

# Innovation Policies and Inventors' Productivity: Evidence from Global AI Initiatives\*

Leo Liu

UTS Business School, University of Technology Sydney, [leo.liu@uts.edu.au](mailto:leo.liu@uts.edu.au)

Fariborz Moshirian

UNSW Business School, University of New South Wales, [f.moshirian@unsw.edu.au](mailto:f.moshirian@unsw.edu.au)

Vikram Nanda<sup>†</sup>

Naveen Jindal School of Business, University of Texas at Dallas, [vikram.nanda@utdallas.edu](mailto:vikram.nanda@utdallas.edu)

Sheng Xu

UNSW Business School, University of New South Wales, [sheng.xu3@unsw.edu.au](mailto:sheng.xu3@unsw.edu.au)

---

**Abstract.** Governments have developed policies and increased investment in AI technology, recognizing its potential as a future engine of economic growth. In this study, we construct two novel datasets of global AI patents and AI innovation policies across 42 countries, to examine, for the first time, the effect of these policies on the productivity of AI scientists. We find that these “mission-oriented” innovation policies are effective in incentivizing scientists to pursue more risky and exploratory AI inventions that the private sector may be unwilling to undertake. While such projects are essential for successful breakthroughs, we find that short-term productivity (measured by the quantity and average quality of innovations) decreases in the post-policy period. This pattern is consistent with the J-curve dynamics commonly observed during the transition to a new innovation paradigm. Our results highlight the transitional risks associated with targeted government support for innovation, as well as evidence that it incentivizes more novel projects with long-run technological potential.

**Keywords:** Innovation, Patents, Artificial Intelligence, Mission-oriented Policy, Scientists, Inventors

**JEL:** O30, O31, O38, I28, G38

---

\*We thank Zahi Ben-David, Michael Evans, Caroline Flammer, Cesare Fracassi, Kristoffer Glover, Xian Gu (discussant), Shenje Hsieh (discussant), Ryoonee Kim (discussant), Alan Kwan, Hao Liang, Xinwei Li, Jing Lu, Ronald Masulis, Christopher Reilly, Janghoon Shon, Sascha Steffen, Luis Felipe Varas Greene, Pingle Wang, Sumudu Watugala, Ting Xu, Yexiao Xu, Yuchen Xu, Bernard Yeung, Harold Zhang, Xiaoyan Zhang, and Dexin Zhou, as well as other seminar and conference participants at the University of Texas at Dallas, the University of Technology Sydney, the ABFER–JFDS Conference on Capital Market Research in the Era of AI 2025, CAFM 2025, and the Global AI Finance Conference 2025, for helpful comments and suggestions. All errors are our own.

<sup>†</sup>Corresponding Author

# 1 Introduction

A recognition that AI could be the next engine of economic growth has prompted governments around the globe to develop policies and increase investment in AI. According to the OECD AI Policy Observatory in 2025,<sup>1</sup> over 1,000 AI policy initiatives have been adopted across 69 countries, reflecting a global consensus that AI is a strategic priority for innovation policies and public governance (OECD.AI 2025). As part of the global trend toward national AI strategies, several major economies have substantially expanded their AI policy commitments over the past two decades, accompanied by significant government investment in AI research and development. A recent example is the National AI Action Plan of the U.S., which involves multi-billion-dollar public investment in firms like OpenAI and NVIDIA to maintain AI leadership. Similarly, the Chinese government has invested heavily in AI since 2017, as part of their signature “whole-of-nation” approach. However, despite the prevalence of policies to promote AI innovation, the effect of these policies is not well understood.

Innovation policies are an important research topic with a substantial literature. However, existing studies offer mixed findings on the impact of government policies (for a review, see Bloom, Van Reenen, and Williams 2019), with much of the evidence based on untargeted tax subsidies for all eligible R&D expenses or on immigration rules designed to attract STEM professionals. In contrast, the current wave of AI policies is best characterized as “mission-oriented,” involving targeted support for high-risk and potentially high-reward projects deemed most promising. For example, the US’s National AI R&D Strategic Plan (2016) states that the goal is to fund long-term, high-risk, high-reward research that private industry is unlikely to undertake.<sup>2</sup>

A rationale behind using targeted support is that market frictions arise such that innovators are forced to internalize the costs and risks of innovation, leading to the neglect of promising projects (Howell 2017). This is particularly true for AI since it involves high upfront costs, such as hardware

---

<sup>1</sup>See, <https://oecd.ai/>

<sup>2</sup>Similarly, other countries have also set ambitious goals in these policies: the United Kingdom’s National AI Strategy (2021) aims to transition the country from an AI nation to an *AI Superpower* over the next ten years, securing its status as a global science and technology superpower. South Korea’s National Strategy for Artificial Intelligence (2019) aims to become a *Top 3 Digital Competitor* in the world by 2030 and lead the global market in AI semiconductor chips. Russia’s National Strategy for the Development of Artificial Intelligence (2019) seeks to ensure national sovereignty and leadership, framed by the declaration that *whoever becomes the leader in this sphere will become the ruler of the world*. The United Arab Emirates’ Strategy for Artificial Intelligence (2017) aims to establish the nation as a global *testbed* for AI innovation, with a vision for 100% reliance on AI for government services and data analysis by 2031. Singapore’s National AI Strategy 2.0 (2023) seeks to develop the country into a *Global Hub* for AI solutions, prioritizing the highest density of AI adoption per capita to ensure national competitiveness.

(e.g., GPU clusters) and human capital. It is also risky because AI has winner-takes-all dynamics. An inventor might expend significant effort but receive few rewards if the product is not the very best. For example, only leading firms in large language models make profits because users only choose the best models. Risk-taking by innovators is required because, in order to achieve the best, one must pursue novel ideas and embrace failure to reach breakthrough solutions. In this regard, an important empirical finding of our study is that national AI policies tend to be successful in fostering inventors to undertake novel, high-risk, high-reward AI projects.

We begin our paper by describing national AI policies and innovation in detail, based on novel datasets constructed for the study. First, we observe that AI policies are often continued following the initial implementation. For example, the US initiated the National Robotics Initiative in 2011; subsequently, it introduced several initiatives, including the BRAIN initiative in 2013, nanotechnology in 2014, strategic computing, and big data research in 2016. This trajectory continued with the AI R&D Strategic Plan in 2016 and then to the 2025 Stargate project. Likewise, China initiated the “New Generation AI Plan” in 2017, subsequently launching several sub-policies under this broader framework. These initiatives often feature long-term goals, such as becoming “world-leading” by 2030.

These policies employ a diverse set of instruments, such as project grants for public research, fellowships and postgraduate loans and scholarships, the establishment of networking and collaborative platforms, and investments in AI computing and research infrastructure. These tools sets them apart from the traditionally studied the traditional R&D tax incentives and grant programs emphasized in the prior literature. This has a general differences among countries, the United States and the United Kingdom, among other tools, place greater emphasis on project grants for public research and on fellowships and postgraduate funding directed at researchers in public research institutions, China focus on funding firms directly and leads VC money and focus on infrastructure building. Japan and Germany devote more policy effort to networking and collaborative platforms as well as investments in AI computing and research infrastructure. European countries tend to place greater emphasis on regulatory-and guidance-based approach-es to promote ethical and trustworthy AI adoption. e.g. France’s National AI Strategy in 2018 focus on human-centric and responsible AI development.

Regarding AI innovation, we observe a shift in AI leadership over the past decades (see [Figure 1](#)). The 1990s were characterized by Japan’s preeminence, where it accounted for nearly half of global AI patent filings, reflecting its early industrial automation capabilities. However, the post-2010 era

marks a distinct paradigm shift, mainly driven by China’s exponential trajectory. China’s share expanded from a negligible fraction in the early 2000s to a dominant majority by 2021–2022, supplanting traditional leaders such as the United States and Japan. Meanwhile, the United States has maintained a resilient but proportionally shrinking presence, suggesting that while western innovation remains active, the sheer volume of AI innovation has increasingly tilted toward East Asia. This observation largely motivates our international investigation, given that most existing studies on innovation policies focus on western economies.

Our empirical analysis focuses on an inventor-level analysis, as inventors are arguably the cornerstone of innovation (Lucas 1988, Moretti 2021). A second important reason is that a country-level analysis would inherently capture the country’s aggregate trajectory such as the general rise of China in AI. By employing an inventor-level analysis and controlling for inventor fixed effects, we can exploit within-inventor variation to examine individual activity. This approach allows us to better examine how inventors respond to incentives rather than capture national trajectories. Furthermore, this strategy largely addresses concerns about confounding factors at both the country and individual levels, such as cultural norms, race, or innate ability.

To characterize high-risk and potentially high-reward projects, we follow extant innovation literature that novel innovation efforts are risky but have a high potential to achieve breakthrough (March 1991, Ahuja and Lampert 2001, Fleming 2001, Uzzi, Mukherjee, Stringer, and Jones 2013). We follow Kelly, Papanikolaou, Seru, and Taddy (2021) by utilizing textual similarity. Specifically, we employ the similarity of a focal AI patent to its cohort of AI patents (serving as a proxy for a set path for innovation) as a measure of novelty; the lower the focal patent’s similarity to its cohort, the more novel it is considered and the more it breaks from that established path. This captures the insight from Kelly, Papanikolaou, Seru, and Taddy (2021) that novelty is generated through exploratory efforts and represents a “leap” from existing knowledge or path. Subsequent analysis validates this novelty measure by showing that novel patents (specifically, AI patents with zero similar cohorts) are inherently risky. Using forward citations — a commonly used measure for patent quality — we observe that these patents exhibit highly volatile and polarized quality: while they possess greater potential to become breakthrough inventions, they are also significantly more likely to fail, receiving zero future citations. Furthermore, these projects are associated with longer development times.

We employ a stacked difference-in-differences (DID) design to analyze the novel AI patents filed by inventors subject to national AI policies. A key finding, noted above, is that the share and

number of novel patents — those patents break from existing set path — increase among treated inventors following policy implementation. However, short-term productivity measures, including the number of patents, average forward citations, and quality-weighted counts, suggest a decline in innovation output. This pattern is consistent with the “J-curve” dynamics frequently observed in innovation trajectories, where long-term success follows an initial, resource-intensive exploratory period (Brynjolfsson, Rock, and Syverson 2021, Cong, Lu, Shi, and Zhu 2024). We observe a recovery in productivity three years following policy implementation, particularly for breakthrough patents; the trajectory broadly follows a J-curve. It is important to note that this decline in output is a relative one, as estimated within a difference-in-differences (DID) framework. The absolute level of productivity in policy-adopting countries continues to increase, but at a slower rate than in control countries without such policies.

Our findings are robust across numerous specifications, including use of sub-samples that exclude major economies such as the US and China. Furthermore, our validity checks show little evidence of confounding pre-existing trends, suggesting that our results can not be explained by an industry-cycle story — specifically, that governments implement policy in response to an observable pre-existing decline in productivity. For such a confounding factor to explain our results, it would require that the unobserved shock occur simultaneously with the policy; in other words, the government would need to perfectly predict the timing of the productivity decline and implement policies.

We argue this is unlikely for two reasons. First, policy formulation is subject to administrative frictions and significant lags; launching distinct policies at the exact moment of a productivity decline is implausible. Second, qualitative evidence from policy documents do not support this “predictive decline” hypothesis. In general, countries are motivated by the potential of AI and a desire to capitalize on it, rather than by a need to correct a slump. For instance, the US explicitly states in its National AI R&D Strategic Plan that AI is expected to create new markets and improve the efficiency of existing industries, including manufacturing, logistics, and finance; the goal is to invest in growth. Similarly, China’s Next Generation Artificial Intelligence Development Plan characterizes AI as a “new engine of economic development” and stresses seizing first-mover advantages and global leadership, despite explicitly noting China’s already strong AI research and patenting performance. This forward-looking orientation extends beyond the United States and China. Saudi Arabia and the United Arab Emirates position AI as a central pillar of economic

diversification away from petroleum dependence under the Vision 2030 framework, explicitly describing AI as the “new oil”.

We also examine the heterogeneous effects of policies at the individual country level by estimating a separate difference-in-differences regression for each country, utilizing its not-yet-treated/never treated cohorts as the control group. We observe a negative correlation between productivity and novelty: countries that experienced positive productivity gains tended to produce fewer novel inventions, while those with declining productivity generally produced more novel inventions. This result provides strong support for our central argument regarding the trade-off between innovation output and novelty. The difference in outcomes could, in part, be the result of policies that place a relatively greater or lesser emphasis on pursuing novel projects. However, we do not claim this to be the sole explanation for the short-term productivity decline, as we observe that many European countries experienced decreases in both innovation novelty and productivity. We find that declines in European countries could be driven by regulation- and guidance-based policy instruments designed to promote ethical and trustworthy AI adoption rather than to directly constrain AI development. Nevertheless, despite their AI-promoting intent, additional regulation or guidance of standard may constrain innovative flexibility. This observation is in line with recent U.S. deregulation initiatives across various states aimed at fostering AI development.<sup>3</sup>

As our treatment is at the country level, we use US AI patents to provide more micro-level evidence, given the rich data available for US patents. We can identify with certainty which AI patents are funded by the government by examining the “government interests” reported in the patent documents. Our findings confirm that the US policy goal targets high-risk projects, as measured by novel AI patents; furthermore, most of these funded projects do not ultimately result in breakthroughs, suggesting a role for the government in embracing failures by supporting projects that the private sector may be unwilling to undertake.

Finally, we examine the welfare effects by focusing on 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 originates from researchers in other countries post-policy. This suggests that while the absolute volume of spillovers may decrease, policies increase the relative share of cross-border knowledge exchange.

---

<sup>3</sup>See Executive Order 14179 & Executive Order “Ensuring a National Policy Framework for AI”: Federal mandates that explicitly challenge state-level AI regulations as “onerous” barriers that “thwart the imperative” of U.S. innovation.

Our work offers at least three main contributions to the literature. First, to the best of our knowledge, we are the first to study the relationship between AI policies and AI innovation. The emerging literature on AI has focused mainly on how AI affects outcomes, such as firm value and growth (Eisfeldt, Schubert, and Zhang 2023, Babina, Fedyk, He, and Hodson 2024), labour (Acemoglu 2025), asset prices (Kelly, Kuznetsov, Malamud, and Xu 2025), productivity (Jiang, Park, Xiao, and Zhang 2025), and how exposure to automation affects corporate innovation (Cong, Lu, Shi, and Zhu 2024). We depart from these streams of literature by shifting the focus to the determinants of AI innovation. Specifically, we examine the critical role of government policies in accelerating the pace of AI innovation.

Second, we contribute to the broader literature on innovation policies. Most existing research focuses on R&D tax credits and immigration rules (Bloom, Van Reenen, and Williams 2019), while several studies examine grant programs (e.g. Howell 2017); however, this literature largely lacks international evidence. Furthermore, studies on national-level innovation policies that target specific sectors remain sparse. In particular, empirical evidence on “mission-oriented” policies is limited (Bloom, Van Reenen, and Williams 2019). Early evidence regarding mission-driven initiatives, such as the Apollo project and DARPA, suffers from small sample sizes and unclear counterfactuals. We contribute to this scarce literature by providing a large sample and systematic international evidence.

Third, we construct and introduce two new datasets valuable for future research. Using machine learning techniques, we build a global AI patent dataset that extends the USPTO’s AI patent classification system across multiple jurisdictions (Arora, Belenzon, Pataconi, and Suh 2020). Given that AI development is a global phenomenon, this comprehensive dataset is essential for international studies of AI innovation. The second is a refined dataset of national AI policies, built upon the OECD’s AI Policy Observatory. Using large language models with human verification, our curation process addresses several limitations of the OECD’s data and provides a more reliable policy classification.

The remainder of the paper is organized as follows. [Section 2](#) provides a review of the literature on innovation policy to contextualize our contribution. [Section 3](#) discusses the data construction and provides a descriptive analysis of AI innovation and policy. [Section 4](#) outlines the empirical methodology, reports the main results on AI policy and inventor productivity, and validates our novelty measure. [Section 5](#) provides additional results at the aggregate level and investigates the sources of innovation novelty and international knowledge spillover. [Section 6](#) concludes.

[Appendix A](#) provides additional information on innovation data and policy. Lastly, we have furnished an [Internet Appendix B](#) with additional robustness checks and results, and [Appendix C](#) presents a simple general equilibrium framework to rationalize the policy effects.

## 2 Literature Review on Innovation Policies

### 2.1 Why We Need Innovation Policies

Innovation is widely regarded as the primary engine of long-run economic growth ([Schumpeter 1947](#), [Solow 1957](#), [Romer 1990](#), [Kogan, Papanikolaou, Seru, and Stoffman 2017](#), [Acemoglu and Restrepo 2018](#)). However, it is broadly acknowledged that markets systematically underinvest in innovation due to two fundamental failures: knowledge spillovers and financial frictions ([Nelson 1959](#), [Arrow 1962](#), [Hall and Lerner 2010](#), [Howell 2024](#)).

The first failure arises because R&D generates positive externalities that the innovating firm cannot fully appropriate. This wedge between private and social returns is substantial: [Lucking, Bloom, and Van Reenen \(2018\)](#) estimate social returns at roughly 60 % compared to private returns of 15 %, a magnitude supported by earlier studies ([Hall 1993a](#), [Adams and Griliches 1998](#)). This divergence creates a clear rationale for public intervention to align private incentives with social welfare.

The second failure stems from capital market imperfections. Innovation projects are inherently risky, intangible, and opaque, making them difficult to finance ([Hall and Lerner 2010](#), [Howell 2017](#)). This is because entrepreneurs cannot fully disclose ideas without risking appropriation, and because deep-tech projects often lack collateral, external financiers demand high premiums or verifiable short-term outcomes ([Lerner 2009](#), [Nanda and Nicholas 2014](#)). This bias against exploration is particularly severe for young firms and radical, long-horizon technologies ([Manso 2011](#), [Haltiwanger, Jarmin, Kulick, and Miranda 2016](#), [Acemoglu and Restrepo 2018](#)).

These failures are most acute in “mission-critical” research, such as Artificial Intelligence (AI). As emphasized by [Mazzucato \(2013, 2021\)](#), such activities require patient, failure-tolerant capital that private markets rarely provide. Consequently, the role of government is not merely to increase the volume of R&D, but to correct the systematic bias against high-risk, socially valuable exploration ([Bloom, Van Reenen, and Williams 2019](#), [Howell 2024](#)).

## 2.2 Non-targeted vs. Targeted Policy

To address these failures, advanced economies have historically relied on broad-based fiscal incentives, particularly R&D tax credits. These “non-targeted” instruments reduce the marginal cost of R&D for all eligible firms, leaving specific allocation decisions to the market (Hall 1993b, Bloom, Griffith, and Van Reenen 2002). The empirical evidence suggests these tools are effective at raising aggregate spending. For instance, Bloom, Griffith, and Van Reenen (2002) find that a 10 % reduction in the tax price of R&D increases spending by at least 10 % in the long run. Furthermore, tax credits generate positive spillovers for rivals and suppliers (Bloom, Schankerman, and Van Reenen 2013, Babina and Howell 2024).

However, innovation policies face an inherent “targeting tension” (Howell 2024). While tax credits avoid the risk of governments “picking winners”, they tend to subsidize incremental R&D that firms would often conduct anyway, failing to incentivize the highest-risk, highest-spillover projects (David, Hall, and Toole 2000). In contrast, “targeted” instruments—such as direct grants, procurement, and mission-oriented strategies—can channel resources toward specific technological frontiers (Azoulay, Fuchs, Goldstein, and Kearney 2019, Hausman 2022).

