011. Growing out of trouble? corporate responses to liability risk. *The Review of Financial Studies* 24:2781–821.

Gunning, D., E. Vorm, Y. Wang, and M. Turek. 2021. Darpa's explainable ai (xai) program: A retrospective. *Authorea Preprints* .

Hall, B. H., A. B. Jaffe, and M. Trajtenberg. 2005. Market value and patent citations. *The RAND Journal of Economics* 36:16–38.

Howell, S. T. 2024. Government intervention in innovation. *Annual Review of Financial Economics* 16.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy. 2021. Measuring technological innovation over the long run. *American Economic Review: Insights* 3:303–20.

Lerner, J. 2009. *Boulevard of broken dreams: why public efforts to boost entrepreneurship and venture capital have failed–and what to do about it.* Princeton University Press.

Lucas, R. E. 1988. On the mechanics of economic development. *Journal of Monetary Economics* 22:3–42. ISSN 0304-3932. doi:10.1016/0304-3932(88)90168-7.

Miller, T. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence* 267:1–38.

Moretti, E. 2021. The effect of high-tech clusters on the productivity of top inventors. *American Economic Review* 111:3328–75–.

Romer, P. M. 1990. Endogenous technological change. *Journal of political Economy* 98:S71–S102.

Sant'Anna, P. H., and J. Zhao. 2020. Doubly robust difference-in-differences estimators. *Journal of Econometrics* 219:101–22.

8 Figures and Tables

Figure 1
Word Cloud for National AI Strategies and Policies

This figure shows the word cloud of descriptions for National AI-related Strategies and Policies from OECD AI Policy Database.

Figure 2
Time Variation of Percentage of AI Patent Filings

This figure shows the time variation of percentage of AI patent filings across 42 countries from 1990 to 2022.

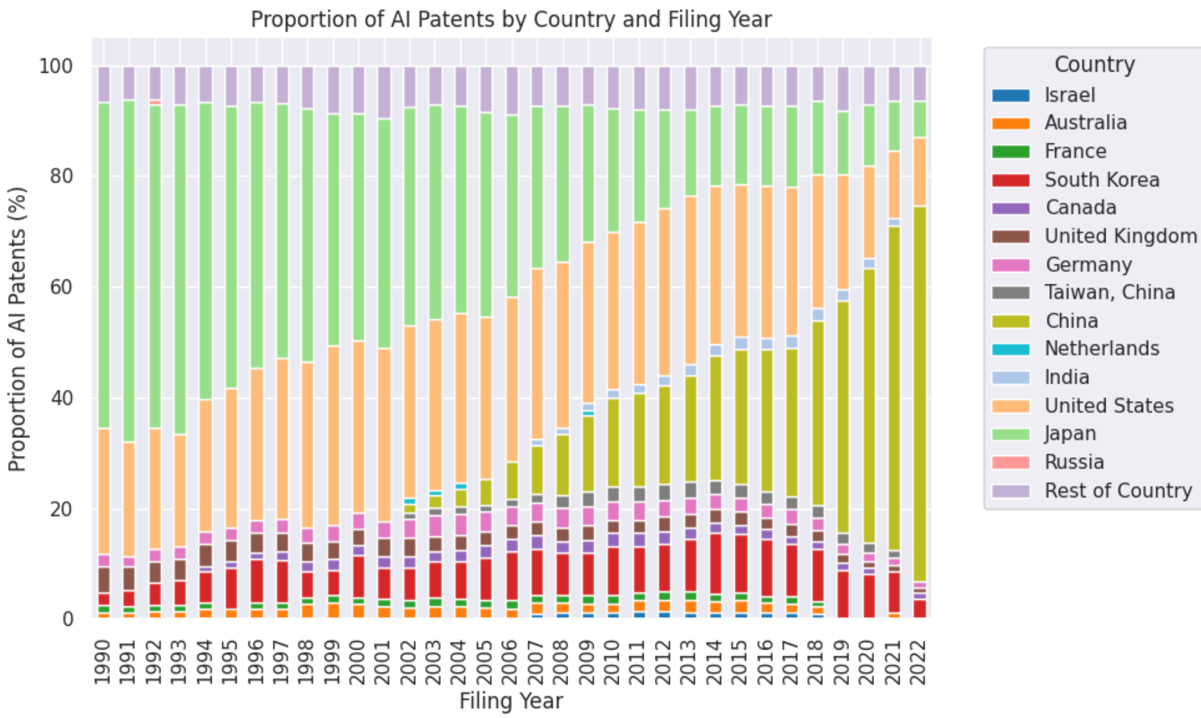

Figure 3
Timeline of USA's Initiatives and Policy Events Related to AI

This figure shows the timeline of the U.S. initiatives and policy events related to AI from 2011 to 2019.

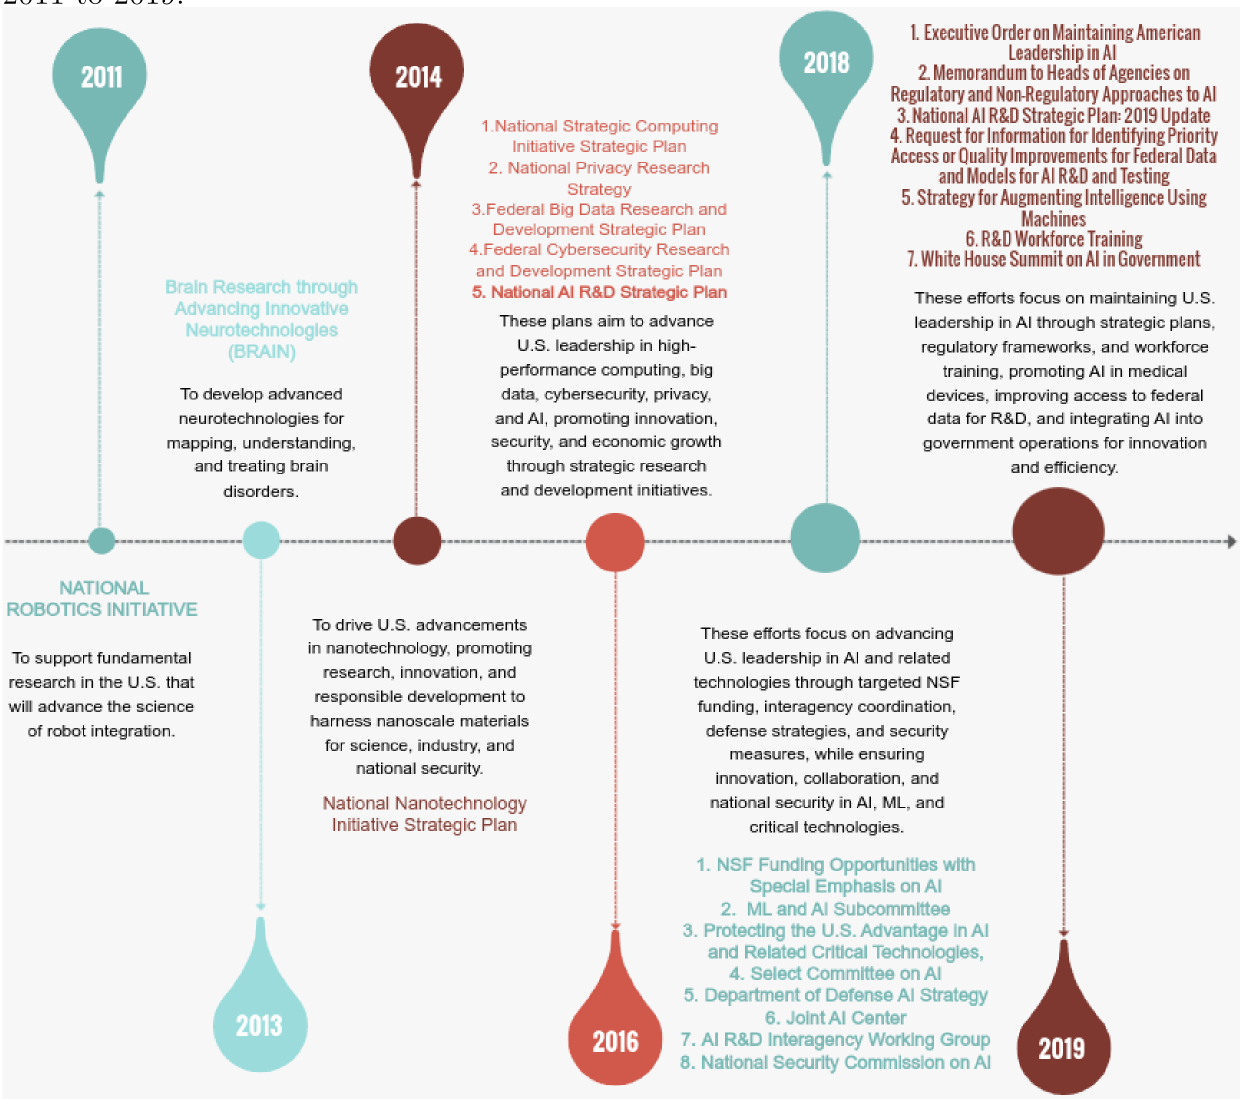

Figure 4
Time Variation of AI Patent Quality & Productivity

This figure reports two subfigures. The figure 4a presents the pretrend analysis for our tests. The specification follow the methodology from Callaway and Sant'Anna (2021) and the treatment is separated by the timing of treatment from $\leq t-8$ to $\geq t+8$. The figure 4b shows the trend in the average number of AI patent annual filings per inventor between 2006 and 2019. The red line indicates the average for inventors in countries that had implemented national AI policies by year t. The blue line represents the control countries that had not.