Recent literature highlights the unique value of these targeted interventions. For instance, the US “DARPA (originally ARPA) model” for funding high-risk, high-reward, breakthrough R&D, characterized by vision-driven programs, has proven successful by combining autonomy with a high tolerance for failure, enabling the pursuit of transformative breakthroughs (Azoulay, Fuchs, Goldstein, and Kearney 2019, Gans 2019). Similarly, U.S. “Small Business Innovation Research” (SBIR) grants have been shown to catalyse growth and patenting in young firms (Howell 2017). These targeted policies are premised on the belief that the government has a distinct role in supporting open-ended research that private actors deem too risky (Akcigit, Hanley, and Stantcheva 2022).

In our study, we focus on the surge in “mission-oriented” policies dealing with AI. As highlighted in the OECD AI Policy Observatory<sup>4</sup> in 2025 and other related OECD studies (i.e. Larrue 2021, Galindo, Perset, and Sheeka 2021), a large number of governments have launched comprehensive strategies to lead in AI. Despite their prevalence, mission-oriented policies remain under-studied compared to tax credits because they are heterogeneous and lack obvious counterfactuals (Bloom, Van Reenen, and Williams 2019). Early evidence on mission-driven initiatives suffers from small

---

<sup>4</sup>See, <https://oecd.ai/>

sample sizes and unclear counterfactual. This study is the first to examine in detail the relationship between AI policies and AI innovation; it also provides, for the first time, a large and systematic sample of international evidence on the impact of innovation policies on scientists' productivity in advancing AI technology.

Since "mission-oriented" policies are anticipated to encourage a pivot toward radical innovation, one should expect a change in the distribution of innovation outcomes: a higher variance (more failures and more breakthroughs) and potentially longer development timelines. Consequently, a policy that successfully induces risk-taking may mechanically result in a short-run decline in average inventor productivity (measured by patent counts), as innovators shift focus from "low-hanging fruit" to high-risk and high-reward projects. As we discuss, our evidence on the nature of innovations being conducted and the short-term decline in innovation is consistent with a number of past and current studies (Brynjolfsson, Rock, and Syverson 2021, Griliches 2009, Jones 2009) that demonstrate such a J-curve process. In addition, a related recent study by Cong, Lu, Shi, and Zhu (2024) finds that innovation shifts prompted by automation exhibit a J-curve, which can explain part of the puzzle of declining research productivity.

### 3 Data

#### 3.1 Constructing Global AI Patent Database

To construct our database, we begin by developing a methodology to identify AI-related patent filings globally. We focus on US AI patent filings using the Artificial Intelligence Patent Dataset (AIPD),<sup>5</sup> as provided by the USPTO (Giczy, Pairolo, and Toole 2022).<sup>6</sup> We classify AI patents in accordance with the USPTO's guidelines.<sup>7</sup>

Next, we utilize US AI patents from the AIPD to identify international AI patent applications using semantic text vectors and a machine learning technique called Support Vector Machine (SVM). We use text vectors<sup>8</sup> provided by Google Patents to train an SVM using labeled AI patents

---

<sup>5</sup>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 pertain to 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.

<sup>6</sup>Note that this paper is written by Chief Economists at the USPTO.

<sup>7</sup>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 classified as not 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.

<sup>8</sup>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.

from the USPTO. The trained SVM model is then used to identify AI patents in other patent offices.

It is important to note that this classification is not US-specific because inventors around the globe file patents in the USPTO. In recent years, for instance, more than half of the inventors filing in the US were domiciled in other countries, implying that the classifier is capable of detecting AI inventions globally for our purpose.

We proceed by categorizing 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, over the 2006 to 2019 sample period, we identified 1,568,818 AI patent filings with unique family IDs, including 455,461 filed by inventors from the United States.

### 3.1.1 Descriptive Statistics of AI Patent Applications

A detailed country-level breakdown of AI patent applications, classified by the lead inventor's address over the period 2006–2019, is presented in [Table B1](#). As indicated, United States, Japan, China, and South Korea are among the most prolific over this period, with approximately 455,461, 370,089, 302,758, and 152,437 applications, respectively. [Figure 1](#) illustrates the time series variation in AI patent applications for 42 countries from 1990 to 2022. The figure highlights the divergent trend between China and United States, with China surpassing United States in 2018 in the number of AI patent applications, accounting for approximately 50% of global applications in 2019, compared to 20% for the United States. Historically, Japan was the leading country in AI patent applications for much of the twentieth century, but experienced a decline from 1990 to 2022, with United States surpassing it around 2002.

The domain-specific distribution of AI patent applications, categorized into eight classifications by the AI Patent Database (AIPD), is provided in [Table B1](#). As indicated, the bulk of patent filings from United States (32.55%), China (30.70%), South Korea (27.50%), Germany (37.86%) and United Kingdom (36.03%) fall under “Planning and Control.”<sup>9</sup> By contrast, Japan, Taiwan and Netherlands primarily focus on “Vision” technologies, which account for 34%, 35% and 30% of their respective applications.

---

<sup>9</sup>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.

### 3.2 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 regarded as the “cornerstone of innovation”.<sup>10</sup> This focus aligns with the nature of the policies studied that typically support AI research and, thereby, directly impact AI researchers. We use the dataset of identified global AI patent applications 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 a citation-weighted count that is the sum of one plus forward citations<sup>11</sup> ([Trajtenberg 1990](#)) received by inventor  $i$ 's patents filed in year  $t$ . Citations are counted up to the end of 2024.

Patent citation data are subject to truncation issues that are a well recognized 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 accumulate forward citations. The presumption is that 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 reported here is smaller than that in our regression tables since we use “stacked” regressions, as detailed in the methodology section. Variable definitions are provided in [Table A1](#).

### 3.3 National AI-Related Strategies and Policies Data

Our primary data source for analyzing national AI strategies and policies is the OECD AI Policy Database.<sup>12</sup> This comprehensive dataset provides detailed information on AI-related policies across countries, including average annual budgets, policy types, and implementation periods.

---

<sup>10</sup>AI innovation is no exception. 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.

<sup>11</sup>Forward citations refer to the number of times a focal patent is cited by subsequent patents.

<sup>12</sup>See, <https://oecd.ai>

The original dataset contains over 1,000 unique policies from 69 countries, spanning the period from 1969 to 2024. Given that some entries are general innovation policies rather than specifically AI-focused, we utilized ChatGPT-4o (see [Table A2](#) for the exact prompt) to analyze each policy description and identify those directly related to AI. To ensure robust classification, each description was evaluated ten times; we retained only those policies confirmed as both AI-related and national in scope in at least seven out of ten evaluations. Notably, all policies in our final sample are AI-promoting in nature, though some incorporate regulatory instruments that provide guidelines on data access and ethics. These instruments aim to facilitate AI development while ensuring it proceeds in a safe and ethical manner.

This process yielded a refined dataset of 120 AI-specific policies from a sample of 42 countries and regions, spanning 2011 to 2019. 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.”

### 3.3.1 Description of AI Policies

Governments worldwide have, particularly in the 21st century, experienced a surge in the launch of national AI strategies and related initiatives. Our analysis includes 188 national AI-related policies from a sample of 42 countries over the period 2006-2019. These policies are national-level, targeted—often described as “mission-oriented” ([Larrue 2021](#))—interventions that aim to support high-risk, potentially high-reward AI research. [Figure 2](#) presents a word cloud generated from policy descriptions obtained from the OECD AI Policy Database,<sup>13</sup> highlighting key terms such as “artificial intelligence”, “development”, “research”, “government”, and “national”. The prevalence of these terms suggest policies in our data are R&D policies led by government and especially national government of AI policy.

Additional quantitative information on these policies is reported in [Table A3](#), summarizing average annual budgets. Notably, the four countries with the largest public AI investment budgets are Saudi Arabia, the United Kingdom, China, and the United States, with average annual

---

<sup>13</sup><https://oecd.ai>

allocations of approximately USD 6,000 million, USD 4,165 million, USD 2,640 million, and USD 1,678 million, respectively. Importantly, the majority of countries in our sample explicitly target high-risk, potentially high-reward AI research. Table A4 provides illustrative examples of major national AI initiatives, particularly from countries with substantial public investment. These examples highlight the mission-oriented nature of AI policy design, which emphasizes exploratory projects characterized by high technological uncertainty, long development horizons, and uncertain payoffs. For instance, China’s 2017 National New Generation Artificial Intelligence Development Plan<sup>14</sup> explicitly calls for “exploratory research where there is no consensus”, while the U.S. National Artificial Intelligence Research and Development Strategic Plan<sup>15</sup> emphasizes support for “long-term, high-risk, high-reward research”. Together, these policy statements illustrate a shared objective of fostering transformative, rather than incremental, innovation in AI.

Beyond their emphasis on high-risk, potentially high-reward research, these policies often involve a diverse range of policy instrument and target entities. They engage both the public and private sectors, distinguishing them from the more narrowly studied instruments such as R&D tax incentives and conventional grant programs. Table A5 summarizes the policy instruments and target entities associated with national AI initiatives across all countries in the OECD AI Policy Database. Overall, these policies support AI development through methods like project grants for public research, fellowships and postgraduate loans and scholarships, the creation of networking and collaborative platforms, and investments in AI computing and research infrastructure. In terms of target entities, national AI policies predominantly focus on actors in the public research system—such as researchers in public research institutes and higher education institutions—while also extending support to private-sector participants ranging from large firms to startups. Specifically, the United States and the United Kingdom place greater weight on project grants for public research and on fellowships and postgraduate funding directed at researchers in public research institutes. Japan and Germany direct greater policy effort toward networking and collaborative platforms and investments in AI computing and research infrastructure, across public and private research institutions.

Thirdly, we observe that AI policies are often continued following the initial implementation. For example, Figure 3 outlines the evolution of major U.S. AI-related policies. the United States

---

<sup>14</sup>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/>

<sup>15</sup>See National Artificial Intelligence Research and Development Strategic Plan (2016) available at [https://www.nitrd.gov/pubs/national\\_ai\\_rd\\_strategic\\_plan.pdf](https://www.nitrd.gov/pubs/national_ai_rd_strategic_plan.pdf)

launched the National Robotics Initiative in 2011,<sup>16</sup> and subsequently introduced a sequence of related initiatives, including the BRAIN Initiative in 2013, nanotechnology initiatives in 2014, strategic computing, and big data research programs in 2016. This policy trajectory continued with the release of the National AI R&D Strategic Plan in 2016 and extends to more recent efforts such as the 2025 Stargate project. Similarly, China introduced the “Internet+” Artificial Intelligence Three-Year Action and Implementation Plan in 2016, followed by the “New Generation Artificial Intelligence Development Plan” in 2017, and subsequently rolled out multiple complementary sub-policies under this broader strategic framework. These initiatives are characterized by long-term policy objectives, including explicit goals of achieving global leadership in AI by 2030.

In general, governments are motivated by the potential of AI and a desire to capitalize on it, rather than by a need to correct a productivity slump. For instance, the United States states in its National AI R&D Strategic Plan that AI is expected to create new markets and enhance the efficiency of existing industries—including manufacturing, logistics, and finance—signaling a growth-oriented investment rationale. Similarly, China’s Next Generation Artificial Intelligence Development Plan characterizes AI as a “new engine of economic development” and stresses seizing first-mover advantages and global leadership, despite explicitly noting China’s already strong AI research and patenting performance. This forward-looking orientation extends beyond the United States and China. The United Kingdom’s UK National AI Strategy, Russia’s National Strategy for the Development of Artificial Intelligence, and South Korea’s National Strategy for Artificial Intelligence similarly frame AI as a strategic engine for new industry creation, global competitiveness, and long-run economic transformation. Taken together, these policy narratives indicate that national AI initiatives are designed not to stabilize declining productivity, but to pursue high-risk, high-reward technological leadership through market-creating innovation.

A primary objective of our paper is to evaluate the impact of these national AI policies. Consequently, our analysis focuses on estimating the *Average Treatment Effects (ATE)* across the sample to develop a broad understanding of their overall influence. However, it is evident that there are substantial differences in scope, design, and implementation of these policies across countries.

---

<sup>16</sup>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 recognized as one of the earliest strategic R&D initiatives relevant to AI. See National Artificial Intelligence Research and Development Strategic Plan (2016), available at [https://www.nitrd.gov/pubs/national\\_ai\\_rd\\_strategic\\_plan.pdf](https://www.nitrd.gov/pubs/national_ai_rd_strategic_plan.pdf); relevant discussion appears on pages 6–7. We also document a significant increase in AI-specific research funding provided by the National Science Foundation (NSF) after the NRI’s launch. [Figure A1](#) shows that the proportion of NSF awards dedicated to AI research rose from approximately 7% in 2010 to around 10% in 2012.

In later sections, we complement the aggregate analysis by examining the effect of policies at the individual country level, where feasible.

## 4 Empirical Method and Baseline Results

### 4.1 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](#)):

$$Y_{ijct} = \beta_1 AI Policy_{jct} + \phi_{ic} + \lambda_{tc} + \epsilon_{icjt} \quad (4.1)$$

In this model,  $Y_{ijct}$  represents one of several dependent variables of interest for inventor  $i$  of country  $j$  and cohort  $c$  in year  $t$ . We stack  $t - 5$  to  $t + 8$  for each cohort, where 8 years is the maximum allowed in our data. The treatment indicator variable  $AI Policy$  equals 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 continued support for AI and enact follow-on policies.

In the stacked difference-in-differences (DID) approach, each treated country is paired with control countries that are either “not-yet-treated” 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-by-cohort fixed effects ( $\phi_{ic}$ ) to absorb all time-invariant unobserved heterogeneity at the individual level, such as gender, race, and innate ability. Given the low rate of cross-border migration in our sample, these fixed effects subsume country fixed effects, and it controls for time-invariant national characteristics such as culture and initial levels of economic

development. We also include cohort-by-year fixed effects ( $\lambda_{tc}$ ) to account for common time shocks specific to each stacked cohort. Standard errors are clustered at the country level, following [Bertrand, Duflo, and Mullainathan \(2004\)](#). This level of clustering is essential because, while the unit of analysis is at the inventor-year level, treatment is assigned at the country level. Clustering ensures that our inference accounts for within-country serial correlation and that test statistics are computed based on the number of countries rather than the raw number of observations. Finally,  $\beta_1$  serves as the difference-in-differences estimator; note that these fixed effects absorb the standalone effects for the treatment group and the post-period indicator.

We also estimate [Callaway and Sant’Anna \(2021\)](#) difference-in-differences (DID) regressions for robustness, which represent a structural approach to addressing bias in TWFE as an alternative to pairing treated units with clean controls. We utilize stacked DID as our main specification due to its flexibility. The [Callaway and Sant’Anna \(2021\)](#) method requires inverse probability weighting implemented via ordinary least squares (OLS); however, the nature of our innovation data often necessitates a count data model. Consequently, we estimate the [Callaway and Sant’Anna \(2021\)](#) model using the logarithm of the count-based dependent variable. While the results are consistent with our main findings, their interpretation requires certain caveats due to potential biases arising from the log transformation of the dependent variable ([Cohn, Liu, and Wardlaw 2022](#)).

## 4.2 Baseline Results

### 4.2.1 AI policies, Patent Novelty and Inventor Productivity

As discussed earlier, “mission-oriented” AI policies are often designed to foster high-risk, high-cost innovations that the private sector would not otherwise undertake ([Mazzucato and Semieniuk 2018](#)). To investigate whether this goal is achieved and to characterize these high-risk, high-reward projects, we draw on the extensive innovation literature, which posits that projects with higher impacts are those novel, exploratory efforts that deviate from an existing path, often described as path-breaking” ([March 1991](#), [Ahuja and Lampert 2001](#), [Fleming 2001](#), [Uzzi, Mukherjee, Stringer, and Jones 2013](#)). These projects represent a “leap” from existing knowledge or trajectories. We follow recent developments in natural language processing to measure innovation steps, specifically the approach of [Kelly, Papanikolaou, Seru, and Taddy \(2021\)](#), who posit that focal patents that deviate from past patents—exhibiting low textual similarity—tend to represent exploratory projects. In a similar spirit, we argue that if a patent deviates (is textually dissimilar) from a cohort of patents representing the normal innovation path, it represents a novel path-breaking effort. We conduct

robustness tests using only past patents and find similar results in [Table B2](#). However, comparing with only past patents may capture technological field specific trends; for instance, in a fast moving AI field, patents may exhibit low similarity from the past without being genuinely path-deviating in their own right. Finally, it is important to note that these patents may not succeed ex-post; in fact, the innovation literature generally finds that radical attempts tend to be trial-and-error efforts that mostly fail (e.g. [March 1991](#), [Fleming 2001](#)).

To assess the novelty of an AI patent, we count the number of patents filed in the same year (the cohort) that are highly similar to the focal patent, as defined by the Google Patent Research Database using their text embeddings.<sup>17</sup> The lower this count, the more novel the patent is considered, with a value of zero representing the highest degree of novelty in our data. In this framework, other patents filed at a similar time serve as a proxy for the existing innovation path; those without similar competing patents (henceforth novel patents) are those that deviate from this set path. Accordingly, we construct two inventor-level measures: the novel patent share (number of novel patents filed over total number of patents filed) and the novel patent count (the absolute number of novel patents filed).

Using these two measures, we find a shift among treated inventors towards filing novel patents. In [Table 2](#), we regress the share of novel patents and novel patent counts for inventor  $i$  in year  $t$  on the indicator of AI policy implementation. In [Table 2](#), we regress the share of novel patents and the number of novel patents for inventor  $i$  in year  $t$  on an indicator for AI policy implementation. Following policy implementation, the share of novel patents increases by 2.8 percentage points, corresponding to an approximately 11% increase relative to the sample mean reported in Column 1 of Panel A. Column 2 of Panel A further indicates that the number of novel AI patents filed after the policy increases by 3.4%. While these magnitudes may seem modest, they should be interpreted in light of the inherent difficulty of path-breaking innovation. Because generating and realizing genuinely novel ideas is a high-risk, low-probability endeavor, those efforts that eventually become patentable may represent only a small proportion of total exploratory attempts. This mirrors the experience of academic research: just as not every ambitious project results in a published paper, not every exploratory effort in AI leads to a successful patent.

In Panel B, we examine total patent counts and patent qualities as measures for inventor productivity. Column 1 indicates an approximately 9% ( $e^{-0.095} - 1$ ) decrease in total patent counts. This stands in contrast to Panel A, Column 2; where total patent counts decrease, however, novel

---

<sup>17</sup>These embeddings are designed to capture the semantic meaning of the text, not only textual similarity.

patents increase both in absolute terms and as a share of the total (Panel A, Column 1), suggesting a shift focus toward novel patents. In Column 2, we use citation-weighted counts ( $= \sum_i^P 1 + C_i$ ) (Trajtenberg 1990) to examine quality-weighted metrics and find a 37.5% ( $e^{-0.471} - 1$ ) decline. This is not purely driven by the decrease in raw counts, as the average forward citations also decrease by 30% ( $e^{-0.355} - 1$ , in column 3). These results suggest that despite the shift toward novelty, even those efforts that eventually prove patentable are not on average high quality. In subsequent sections, we show that novel innovations exhibit polarized quality: while they are more likely to achieve breakthroughs, they are simultaneously more likely to fail. This is not surprising, as achieving a breakthrough requires a process of trial and error and experimentation that embraces failure (Tian and Wang 2014, Manso 2011, Ahuja and Lampert 2001). This phenomenon is especially pronounced in the short term, as many studies demonstrate J-curved dynamics in AI adoption and innovation Brynjolfsson, Rock, and Syverson (2021), Cong, Lu, Shi, and Zhu (2024).