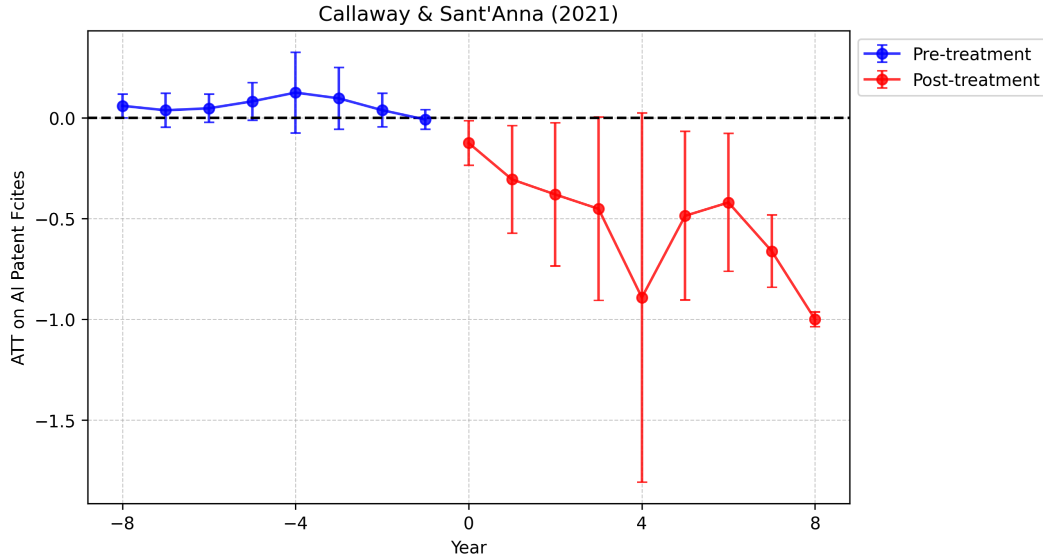

(a) Pre-trend analysis of AI patent forward citations by Methodology from Callaway and Sant'Anna (2021)

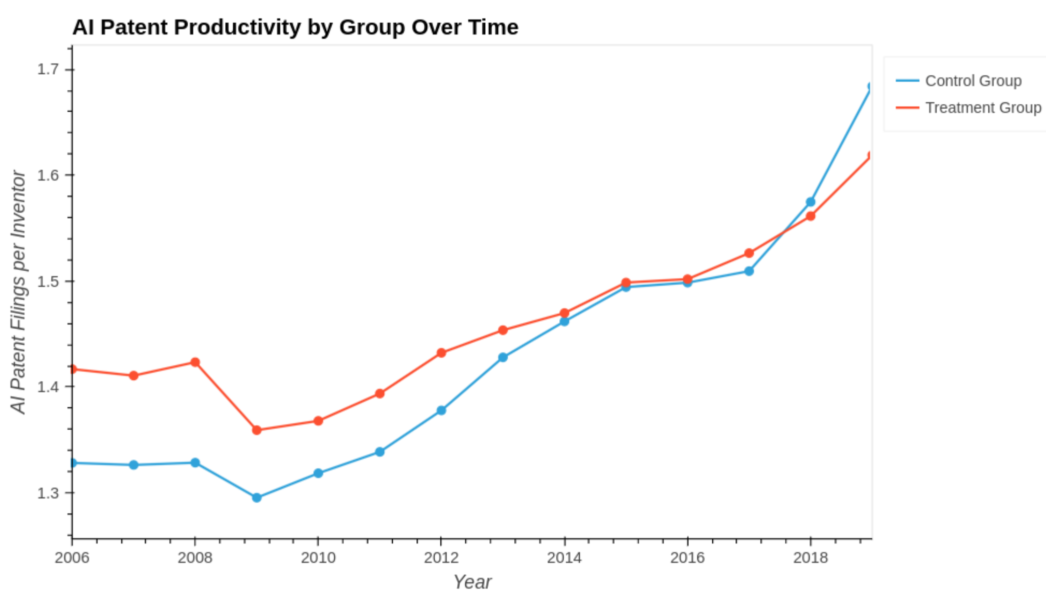

(b) Time Variation of AI Patent Productivity by Group

Figure 5
Time Series Variation of Percentage of Award Amount in AI from NSF

This figure plots the time-series variation of percentage of award amount in AI from National Science Foundation (NSF) from 2000 to 2023.

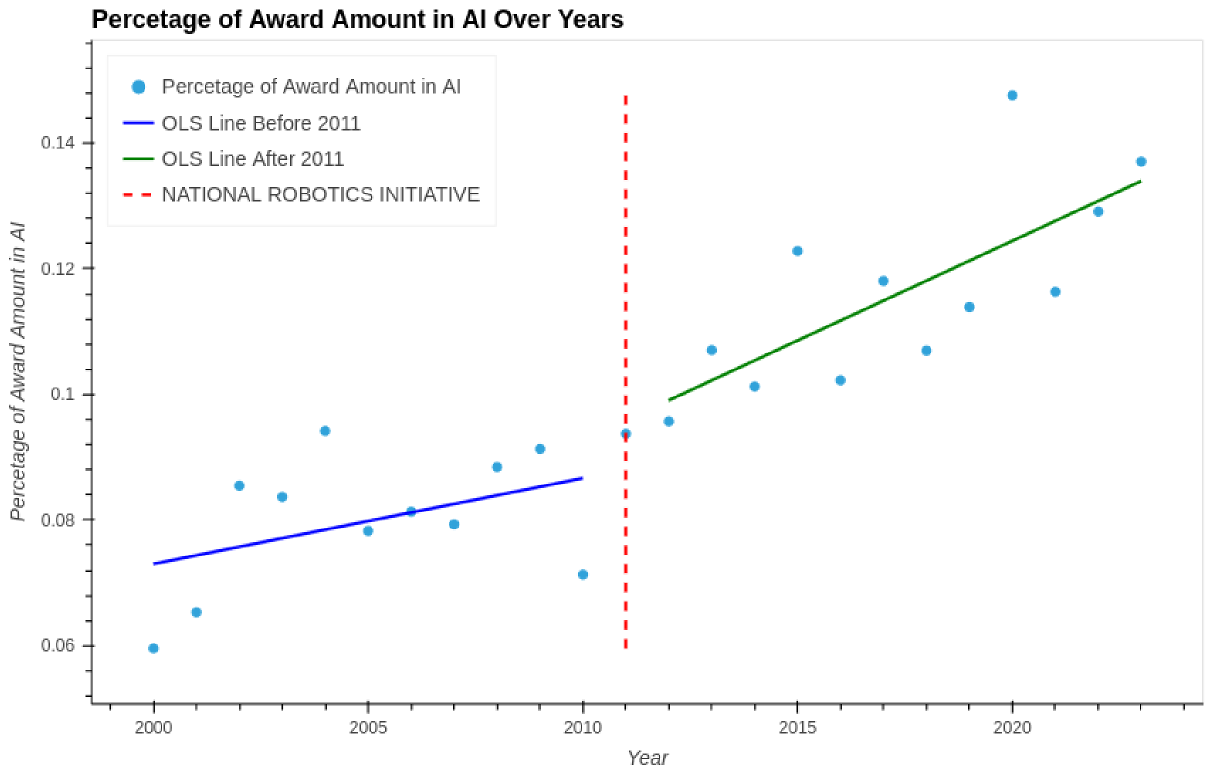

Table 1
Inventor-Level Panel Data Descriptive Statistics

This table reports the overview statistics from 2006 to 2019. The definitions of all variables are reported in the Appendix.

	N	Mean	Std	25%	Median	75%
AI Patent Counts	2,960,410	1.47	1.91	1.00	1.00	1.00
AI Patent Forward Cites	2,960,410	7.24	28.66	0.00	2.00	6.00
Avg AI Forward Cites	2,960,410	4.68	16.53	0.00	1.00	4.50
AI Patent Backward Cites	2,960,410	20.84	168.24	0.00	5.00	12.00
Avg AI Backward Cites	2,960,410	12.83	65.12	0.00	4.00	9.00
Count of Competing AI Patents (Negative Novelty)	2,960,410	3.69	10.01	1.00	2.00	4.00
Inentor AI Domains	2,960,410	1.19	0.57	1.00	1.00	1.00
Inventor AI Domains HHI	2,960,410	0.93	0.18	1.00	1.00	1.00

Table 2
National AI-Related Policy and AI Innovation

This table reports the coefficients from Poisson and OLS regressions of the AI innovation inventor-level panel on an indicator for national AI-related policies. Panel A presents stacked Difference-in-Differences (DID) estimates, controlling for inventor-by-cohort and year-by-cohort fixed effects, following Gormley and Matsa (2011); Baker, Larcker, and Wang (2022). Panel B reports estimates from the structural DID approach proposed by Callaway and Sant'Anna (2021). The AI policy indicator equals one if the country has implemented national AI-related policies. The dependent variables include AI patent filing counts, forward citations and average forward citations. The sample consists of inventor-year observations from five years prior to, and up to ten years following, the implementation of national AI-related policies across countries between 2006 and 2019. For all regressions, robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

Panel A: Stacked DID (Poisson)

	AI Patent Counts	AI Patent Forward Cites	Avg AI Forward Cites
	(1)	(2)	(3)
DID(AI Policy)	-0.095***	-0.471*	-0.355*
	(-2.58)	(-1.93)	(-1.73)
Year × Cohort FE	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes
Obs.	3,900,297	3,515,936	3,515,936

Panel B: Methodology of Callaway and Sant'Anna (2021) (OLS)

	ln(1+AI Patent Forward Cites)		
	ATT on Treated	ATT by Calendar Period	ATT by Group
	(1)	(2)	(3)
DID(AI Policy)	-0.282**	-0.388**	-0.182***
	(-2.57)	(-2.57)	(-2.85)
Obs.	695,738	695,738	695,738

Table 3
National AI-Related Policy and Aggregate AI Innovation (Country Level)

This table reports the impact of AI policy on AI innovation at country level and pretrend analysis. The Panel A reports coefficients from country-panel Poisson regressions of countries' AI innovation on an indicator for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. The dependent variables are AI patent filings counts, forward citations, Inventor Counts. Inventors are classified annually as high-impact if their AI patents rank in the top 10% by forward citations over the previous five years. The high-impact inventor share is calculated as the number of high-impact inventors divided by the total number of unique inventors filing AI patents in that year; the remainder constitutes the low-impact inventor share. The data include country-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. The Panel B presents the pretrend analysis for our tests. All specifications follow the main table but the treatment is separated by the timing of treatment from $\leq t-3$ to $\geq t+3$. The robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

Panel A — Impact of AI Policy on AI Innovation and Inventor

	Poisson	Poisson	Poisson	OLS
	AI Patent Counts	AI Forward Cites	Inventor Counts	Low-Impact Inventor Share
	(1)	(2)	(3)	(4)
DID(AI Policy)	-0.267**	-0.434*	-0.292***	0.090
	(-2.29)	(-1.67)	(-4.41)	(0.13)
Year × Cohort FE	Yes	Yes	Yes	Yes
Country × Cohort FE	Yes	Yes	Yes	Yes
Obs.	2,712	2,712	2,712	2,712
R^2				0.86

Panel B — Pretrend Analysis

	Poisson	Poisson	Poisson	OLS
	AI Patent Counts	AI Forward Cites	Inventor Counts	Low-Impact Inventor Share
	(1)	(2)	(3)	(4)
t ≤ -3	0.058	-0.005	0.091	0.324
	(0.95)	(-0.04)	(1.45)	(0.64)
t = -2	0.052	0.028	0.060**	0.147
	(1.60)	(0.44)	(2.24)	(0.43)
t = 0	-0.077	-0.134	-0.092***	-0.028
	(-1.52)	(-1.16)	(-2.87)	(-0.11)
t = 1	-0.206*	-0.404**	-0.193***	-0.438
	(-1.93)	(-2.10)	(-3.61)	(-0.95)
t = 2	-0.322**	-0.525*	-0.282***	-0.283
	(-2.06)	(-1.82)	(-3.36)	(-0.40)
t ≥ -3	-0.360**	-0.628*	-0.434***	1.454
	(-2.27)	(-1.87)	(-4.42)	(0.87)
Year × Cohort FE	Yes	Yes	Yes	Yes
Country × Cohort FE	Yes	Yes	Yes	Yes
Obs.	2,712	2,712	2,712	2,712
R^2				0.86

Table 4
National AI-Related Policy and AI Innovation Novelty

This table reports coefficients of Poisson regressions of the AI innovation inventor-panel on an indicator for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. The AI policy indicator equals one if the country has implemented national AI-related policies. The dependent variables are backward citations, average backward citations and the number of similar AI patent filings in the year preceding the current observation year. The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. For all regressions, the robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively. The sample period is from 2006 to 2019.

	Poisson AI Patent Backward Cites	Poisson Avg AI Backward Cites	Poisson AI Novelty
	(1)	(2)	(3)
DID(AI Policy)	-0.112*	-0.041*	0.209***
	(-1.78)	(-1.87)	(2.61)
Year × Cohort FE	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes
Obs.	3,675,748	3,675,748	3,743,038

Table 5
Risk and Return Associated with Novelty

This table presents a patent-level analysis of AI novelty. Panel A reports the relationship between AI novelty—measured within one year prior to the focal patent's filing year—and two outcomes: low-impact AI patents (defined as those receiving zero forward citations) and patent development duration. Panel B examines the association between AI novelty and forward citation counts, as well as the likelihood of a patent being classified as high-impact. The key independent variable is AI Novelty, defined as the number of competing AI patent filings within the preceding year. Dependent variables include indicators for low-impact AI patents, forward citation counts, the natural logarithm of development duration, and binary indicators for high-impact AI patents based on various citation thresholds. A high-impact (or low-impact) AI patent is defined as a dummy variable equal to one if the patent falls above (or below) a specified citation threshold for its filing year, and zero otherwise. Patent development duration is calculated as the number of days between the latest filing date among the patent's backward citations and the filing date of the focal AI patent. For all regressions, the robust standard errors are clustered at the IPC4 level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively. The sample period is from 2006 to 2019.

Panel A: Risk of Novel AI Patent

	OLS	OLS
	Low-Impact AI Patent (Forward Cites=0)	ln(Patent Development Duration)
	(1)	(2)
AI Novelty	0.003**	0.025***
	(2.58)	(9.15)
IPC4 FE	Yes	Yes
Year FE	Yes	Yes
Country FE	Yes	Yes
Obs.	2,139,308	1,399,212
R^2	0.08	0.10

Panel B: Reward of Novel AI Patent

	Poisson	OLS	OLS	OLS	OLS	OLS
	Forward Cites	Top 0.01%	Top 0.05%	Top 0.1%	Top 1%	Top 10%
	(1)	(2)	(3)	(4)	(5)	(6)
AI Novelty	0.037***	0.006***	0.011***	0.047***	0.080***	0.277***
	(12.05)	(4.78)	(6.43)	(8.10)	(6.82)	(6.94)
IPC4 FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	2,139,233	2,139,308	2,139,308	2,139,308	2,139,308	2,139,308
R^2		0.00	0.00	0.01	0.01	0.05

Table 6
Risk Associated with Novelty

This table presents the variance of forward citations among groups of AI patent filings, classified by the number of competing patent filings relative to a focal AI patent filing in the prior one year from Panel A. The Panel B reports F-tests to compare the differences between these groups.