It is important to interpret the results 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 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 B1, 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 than for those in the control countries.

One might hypothesize that the decline in productivity is driven simply by policy ineffectiveness. However, if the policy were simply ineffective, one would expect a null effect rather than the *decline* observed in our results. Alternatively, one could argue that the decline stems from policy-induced moral hazard, in which an influx of funding leads to reduced effort. Under such a scenario, one would expect to observe a drop in total output alongside a shift toward “tick-the-box,” incremental innovations (which are less novel) as inventors attempt to maintain their professional standing with minimal risk. Our findings contradict this hypothesis: novelty and exploratory work have actually increased.

## 4.3 Robustness checks

### 4.3.1 Parallel Trends

The key identifying assumption in 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, we plot the coefficients on the interaction of the treatment indicator with time periods from  $\leq t - 3$  to  $\geq t + 3$  in [Figure 4](#).

Since the DID regression satisfies parallel trends, our results cannot be explained by an industry-cycle story — specifically, that governments implement policy in response to an observable pre-existing decline in productivity. For such a confounding factor to explain our results, the unobserved shock would need to occur simultaneously with the policy; in other words, the government would need to perfectly predict the timing of the productivity decline and implement policies.

We argue that this is unlikely for two reasons. First, policy formulation is subject to administrative frictions and significant lags; launching distinct policies at the exact moment of a productivity decline is impractical. Second, qualitative evidence from policy documents does not support the "predictive decline" hypothesis. In general, countries are motivated by the potential of AI and a desire to capitalize on it, rather than by a need to correct a slump. For instance, the US explicitly states in its National AI R&D Strategic Plan that AI is expected to create new markets and improve the efficiency of existing industries, including manufacturing, logistics, and finance; the goal is to invest in growth. Similarly, China's New Generation AI Plan views AI as a "leapfrog technology" that can help the nation bypass traditional stages of industrial development and overtake Western competitors. Saudi Arabia and the UAE likewise view AI as the primary vehicle for diversifying their economies away from petroleum dependence under the Vision 2030 framework, operating on the stated premise that AI represents the "new oil."

We conduct a comprehensive set of robustness checks. First, our results are robust to excluding specific countries, such as the US, China, or both in [Table B3](#). Second, in [Table B4](#), 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. Third, in [Table B5](#), 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 B6](#), we estimate an OLS regression using the AI policy indicator, with the dependent variable being the log-transformed count of AI patents. Across all specifications, the results remain qualitatively consistent with our baseline findings.

### **4.3.2 Using Callaway and Sant'Anna (2021) DID model**

In [Table B7](#), we use the methodology from ([Callaway and Sant'Anna 2021](#)) (also known as CSDID) to address potential biases inherent in standard TWFE models. As discussed earlier,

CSDID requires OLS estimation; consequently, we cannot estimate a count model for most of our results. Instead, we use log-transformed count variables, which may introduce potential biases [Cohn, Liu, and Wardlaw \(2022\)](#). Nevertheless, the results are qualitatively consistent with our primary findings. Across all columns, the introduction of AI policies is associated with increased novelty and a decline in both the quality and quantity of AI patents.

### **4.3.3 Placebo tests 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, thereby biasing our results. To lend further confidence to the link between AI policy and AI innovation output, we conduct a placebo test using non-AI patents as the outcome variable. The premise is that if our results were spurious, for example, if AI policies were merely reflecting a national-level trajectory and correlated with innovation, we should observe similar results for non-AI patents. If national AI policies primarily affect AI innovation while leaving non-AI patenting behavior relatively unchanged, this result lends greater confidence to our main findings.

The results are presented in [Table B8](#), which reports the policy effects on non-AI patent counts, citation-weighted counts, and average forward citations. We find that most of these results yield negative coefficients that are much smaller in magnitude than those for AI patents and are also statistically insignificant. Some negative coefficients are expected because a productivity decline in AI could spillover to non-AI as AI innovations are general purpose technologies that spur innovation and applications in other fields ([Babina, Fedyk, He, and Hodson 2024](#), [Agrawal, Gans, and Goldfarb 2022](#), [Crafts 2021](#), [Cong, Lu, Shi, and Zhu 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.

### **4.3.4 Placebo tests using Simulation**

One concern is that our results may arise randomly rather than reflecting a genuine statistically significant relationship. To address this concern, we assess the probability of observing the same productivity decline when the timing of treatment across countries is randomly assigned.

We conduct placebo simulations using the same stacked difference-in-differences (DID) specification. 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. As shown in [Figure B2](#), the simulated distribution centers around zero, whereas the actual estimate lies far in the left tail. Only about 3% of the simulations yield a negative and statistically significant effect on productivity comparable to our baseline results, indicating that the observed findings are highly unlikely to have occurred by chance.

### 4.3.5 Disclosure Story

The observed decline in AI patents may reflect a disclosure story rather than a genuine loss of productivity. As AI becomes strategically vital, firms and governments may choose to withhold or delay patenting — which discloses technological details to the public — particularly for sensitive technologies to protect national security. In this case, national AI policies would incentivize secrecy and reduce observable patent counts without an actual decline in innovation.

To test this mechanism, we examine AI-military patents, which are most susceptible to non-disclosure for security reasons. We identify these military patents based on CPC codes spanning weapons and ammunition (F41–F42), military transport (B63G, B64C–D), sensors and target tracking (G01S, G05D), and secure communications (H01Q, H04B–W).<sup>18</sup>

The results reported in [Table B9](#) do not support the disclosure hypothesis. We find no significant change in either the volume or quality of AI-military patent filings following the implementation of national AI policies, suggesting that strategic withholding is unlikely to be the main driver of our results.

## 4.4 Long-term Effects

Our main premise is that a focus on technological novelty induces a transitory decline in measured productivity. This “short-term pain” is a standard feature of the innovation trajectory when economy transits into new technological spaces. In the long run, this process is expected to follow the “Productivity J-curve” ([Brynjolfsson, Rock, and Syverson 2021](#)).

To test this, we examine long-term effects by extending our DID estimator to  $t + 8$ , the maximum horizon permitted by our data. In [Figure 5](#), we observe a systematic recovery in productivity across AI-related patent metrics. Most notably, breakthrough patents (those in the top 1% of the citation

---

<sup>18</sup>Detailed CPC classifications include: firearms and missiles (F41–F42); warships, military aircraft, and specialized vehicles (B); radar, navigation, and simulators (G); and electronic warfare or cybersecurity (H).

distribution) exhibit the most pronounced recovery, becoming positive and statistically significant by  $t + 8$ . Meanwhile, patent counts remain negative and are the least recovered. This dynamics suggests that inventors undergo an initial exploratory period (March 1991) characterized by high-variance experimentation and a high failure rate. Once a “good path” is identified, the R&D process shifts into a “deep dive” phase where quality outweighs quantity; in this stage, resources are concentrated on refining the most promising trajectories to achieve breakthroughs (Weitzman 1979).

However, an 8-year horizon may still be insufficient to capture the full realization of such shifts. For instance, U.S. National Robotics Initiative started in 2011, yet at the time of writing (2025), we yet to see major commercial success of robots. As non-truncated innovation and commercialization data become available, future research may confirm that this J-curve trajectory eventually yields even more substantial productivity boosts over longer horizon.

#### 4.5 Novel Innovation is high-risk and high-reward

Our measure of novelty is new to the current literature, as we leverage recent methodological advancements in Natural Language Processing (NLP) and Large Language Models (LLM) alongside recent developments in the economics of innovation (Kelly, Papanikolaou, Seru, and Taddy 2021). A potential concern is whether our measure of novelty truly captures high-risk, high-reward projects aimed at achieving breakthroughs. In this section, we address this directly.

We conduct a patent-level analysis to assess the risk–return trade-off associated with novel inventions. To do so, we define three measures. We capture the risk of failure by identifying “low-impact” patents as those receiving zero forward citations. We proxy for development costs 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.<sup>19</sup> On the return side, we measure patent quality using citations (Trajtenberg 1990, Hall, Jaffe, and Trajtenberg 2005). While other popular alternatives include the market-value-based measures from Kogan, Papanikolaou, Seru, and Stoffman (2017), these are restricted to US public firms; we cannot use them because our AI patent dataset includes both foreign entities and non-public firms. Furthermore, Kelly, Papanikolaou, Seru, and Taddy (2021) also define breakthroughs using textual similarity, but we

---

<sup>19</sup>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.

refrain from using that approach because our own novelty measures could mechanically correlate with it. Instead, we rely exclusively on citations, which is a robust and widely utilized measure of impact.

Panel A of [Table 3](#) shows that novel patents exhibit a polarized quality distribution. They are approximately 30% more likely to be low-impact (i.e., receiving zero citations) and are associated with 5% longer development times. However, as Column 1 of Panel B shows, novel patents are of higher quality on average—exhibiting modest a 1.3% higher citation count. This occurs despite their higher probability of failure because they are also disproportionately likely to achieve breakthrough (i.e., extremely high-impact) innovations. We construct a series of binary indicators for high-impact patents based on increasingly stringent citation thresholds (e.g., top 10%, top 1%, top 0.5%, and top 0.05% and top 0.01%). Panel B of [Table 3](#) 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, novel patents are 1.21 percentage points (a 12.1% increase relative to the sample mean) more likely to rank among the top 10% of high-impact AI patents, and are 0.098 percentage points more likely to fall within the top 0.1%, corresponding to a 98% increase relative to the sample mean.

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, demonstrating that novel patents do indeed exhibit a polarized quality distribution. As shown in [Table 4](#), 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 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 that of Group 5 (which includes patents with more than three competing prior-year filings). This suggests that more novel patents are associated with greater uncertainty about their future impact, highlighting the inherent risk of pursuing novelty. Panel B confirms these differences statistically through F-tests, all of which indicate that Group 1 has significantly higher variance than other groups, a result that is statistically significant at the 1% level.

Taken together, these findings suggest that while novel AI patents are more likely to fail—often becoming low-impact patents with no forward citations—and require longer development periods, they are also disproportionately likely to become radical innovations. These results validate that our novelty measure successfully captures high-risk and high-return projects.

## 4.6 Individual Country Results

As discussed earlier, countries implement policies using different sets of tools, and their scope can vary quite substantially. Therefore, we examine the effects of individual countries' policies. We follow baseline but executed separately for each country, with each treated country paired with a control group of never-treated and not-yet-treated countries, we plot the coefficients for novel patent shares and total patent counts. These results are illustrated in [Figure 6](#), with the corresponding tabulated results provided in [Table B10](#). The x-axis represents the coefficients for a country's patent count, while the y-axis represents the coefficients for a country's novel patent share. The results are noteworthy: the correlation between these two sets of coefficients is -41%. This suggests that countries exhibiting a higher share of novel patents post-policy are more likely to experience a decline in total patent output, lending support to our hypothesis that the shift toward novel innovation contributes to the observed decline in productivity. However, we do not claim that this is the exclusive explanation. Notably, certain European countries (in the bottom left of the graph) experience declines in both novelty and patent counts. We explore this in the subsequent section.

## 4.7 Heterogeneous Effects

### 4.7.1 Policy Types

As discussed earlier, countries used a diverse set of tools (see [Table A5](#)) to promote innovation, and there may be large, heterogeneous effects across different policy types. We examine this heterogeneity along several dimensions. First, we follow the OECD classifications<sup>20</sup> to examine different modes of support, with results reported in [Table 5](#).

We find that governance-related policies (those providing strategic direction and oversight) have little effect on the outcomes examined. Conversely, financial support (direct investment in the form of grants or firm-level capital) and AI enablers (infrastructure and ecosystem development) yield results in line with our baseline: increased novelty alongside declining patent counts. However,

---

<sup>20</sup>There are four categories: governance, financial support, guidance & regulation, and AI enablers & other incentives. Each of which contains several subtypes. 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.

guidance and regulation policies (those establishing legal and ethical standards) exert negative effects on both novelty and total patent output. It is important to note that the guidance- and regulation-based policies in our sample are intended to promote ethical and trustworthy AI adoption and development, rather than to directly constrain AI development.

Instead, we suspect this may be due to the fact that excessive standardized guidance — even when intended to promote AI — may render AI development more difficult and “handcuff” exploratory freedom, leading to undesired outcomes. This view is consistent with recent U.S. policy shifts towards deregulating<sup>21</sup> the AI field to achieve technological success. However, further research is required to investigate how AI regulations specifically affect outcomes and to test the related mechanisms. We observe that these regulations are mostly enacted in European countries, which may explain the country-level results discussed in the previous section, where several European nations experienced declines in both novelty and patent output.

#### **4.7.2 Target Entities**

Using available OECD data, we classify targeted entities into specific size buckets for private sector investments: small firms (<10 employees), medium firms (10–249 employees), and large firms (250+ employees). Additionally, we examine policies targeting public sector: individual researchers, students, and teachers<sup>22</sup> and research and education organizations.<sup>23</sup> We study the heterogeneous effects of these policies to address two questions: 1) whether investment in the public versus private sector yields different outcomes, and 2) whether supporting different types of entities yields different results. Existing literature suggests that firms vary considerably by size in their innovation strategies; while small firms are often the primary source of exploratory venturing efforts, larger firms may face significant organizational inflexibility and regulatory frictions that hinder breakthrough research (Aghion, Bergeaud, and Van Reenen 2023).

---

<sup>21</sup>For example, the 2025 shift in U.S. AI policy represents a decisive pivot toward an “innovation-first” framework that prioritizes American technological dominance by dismantling centralized federal oversight and preempting restrictive state-level laws. By rescinding the Biden-era reporting requirements for powerful foundation models through Executive Order 14179 and establishing a DOJ Litigation Task Force to challenge state regulations, the administration has replaced a risk-mitigation approach with one centered on unfettered development. This strategy is further solidified by the “America’s AI Action Plan,” which fast-tracks infrastructure build-outs for data centers and mandates that federal agencies purge “ideological” or “non-technical” criteria from their standards. Collectively, these actions aim to lower compliance costs for developers, foster an open-source ecosystem, and ensure that a “patchwork” of state regulations does not slow the rapid deployment of American AI technology.

<sup>22</sup>This category includes established researchers, PhD students, teachers, postdocs, other early-career researchers, as well as secondary education, undergraduate, and master’s students.

<sup>23</sup>This includes public research institutes, higher education institutions.

Table 6 reports these results, which are nuanced. Larger firms experience declines in both novelty and total patent output. This is likely driven by the governance and regulatory mechanisms discussed earlier, as larger firms face higher regulatory compliance costs post-AI policy that may constrain their innovative capacity. Conversely, we find that small firms experience lower novelty but higher patent output. This result is consistent with [Howell, Rathje, Van Reenen, and Wong \(2025\)](#), which shows that policy support for new ventures can have a positive effect on their growth and scaling. However, further research could fully disentangle how AI-specific guidance and regulations affect these firm-level outcomes.

Importantly, we find that support for the public sector—whether directed at research institutes or individual researchers—is associated with an increase in novelty and a decline in patent output, a finding that aligns with our main argument. This result is intuitive: public entities often conduct basic science that is distant from commercialization and thus more long-term in nature. Many such foundational innovations are unlikely to be pursued by the private sector alone, necessitating government intervention. This is particularly vital for the AI field, which remains deeply rooted in basic research. For example, ChatGPT is built on the research in [Radford and Narasimhan \(2018\)](#), and modern Large Language Models (LLMs) are built on the Transformer architecture introduced in the foundational “Attention Is All You Need” a study by [Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser, and Polosukhin \(2017\)](#). In untabulated results, we find that AI patents cite academic papers more frequently than non-AI patents.

#### **4.8 Do governments really target novelty innovation and embrace failure?**

We posit that “mission-oriented” policies specifically promote novel innovation. While our aggregate results show an increase in novelty post-policy, more direct evidence requires observing whether the funded projects are themselves actually more novel. Establishing such micro-level evidence is challenging at a global scale because, first, the exact beneficiaries of these policies are difficult to identify, given the multifaceted nature of AI support (e.g., infrastructure and ecosystem building benefit a large number of projects but are difficult to track individually). The second-best approach is to examine grant-level data; however, this represents only a small subset of policy instruments and is not available in many countries.

To provide at least some micro-level evidence, we focus on US patents, where patent law requires the disclosure of a “government interest” in patent documents, enabling us to identify subsidized

inventions receiving grants. This allows us to observe whether government-funded patents exhibit higher levels of novelty during the post-AI policy period (2011–2019).

In [Table 7](#), column 1, we find that government-funded patents are 17.6% more likely to be novel than non-subsidized patents, representing 62% of the unconditional mean. Although these grants represent only a subset of the policy tools, the result provides empirical support for the hypothesis that mission-oriented funding encourages exploratory research at the technological frontier.<sup>24</sup>

A natural question that follows is how effectively government agencies can identify and target breakthroughs. While agencies can use novelty as a screening criterion, the ex-post quality of innovation is unobservable during project selection. Interestingly, we find that the novel patents selected for US government support are centered in the distribution of quality. The average quality of funded projects is not higher than that of unfunded projects (Column 2). In fact, subsidized projects are both less likely to result in “failures” (Column 3) and less likely to become “breakthroughs” (top 1%, Column 5), clustering instead in the “marginal success” category (top 10%, Column 4).

This distribution contrasts with our earlier finding in [Table 3](#), which shows that novel patents typically have polarized outcomes, characterized by higher probabilities of both failure and breakthrough. This divergence suggests that US grant mechanisms may select a specific profile of novel projects that differs from the broader set of projects targeted by global policies. Nevertheless, the negative coefficient on “breakthrough” status (top 1%) indicates that many funded projects do not reach the extreme upper tail of impact. This result aligns with the theoretical mandate of public R&D: to address market failures by funding projects that the private sector is unwilling to undertake. If government-funded projects were consistently successful, it would suggest that the private sector would likely have funded them anyway. Such a scenario would imply that public funding is crowding out private investment by subsidizing projects that the market would have supported on its own.

Ultimately, our finding confirms that AI support in the form of US grants targets “high-risk” (novel) research, as stated in official policy documents. This micro-level evidence provides support for our premise.

---

<sup>24</sup>While we cannot link every grant to a specific policy, we observe that the reporting agencies are partners of the Networking and Information Technology Research and Development (NITRD) Program, which coordinates the federal AI initiatives central to our study.

## 5 Additional Findings

### 5.1 Aggregate-Level Analysis

An important question is whether micro-level declines in productivity translate into a broader decline in national innovation output. To address this, we examine the impact of national AI policies on country-level aggregate innovation. As discussed earlier, the country level is more sensitive to a country's own trajectory or time-varying country characteristics; nevertheless, it is useful to observe the aggregate results.

The Poisson estimates, in Panel A of [Table B11](#), indicate 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 shows that this decline is not driven by pre-existing downward trends, as there is no evidence of a negative trajectory prior to the policy's implementation. 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 3, this relationship appears to be driven by selection rather than causation. Specifically, the Poisson estimate of  $-0.292$  in Column 3 of Panel A implies a 25.32% ( $e^{-0.292} - 1$ ) reduction in the probability of having an additional AI inventor following policy adoption. Moreover, Column 3 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.

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.