Panel A: Forward Citation Variance by Novelty	Variance (Forward Cites)	N
Competing Patent Filing Count = 0 (Group 1, Most Novel)	2232.45	1,175,474
Competing Patent Filing Count = 1 (Group 2)	884.94	1,035,884
Competing Patent Filing Count = 2 (Group 3)	595.13	715,713
Competing Patent Filing Count = 3 (Group 4)	375.90	431,680
Competing Patent Filing Count > 3 (Group 5, Least Novel)	312.66	524,160

Panel B: F-Tests Comparing Variance Across Novelty Groups	F-statistic	P-value	Larger Variance Group
F-test (Group 1 vs Group 2)	2.52	0.00	Group 1
F-test (Group 1 vs Group 3)	3.75	0.00	Group 1
F-test (Group 1 vs Group 4)	5.94	0.00	Group 1
F-test (Group 1 vs Group 5)	7.95	0.00	Group 1
F-test (Group 1 vs Non-Group 1)	3.61	0.00	Group 1

Table 7
National AI-Related Policy, Innovation Novelty, and High-Impact Inventors

This table reports coefficients of Poisson regressions of the AI innovation inventor-panel on an indicator for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. Inventors are classified annually as high-impact based on their AI patent forward citations over the previous five years. The high-impact inventor indicator is set to 1 if the inventors rank in the top 10% within their countries. The AI policy indicator equals one if the country has implemented national AI-related policies. Te dependent variables are backward citations, average backward citations and AI novelty (the number of competing AI patent filings in the year preceding the current observation year). The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. For all regressions, the robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

	Poisson AI Patent Backward Cites	Poisson Avg AI Backward Cites	Poisson AI Novelty
	(1)	(2)	(3)
DID(AI Policy)	-0.014	-0.012	0.129**
	(-0.29)	(-0.54)	(2.11)
High-Impact Inventor(Top 10%)	0.340***	0.101***	-0.369***
	(5.36)	(5.86)	(-5.31)
DID(AI Policy)×High-Impact Inventor (Top 10%)	-0.188***	-0.080***	0.150***
	(-6.09)	(-7.48)	(5.34)
Year × Cohort FE	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes
Obs.	3,675,748	3,675,748	3,743,038

Table 8
National AI-Related Policy and High-Impact Inventors' AI Innovation

This table reports coefficients from inventor-panel Poisson regressions of inventors' AI innovation on an interaction terms of high productivity inventor indicator and national AI-related policies indicator, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. Inventors are classified annually as high-impact based on their AI patent forward citations over the previous five years. The high-impact inventor indicator is set to 1 if the inventors rank in the top 10% within their countries. The national AI-related policies indicator is set to 1 if the country has implemented national AI-related policies. The dependent variables are AI patent filing counts, forward citations, average forward citations, and average forward citations. The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. For all regressions, robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

	Poisson AI Patent Counts	Poisson AI Patent Forward Cites	Poisson Avg AI Forward Cites
	(1)	(2)	(3)
DID(AI Policy)	-0.019	-0.180	-0.147
	(-0.94)	(-0.91)	(-0.88)
High-Impact Inventor(Top 10%)	0.311***	0.937***	0.669***
	(4.38)	(6.29)	(7.43)
DID(AI Policy)×High-Impact Inventor(Top 10%)	-0.160***	-0.395***	-0.386***
	(-6.50)	(-13.28)	(-4.26)
Year × Cohort FE	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes
Obs.	3,900,297	3,515,936	3,515,936

Table 9
Government Failure to Target Radical AI innovation

This table presents the correlation between government-funded patents and AI patent quality based on a patent-level analysis. The independent variable is government-funded project, which equals to 1 if the project receive government grant. The dependent variables are high-impact AI patents, low-impact AI patents and Fcites. High-impact (Low-impact) AI patent is a dummy variable equal to 1 if the patent falls within the high (low) threshold for its filing year, and 0 otherwise. All regressions include IPC4 fixed effects, country fixed effects, and year fixed effects. Robust standard errors are clustered at the IPC4 level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively. The sample period is from 2006 to 2019.

Panel A: Government Funding and Patent Quality

	OLS	OLS	OLS	OLS	Poisson	Poisson
	Low-Impact AI Patent (Forward Cites = 0)	Top 10%	Top 5%	Top 1%	Forward Cites	AI Novelty
	(1)	(2)	(3)	(4)	(5)	(6)
Gov Funded Project	-0.044***	0.026***	0.010*	-0.003	0.061	0.474***
	(-8.25)	(3.50)	(1.91)	(-1.37)	(1.41)	(12.33)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
IPC4 FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	504,228	504,228	504,228	504,228	504,215	504,202
R^2	0.17	0.04	0.03	0.01		

Panel B: Government Funding and Patent Quality after National Robotics Initiative

	OLS	OLS	OLS	OLS	Poisson	Poisson
	Low-Impact AI Patent (Forward Cites = 0)	Top 10%	Top 5%	Top 1%	Forward Cites	AI Novelty
	(1)	(2)	(3)	(4)	(5)	(6)
Gov Funded Project	-0.052***	0.060***	0.030***	0.003	0.143***	0.478***
	(-7.12)	(4.53)	(3.06)	(0.75)	(3.55)	(7.72)
Gov Funded Project× post2011	0.012	-0.050***	-0.029***	-0.008**	-0.200***	0.006
	(1.45)	(-4.69)	(-3.69)	(-2.21)	(-5.85)	(0.11)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
IPC4 FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	504,228	504,228	504,228	504,228	504,215	504,202
R^2	0.17	0.04	0.03	0.01		

Table 10
Alternative Explanation: National AI-Related Policy and AI Military Patent

This table reports coefficients of Poisson regressions of the AI innovation inventor-panel on an indicator for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. The AI policy indicator is set to 1 if the country has implemented national AI-related policies. The dependent variables are AI military patent counts, AI military patent forward citations, AI military patent backward citations. We identify AI military patents based on CPC codes related to weapons, defense systems, and military applications: F41-F42 (Weapons & Ammunition): includes firearms (F41A-C, F41F-G, F41H-J), ammunition, missiles, fuzes, and blasting technologies (F42B-D). B (Military Transport & Vehicles): covers warships, submarines (B63G), military aircraft (B64C-D), and specialized military vehicles (B60P). G (Sensors & Control): encompasses radar, navigation, target tracking (G01S), vehicle control (G05D), data processing (G06F), alarm systems (G08B), and simulators (G09B). H (Military Communication): includes antennas (H01Q), secure communication, electronic warfare, cybersecurity, and military networks (H04B, H04L, H04W). The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. For all regressions, the robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

	Poisson	Poisson	Poisson	Poisson	Poisson
	AI Military Patent Counts	AI Military Patent Forward Cites	AI Military Patent Backward Cites	Avg AI Military Forward Cites	Avg AI Military Backward Cites
	(1)	(2)	(3)	(4)	(5)
DID(AI Policy)	-0.017	-0.197	-0.024	-0.178	-0.038
	(-0.88)	(-1.47)	(-0.92)	(-1.61)	(-1.58)
Year × Cohort FE	Yes	Yes	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes	Yes	Yes
Obs.	1,432,195	1,290,901	1,366,487	1,290,901	1,366,487

Table 11
National AI-Related Policy and Cross-Country Spillovers

This table reports coefficients from inventor-panel OLS regressions of inventors' fraction of cross country citations on an indicator for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. The AI policy indicator is set to 1 if the country has implemented national AI-related policies. The dependent variables are percentage of cross country forward citations and backward citations. The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. For all regressions, the robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

	OLS Cross Country Backward Cites.%	OLS Cross Country Forward Cites.%
	(1)	(2)
DID(AI Policy)	4.662**	5.763**
	(2.50)	(2.32)
Year × Cohort FE	Yes	Yes
Inventor × Cohort FE	Yes	Yes
Obs.	2,668,603	2,308,300
R^2	0.74	0.65

Appendix A. Additional Tables and Figures

Table B1
Variable Definitions

Variable Name	Abbrev.	Description
Innovation Variables (Source: Google Patents Public Data and AIPD)		
AI Patent Counts		Number of AI patent in a given year
AI Patent Forward Citations		Number of Forward citations of AI patent in a given year with null value replaced by zero in a given year
AI Patent Backward Citations		Number of backward citations of AI patent in a given year with null value replaced by zero
Average AI Patent forward citations	Avg AI Forward Cites	Number of forward citations of AI patent in a given year divided by Number of AI patent
Average AI Patent Backward Citations	Avg AI backward Cites	Number of backward citations of AI patent in a given year divided by Number of AI patent
Cross Country Forward Cites(%)		External forward citation is divided by total forward citations in a given year. a forward citation is classified as external if the country of the lead inventor of the citing filing does not match that of the cited filing.
Cross Country Backward Cites(%)		External backward citation is divided by total backward citations in a given year. a backward citation is classified as external if the country of the lead inventor of the citing filing does not match that of the cited filing.
Count of Competing AI Patents (Negative Novelty)		The number of similar AI patents filed in the same year as the AI patent with null value replaced by zero
Inventor AI Domains		The number of distinct AI Patent domain of AI patents in a given year
Inventor AI Domains HHI		Squaring the share of AI patents in each AI domain for an individual inventor, and then summing these squared shares
National AI-related Stratgies & Policies Variables (Source: OECD AI)		
National AI policies	DID(AI Policy)	The treat-post dummy variable assigning a value of 1 during the periods they are implemented and 0 otherwise.
Governance	DID(Governance)	The treat-post dummy variable assigning a value of 1 during the periods they are implemented policy focus on governance and 0 otherwise.
Financial Support	DID(Financial Support)	The treat-post dummy variable assigning a value of 1 during the periods they are implemented policy focus on financial support and 0 otherwise.
Guidance and Regulation	DID((Guidance and Regulation)	The treat-post dummy variable assigning a value of 1 during the periods they are implemented policy focus on guidance and Regulation and 0 otherwise.
AI Enablers and Other Incentives	DID(AI Enablers and Other Incentives)	The treat-post dummy variable assigning a value of 1 during the periods they are implemented policy focus on AI enablers and other incentives and 0 otherwise.

Table B2
Summary Statistics for AI Policy Budgets and AI Innovation by Country

This table reports the total AI policy budget, the budget allocated to specific policy categories, and the number of AI patent filings and AI-related scientific publications by country over the sample period from 2006 to 2019.

Country	Average Annual Budget ($M)	Governance ($M)	Financial Support ($M)	Guidance and Regulation ($M)	AI Enablers and Other Incentives ($M)	No. of AI Patents	No. of AI Inventors	No. of AI Academic Publications
Saudi Arabia	6000.00	5333.33	0.00	0.00	666.67	1309	1860	17249
United Kingdom	4165.53	3524.14	391.48	119.46	130.46	37015	41956	199070
China	2640.00	2640.00	0.00	0.00	0.00	370089	439440	1228565
United States	1678.00	423.00	1235.00	20.00	0.00	455461	525833	734519
France	1278.00	442.31	238.46	184.62	412.62	22366	27556	188607
Singapore	1140.00	645.00	310.00	0.00	185.00	4037	5476	37429
Korea, Republic of	1005.00	732.86	132.32	42.86	96.96	152437	146702	96457
Russian Federation	920.00	620.00	300.00	0.00	0.00	13430	28766	72042
Japan	790.50	610.50	37.50	0.00	142.50	302758	226842	180732
Czechia	740.00	356.67	13.33	6.67	363.33	1676	1934	26913
Netherlands	710.00	92.50	0.00	0.00	617.50	13121	16890	60445
Germany	499.00	335.00	112.33	0.00	51.67	46881	58816	213504
Israel	300.00	0.00	300.00	0.00	0.00	19622	19387	24341
Spain	275.00	275.00	0.00	0.00	0.00	7487	15251	98225
Switzerland	250.00	250.00	0.00	0.00	0.00	9135	10485	44514
Finland	190.50	105.00	35.00	50.00	0.50	7823	8017	22985
Australia	105.00	30.08	57.17	0.58	17.17	28060	53372	92894
Norway	85.00	0.00	75.00	10.00	0.00	2878	3677	19194
Denmark	85.00	21.43	10.71	42.14	10.71	4106	5139	21323
Sweden	82.50	1.00	75.00	3.00	3.50	12857	12123	34752
Mexico	60.00	40.00	0.00	0.00	20.00	1924	4006	28384
India	55.00	35.00	10.00	0.00	10.00	30156	35724	185571
Italy	51.00	16.00	17.50	0.00	17.50	10514	14436	115695
Canada	36.00	35.25	0.00	0.75	0.00	31671	33925	118679
Ireland	30.00	20.00	0.00	0.00	10.00	4600	4534	16194
Egypt	30.00	22.50	0.00	0.00	7.50	364	0	18680
Argentina	25.00	25.00	0.00	0.00	0.00	456	573	9287
Brazil	20.50	20.50	0.00	0.00	0.00	2627	5202	72016
Portugal	19.00	14.50	3.50	1.00	0.00	617	1155	28926
Austria	15.00	6.50	0.00	8.50	0.00	3333	4243	33096
Türkiye	14.50	6.25	0.00	0.00	8.25	1525	2338	42413
South Africa	10.00	10.00	0.00	0.00	0.00	940	1194	11177
Estonia	7.00	7.00	0.00	0.00	0.00	278	0	3337
Hungary	1.00	0.50	0.00	0.50	0.00	1092	1522	12342
Romania	1.00	1.00	0.00	0.00	0.00	894	1184	21746
Luxembourg	0.50	0.33	0.00	0.17	0.00	370	0	4334
Belgium	0.50	0.25	0.12	0.12	0.00	5016	5659	37536
New Zealand	0.00	0.00	0.00	0.00	0.00	1446	2165	13271
United Arab Emirates	0.00	0.00	0.00	0.00	0.00	369	0	7128

Table B3
Summary Statistics of Domains for AI Patent Filings

This table reports the overview statistics of domains for AI patent filings by country from 2006 to 2019.

Country	No. of AI Patent Filings	Vision(%)	AI Hardware(%)	Evolutionary Computation(%)	Machine Learning(%)	Planning and Control(%)	Natural Language Processing(%)	Speech(%)	Knowledge Processing(%)	Primary Filings Focus
United States	455461	16.72	23.27	3.53	3.23	32.55	8.19	3.09	9.42	Planning and Control
China	370089	24.08	20.26	3.49	2.94	30.70	6.92	3.30	8.31	Planning and Control
Japan	302758	33.99	16.30	6.04	1.33	25.70	6.31	2.56	7.78	Vision
Korea, Republic of	152437	26.10	21.34	3.91	1.85	27.50	8.00	3.63	7.66	Planning and Control
Germany	46881	22.27	20.12	3.65	2.22	37.86	4.28	2.95	6.64	Planning and Control
Taiwan, Province of China	38989	34.98	24.46	2.80	1.63	22.82	5.55	2.95	4.82	Vision
United Kingdom	37015	16.70	22.27	3.61	2.83	36.03	6.13	3.26	9.16	Planning and Control
Canada	31671	16.70	20.98	3.83	2.76	32.53	9.34	3.36	10.49	Planning and Control
India	30156	12.10	24.86	2.74	3.28	31.84	11.47	2.41	11.30	Planning and Control
Australia	28060	16.32	15.98	6.53	2.40	39.63	6.73	3.57	8.84	Planning and Control
France	22366	24.27	21.09	3.98	3.61	31.36	5.42	2.23	8.04	Planning and Control
Israel	19622	21.51	27.66	2.92	3.52	24.04	8.10	2.66	9.60	AI Hardware
Russian Federation	13430	24.47	20.69	4.44	2.79	32.21	6.40	2.28	6.73	Planning and Control
Netherlands	13121	29.92	18.32	4.73	2.81	29.39	4.38	3.09	7.35	Vision
Sweden	12857	17.99	16.43	4.32	2.61	37.48	4.50	4.10	12.56	Planning and Control
Italy	10514	16.92	19.75	3.73	2.38	43.99	4.41	1.76	7.06	Planning and Control
Switzerland	9135	20.43	20.76	4.32	4.07	33.05	6.56	3.89	6.93	Planning and Control
Finland	7823	18.37	18.75	3.80	2.94	30.13	6.14	5.57	14.30	Planning and Control
Spain	7487	19.95	18.93	5.82	2.83	36.26	5.41	2.60	8.19	Planning and Control
Belgium	5016	21.09	23.68	5.52	2.87	31.74	4.78	2.61	7.70	Planning and Control
Ireland	4600	13.33	20.35	2.61	4.00	34.83	12.72	2.39	9.78	Planning and Control
Denmark	4106	12.52	14.34	6.41	2.29	44.20	4.48	9.99	5.77	Planning and Control
Singapore	4037	20.83	24.85	2.82	3.81	31.04	5.92	3.32	7.41	Planning and Control
Austria	3333	21.84	23.13	3.96	2.73	36.45	3.75	2.88	5.25	Planning and Control
Norway	2878	18.31	20.95	2.85	2.36	40.34	5.84	2.50	6.85	Planning and Control
Brazil	2627	13.06	18.20	5.41	3.81	41.07	8.60	2.17	7.69	Planning and Control
Poland	2015	18.56	21.44	2.83	2.78	32.75	6.60	5.51	9.53	Planning and Control
Mexico	1924	14.92	15.44	6.24	2.70	44.33	6.39	3.17	6.81	Planning and Control
Czechia	1676	14.62	21.78	3.04	4.18	37.23	7.40	2.15	9.61	Planning and Control
Türkiye	1525	23.74	16.66	3.34	2.23	34.23	5.64	5.05	9.11	Planning and Control
Malaysia	1460	20.62	30.41	3.56	2.74	26.85	5.34	3.22	7.26	AI Hardware
New Zealand	1446	17.50	17.63	6.29	2.42	37.48	6.78	3.32	8.58	Planning and Control
Saudi Arabia	1309	14.21	20.02	3.82	7.94	38.81	6.80	0.69	7.72	Planning and Control
Hong Kong	1109	29.49	19.39	3.16	3.97	27.23	7.48	3.25	6.04	Vision
Hungary	1092	19.23	18.86	3.85	3.66	30.86	4.67	1.92	16.94	Planning and Control
South Africa	940	12.77	20.74	6.06	2.98	39.89	5.74	2.34	9.47	Planning and Control
Romania	894	25.17	25.73	2.24	4.59	22.93	8.50	1.68	9.17	AI Hardware
Ukraine	661	26.17	24.96	4.39	2.27	24.66	6.51	2.57	8.47	Vision
Portugal	617	17.99	27.39	4.38	4.54	28.85	4.38	2.59	9.89	Planning and Control
Greece	597	17.09	24.62	3.69	4.86	28.48	6.70	5.03	9.55	Planning and Control
Bulgaria	473	8.67	34.04	1.69	1.69	35.10	7.61	1.69	9.51	Planning and Control
Argentina	456	14.91	24.34	4.82	5.48	33.33	7.89	2.41	6.80	Planning and Control