### 5.2 The Pursuit of Novelty via Domain Specialization

Novelty in innovation typically arises in two ways: the recombination of existing knowledge across diverse fields or a "deep dive" into a specific technological domain ([Verhoeven et al. 2016](#)). We investigate how national AI policies influence these paths by examining whether inventors concentrate their efforts on a narrower range of technological domains.

We employ two measures of domain specialization: (1) the number of unique AI patent domains per inventor and (2) a Herfindahl-Hirschman Index (HHI) based on classifications<sup>25</sup> from the AI Patent Database (AIPD). As shown in [Table 8](#), column 1 shows a statistically significant decline in the number of AI patent domains per inventor post-policy. The coefficient of -0.058 represents approximately a 5% reduction relative to the sample mean of 1.19. Correspondingly, column 2 shows a significant increase in the HHI, indicating a higher concentration of inventive activity.

Together, these results suggest that treated inventors have narrowed their technological focus. This shift is consistent with a strategic move towards exploratory innovation within specialized domains rather than cross-field combinations.

### 5.3 Knowledge Spillover and International Collaboration

Standard economic theory suggests that in the absence of market failures, investment decisions are best left to private actors. The primary justification for government intervention is the existence of knowledge spillovers — the “promiscuous” nature of ideas ([Bloom, Van Reenen, and Williams 2019](#)). Because innovators often cannot capture the full social value of their inventions, they may under-invest in R&D. Thus, it is critical to assess how AI policies affect knowledge diffusion.

While our baseline estimates indicate a decline in overall citations (a proxy for spillovers), knowledge diffusion may also manifest through international collaboration. To test this, we examine changes in the share of cross-country backward and forward citations following policy implementation. As reported in [Table 9](#), OLS estimates in columns 1 and 2 show increases of 4.8% and 5.3% in the share of cross-country backward and forward citations, respectively. These findings suggest that while domestic spillovers may have slowed, these policies are associated with a greater reliance on international knowledge and a broader global reach for domestic innovation.

## 6 Conclusion

As the strategic importance of Artificial Intelligence grows, governments have increasingly implemented targeted policies and investments to secure a competitive edge. Utilizing two novel datasets, this paper provides a first look at the effects of these national AI-related policies on innovation. Our findings suggest that these policies are largely achieving their stated objectives

---

<sup>25</sup>There are 8 classifications: Machine Learning, Computer Vision, Natural Language Processing, Speech, Evolutionary Computation, Knowledge Processing, Planning/Control, AI Hardware.

by encouraging a shift toward exploratory and novel innovation. While these activities are characterized as high-risk and high-reward, the transition appears to induce short-term productivity declines in both the quantity and quality of AI patenting. This phenomenon is consistent with the “Productivity J-curve” typically observed during major technological transitions (Brynjolfsson, Rock, and Syverson 2021, Cong, Lu, Shi, and Zhu 2024), where the reorganization of resources and the learning costs associated with new paradigms lead to a temporary dip in output before long-term gains are realized.

We acknowledge, however, that a shift toward novelty may not be the sole explanation for the observed productivity declines. Notably, several European countries in our sample experienced simultaneous declines in both productivity and novelty following policy implementation. Future research is needed to “zoom in” on these countries and explore potential mechanisms. Additionally, as longer time-series data become available, future studies can examine the long-run dynamics of inventor productivity to see whether these short-term declines eventually give way to gains.

Furthermore, the current wave of AI policies utilizes a diverse array of instruments, such as infrastructure development, ecosystem networking, and public-private partnerships, and many of which remain under-studied. Future research could attempt to isolate the effects of these individual policy tools. Finally, while our analysis focuses on innovation metrics, future studies could examine the effects of these national policies on other outcomes of interest, such as productivity, labor market dynamics, and international trade.

## References

- Acemoglu D (2025) The simple macroeconomics of AI. *Economic Policy* 40(121):13–58.
- Acemoglu D, Akcigit U, Alp H, Bloom N, Kerr W (2018) Innovation, Reallocation, and Growth. *American Economic Review* 108(11):3450–91.
- Acemoglu D, Restrepo P (2018) The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review* 108(6):1488–1542.
- Adams JD, Griliches Z (1998) Research productivity in a system of universities. *Annales d'Économie et de Statistique* 49/50:127–162.
- Aghion P, Bergeaud A, Van Reenen J (2023) The Impact of Regulation on Innovation. *American Economic Review* 113(11):2894–2936.
- Aghion P, Howitt P (1992) A model of growth through creative destruction. *Econometrica* 60:323–351.
- Agrawal A, Gans J, Goldfarb A (2022) *Prediction machines, updated and expanded: The simple economics of artificial intelligence* (Harvard Business Press).
- Ahuja G, Lampert CM (2001) Entrepreneurship in the large corporation: A longitudinal study of how established firms create breakthrough inventions. *Strategic Management Journal* 22(6-7):521–543.
- Akcigit U, Hanley D, Stantcheva S (2022) Optimal taxation and R&D policies. *Econometrica* 90(2):645–684.
- Akcigit U, Kerr WR (2018) Growth through Heterogeneous Innovations. *Journal of Political Economy* 126(4):1374–1443.
- Arora A, Belenzon S, Pataconi A, Suh J (2020) The changing structure of American innovation: Some cautionary remarks for economic growth. *Innovation Policy and the Economy* 20(1):39–93.
- Arrow KJ (1962) Economic welfare and the allocation of resources for invention. *Readings in industrial economics: Volume two: Private enterprise and state intervention*, 219–236 (Springer).
- Atkeson A, Burstein A (2019) Aggregate Implications of Innovation Policy. *Journal of Political Economy* 127(6):2625–2683.
- Azoulay P, Fuchs E, Goldstein AP, Kearney M (2019) Funding breakthrough research: promises and challenges of the “ARPA Model”. *Innovation Policy and the Economy* 19(1):69–96.
- Babina T, Fedyk A, He A, Hodson J (2024) Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151:103745.
- Babina T, Howell ST (2024) Entrepreneurial spillovers from corporate R&D. *Journal of Labor Economics* 42(2):469–509.
- Baker AC, Larcker DF, Wang CC (2022) How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics* 144(2):370–395.
- Bertrand M, Duflo E, Mullainathan S (2004) How Much Should We Trust Differences-In-Differences Estimates? *The Quarterly Journal of Economics* 119(1):249–275.
- Bloom N, Griffith R, Van Reenen J (2002) Do R&D tax credits work? Evidence from a panel of countries 1979–1997. *Journal of Public Economics* 85(1):1–31.
- Bloom N, Schankerman M, Van Reenen J (2013) Identifying technology spillovers and product market rivalry. *Econometrica* 81(4):1347–1393.
- Bloom N, Van Reenen J, Williams H (2019) A Toolkit of Policies to Promote Innovation. *Journal of Economic Perspectives* 33(3):163–84.
- Brynjolfsson E, Rock D, Syverson C (2021) The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *American Economic Journal: Macroeconomics* 13(1):333–372.

- Callaway B, Sant'Anna PH (2021) Difference-in-differences with multiple time periods. *Journal of econometrics* 225(2):200–230.
- Canay IA, Santos A, Shaikh AM (2021) The wild bootstrap with a “small” number of “large” clusters. *Review of Economics and Statistics* 103(2):346–363.
- Cengiz D, Dube A, Lindner A, Zipperer B (2019) The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics* 134(3):1405–1454.
- Cohn JB, Liu Z, Wardlaw MI (2022) Count (and count-like) data in finance. *Journal of Financial Economics* 146(2):529–551.
- Cong LW, Lu Y, Shi H, Zhu W (2024) Automation-Induced Innovation Shift. *Working paper* .
- Crafts N (2021) Artificial intelligence as a general-purpose technology: an historical perspective. *Oxford Review of Economic Policy* 37(3):521–536.
- David PA, Hall BH, Toole AA (2000) Is public R&D a complement or substitute for private R&D? A review of the econometric evidence. *Research policy* 29(4-5):497–529.
- Deshpande M, Li Y (2019) Who is screened out? Application costs and the targeting of disability programs. *American Economic Journal: Economic Policy* 11(4):213–248.
- Eisfeldt AL, Schubert G, Zhang MB (2023) Generative AI and firm values. Technical report, National Bureau of Economic Research.
- Fleming L (2001) Recombinant uncertainty: Mating games on technological landscapes. *Management Science* 47(1):117–132.
- Galindo L, Perset K, Sheeka F (2021) An overview of national AI strategies and policies. *OECD Going Digital Toolkit Notes* (14).
- Gans J (2019) Science policy: An ARPA for science. *Science* 363(6425):1103–1105.
- Giczy AV, Pairolero NA, Toole AA (2022) Identifying artificial intelligence (AI) invention: a novel AI patent dataset. *Journal of Technology Transfer* 47:476–505.
- Gormley TA, Matsa DA (2011) Growing out of trouble? Corporate responses to liability risk. *Review of Financial Studies* 24(8):2781–2821.
- Griliches Z (2009) Productivity, R&D, and the data constraint. Neef D, Siesfeld GA, Cefola J, eds., *The Economic Impact of Knowledge*, 253–274 (Butterworth-Heinemann).
- Hall BH (1993a) Industrial research during the 1980s: Did the rate of return fall? *Brookings Papers on Economic Activity: Microeconomics* 289–343.
- Hall BH (1993b) R&D tax policy during the 1980s: success or failure? *Tax Policy and the Economy* 7:1–35.
- Hall BH, Jaffe AB, Trajtenberg M (2005) Market Value and Patent Citations. *The RAND Journal of Economics* 36(1):16–38.
- Hall BH, Lerner J (2010) The financing of R&D and innovation. *Handbook of the Economics of Innovation*, volume 1, 609–639 (Elsevier).
- Haltiwanger J, Jarmin RS, Kulick R, Miranda J (2016) High growth young firms: contribution to job, output, and productivity growth. *Measuring entrepreneurial businesses: Current knowledge and challenges*, 11–62 (University of Chicago Press).
- Hausman N (2022) Public R&D investments and private-sector innovation: Evidence from NIH funding. *American Economic Journal: Applied Economics* .
- Howell ST (2017) Financing Innovation: Evidence from R&D Grants. *American Economic Review* 107(4):1136–1164.
- Howell ST (2024) Government intervention in innovation. *Annual Review of Financial Economics* 16.

- Howell ST, Rathje J, Van Reenen J, Wong J (2025) Opening Up Military Innovation: Causal Effects of Reforms to US Defense Research. *Journal of Political Economy* 133(11):3605–3651.
- Jiang W, Park J, Xiao RJ, Zhang S (2025) AI and the Extended Workday: Productivity, Contracting Efficiency, and Distribution of Rents. Technical report, National Bureau of Economic Research.
- Jones BF (2009) The Burden of Knowledge and the “Death of the Renaissance Man”: Is Innovation Getting Harder? *Review of Economic Studies* 76(1):283–317.
- Kelly B, Papanikolaou D, Seru A, Taddy M (2021) Measuring technological innovation over the long run. *American Economic Review: Insights* 3(3):303–320.
- Kelly BT, Kuznetsov B, Malamud S, Xu TA (2025) Artificial intelligence asset pricing models. Working Paper 33351, National Bureau of Economic Research.
- Kogan L, Papanikolaou D, Seru A, Stoffman N (2017) Technological Innovation, Resource Allocation, and Growth. *Quarterly Journal of Economics* 132(2):665–712.
- Larrue P (2021) The design and implementation of mission-oriented innovation policies. *OECD Science, Technology and Industry Policy Papers* .
- Lerner J (2009) *Boulevard of Broken Dreams: Why Public Efforts to Boost Entrepreneurship and Venture Capital Have Failed* (Princeton University Press).
- Lucas RE (1988) On the mechanics of economic development. *Journal of Monetary Economics* 22(1):3–42.
- Lucking B, Bloom N, Van Reenen J (2018) Have R&D spillovers changed? Technical report, National Bureau of Economic Research.
- Manso G (2011) Motivating innovation. *Journal of Finance* 66(5):1823–1860.
- March JG (1991) Exploration and exploitation in organizational learning. *Organization Science* 2(1):71–87.
- Mazzucato M (2013) *The Entrepreneurial State: Debunking Public vs. Private Sector Myths* (Anthem Press).
- Mazzucato M (2021) *Mission Economy: A Moonshot Guide to Changing Capitalism*.
- Mazzucato M, Semieniuk G (2018) Mission-oriented innovation policy: Challenges and opportunities. *Industrial and Corporate Change* 27(5):803–815.
- Moretti E (2021) The Effect of High-Tech Clusters on the Productivity of Top Inventors. *American Economic Review* 111(10):3328–75.
- Nanda R, Nicholas T (2014) Did bank distress stifle innovation during the Great Depression? *Journal of Financial Economics* 114(2):273–292.
- Nelson RR (1959) The simple economics of basic scientific research. *Journal of political economy* 67(3):297–306.
- OECD AI (2025) OECD.AI Policy Navigator. <https://oecd.ai/dashboards>.
- Radford A, Narasimhan K (2018) Improving Language Understanding by Generative Pre-Training. *OPENAI*.
- Romer PM (1990) Endogenous technological change. *Journal of political Economy* 98(5, Part 2):S71–S102.
- Sant’Anna PH, Zhao J (2020) Doubly robust difference-in-differences estimators. *Journal of Econometrics* 219(1):101–122.
- Schumpeter JA (1947) The creative response in economic history. *Journal of Economic History* 7(2):149–159.
- Solow RM (1957) Technical change and the aggregate production function. *Review of Economics and Statistics* 39(3):312–320.
- Tian X, Wang TY (2014) Tolerance for failure and corporate innovation. *Review of Financial Studies* 27(1):211–255.

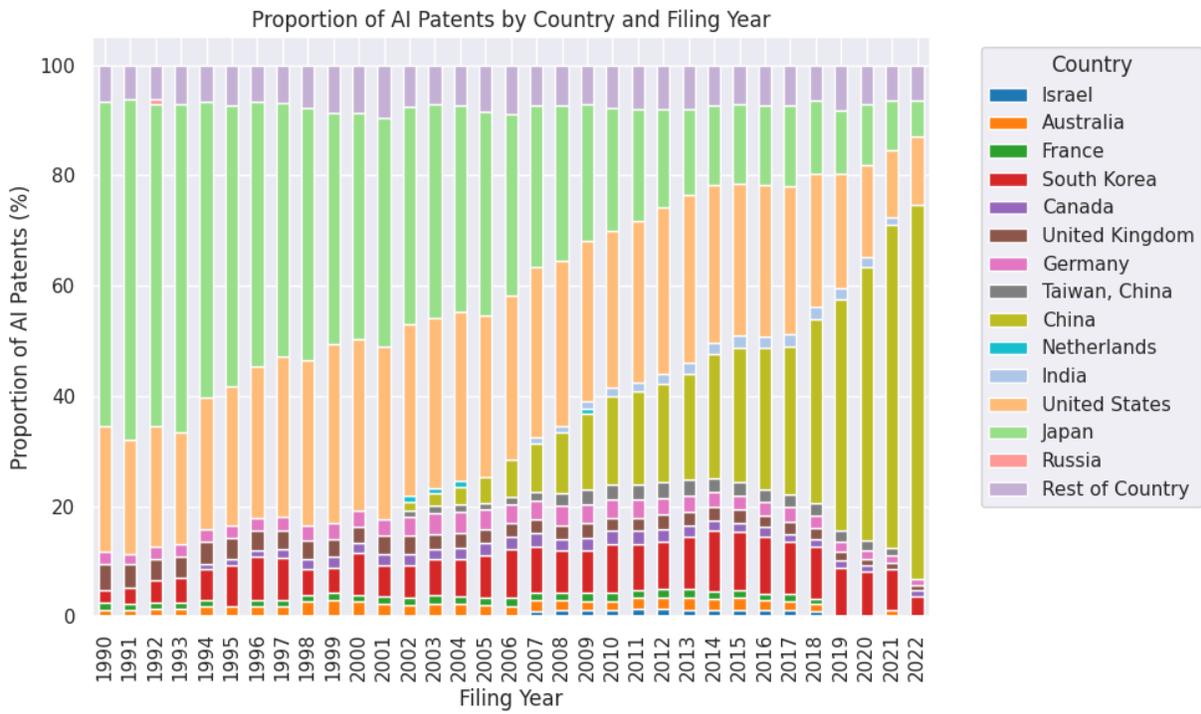
- Trajtenberg M (1990) A penny for your quotes: Patent citations and the value of innovations. *RAND Journal of Economics* 21(1):172–187.
- Uzzi B, Mukherjee S, Stringer M, Jones B (2013) Atypical combinations and scientific impact. *Science* 342(6157):468–472.
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30.
- Verhoeven D, Bakker J, Veugelers R (2016) Measuring technological novelty with patent-based indicators. *Research Policy* 45(3):707–723.
- Weitzman ML (1979) Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society* 641–654.