Table B4
Model Parameters and Variables

Variable/Parameter	Definition
p_i	Inventor i's idiosyncratic probability of success on a radical project. Represents innate talent.
$F(p)$	The cumulative distribution function (CDF) of inventor talent p_i over the support $[0, \overline{p}]$.
V_R	The private value (e.g., capitalized patent profits) of a successful radical innovation.
V_S	The total social value of a successful radical innovation, where $V_S > V_R$.
w_t	The economy-wide wage, representing the opportunity cost of pursuing innovation.
c	The private cost of undertaking a radical innovation project.
G	A government grant (subsidy) provided for pursuing a radical innovation project.
p^*	The equilibrium cutoff talent level for choosing radical innovation.
ψ^*	The fraction of the inventor population that chooses to pursue radical innovation.
A_t	The aggregate stock of technology in the economy at time t.
δ	The small, certain productivity gain from a successful incremental innovation.
γ	The large productivity gain from a successful radical innovation, with $\gamma \gg \delta$.
p_H	The talent threshold that defines a "high-quality" project.
α	Type I Error (False Negative) probability: the government denies funding to a high-quality project ($p_i \geq p_H$).
β	Type II Error (False Positive) probability: the government provides funding to a low-quality project ($p_i < p_H$).

Table B5
Prompt used to classify whether a policy qualifies as a national AI policy

This table presents the Python prompt used with ChatGPT GPT-4o to assess whether a given policy pertains to Artificial Intelligence (AI). Each policy was evaluated ten times, and we retained only those classified as (1) AI-related, (2) national in scope, and (3) applicable to all entity groups in at least seven out of ten evaluations. Criteria (2) and (3) ensure consistency with our dependent variable, which includes all inventors.

Assess whether the following policy is a national policy related to Artificial Intelligence (AI)

based on its description, background information, and objectives, as well as other reliable internet resources.

Consider whether the policy suggests a nationwide implementation related to AI by the central government.

Description: {Description}

Background: {Background if pd.notna(Background) else 'No additional background provided.'}

Objectives: {Objective}

Respond with **'Yes'** or **'No'** to indicate whether it is a national AI policy, followed by a brief explanation

incorporating general internet knowledge where relevant.

Example: Yes: The policy outlines national initiatives for AI development across all government departments.

Table B6
Keywords Describing National AI Strategies and Policies

This table shows the keywords describing national AI-related policies in thirteen principal countries, from 2006 to 2019.

Country	Keywords
Australia	Provide financial incentives, Promote public-private partnerships, Foster AI education and workforce training, Support AI startups, Develop national AI strategies, Establish AI regulatory frameworks, Improve infrastructure, Guide corporate investment, Boost internal R&D, Enhance AI governance, Fund research missions, Support innovative collaborative technologies, Increase transparency and integrity, Encourage adoption and use of AI technologies, Co-fund collaborative projects, Provide access to emerging technologies, Develop ethical AI frameworks, Build international AI relationships, Strengthen AI capability and adoption, Attract and train AI specialists
Canada	risk-based approach, Algorithmic Impact Assessment, regulatory frameworks, public-facing automated decision-making systems, Pan-Canadian Artificial Intelligence Strategy, investments in talent, research capacity, commercialization, standardization, interconnected nodes of scientific excellence, global thought leadership, support national research community, foster co-operation between AI research centres and industry, multi-stakeholder and multi-disciplinary expertise, AI Source List, expedite procurement, establish pre-qualified list of suppliers
China	Enhancing AI hardware capacity, Strengthen platform ecosystems, Develop research and industrial leadership, Build technology parks for AI research, Establish AI governance and regulatory frameworks, Promote AI education and workforce training, Foster international cooperation in AI, Support AI startups and innovation centers, Optimize university infrastructure for AI, Set national standards for AI technologies, Increase R&D spending in AI, Develop national AI strategies, Cultivate AI talent development, Provide financial incentives for AI development, Promote public-private partnerships in AI, Enhance capabilities of data infrastructure, Regulate algorithm security and development, Standard
France	Enhance research ecosystem, Accelerate AI dissemination, Encourage trustworthy AI development, Develop AI educational excellence, Foster interdisciplinary AI research, Support AI startups and SMEs, Create prestigious AI research institutes, Offer high-level AI scientific training, Establish collaborative research platforms, Develop national AI strategies, Promote public-private partnerships, Provide financial incentives, Improve infrastructure (e.g., supercomputers, data hubs), Foster AI education and workforce training, Support the digitalisation of SMEs, Develop embedded and frugal AI solutions, Enhance AI governance, Develop sovereign AI frameworks
Germany	Develop national AI strategies, Establish AI regulatory frameworks, Foster AI education and workforce training, Promote public-private partnerships, Provide financial incentives, Enhance AI governance, Support AI startups, Improve infrastructure, Guide corporate investment, Boost internal R&D, Create ethical rules for AI, Build AI research ecosystem, Develop technology transfer initiatives, Fund AI applications for environmental benefits, Integrate AI in societal sectors, Support demographic and ecological transformation, Foster international research collaboration, Enhance participatory technology design, Develop AI platforms for universal use, Harmonize legislation for AI systems
India	Foster talent and build a community, Encourage research and reduce R&D costs, Reduce import dependency, Position as a provider of end-to-end solutions, Create a policy and legal framework, Issue concrete recommendations for government, industry, and research programs, Standardisation of Artificial Intelligence technology, Establish the National Program on AI, Undertake exploratory proof-of-concept AI projects, Craft a national strategy for AI ecosystem, **Collaborate with experts
Israel	Develop national AI strategies, Improve infrastructure, Support AI startups, Provide financial incentives, Encourage entrepreneurship, **
Japan	Construct high-speed network, Promote R&D for innovative fundamental technology, Accelerate Joint AI R&D with Industries, Academia, and Governments, Establish new governance models, Promote appropriate and proactive implementation of AI
Korea	Provide AI technology development infrastructure, Develop and promote AI R&D strategies, Establish intelligent network, Reform R&D system, Promote cross-ministerial collaboration, Enhance interactions with the general public, Foster data and AI convergence, Establish ethical standards for AI, Build a foundation for social deliberation on AI ethics, Secure world-class AI technology, Promote active use of the data value chain, Create economic surplus through AI
Netherlands	Promote public-private partnerships, Foster AI education and workforce training, Develop national AI strategies, Provide financial incentives, Enhance AI governance
Sweden	Elevate competence about AI, Promote education and competence development, Deliver policy proposals, Adapt regulatory frameworks, Accelerate policy development, Formulate a national roadmap, Recommendations on investment and policy development, Map the use of AI and big data analysis, Accelerate applied AI research and innovations, Develop and provide infrastructure for data management, Support industry-science cooperation, Increase undergraduate AI education
United Kingdom	Develop national AI strategies, Enhance AI governance, Foster AI education and workforce training, Support AI startups, Establish AI regulatory frameworks, Provide financial incentives
United States	advance AI research, develop regulatory frameworks, promote public-private partnerships, foster AI education and workforce training, support AI startups, enhance AI governance, provide financial incentives, develop national AI strategies, guide corporate investment, boost internal R&D, ensure AI safety and security, accelerate AI adoption, cultivate AI workforce, establish AI ethical guidelines, expand AI R&D investments, promote AI in national security, develop AI standards and benchmarks, increase access to AI resources

Table B7

Examples of National AI Strategies and Policies

This table shows examples of national AI-related strategies and policies with the policy name, country, start year and description from OECD AI Policy Database.