## 7 Figures and Tables

**Figure 1**

### **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.

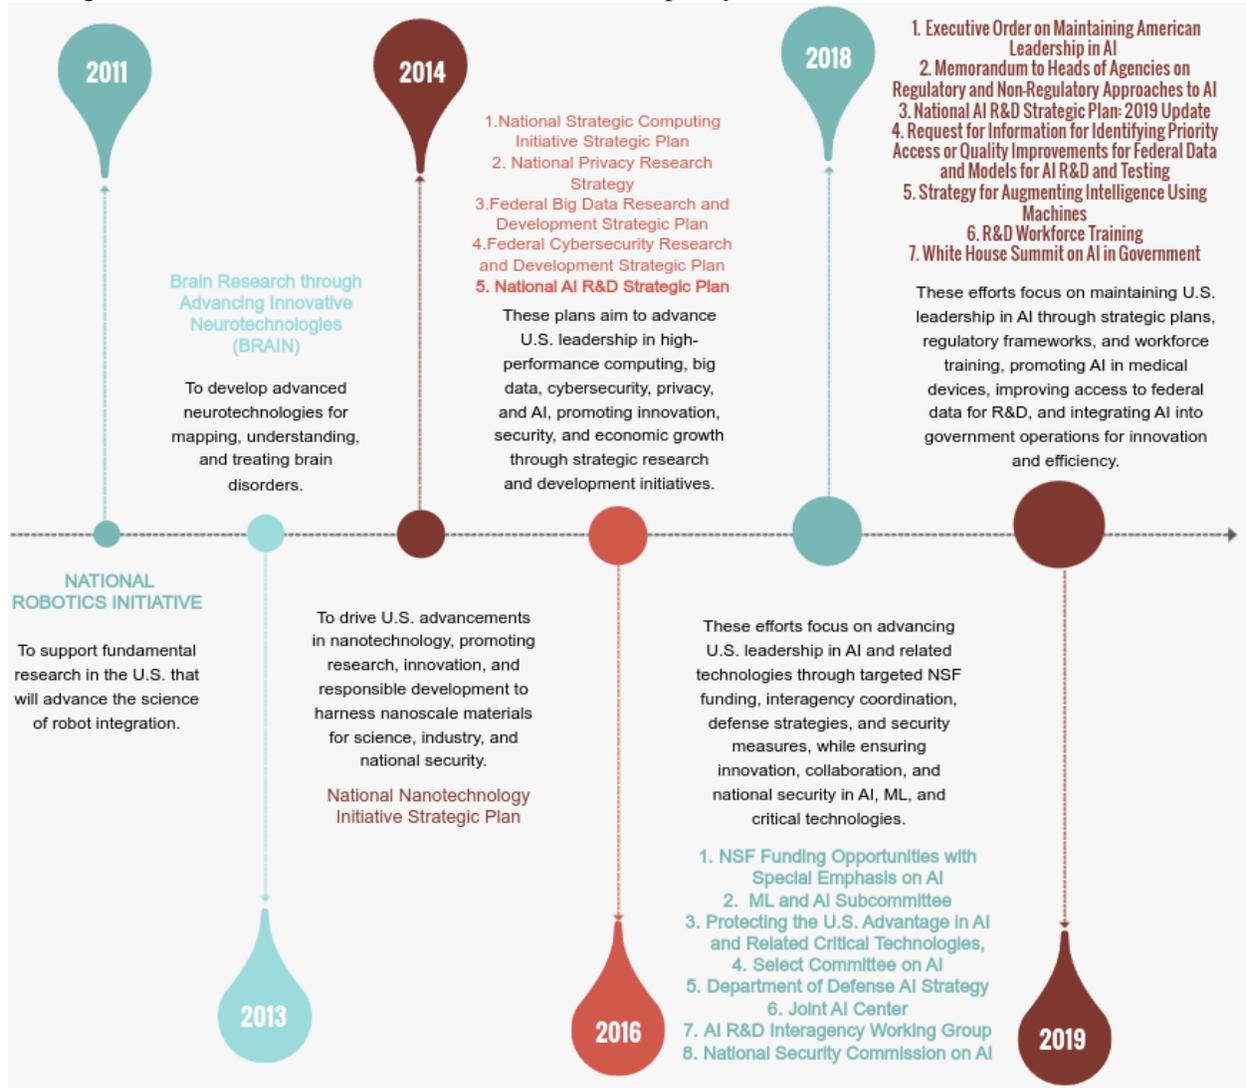




**Figure 3**

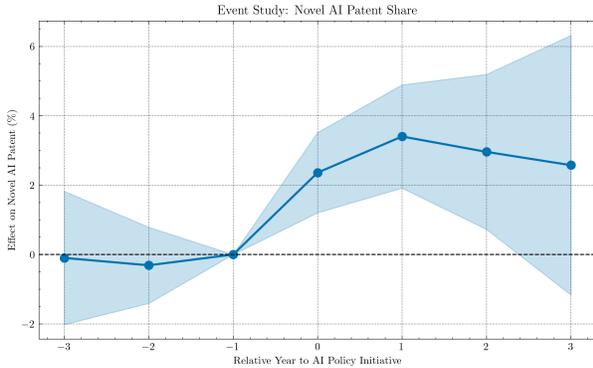
**Timeline of USA’s Initiatives and Policy Events Related to AI**

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

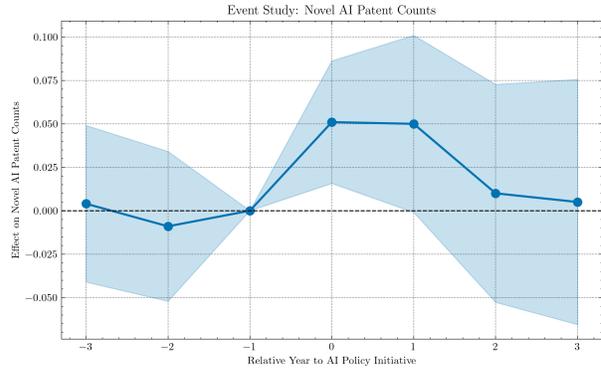


**Figure 4**  
**Parallel Trends**

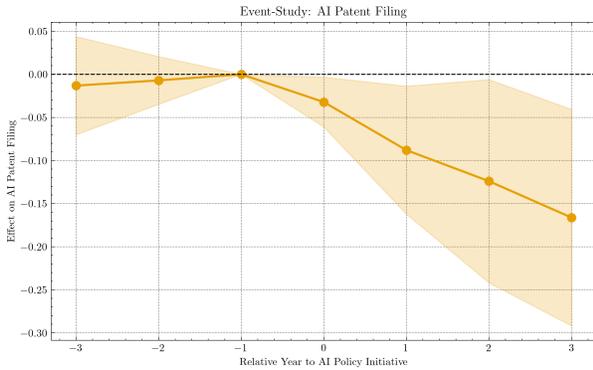
This figure plots the dynamic difference-in-differences coefficients to test the parallel trends assumption and examine the evolution of treatment effects over time. Panel (a) and (b) reports the results for novel AI patents in share and count. Panels (c), (d) and (e) report the results for AI patent counts and citation-weighted counts and average citations, respectively. The specification follows the baseline regression but replaces the single post-treatment indicator with a set of event-time dummies ranging from  $t - 3$  to  $t + 3$ , where  $t = 0$  denotes the year of policy implementation. The reference year is set to  $t - 1$ . Shaded areas represent 95% confidence intervals.



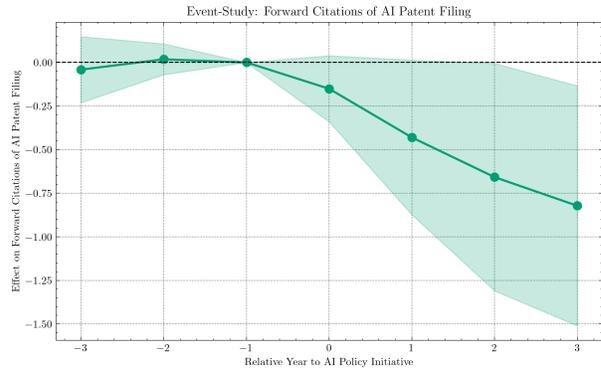
**(a) Patent novelty (share)**



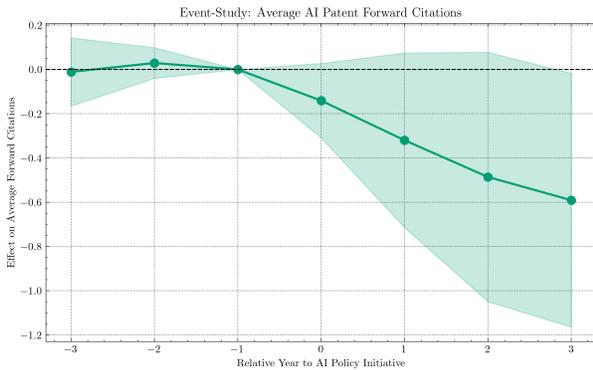
**(b) Patent novelty (count)**



**(c) Patent counts**



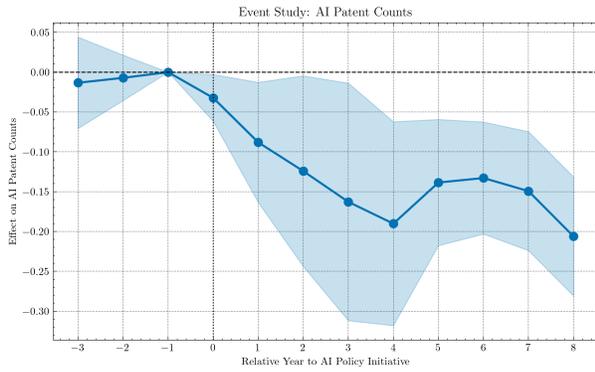
**(d) Citation-weighted counts**



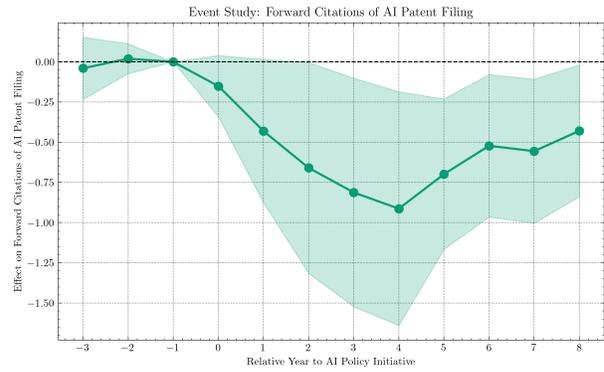
**(e) Average citations**

**Figure 5**  
**Long-term Effects**

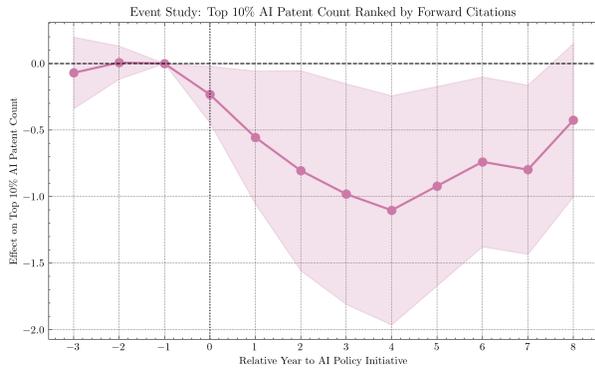
This figure plots the dynamic difference-in-differences estimator to assess long-term effects. Panel (a) and (b) report the results for AI patent counts and citation-weighted counts. Panel (c) reports the top 10% patent indicator ranked by citations, and (d) reports the top 1% patent indicator ranked by citations. The specification follows the baseline regression but replaces the single post-treatment indicator with a set of event-time dummies ranging from  $t - 3$  to  $t + 8$ , where  $t = 0$  denotes the year of policy implementation. The reference year is set to  $t - 1$ . Shaded areas represent 95% confidence intervals.



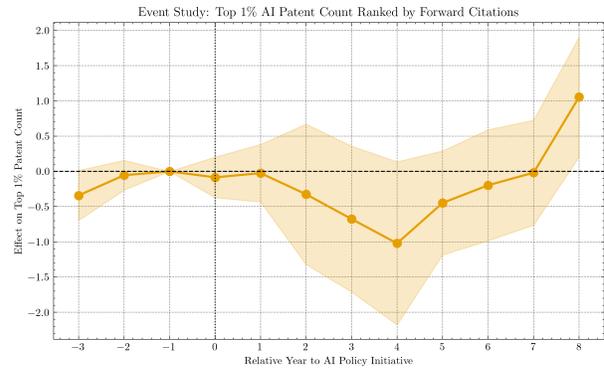
**(a) Patent counts**



**(b) Citation-weighted counts**



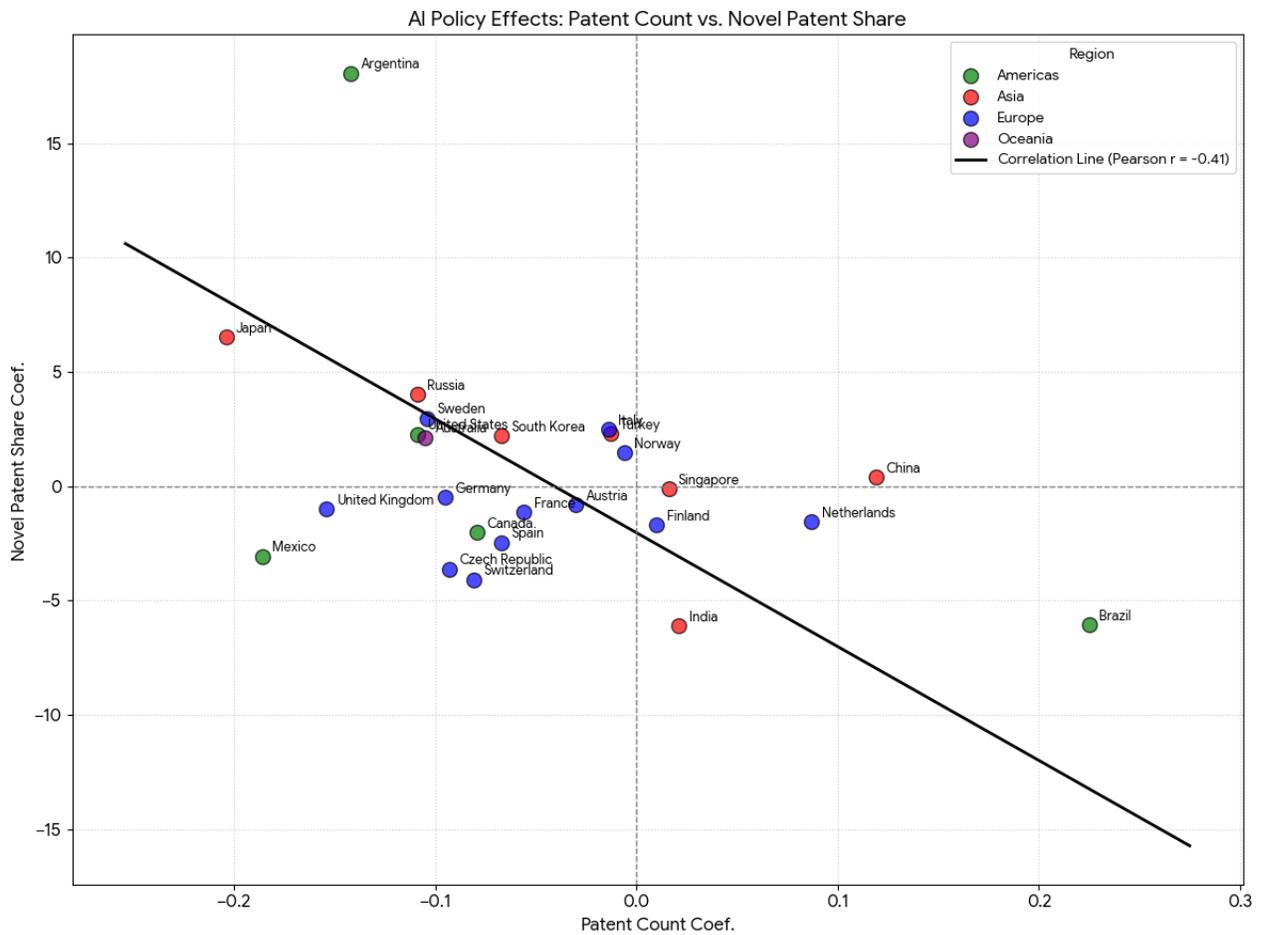
**(c) Top 10% patent counts ranked by citations**



**(d) Top 1% patent counts ranked by citations**

**Figure 6**  
**Individual Country Results**

This figure presents the coefficients derived from the estimation of individual countries' novel AI patent shares and total AI patent counts. The empirical setup follows the baseline specification by estimating a stacked difference-in-differences (DID) regression for each treated country, where the control group consists of never-treated and not-yet-treated countries.



**Table 1**  
**Inventor-Level Panel Data Descriptive Statistics**

This table reports descriptive statistics for inventor-level panel data about AI patent activity from 2006 to 2019. Panel A presents summary statistics for the full inventor-level panel prior to stacking for the difference-in-differences analysis. *Novel AI Patent Share* is defined as the number of novel AI patents—those with zero competing (textually similar) patents—divided by the total number of AI patents filed by an inventor in a given year. *Novel AI Patent Count* measures the total number of novel AI patents filed by an inventor in a given year. *AI Patent Counts* measure the total number of AI patents filed by an inventor in a given year. *AI Patent Forward Citations* measure the number of subsequent patents that cite an inventor’s AI patents, serving as a proxy for technological importance and spillover. *Average AI Patent Forward Citations* are calculated as the total number of forward citations divided by the inventor’s total AI patent filings in that year. *Inventor AI Domains* count the number of distinct AI technology domains in which an inventor files patents in a given year. *Inventor AI Domains HHI* measures the concentration of an inventor’s AI patenting activity across these domains. Panel B reports corresponding summary statistics for inventors located in the 5 major AI-patenting countries. Full variable definitions are provided in the Appendix.

	N	Mean	std. dev	25%	50%	75%
<b>Panel A: Inventor-Level Panel Data</b>						
Novel AI Patents Share (%)	2,960,410	26.27	42.18	0.00	0.00	50.00
Novel AI Patent Counts	2,960,410	0.36	0.74	0.00	0.00	1.00
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
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
<b>Panel B: Inventor-Level Panel Data by Country: Five Main Countries</b>						
<b>United States</b>						
Novel AI Patents Share (%)	855,901	34.11	45.35	0.00	0.00	100.00
Novel AI Patent Counts	855,901	0.48	0.92	0.00	0.00	1.00
AI Patent Counts	855,901	1.48	1.66	1.00	1.00	1.00
AI Patent Forward Cites	855,901	11.52	46.87	0.00	2.00	8.00
<b>China</b>						
Novel AI Patents Share (%)	683,552	17.73	36.30	0.00	0.00	0.00
Novel AI Patent Counts	683,552	0.26	0.65	0.00	0.00	0.00
AI Patent Counts	683,552	1.64	3.05	1.00	1.00	1.00
AI Patent Forward Cites	683,552	7.95	19.34	0.00	3.00	8.00
<b>Japan</b>						
Novel AI Patents Share (%)	466,581	19.53	37.55	0.00	0.00	0.00
Novel AI Patent Counts	466,581	0.27	0.59	0.00	0.00	0.00
AI Patent Counts	466,581	1.45	1.20	1.00	1.00	1.00
AI Patent Forward Cites	466,581	5.28	11.38	0.00	2.00	6.00
<b>South Korea</b>						
Novel AI Patents Share (%)	272,154	22.09	39.08	0.00	0.00	33.33
Novel AI Patent Counts	272,154	0.32	0.76	0.00	0.00	1.00
AI Patent Counts	272,154	1.50	1.52	1.00	1.00	1.00
AI Patent Forward Cites	272,154	3.89	9.54	0.00	2.00	4.00
<b>United Kingdom</b>						
Novel AI Patents Share (%)	57,226	33.62	45.94	0.00	0.00	100.00
Novel AI Patent Counts	57,226	0.40	0.62	0.00	0.00	1.00
AI Patent Counts	57,226	1.26	0.86	1.00	1.00	1.00
AI Patent Forward Cites	57,226 <sup>41</sup>	4.88	21.50	0.00	0.00	4.00

**Table 2****National AI-Related Policy and Inventor's Innovation Novelty and Productivity**

This table reports the inventor-year level analysis on AI policy on AI innovation. Both panels report stacked Difference-in-Differences (DID) estimates, controlling for inventor-by-cohort and year-by-cohort fixed effects. The AI policy indicator takes the value of one if inventors are subject to national AI-related policies in the given year. Panel A reports results on AI Patent Novelty: Column (1) reports the results on the share of novel patents, while Column (2) reports the results on the absolute number of novel patents. Novel patents are defined as those that do not have competing (textually similar) patents filed within the same year, as detailed in [Section 4.2](#). Panel B reports results on total patent counts (Column 1), citation-weighted counts (Column 2), and average citations (Column 3). Robust standard errors are clustered at the country level, with t-statistics reported in parentheses. \*\*\*, \*\*, and \* represent significance levels at the 1%, 5%, and 10% levels, respectively. The sample period is from 2006 to 2019.

**Panel A: AI Innovation Novelty**

	OLS Novel AI Patents Share (%) (1)	Poisson Novel AI Patent Counts (2)
AI Policy	2.834** (2.49)	0.034** (2.07)
Year × Cohort FE	Yes	Yes
Inventor × Cohort FE	Yes	Yes
Obs.	3,900,297	2,461,111
$R^2$	0.44	

**Panel B: AI Innovation Quantity and Quality**

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

**Table 3**

**Risk and Return Associated with Novelty**

This table presents a patent-level analysis of AI patent novelty. The dependent variables in Panel A are low-impact AI patents (defined as those receiving zero forward citations) and patent development duration, which 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. The dependent variables in Panel B are forward citation counts and a set of indicator variables denoting whether the AI patent is ranked at the top by citations in a given year. Columns (2)–(5) of Panel B use indicators for the top 0.01%, 0.1%, 1%, and 10% AI patents. For reporting purposes, we multiplied the coefficients by 10,000, 1,000, 100, and 10, respectively, to report the effects relative to the unconditional means. The independent variable of interest is the novel patent indicator, which is defined as those patents with zero competing (textually similar) filings in a given year. Robust standard errors are clustered at the IPC4 technology class level, and t-statistics are reported in parentheses. \*\*\*, \*\*, and \* represent significance levels at 1%, 5%, and 10%, respectively. The sample period is from 2011 to 2019, following the first AI policy.

**Panel A: Risk of Novel AI Patent**

	OLS	
	Low-Impact AI Patent (Forward Cites = 0)	ln(Patent Development Duration)
	(1)	(2)
Novel AI Patent	0.008** (2.13)	0.123*** (10.32)
IPC4 × Year × Country FE	Yes	Yes
Obs.	1,551,917	1,059,153
R <sup>2</sup>	0.11	0.15

**Panel B: Reward of Novel AI Patent**

	Forward Cites	Top 0.01%	Top 0.1%	Top 1%	Top 10%
	(1)	(2)	(3)	(4)	(5)
Novel AI Patent	0.176*** (10.44)	1.951** (2.53)	0.984*** (5.94)	0.510*** (8.06)	0.121*** (6.64)
IPC4 × Year × Country FE	Yes	Yes	Yes	Yes	Yes
Obs.	1,533,160	1,551,917	1,551,917	1,551,917	1,551,917
R <sup>2</sup>		0.01	0.01	0.02	0.07

**Table 4**  
**Risk Associated with Novelty**

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

<b>Panel A: Forward Citation Variance by Novelty</b>	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

<b>Panel B: F-Tests Comparing Variance Across Novelty Groups</b>	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 5**  
**Heterogeneous Effect of Policy Types on AI Innovation**

This table reports the heterogeneous effects of policies across different types. The regression setup follows the baseline specification but disaggregates all policies into distinct categories as specified by the OECD: governance, guidance and regulation, financial support, and AI enablers and other incentives. Governance policies provide strategic direction and institutional oversight, seeking to align 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 are those establishing legal and ethical standards for AI development. Finally, AI enablers and other incentives focus on building long-term innovation capacity through investments in infrastructure, human capital, and collaborative ecosystems. The dependent variables are novel AI patent shares, AI patent filings counts and forward citations. Robust standard errors are clustered at the country level, and t-statistics are reported in parentheses. \*\*\*, \*\*, and \* represent significance levels at 1%, 5%, and 10%, respectively. The sample period is from 2006 to 2019.

	OLS			Poisson			Poisson					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Governance	0.446 (0.43)				-0.019 (-1.48)				-0.060 (-0.65)			
Financial Support		2.402** (2.53)				-0.093*** (-2.73)				-0.473* (-1.91)		
Guidance and Regulation			-2.277** (-2.27)				-0.084*** (-5.00)				-0.271*** (-2.95)	
AI Enablers and Other Incentives				3.622** (2.61)				-0.153*** (-3.22)				-0.565** (-1.99)
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	3,900,297	3,900,297	3,900,297	3,900,297	3,515,936	3,515,936	3,515,936	3,515,936
R <sup>2</sup>	0.44	0.44	0.44	0.44								

**Table 6**  
**Heterogeneous Effect of Target Groups on AI Innovation**

This table reports stacked DID analysis on the heterogeneous effects of AI policies on innovation, categorized by their respective target groups. The regression setup follows the baseline specification but disaggregates all policies into distinct indicators targeting different entities. In particular, policies targeting the private sector are classified by firm size: small firms (fewer than 10 employees), medium-sized firms (10–249 employees), and large firms (250 or more employees). Additionally, two indicators are included for policies targeting the public sector: one for individuals, such as researchers, students, and teachers, and another for research and educational institutes. The dependent variables are novel AI patent shares, AI patent filing counts and forward citations. Robust standard errors are clustered at the country level, and t-statistics are reported in parentheses. \*\*\*, \*\*, and \* represent significance levels at 1%, 5%, and 10%, respectively. The sample period spans from 2006 to 2019.

	OLS			Poisson			Poisson								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Large Firms	-2.868** (-2.03)					-0.010 (-0.47)					-0.209** (-2.13)				
Medium Firms		-0.056 (-0.04)					-0.056*** (-2.72)					0.029 (0.24)			
Small Firms			-1.998* (-1.83)					0.093*** (4.46)					0.123*** (3.23)		
Research & Education Institutes				2.988** (2.37)					-0.138*** (-2.90)					-0.585** (-2.13)	
Individual Researchers, Students & Teachers					3.172** (2.57)					-0.141*** (-2.94)					-0.587** (-2.12)
Year × Cohort FE	Yes	Yes	Yes	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	Yes	Yes	Yes
Obs.	3,900,297	3,900,297	3,900,297	3,900,297	3,900,297	3,900,297	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
R <sup>2</sup>	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44	0.44

**Table 7****Governments Target Novel AI innovation (US AI patents)**

This table presents a patent-level analysis regarding government-funded patents and AI patent quality. The sample is restricted to U.S. patents only. The independent variable of interest is the government-funded AI project indicator, which equals one if the project received a government grant and zero otherwise. The dependent variables consist of novel AI patents, forward citations, low-impact AI patents, and high-impact AI patents. Specifically, a novel AI patent is an indicator variable equal to one if a patent has zero competing (textually similar) patents in its filing year, and zero otherwise. A low-impact AI patent is an indicator variable equal to one if a patent receives zero forward citations. A high-impact AI patent is defined as an indicator variable equal to one if the patent is ranked at the top by citations in a given year. We multiply the coefficients by 10 and 100 for top 10% and 1% patents, respectively, to examine the effects relative to the unconditional means. All regressions include IPC4  $\times$  year fixed effects. Robust standard errors are clustered at the IPC4 technology class level, and t-statistics are reported in parentheses. \*\*\*, \*\*, and \* represent significance levels at 1%, 5%, and 10%, respectively. The sample period is from 2011 to 2019, following the first AI policy.

	OLS Novel AI Patent (1)	Poisson Forward Cites (2)	OLS Low-Impact AI Patent (3)	OLS Top 10% (4)	OLS Top 1% (5)
Gov Funded Project	0.176*** (17.29)	-0.043 (-0.76)	-0.043*** (-7.05)	0.099* (1.75)	-0.503*** (-3.04)
Year $\times$ IPC4 FE	Yes	Yes	Yes	Yes	Yes
Obs.	363,897	362,980	363,897	363,897	363,897
$R^2$	0.05		0.13	0.04	0.01

**Table 8****National AI-Related Policy and Domains of AI Innovation**

This table reports an stacked DID analysis regarding national AI-related policies and technological domain specialization. The AI policy indicator is set to one if the inventors are subject to national AI-related policies in a given year and zero otherwise. The dependent variables are the inventor's number of AI domains using AIPD definitions and the AI patent Herfindahl-Hirschman Index (HHI), which measures the concentration of patenting activity across AI domains. Robust standard errors are clustered at the country level, and t-statistics are reported in parentheses. \*\*\*, \*\*, and \* represent significance levels at 1%, 5%, and 10%, respectively. The sample period spans from 2006 to 2019.

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

**Table 9****National AI-Related Policy and Cross-Country Spillovers**

This table reports a stacked difference-in-differences (DID) analysis on national AI-related policies and international knowledge spillovers. The AI policy indicator is an indicator variable taking the value of one if the inventors are subject to national AI-related policies in a given year and zero otherwise. The dependent variables represent the intensity of cross-border knowledge flows, specifically the percentage of cross-country forward citations and the percentage of cross-country backward citations. Robust standard errors are clustered at the country level, and t-statistics are reported in parentheses. \*\*\*, \*\*, and \* represent significance levels at 1%, 5%, and 10%, respectively. The sample period spans from 2006 to 2019.

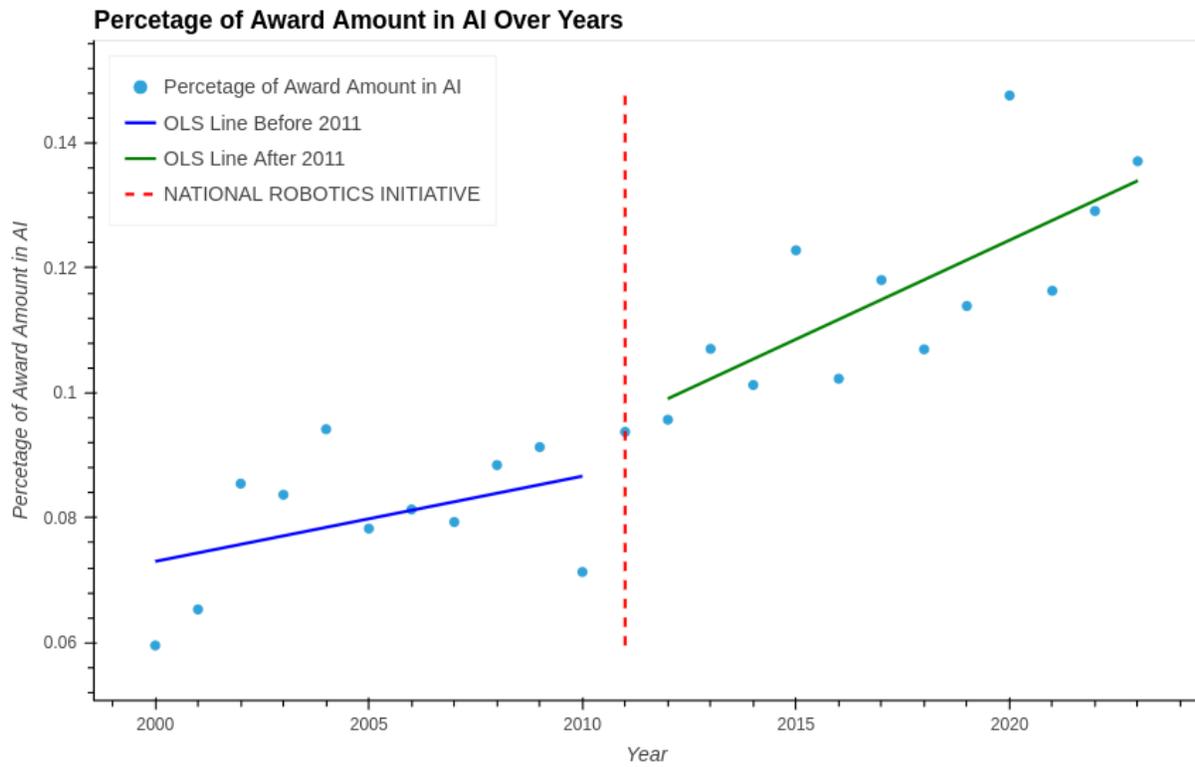
	OLS Cross Country Backward Cites.% (1)	OLS Cross Country Forward Cites.% (2)
AI Policy	4.662** (2.50)	5.763** (2.32)
Year × Cohort FE	Yes	Yes
Inventor × Cohort FE	Yes	Yes
Obs.	2,668,603	2,308,300
$R^2$	0.74	0.65

## A Appendix Figures and Tables

**Figure A1**

### **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 A1**  
**Variable Definitions**

Variable Name	Abbrev.	Description
<b>Innovation Variables (Source: Google Patents Public Data and AIPD)</b>		
Novel AI Patents Share		The number of novel AI patents—defined as patents with zero competing (textually similar) patents—divided by the total number of AI patents filed by an inventor in a given year.
Novel AI Patent Counts		The number of novel AI patents—defined as those with zero competing (textually similar) patents—filed by an inventor in a given year.
AI Patent Counts		Number of AI patent filed 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.
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
<b>National AI-related Strategies &amp; Policies Variables (Source: OECD AI Policy Database)</b>		
National AI policies	AI Policy	The treat-post dummy variable assigning a value of 1 during the periods they are implemented and 0 otherwise.
Governance		The treat-post dummy variable equals one during periods in which a policy focusing on governance is implemented, and zero otherwise. Governance-related policies are defined as providing strategic direction and institutional oversight, aligning AI development with public priorities through coordination and state-led initiatives.
Financial Support		The treat-post dummy variable equals one during periods in which a policy focusing on financial support is implemented, and zero otherwise. Financial support policies are defined as providing direct subsidies for innovation inputs, including public R&D funding and capital investment.
Guidance and Regulation		The treat-post dummy variable equals one during periods in which a policy focusing on guidance and regulation is implemented, and zero otherwise. Guidance and regulation policies are defined as managing externalities by establishing legal and ethical standards for AI development.
AI Enablers and Other Incentives		The treat-post dummy variable equals one during periods in which a policy focusing on AI enablers and other incentives is implemented, and zero otherwise. AI enablers and other incentives focus on building long-term innovation capacity through investments in infrastructure, human capital, and collaborative ecosystems.
Large Firms		The treat-post dummy variable equals one during periods in which a policy targeting large firms—defined as firms with more than 250 employees—is implemented, and zero otherwise.
Medium Firms		The treat-post dummy variable equals one during periods in which a policy targeting large firms—defined as firms with 10–249 employees—is implemented, and zero otherwise.
Small Firms		The treat-post dummy variable equals one during periods in which a policy targeting large firms—defined as firms with fewer than 10 employees—is implemented, and zero otherwise.
Research & Education Institutes		The treat-post dummy variable equals one during periods in which a policy targeting Research & Education Institutes—defined as public research institutes and higher education institutions—is implemented, and zero otherwise.
Individual researchers, students, and teachers		The treat-post dummy variable equals one during periods in which a policy targeting individual researchers, students, and teachers—defined to include established researchers, PhD students, teachers, postdoctoral and other early-career researchers, as well as secondary education, undergraduate, and master’s students—is implemented, and zero otherwise.

## Table A2

### Prompt used to classify whether a policy qualifies as a national AI policy

This table presents the Python prompt utilized with ChatGPT GPT-4o to determine whether a given policy pertains to Artificial Intelligence (AI). To account for model variability, each policy description was evaluated ten times; we retained only those classified as “AI-related” and “national in scope” in at least seven out of ten evaluations. Input data and code originating from our policy datasets and generated on our end are indicated in italics.

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 Background is not null 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 A3****Summary Statistics of AI Policy Budgets, AI Innovation, and AI Inventors by Country**

This table reports the average annual AI policy budget, the number of AI patent filings and AI inventors by country over the sample period from 2006 to 2019.

Country	Average Annual Budget (\$M)	No. of AI Patents	No. of AI Inventors
Saudi Arabia	6000.00	1309	1860
United Kingdom	4165.53	37015	41956
China	2640.00	370089	439440
United States	1678.00	455461	525833
France	1278.00	22366	27556
Singapore	1140.00	4037	5476
Korea, Republic of	1005.00	152437	146702
Russian Federation	920.00	13430	28766
Japan	790.50	302758	226842
Czechia	740.00	1676	1934
Netherlands	710.00	13121	16890
Germany	499.00	46881	58816
Israel	300.00	19622	19387
Spain	275.00	7487	15251
Switzerland	250.00	9135	10485
Finland	190.50	7823	8017
Australia	105.00	28060	53372
Norway	85.00	2878	3677
Denmark	85.00	4106	5139
Sweden	82.50	12857	12123
Mexico	60.00	1924	4006
India	55.00	30156	35724
Italy	51.00	10514	14436
Canada	36.00	31671	33925
Ireland	30.00	4600	4534
Egypt	30.00	364	0
Argentina	25.00	456	573
Brazil	20.50	2627	5202
Portugal	19.00	617	1155
Austria	15.00	3333	4243
Türkiye	14.50	1525	2338
South Africa	10.00	940	1194
Estonia	7.00	278	0
Hungary	1.00	1092	1522
Romania	1.00	894	1184
Luxembourg	0.50	370	0
Belgium	0.50	5016	5659
New Zealand	0.00	1446	2165
United Arab Emirates	0.00	369	0

**Table A4**  
**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.

<b>Policy Name</b>	<b>Country</b>	<b>Start Year</b>	<b>Description</b>
National AI R&D Plan	United States	2016	The ultimate goal is to produce <b>new</b> AI knowledge and technologies that generate broad societal benefits. To achieve this, the AI R&D Strategic Plan emphasizes <b>federally funded</b> research investments in areas with <b>substantial long-term potential</b> , recognizing that sustained support for <b>high-risk</b> projects can yield transformative, <b>high-reward</b> outcomes.
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 <b>boost scientific research</b> , economic growth, and national security. It strives to set technical standards, simplify the process for AI testing and deployment, and support the <b>growth of new AI-focused industries</b> and the broader adoption of AI in existing sectors.
National New Generation AI Plan	China	2017	A comprehensive plan to position China as a global AI leader by 2030, setting ambitious goals to <b>promote strongly original and exploratory</b> research, encourage scientific address frontier challenges in AI, and foster the creation of novel theories and discoveries.
“Internet+” AI Three-Year Action and Implementation Plan	China	2016	Carry out research on <b>cutting-edge theories</b> and technologies in brain-inspired intelligence fields, such as brain-inspired neural computing systems and brain-inspired information processing.
Alan Turing Institute	United Kingdom	2015	<b>Advance AI research</b> , apply it to <b>real-world challenges</b> , promote ethical AI innovation, and train future talent.
Pan-Canadian AI Strategy	Canada	2017	Focuses on attracting/retaining AI talent, <b>boosting research</b> , and promoting commercialization/standardization; includes a CAD \$125 million investment.
AI Research Programme	Singapore	2018	Funds <b>fundamental AI research</b> , encourages collaboration, and develops local AI talent.
Saudi Data and AI Authority	Saudi Arabia	2019	Drive Saudi Arabia’s data and AI agenda, aiming to position the Kingdom as a <b>leading data-driven</b> economy while supporting G20 AI Principles.

**Table A5**

**National AI Policies: Support Instruments and Target Entities**

This table reports the support instruments and target entities of national AI policies across various countries. It details the types of instruments employed (e.g., funding, regulatory frameworks) and distinguishes between public and private sector entities that are the focus of these policies. Values in parentheses indicate the number of policies that feature the specified instrument or target.