Policy Name	Country	Start Year	Description
NATIONAL AI R&D STRATEGIC PLAN	United States	2016	The purpose of this Plan is to convey a clear set of R&D priorities that address strategic research goals, focus Federal investments on those areas in which industry is unlikely to invest, and address the need to expand and sustain the pipeline of AI R&D talent.
EXECUTIVE ORDER ON MAINTAINING AMERICAN LEADERSHIP IN AI	United States	2019	This order seeks to spark significant advancements in AI within the federal government, industry, and academia to boost scientific research, economic growth, and national security. It strives to set technical standards, simplify the process for AI testing and deployment, and support the growth of new AI-focused industries and the broader adoption of AI in existing sectors.
AI RESEARCH PROGRAMME	Singapore	2018	AI Research Programme seeds high quality research efforts aimed at developing fundamental AI novel techniques, algorithms and adjacent technologies that will eventually significantly contribute to the other pillars of AI Singapore. The AI Singapore Research Programme will also encourage national research collaborations and nurture local AI talents.
HIGH PERFORMANCE COMPUTING INFRASTRUCTURE PROJECT	Japan	2012	Supercomputer "Fugaku" as a central core, construct high-speed network among universities and public research sectors for various uses to create outstanding results in various fields such as AI, data science, medicine, climate, space, disaster defense.
NATIONAL NEW GENERATION AI PLAN	China	2017	The Plan involves initiatives and goals for R&D, industrialization, talent development, education and skills acquisition, standard setting and regulations, ethical norms, and security. To make China's AI industry "in-line" with competitors by 2020. To reach "world-leading" in some AI fields by 2025. To become the "primary" center for AI innovation by 2030. By 2030, the Chinese government aims to cultivate an AI industry worth RMB 1 trillion (EUR 130 billion), with related industries worth RMB 10 trillion (EUR 1300 billion) . In addition, the Chinese government has also partnered with national tech companies to develop research and industrial leadership in specific fields of AI and will build a USD 2.1 billion (EUR 1.8 billion) technology park for AI research in Beijing.
PAN-CANADIAN AI STRATEGY	Canada	2017	The Government of Canada is providing the Canadian Institute for Advanced Research (CIFAR) with a contribution of $125 million (in Canadian dollars) to launch a Pan-Canadian Artificial Intelligence Strategy to retain and attract top academic talent, and to increase the number of post-graduate trainees and researchers studying AI. Strategy focused on investments in talent, research capacity, commercialization and standardization of (generative) AI.

Table B8
Pretrend Analysis

This table presents the pretrend analysis for our tests. All specifications follow the main table but the treatment is separated by the timing of treatment from $\leq t-3$ to $\geq t+3$.

	AI Patent Counts	AI Patent Forward Cites	Avg AI Forward Cites
	(1)	(2)	(3)
t $\leq$ -3	-0.013	-0.041	-0.011
	(-0.45)	(-0.43)	(-0.14)
t = -2	-0.007	0.018	0.029
	(-0.50)	(0.40)	(0.82)
t = 0	-0.032**	-0.151	-0.141
	(-2.18)	(-1.56)	(-1.64)
t = 1	-0.088**	-0.430*	-0.320
	(-2.32)	(-1.90)	(-1.59)
t = 2	-0.124**	-0.658**	-0.486*
	(-2.06)	(-1.98)	(-1.69)
t $\geq$ -3	-0.166***	-0.822**	-0.591**
	(-2.59)	(-2.34)	(-2.02)
Year $\times$ Cohort FE	Yes	Yes	Yes
Inventor $\times$ Cohort FE	Yes	Yes	Yes
Obs.	3,900,297	3,515,936	3,515,936

Table B9
National AI-Related Policy and AI Innovation (Poisson Regressions Controlling GDP)

This table reports coefficients from inventor-panel Poisson regressions of inventors' AI innovation on an indicator for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. The indicator of AI policy is set to 1 if the country has implemented national AI-related policies. The dependent variables are AI patent filings counts, forward citations, and average forward citations. The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. For all regressions, we include control variables ln(GDP per Capita). The robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

	Poisson AI Patent Counts	Poisson AI Patent Forward Cites	Poisson Avg AI Forward Cites
	(1)	(2)	(3)
DID(AI Policy)	-0.064***	-0.197**	-0.157**
	(-3.70)	(-2.25)	(-2.06)
ln(GDP per capita)	1.057***	4.601***	3.776***
	(6.22)	(9.35)	(7.07)
Year × Cohort FE	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes
Obs.	3,900,297	3,515,936	3,515,936

Table B10
National AI-Related Policy's ln(1+Budget) and AI Innovation (Poisson)

This table reports coefficients from inventor-panel Poisson regressions of inventors' AI innovation on budgets for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. Total budget is ln(1+ Total Budgets) for the policy of the country. The dependent variables are AI patent filings counts, forward citations and average forward citations. The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. For all regressions, the robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

	Poisson AI Patent Counts	Poisson AI Patent Forward Cites	Poisson Avg AI Forward Cites
	(1)	(2)	(3)
ln(1+Total Budget)	-0.013**	-0.074*	-0.058
	(-2.56)	(-1.71)	(-1.62)
Year × Cohort FE	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes
Obs.	3,900,297	3,515,936	3,515,936

Table B11
National AI-Related Policy Indicator and AI Innovation (OLS)

This table reports coefficients from inventor-panel OLS regressions of inventor's AI innovation on an indicator for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects.The indicator of AI policy is set to 1 if the country has implemented national AI-related policies. The dependent variables are ln(AI patent filings counts), ln(forward citations) and ln(average forward citations). The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. For all regressions, the robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

	OLS ln(AI Patent Counts)	OLS ln(AI Patent Forward Cites)	OLS ln(Avg AI Forward Cites)
	(1)	(2)	(3)
DID(AI Policy)	-0.066**	-0.301***	-0.232**
	(-2.65)	(-2.73)	(-2.39)
Year × Cohort FE	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes
Obs.	3,900,297	2,308,300	2,308,300
R^2	0.54	0.55	0.55

Table B12
Heterogeneous Effect of National AI-Related Policy on AI Innovation

This table reports coefficients of Poisson regressions of the AI innovation inventor-panel on an indicator for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. The AI policy indicators reflect the implementation of specific categories of AI policies such as governance, guidance & regulation, financial support, and AI enablers & other incentives. The dependent variables are are AI patent filings counts, forward citations and average forward citations. The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. For all regressions, the robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

	Poisson AI Patent Counts				Poisson AI Patent Forward Cites				Poisson Avg AI Forward Cites			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
DID(Governance)	-0.019				-0.060				-0.068			
	(-1.48)				(-0.65)				(-0.81)			
DID(Financial Support)		-0.093^{***}				-0.473^{*}				-0.365^{*}		
		(-2.73)				(-1.91)				(-1.78)		
DID(Guidance and Regulation)			-0.084^{***}				-0.271^{***}				-0.206^{**}	
			(-5.00)				(-2.95)				(-2.38)	
DID(AI Enablers and Other Incentives)				-0.153^{***}				-0.565^{**}				-0.340
				(-3.22)				(-1.99)				(-1.41)
Year × Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	3,900,297	3,900,297	3,900,297	3,900,297	3,515,936	3,515,936	3,515,936	3,515,936	3,515,936	3,515,936	3,515,936	3,515,936

Table B13
Placebo Test: National AI-Related Policy and Non-AI Innovation

This table reports the coefficients from Poisson and OLS regressions of the non-AI innovation inventor-level panel on an indicator for national AI-related policies. Panel A presents stacked Difference-in-Differences (DID) estimates, controlling for inventor-by-cohort and year-by-cohort fixed effects, following Gormley and Matsa (2011); Baker, Larcker, and Wang (2022). Panel B reports estimates from the structural DID approach proposed by Callaway and Sant'Anna (2021). The AI policy indicator equals one if the country has implemented national AI-related policies. The dependent variables include non-AI patent filing counts, forward citations and average forward citations. The sample consists of inventor-year observations from five years prior to, and up to ten years following, the implementation of national AI-related policies across countries between 2006 and 2019. For all regressions, robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

Panel A: Stacked DID (Poisson)

	Non-AI Patent Counts	Non-AI Patent Forward Cites	Avg Non-AI Patent Forward Cites
	(1)	(2)	(3)
DID(AI Policy)	-0.097	-0.370	-0.373*
	(-1.62)	(-1.52)	(-1.70)
Year × Cohort FE	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes
Obs.	21,796,942	21,796,942	21,796,942

Panel B: Methodology from Callaway and Sant'Anna (2021) (OLS)

	ln(1+Non-AI Patent Forward Cites)		
	ATT on Treated	ATT by Calendar Period	ATT by Group
	(1)	(2)	(3)
DID(AI Policy)	-0.108*	-0.145**	-0.058*
	(-1.77)	(-2.00)	(-1.86)
Obs.	4,836,545	4,836,545	4,836,545

Table B14
Robustness Test: National AI-Related Policy and AI Innovation

This table reports coefficients of Poisson regressions of the AI innovation inventor-panel on an indicator for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. The AI policy indicator equals one if the country has implemented national AI-related policies. The dependent variable is AI patent filing counts. The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. Columns (1)–(3) present robustness tests excluding specific groups of inventors from the sample: Column (1) excludes inventors based in the United States, Column (2) excludes inventors based in China, and Column (3) excludes inventors based in both the United States and China. For all regressions, the robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively. The sample period is from 2006 to 2019.

	AI Patent Counts		
	Exclude USA	Exclude China	Exclude USA & China
	(1)	(2)	(3)
DID(AI Policy)	-0.088**	-0.071***	-0.091***
	(-2.13)	(-2.99)	(-4.04)
Year $\times$ Cohort FE	Yes	Yes	Yes
Inventor $\times$ Cohort FE	Yes	Yes	Yes
Obs.	3,295,872	2,923,814	2,319,389

Table B15
Government-Funded Inventor and High-Impact Inventor

This table reports coefficients from inventor-panel OLS regressions of an indicator for government-funded inventor on an indicator for high-impact inventor, year fixed effects, and inventor field fixed effects. The government funding indicator equals 1 if the inventor filed an AI patent supported by government funding in a given year, and 0 otherwise. The high-impact inventor indicator equals 1 if the inventor is among the top 10% (5 % or 1%) of inventors based on forward citations to AI patents over the previous five years. The post-2011 indicator equals 1 for years after 2011, and 0 otherwise. For all regressions, the robust standard errors are clustered at the inventor level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively. The sample period is from 2006 to 2019.

	Government Funded Inventor					
	(1)	(2)	(3)	(4)	(5)	(6)
High-Impact Inventor (Top 10%)	0.007***	0.008***				
	(7.99)	(4.98)				
High-Impact Inventor (Top 10%)$\times$ post-2011		-0.001				
		(-0.52)				
High-Impact Inventor (Top 5%)			0.007***	0.004*		
			(5.57)	(1.89)		
High-Impact Inventor (Top 5%)$\times$ post-2011				0.004*		
				(1.79)		
High-Impact Inventor (Top 1%)					-0.000	-0.004
					(-0.07)	(-1.00)
High-Impact Inventor (Top 1%)$\times$ post-2011						0.006
						(1.20)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	1,046,698	1,046,698	1,046,698	1,046,698	1,046,698	1,046,698
R^2	0.04	0.04	0.04	0.04	0.04	0.04

Table B16
National AI-Related Policy and Uniqueness of AI Innovation

This table reports coefficients from inventor-panel Poisson regressions of the number of competing AI patents on an indicator for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. The AI policy indicator is set to 1 if the country has implemented national AI-related policies. The dependent variables include: (i) the number of competing AI patents filed in the current year and (ii) cumulatively up to the current year. The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. For all regressions, the robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

	AI Novelty (t = 0)	AI Novelty (t < 0)
	(1)	(2)
DID(AI Policy)	0.165**	0.177***
	(2.22)	(3.34)
Year × Cohort FE	Yes	Yes
Inventor × Cohort FE	Yes	Yes
Obs.	3,800,627	3,887,712

Table B17
National AI-Related Policy and Domains of AI Innovation

This table reports coefficients from Poisson and OLS regressions of the inventor-panel of the number of AI patent filings/applications domains on an indicator for national AI-related policies, inventor-by-cohort fixed effects, and year-by-cohort fixed effects. The AI policy indicator is set to 1 if the country has implemented national AI-related policies. The dependent variables are inventor's number of AI patent domains and AI patent HHI. The data include inventor-year observations in the 5 years prior and up to 10 years following the implementation of national AI-related policies in each country from 2006 to 2019. For all regressions, robust standard errors are clustered at the country level, and t-statistics are illustrated in parentheses.***,**,* represent the significance level at 1%, 5% and 10% ,respectively.

	Poisson	OLS
	Inventor AI Domains	Inventor AI Domains HHI
	(1)	(2)
DID(AI Policy)	-0.058**	0.023**
	(-2.29)	(2.43)
Year $\times$ Cohort FE	Yes	Yes
Inventor $\times$ Cohort FE	Yes	Yes
Obs.	3,900,297	3,900,297
R^2		0.49

Figure B1
Placebo Simulation: Distribution of t-Statistics from Randomized Policy Timing

These figures present placebo simulations designed to assess whether the estimated treatment effect on patent filings reported in Panel A of Table 2 could occur by chance. In each iteration, we randomly assign pseudo-policy adoption years across countries and re-estimate the stacked difference-in-differences (DID) model following Gormley and Matsa (2011); Baker, Larcker, and Wang (2022), recording the corresponding t-statistic for the treatment effect.

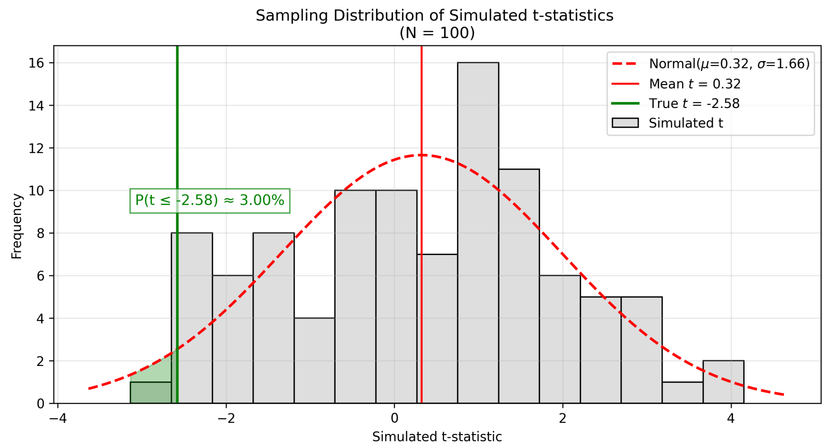

(a) 100 Iterations

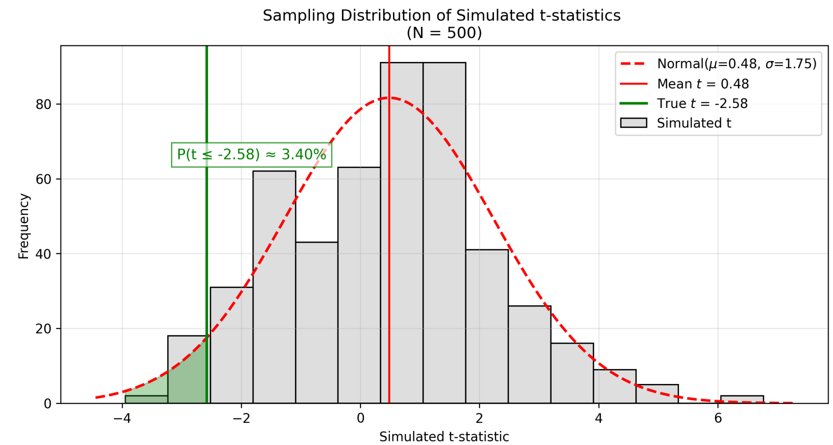

(b) 500 Iterations

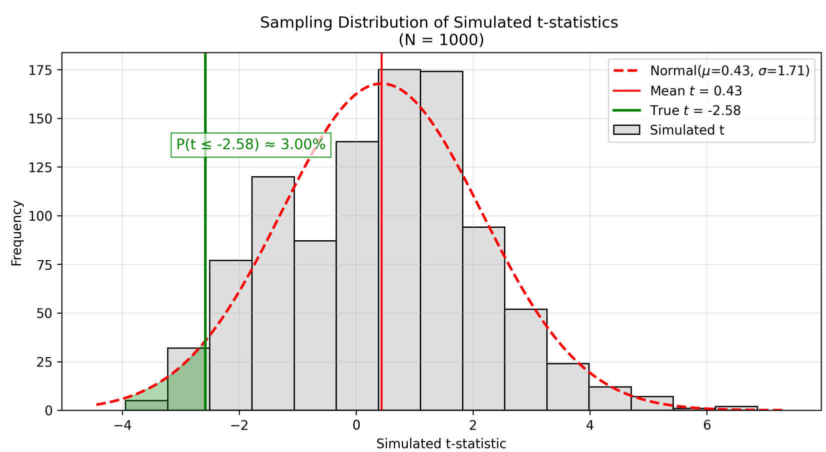

(c) 1000 Iterations

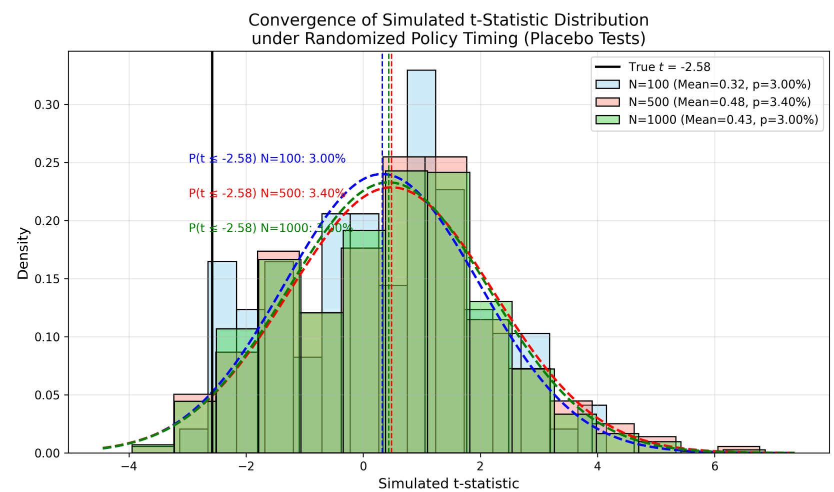

(d) Comparison: 100, 500, and 1000 Iterations