Country	Policy count	Instruments	Target entity	
			Public sector targets	Private sector targets
United States	29	Project grants for public research (11); Centres of excellence grants (2); Fellowships and postgraduate loans and scholarships (1)	PhD students (12); Undergraduate and master students (12); Established researchers (11); Public research institutes (11); Technology transfer offices (4); Postdocs and other early-career researchers (3); Higher education institutes (2)	Firms of any age (8); Firms of any size (8); Private research and development lab (6)
United Kingdom	27	Fellowships and postgraduate loans and scholarships (7); Knowledge transfers and business advisory services (5); Procurement programmes for AI R&D and innovation (3); Project grants for public research (3); Centres of excellence grants (2); Data access and sharing (2); Equity financing (2); Indirect financial support (2); Networking and collaborative platforms (2)	Established researchers (17); Higher education institutes (15); Postdocs and other early-career researchers (12); Public research institutes (9); PhD students (3); Academic societies / academies (2); Teachers (2); Undergraduate and master students (2)	Private research and development lab (9); Firms of any age (8); Firms of any size (8); Incubators, accelerators, hubs (5); Entrepreneurs (3)
Germany	19	Centres of excellence grants (5); Networking and collaborative platforms (5); Project grants for public research (5); Grants for business R&D and innovation (3); Institutional funding for public research (2); AI skills and education (1); Knowledge transfers and business advisory services (1); Science and innovation challenges, prizes and awards (1)	Higher education institutes (9); Public research institutes (9); Academic societies / academies (8); Established researchers (6); Postdocs and other early-career researchers (4); PhD students (1); Undergraduate and master students (1)	Private research and development lab (8); Firms of any size (6); Industry associations (5); SMEs (10 to 249 employees) (4); Established firms (more than 5 years old) (3); Young firms (1 to 5 years old) (3); Entrepreneurs (2); Firms of any age (2); Incubators, accelerators, hubs (2)
Japan	19	Networking and collaborative platforms (14); AI computing and research infrastructure (10); Grants for business R&D and innovation (4); Data access and sharing (2)	Public research institutes (19); Higher education institutes (14); Undergraduate and master students (13); Established researchers (12); PhD students (12); Postdocs and other early-career researchers (12); Academic societies / academies (4); Secondary education students (1); Teachers (1)	Firms of any size (6); Industry associations (4); Entrepreneurs (3); Firms of any age (3); Private research and development lab (3)
Australia	16	Grants for business R&D and innovation (4); Project grants for public research (4); Fellowships and postgraduate loans and scholarships (3); AI computing and research infrastructure (2); Standards and certification for technology development and adoption (1)	Established researchers (6); Public research institutes (6); PhD students (3); Postdocs and other early-career researchers (2); Secondary education students (2); Teachers (2); Higher education institutes (1)	Firms of any age (4); Firms of any size (4); Private research and development lab (2)
Singapore	11	Grants for business R&D and innovation (2); Project grants for public research (2)	Public research institutes (9); Established researchers (3); Higher education institutes (3); Postdocs and other early-career researchers (3); Undergraduate and master students (3); Academic societies / academies (1); PhD students (1); Secondary education students (1); Teachers (1); Technology transfer offices (1)	Firms of any size (6); Private research and development lab (2); Entrepreneurs (1); Firms of any age (1); Incubators, accelerators, hubs (1); Industry associations (1); Private investors (1)

*Continued on next page*

Table A5 – continued from previous page

Country	Policy count	Instruments	Target entity	
			Public sector targets	Private sector targets
Austria	10		Higher education institutes (3)	Large firms (250 or more employees) (4); Multinational enterprises (4); Private investors (4); Private research and development lab (4); Entrepreneurs (3); Firms of any age (3); Firms of any size (3)
Finland	9	Grants for business R&D and innovation (2); Knowledge transfers and business advisory services (2); Procurement programmes for AI R&D and innovation (2); Project grants for public research (2)	Public research institutes (6); Academic societies / academies (3); Established researchers (3); Higher education institutes (3); PhD students (3); Postdocs and other early-career researchers (3); Teachers (3); Technology transfer offices (3); Undergraduate and master students (3)	Firms of any size (4); Entrepreneurs (3); Firms of any age (3); Incubators, accelerators, hubs (3); Industry associations (3); Private investors (3); Private research and development lab (3); Large firms (250 or more employees) (2); Nascent firms (0 to less than 1 year old) (2); SMEs (10 to 249 employees) (2)
South Korea	9	AI computing and research infrastructure (5); Networking and collaborative platforms (5); Data access and sharing (3); Indirect financial support (3); Centres of excellence grants (2); Grants for business R&D and innovation (2); Project grants for public research (2)	Higher education institutes (6); Public research institutes (6); Academic societies / academies (2); Established researchers (2); PhD students (2); Postdocs and other early-career researchers (2); Undergraduate and master students (2)	Entrepreneurs (5); Firms of any age (5); Firms of any size (4); Private research and development lab (4); SMEs (10 to 249 employees) (3); Young firms (1 to 5 years old) (3)
Sweden	9	Networking and collaborative platforms (2); Procurement programmes for AI R&D and innovation (2)	Higher education institutes (3); Public research institutes (2); Established researchers (1)	Firms of any size (3); Entrepreneurs (2); Firms of any age (2); Large firms (250 or more employees) (1); Private investors (1); Private research and development lab (1); SMEs (10 to 249 employees) (1)
Canada	8		Established researchers (3); Higher education institutes (3); PhD students (3); Postdocs and other early-career researchers (3)	
Turkey	8	AI computing and research infrastructure (7); Data access and sharing (1)	Academic societies / academies (7); Established researchers (7); Higher education institutes (7); Public research institutes (7); Technology transfer offices (7)	Firms of any age (7); Firms of any size (7); Incubators, accelerators, hubs (7); Industry associations (7); Private research and development lab (7)
India	7	Centres of excellence grants (5); Networking and collaborative platforms (5)	Established researchers (7); Public research institutes (7); Postdocs and other early-career researchers (5)	Large firms (250 or more employees) (7); Multinational enterprises (7); Private investors (7); SMEs (10 to 249 employees) (7); Firms of any age (5); Firms of any size (5); Incubators, accelerators, hubs (5); Private research and development lab (5); Entrepreneurs (2); Industry associations (2); Micro-enterprises (10 or less employees) (2)
Portugal	7	Project grants for public research (4); AI use in the public sector (2)	Higher education institutes (5); Public research institutes (5)	Private research and development lab (5); Firms of any size (1)
China	6			Firms of any size (6); Firms of any age (3)
France	6	AI computing and research infrastructure (3); Data access and sharing (3); Fellowships and postgraduate loans and scholarships (3); Networking and collaborative platforms (3); Centres of excellence grants (2); Grants for business R&D and innovation (2); Science and innovation challenges, prizes and awards (2); Standards and certification for technology development and adoption (2); AI use in the public sector (1); Knowledge transfers and business advisory services (1)	Established researchers (6); Higher education institutes (4); PhD students (4); Public research institutes (4); Postdocs and other early-career researchers (3); Technology transfer offices (1)	Firms of any size (4); Private research and development lab (4); Firms of any age (3); Entrepreneurs (2); SMEs (10 to 249 employees) (2); Incubators, accelerators, hubs (1)

Continued on next page

Table A5 – continued from previous page

Country	Policy count	Instruments	Target entity	
			Public sector targets	Private sector targets
Netherlands	6	Networking and collaborative platforms (5)	Academic societies / academies (2); Established researchers (2); Higher education institutes (2); PhD students (2); Postdocs and other early-career researchers (2); Secondary education students (2); Teachers (2); Technology transfer offices (2); Undergraduate and master students (2); Public research institutes (1)	Incubators, accelerators, hubs (6); Entrepreneurs (5); Private research and development lab (5); Firms of any age (2); Firms of any size (2); Industry associations (2); Private investors (1); SMEs (10 to 249 employees) (1)
Russia	6	Centres of excellence grants (5); Grants for business R&D and innovation (5)	Higher education institutes (6); Public research institutes (6)	Private research and development lab (6); Private investors (5); SMEs (10 to 249 employees) (5); Large firms (250 or more employees) (1)
Norway	5	Project grants for public research (2)	Established researchers (2); Higher education institutes (2); PhD students (2); Postdocs and other early-career researchers (2); Public research institutes (2)	Firms of any size (5); Firms of any age (3); Private research and development lab (2)
Spain	5		Academic societies / academies (1); Established researchers (1); Higher education institutes (1); PhD students (1); Postdocs and other early-career researchers (1); Public research institutes (1); Teachers (1); Technology transfer offices (1); Undergraduate and master students (1)	Firms of any age (5); SMEs (10 to 249 employees) (4); Firms of any size (1); Industry associations (1); Private research and development lab (1)
Denmark	4	Data access and sharing (1); Project grants for public research (1)		Multinational enterprises (3)
Argentina	3	AI use in the public sector (2)	Established researchers (3); Higher education institutes (3); Public research institutes (3); Postdocs and other early-career researchers (2)	Firms of any size (3); Firms of any age (2); Private research and development lab (2); Entrepreneurs (1); SMEs (10 to 249 employees) (1)
Brazil	3		Public research institutes (3); Established researchers (1); Postdocs and other early-career researchers (1)	Firms of any size (1); Industry associations (1); Private research and development lab (1)
Czech Republic	3	Networking and collaborative platforms (3); AI computing and research infrastructure (1); Centres of excellence grants (1); Data access and sharing (1); Equity financing (1); Fellowships and postgraduate loans and scholarships (1); Knowledge transfers and business advisory services (1); Loans and credits for innovation in firms (1)	Public research institutes (2); Academic societies / academies (1); PhD students (1); Postdocs and other early-career researchers (1)	Entrepreneurs (3); SMEs (10 to 249 employees) (3); Firms of any size (2); Incubators, accelerators, hubs (2); Industry associations (2)
Hungary	3		Established researchers (2); Higher education institutes (2); Public research institutes (2); Technology transfer offices (2)	Firms of any size (3); Incubators, accelerators, hubs (2); Private research and development lab (2)
Italy	3	Centres of excellence grants (1); Science and innovation challenges, prizes and awards (1)	Higher education institutes (3)	
Switzerland	3			
Estonia	2		PhD students (1); Public research institutes (1)	Firms of any age (2); Firms of any size (2)
Mexico	2	Networking and collaborative platforms (2)	Public research institutes (2)	Firms of any age (2); Firms of any size (2); Private research and development lab (2)
Saudi Arabia	2	AI computing and research infrastructure (1); Data access and sharing (1)		Firms of any size (1)
Belgium	1	Fellowships and postgraduate loans and scholarships (1)		

Continued on next page

Table A5 – continued from previous page

Country	Policy count	Instruments	Target entity	
			Public sector targets	Private sector targets
Egypt	1	Networking and collaborative platforms (1)	Established researchers (1); Higher education institutes (1); Public research institutes (1)	Firms of any age (1); Firms of any size (1); Private research and development lab (1)
Ireland	1			
Israel	1	Grants for business R&D and innovation (1)		Entrepreneurs (1); Firms of any age (1); Firms of any size (1); Incubators, accelerators, hubs (1)
Luxembourg	1		Public research institutes (1)	Entrepreneurs (1); Private research and development lab (1)
Romania	1		Established researchers (1); Higher education institutes (1)	Firms of any age (1); Firms of any size (1)
South Africa	1		Established researchers (1); Higher education institutes (1); Postdocs and other early-career researchers (1); Public research institutes (1)	Firms of any age (1); Firms of any size (1); Private research and development lab (1)

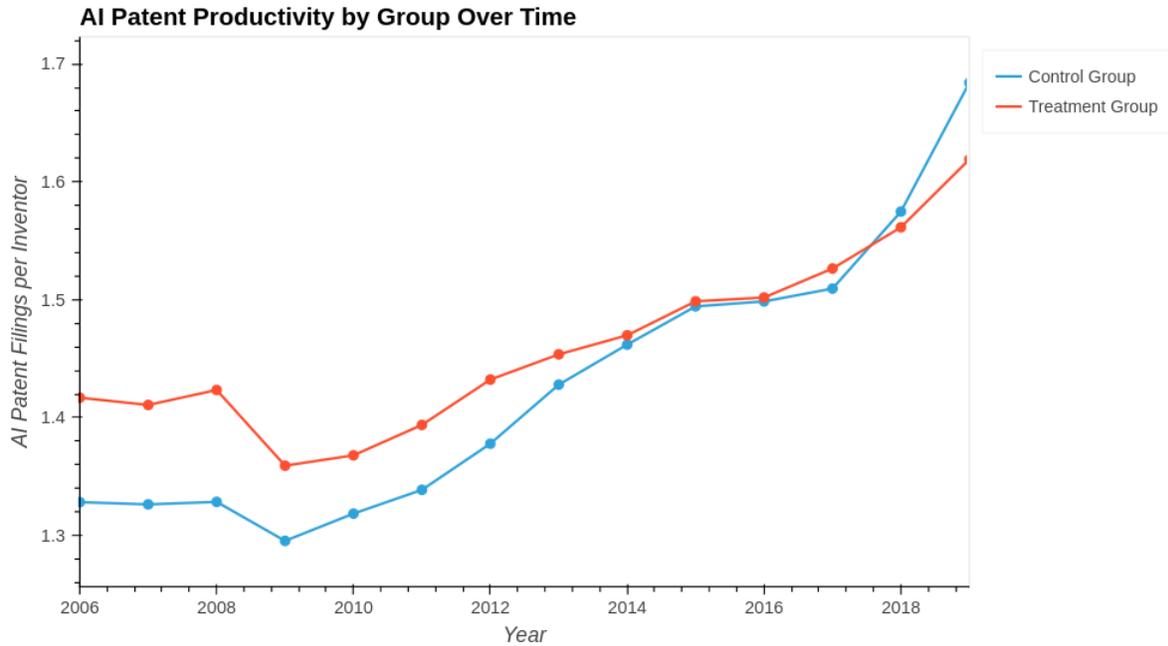
# Internet Appendix to: “Innovation Policies and Inventors’ Productivity”

## B Additional Figures and Tables

Figure B1

### Time Variation of AI Patent Productivity by Group

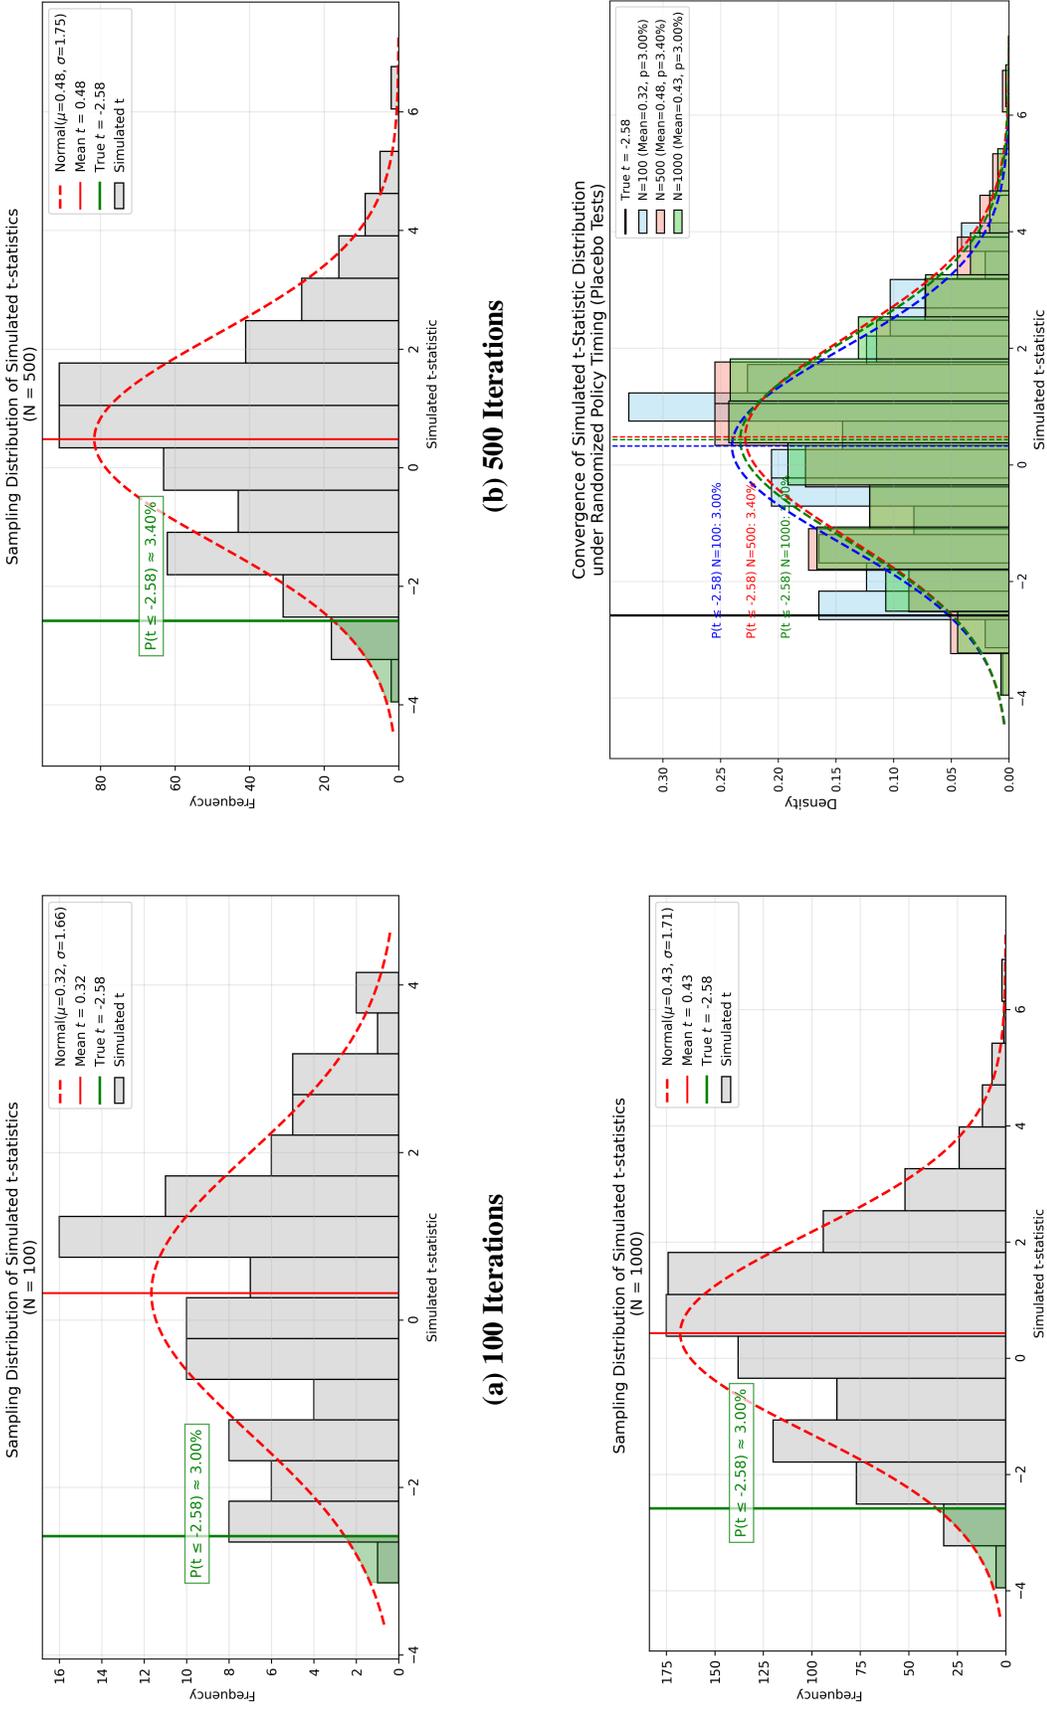
This figure shows the trend in the average number of AI patent annual filings per inventor between 2006 and 2019. The red line indicates the average for inventors in countries that had implemented national AI policies by year  $t$ . The blue line represents the control countries that had not.



**Figure B2**

**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 B 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, recording the corresponding t-statistic for the treatment effect.



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

**(c) 1000 Iterations**

**Table B1**  
**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	2.40	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	2.38	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
Turkiye	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, China	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 B2****Robustness: National AI-Related Policy and Novel Patents**

This table reports results following the specifications of [Table 2](#), but with an alternative approach to defining novel patents. The baseline utilizes the *Novel AI Patent Share* at  $t = 0$ . Here, the *Novel AI Patent Share* at  $t = -1$  is defined as the share of novel AI patents with zero competing (textually similar) patents in the year immediately preceding the filing year. The *Novel AI Patent Share* at  $t < 0$  is defined analogously, but based on the cumulative absence of competing (textually similar) patents up to the filing year.

	OLS	
	Novel AI Patents Share (t=-1) (1)	Novel AI Patents Share(t<0) (2)
AI Policy	3.349** (2.16)	1.743** (2.37)
Year $\times$ Cohort FE	Yes	Yes
Inventor $\times$ Cohort FE	Yes	Yes
Obs.	3,900,297	3,900,297
$R^2$	0.42	0.43

**Table B3****Robustness: National AI-Related Policy and AI Innovation (Excluding US and China)**

This table reports results that follow the baseline specification in [Table 2](#), but excludes observations from the United States and China.

	AI Patent Counts		
	Exclude USA (1)	Exclude China (2)	Exclude USA & China (3)
AI Policy	-0.088** (-2.13)	-0.071*** (-2.99)	-0.091*** (-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 B4****Robustness: National AI-Related Policy and AI Innovation (Controlling for GDP)**

This table reports results mirroring the setup of the baseline regression in [Table 2](#), with the additional inclusion of controls for GDP per capita.

**Panel A: AI Innovation Novelty**

	OLS	Poisson
	Novel AI Patents Share (%)	Novel AI Patent Counts
	(1)	(2)
AI Policy	2.036*** (4.33)	0.031* (1.81)
ln(GDP per capita)	-25.519*** (-13.64)	-0.105 (-1.15)
Year × Cohort FE	Yes	Yes
Inventor × Cohort FE	Yes	Yes
Obs.	3,900,297	2,461,111
$R^2$	0.44	

**Panel B: AI Innovation Quantity and Quality**

	Poisson	Poisson	Poisson
	AI Patent Counts	AI Patent Forward Cites	Avg AI Forward Cites
	(1)	(2)	(3)
AI Policy	-0.064*** (-3.70)	-0.197** (-2.25)	-0.157** (-2.06)
ln(GDP per capita)	1.057*** (6.22)	4.601*** (9.35)	3.776*** (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 B5****Robustness: National AI-Related Policy's and AI Innovation (Dynamic Treatment)**

This table reports results mirroring the setup of the baseline regression in [Table 2](#), but replacing the indicator treatment with the Ln(one plus AI policy budget).

**Panel A: AI Innovation Novelty**

	OLS	Poisson
	Novel AI Patents Share(%)	Novel AI Patent Counts
	(1)	(2)
ln(1+Total Budget)	0.371* (1.78)	0.006* (1.73)
Year × Cohort FE	Yes	Yes
Inventor × Cohort FE	Yes	Yes
Obs.	3,900,297	2,461,111
R <sup>2</sup>	0.44	

**Panel B: AI Innovation Quantity and Quality**

	Poisson	Poisson	Poisson
	AI Patent Counts	AI Patent Forward Cites	Avg AI Forward Cites
	(1)	(2)	(3)
ln(1+Total Budget)	-0.013** (-2.56)	-0.074* (-1.71)	-0.058 (-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 B6****Robustness: National AI-Related Policy Indicator and AI Innovation (OLS)**

This table reports results mirroring the setup of the baseline regression in [Table 2](#), but employs OLS regression with Ln(one plus count data) as the dependent variable.

**Panel A: AI Innovation Novelty**

	OLS Novel AI Patents Share (%) (1)	OLS ln(1+Novel AI Patent Counts) (2)
AI Policy	2.682** (2.66)	0.008* (1.89)
Year × Cohort FE	Yes	Yes
Inventor × Cohort FE	Yes	Yes
Obs.	3,900,297	3,900,297
R <sup>2</sup>	0.44	0.44

**Panel B: AI Innovation Quantity and Quality**

	OLS ln(AI Patent Counts) (1)	OLS ln(1+ AI Patent Forward Cites) (2)	OLS ln(1+Avg AI Forward Cites) (3)
AI Policy	-0.066** (-2.65)	-0.402** (-2.23)	-0.330** (-2.25)
Year × Cohort FE	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes
Obs.	3,900,297	3,900,297	3,900,297
R <sup>2</sup>	0.54	0.56	0.56

**Table B7****Robustness: National AI-Related Policy and AI Innovation (CSDID)**

This table reports estimates using the doubly robust difference-in-differences estimator proposed by [Callaway and Sant'Anna \(2021\)](#). The model specifications follow the baseline approach but are estimated on the inventor-year sample without stacking. The dependent variables are defined as the log-transformed counts of novel AI patents, total AI patent filings, citation-weighted counts, and average forward citations. Standard errors are robust and clustered at the country level; *t*-statistics are reported in parentheses. \*\*\*, \*\*, and \* represent statistical significance at the 1%, 5%, and 10% levels, respectively. The sample period spans 2006 to 2019.

**Panel A: AI Innovation Novelty**

	Novel AI Patents Share (%) (1)	ln(1+Novel AI Patent Counts) (2)
AI Policy	1.032* (1.64)	0.012* (1.94)
Obs.	695,738	695,738

**Panel B: AI Innovation Quantity and Quality**

	ln(AI Patent Counts) (1)	ln(1+ AI Patent Forward Cites) (2)	ln(1+ Avg AI Forward Cites) (3)
AI Policy	-0.000 (-0.01)	-0.176*** (-2.81)	-0.186*** (-3.52)
Obs.	695,738	695,738	695,738

**Table B8****Placebo Test: National AI-Related Policy and Non-AI Innovation**

This table reports results mirroring the setup of Panel B of [Table 2](#), but using a sample of non-AI patents instead of AI patents.

	Poisson Non-AI Patent Counts (1)	Poisson Non-AI Patent Forward Cites (2)	Poisson Avg Non-AI Patent Forward Cites (3)
AI Policy	-0.097 (-1.62)	-0.405 (-1.49)	-0.453 (-1.64)
Year × Cohort FE	Yes	Yes	Yes
Inventor × Cohort FE	Yes	Yes	Yes
Obs.	21,796,942	19,519,519	19,519,519

**Table B9**  
**Robustness: National AI-Related Policy and AI Military Patent**

This table reports results that follow the specifications of the baseline regression in Table 2, but replaces the dependent variables with AI military patent counts, AI military patent forward citations, and 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), which includes firearms (F41A-C, F41F-G, F41H-J), ammunition, missiles, fuzes, and blasting technologies (F42B-D); B (Military Transport & Vehicles), covering warships and submarines (B63G), military aircraft (B64C-D), and specialized military vehicles (B60P); G (Sensors & Control), encompassing radar, navigation, and 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).

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

**Table B10**

**Cross-Country Effects of AI Policy: Novel AI Patent Share and Patent Quantity**

This table reports the underlying data for Figure 6. The empirical tests follow the baseline specification but are estimated separately for individual treated countries, utilizing not-yet-treated and never-treated countries as controls. In place of the baseline’s country-clustered standard errors, we employ the wild bootstrap for inference because individual country-level tests provide too few clusters for standard cluster-robust inference (Canay et al. 2021).

Country	Novel Patent Share		Patent Count		Policy Name	Year
	Coef.	t-stat	Coef.	t-stat		
<b>Panel A: Positive Impact on Novel Patent Share</b>						
<b>Asia</b>						
South Korea	2.193***	(4.14)	-0.067***	(-5.13)	<i>Korea AI Innovation Hub</i>	2017
Japan	6.514***	(22.47)	-0.204***	(-3.57)	<i>HPCI Project</i>	2012
China	0.389	(1.04)	0.119***	(4.40)	<i>“Internet +” AI Three-Year Action Plan</i>	2016
Turkey	2.313	(0.28)	-0.013	(-0.29)	<i>Cloud Computing and Big Data Lab</i>	2013
Russia	4.030**	(2.51)	-0.109***	(-6.58)	<i>National Technology Initiative</i>	2015
<b>Europe</b>						
Sweden	2.970	(1.47)	-0.104***	(-4.33)	<i>AI Competence for Sweden</i>	2018
Norway	1.480	(0.34)	-0.006	(-0.52)	<i>Testing Self-Driving Vehicles Act</i>	2017
Italy	2.496	(0.96)	-0.014	(-0.60)	<i>Italian AI Strategy Report</i>	2018
<b>Americas</b>						
Argentina	18.059	(1.17)	-0.142***	(-5.90)	<i>PNIA</i>	2018
United States	2.281***	(8.18)	-0.109**	(-2.01)	<i>National Robotics Initiative</i>	2011
<b>Oceania</b>						
Australia	2.144*	(1.77)	-0.105***	(-10.52)	<i>Medical Research Future Fund</i>	2015
<b>Panel B: Negative Impact on Novel Patent Share</b>						
<b>Asia</b>						
India	-6.109***	(-6.22)	0.021	(0.93)	<i>CoE-IoT &amp; AI</i>	2015
Singapore	-0.103	(-0.03)	0.016	(1.22)	<i>AI Singapore</i>	2017
<b>Europe</b>						
Netherlands	-1.549	(-1.05)	0.087***	(6.82)	<i>KIA Digitalisation</i>	2016
Finland	-1.704	(-0.83)	0.010	(0.76)	<i>Finland’s AI Programme</i>	2017
United Kingdom	-1.008	(-1.03)	-0.154***	(-7.57)	<i>Alan Turing Institute</i>	2015
Germany	-0.478	(-0.67)	-0.095***	(-3.78)	<i>Strategy for Automated and Connected Driving</i>	2015
Czech Republic	-3.619	(-0.60)	-0.093***	(-4.40)	<i>Digital Czech Republic</i>	2018
Switzerland	-4.089*	(-1.91)	-0.081***	(-5.44)	<i>IDWG AI</i>	2017
Spain	-2.453	(-0.79)	-0.067***	(-7.55)	<i>Language Technology Plan</i>	2016
France	-1.131	(-0.80)	-0.056***	(-4.02)	<i>France IA</i>	2017
Austria	-0.798	(-0.29)	-0.030***	(-3.04)	<i>ACRAI</i>	2017
<b>Americas</b>						
Brazil	-6.042	(-1.41)	0.225***	(10.47)	<i>Brazilian Strategy for Digital Transformation</i>	2018
Mexico	-3.062	(-0.42)	-0.186***	(-8.20)	<i>National AI Agenda</i>	2018
Canada	-2.026*	(-1.75)	-0.079***	(-5.82)	<i>Directive on Automated Decision-Making</i>	2016

**Table B11****National AI-Related Policy and Aggregate AI Innovation (Aggregate Level)**

This table reports a stacked DID analysis of AI policy on AI innovation at the country-year level and the corresponding pretrend analysis. Panel A reports results regressing AI innovation on an indicator for national AI-related policies that turned on starting in the year of the first AI policy. Panel B presents the pretrend analysis for Panel A. Robust standard errors are clustered at the country level, and t-statistics are reported in parentheses. \*\*\*, \*\*, and \* represent significance levels at 1%, 5%, and 10%, respectively. The sample period spans 2006 to 2019.

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

	Poisson AI Patent Counts (1)	Poisson AI Forward Cites (2)	Poisson Inventor Counts (3)
AI Policy	-0.267** (-2.29)	-0.434* (-1.67)	-0.292*** (-4.41)
Year × Cohort FE	Yes	Yes	Yes
Country × Cohort FE	Yes	Yes	Yes
Obs.	2,712	2,712	2,712

**Panel B — Pretrend Analysis**

	Poisson AI Patent Counts (1)	Poisson AI Forward Cites (2)	Poisson Inventor Counts (3)
t ≤ -3	0.058 (0.95)	-0.005 (-0.04)	0.091 (1.45)
t = -2	0.052 (1.60)	0.028 (0.44)	0.060** (2.24)
t = 0	-0.077 (-1.52)	-0.134 (-1.16)	-0.092*** (-2.87)
t = 1	-0.206* (-1.93)	-0.404** (-2.10)	-0.193*** (-3.61)
t = 2	-0.322** (-2.06)	-0.525* (-1.82)	-0.282*** (-3.36)
t ≥ -3	-0.360** (-2.27)	-0.628* (-1.87)	-0.434*** (-4.42)
Year × Cohort FE	Yes	Yes	Yes
Country × Cohort FE	Yes	Yes	Yes
Obs.	2,712	2,712	2,712

## C Conceptual 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 policies 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 is measured by novelty), which offer the potential for significant gains but carry inherent risks, and incremental innovations, which provide safer, more modest returns.<sup>26</sup> Our model incorporates inventors, consumers, firms, and government intervention, to examine how targeted policies shape innovation 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 in innovation activity. 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 of innovation activity between incremental and radical projects. 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.

---

<sup>26</sup>In fact, this way of modeling innovation is not new and has been heavily used in policy and innovation studies (e.g. [Aghion et al. 2023](#), [Atkeson and Burstein 2019](#)).

## C.1 Model Setup

### C.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 represent 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:

$$Y_t = A_t L_Y \tag{C.1}$$

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:

$$w_t = A_t \tag{C.2}$$

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  $\epsilon$  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  $i$  is characterized by a project-specific probability of success,  $p_i$ , when undertaking a radical innovation. This parameter captures heterogeneity in the feasibility of high-risk projects rather than differences in researcher quality per se. We assume  $p_i$  is drawn from a known, continuous distribution  $F(p)$  with support on  $[0, \bar{p}]$ , where  $\bar{p} < 1$ .

*Firms:* A sector of perfectly competitive firms produces a single homogeneous final good,  $Y_t$ . Production requires two inputs: labor,  $L_t$ , and the economy's aggregate stock of technology,  $A_t$ . We assume a linear production function,  $Y_t = A_t L_t$ . 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 C.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.

### C.1.2 Additive Technological Progress

While the model is static, the equilibrium allocation of innovation activity 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  $\delta$  be the small, certain productivity gain from an incremental success, and let  $\gamma$  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:

$$E[A_{t+1}] = A_t + (1 - \psi^*) \cdot \delta + \psi^* \cdot E[p_i | p_i \geq p^*] \cdot \gamma \quad (\text{C.3})$$

This equation makes the government's long-term objective explicit. By increasing the subsidy  $G$ , the government raises  $\psi^*$ , the fraction of inventors who pursue the radical path (as defined in the next section in [Equation C.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.

## C.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.

### C.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:

$$p_i \cdot V_R - c + G \geq w_t \quad (\text{C.4})$$

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:

$$p^* \cdot V_R - c + G = w_t \quad (\text{C.5})$$

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:

$$p^* = \frac{w_t + c - G}{V_R} \quad (\text{C.6})$$

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.

## C.2.2 Equilibrium Allocation of Innovation Activity

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

$$\psi^* = \int_{p^*}^{\bar{p}} f(p) dp = 1 - F(p^*) \quad (\text{C.7})$$

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  $w_t = A_t$ .

## C.2.3 Comparative Statics

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

$$\frac{\partial p^*}{\partial G} = -\frac{1}{V_R} < 0 \quad (\text{C.8})$$

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.

## C.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:

$$E[\text{Quantity}] = (1 - \psi^*) \cdot 1 + \psi^* \cdot E[p_i | p_i \geq p^*] \quad (\text{C.9})$$

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  $\psi^*$  inventors pursuing radical innovation. When policy  $G$  increases,  $\psi^*$  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:

$$E[\text{Quality-weighted Output}] = (1 - \psi^*) \cdot 1 + \psi^* \cdot E[p_i | p_i \geq p^*] \times R \quad (\text{C.10})$$

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.

### C.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 C.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 ( $\gamma$ ) is substantially larger than that from an incremental one ( $\epsilon$ ), the expected gain from a radical attempt ( $p_i \cdot \gamma$ ) can far exceed the certain gain from an incremental project ( $\epsilon$ ), even for projects with a modest probability of success. By inducing more inventors to pursue radical innovation (increasing  $\phi^*$ ), 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.

## C.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:

$$V_S = V_R + V_{Spillover} \quad (C.11)$$

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:

$$p^* \cdot V_R - c \geq w_t - G \quad (C.12)$$

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:

$$p^* \cdot V_S - c \geq w_t \quad (C.13)$$

Substituting the definition of  $V_S$ , we get:

$$p^* \cdot (V_R + V_{Spillover}) - c \geq w_t \quad (C.14)$$

A comparison of the private and social investment criteria highlights 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:

$$p^* \cdot V_R - c < \epsilon \text{ and } p^* \cdot (V_R + V_{Spillover}) - c \geq w_t \quad (\text{C.15})$$

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:

$$G^* = p^* \cdot (V_S - V_R) = p^* \cdot V_{Spillover} \quad (\text{C.16})$$

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](#)).

#### **C.4 Model Extension 2: Imperfect Government Screening and Project Allocation**

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.

### C.4.1 The Screening Mechanism

We assume the government classifies projects as eligible for funding if their underlying success probability exceeds a threshold  $p_H$ . Screening is imperfect and subject to two errors:

*Type I Error (False Negative)*: Projects with  $p_i \geq p_H$  are denied funding with probability  $\alpha$ .

*Type II Error (False Positive)*: Projects with  $p_i < p_H$  are mistakenly funded with probability  $\beta$ .

The parameters  $\alpha$  and  $\beta$  capture screening imprecision; lower values correspond to more accurate project evaluation.

### C.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 project success probabilities and grant status. The decision to proceed depends on two equilibrium cutoffs: one for inventors who receive the grant, denoted  $p_G^*$ , and a higher cutoff for inventors who do not receive funding, denoted  $p_N^*$ . We refer to these as the cutoff with grant ( $p_G^*$ ) and the cutoff without grant ( $p_N^*$ ), respectively.

The total fraction of inventors pursuing radical innovation,  $\psi_{\text{Radical}}$ , is the sum of three groups:

(1) *Funded, Eligible Projects* ( $S_1$ ): Inventors whose projects satisfy the eligibility threshold ( $p_i \geq p_H$ ) and who receive the grant (with probability  $1 - \alpha$ ) choose the radical path. Assuming  $p_H > p_G^*$ , all such projects proceed. The size of this group is

$$S_1 = (1 - \alpha) \int_{p_H}^{\bar{p}} f(p) dp. \quad (\text{C.17})$$

(2) *Unfunded, Eligible Projects* ( $S_2$ ): Inventors whose projects satisfy the eligibility threshold ( $p_i \geq p_H$ ) but are denied funding (with probability  $\alpha$ ) may nevertheless proceed if their success probability exceeds the unfunded cutoff ( $p_i \geq p_N^*$ ). The size of this group is

$$S_2 = \alpha \int_{p_N^*}^{\bar{p}} f(p) dp. \quad (\text{C.18})$$

(3) *Funded, Ineligible Projects* ( $S_3$ ): Inventors whose projects fall below the eligibility threshold ( $p_i < p_H$ ) but are mistakenly awarded a grant (with probability  $\beta$ ) may proceed if their success probability exceeds the subsidized cutoff ( $p_i \geq p_G^*$ ). The size of this group is

$$S_3 = \beta \int_{p_G^*}^{p_H} f(p) dp. \quad (\text{C.19})$$

The total pool of radical innovation is therefore given by  $\psi_{\text{Radical}} = S_1 + S_2 + S_3$ .

An increase in  $\alpha$  (a higher false-negative rate) reduces the number of funded projects that meet the eligibility threshold. As a result, some projects with sufficiently high success probabilities are deterred from pursuing radical innovation and instead revert to incremental activity, leading to a misallocation across innovation activities.

An increase in  $\beta$  (a higher false-positive rate) expands the set of subsidized projects with relatively low success probabilities. This introduces additional low-expected-return experimentation into the radical R&D portfolio, resulting in a misallocation of public funds.

Consequently, increases in either  $\alpha$  or  $\beta$  can contribute to a decline in average observed patent quality by shifting the composition of radical innovation toward projects with lower expected success rates.