The Indian Innovation Puzzle

Namrata Kala, Josh Lerner, and Junxi Liu *

September 9, 2026

Preliminary and Incomplete

Abstract

Despite recent increases in some measures of innovation, India continues to have very modest levels of domestic patenting, contradicting standard views of the importance of property rights in economic development. Using a novel, newly constructed dataset of Indian patent grants, we explore why Indian patenting has been so low. We examine whether an absence of non-patented innovation, differences in industry composition and public attitudes, and a lack of public market reactions to patent awards can explain these differences, and find only modest support for these explanations. We then show evidence of an institutional quality mechanism, i.e., long delays in patent examination. We estimate the causal effects of reduced delays using a natural experiment. Specifically, we show that a 2016 reform by IP India (the patent office) increased examination speeds for technology-office pairs that gained access to more same-field examiner capacity. In response, the mixture of applicants shifted away from applicants with very low success rates and towards high-quality filers. This improved institutional capacity also enhanced other innovation outcomes, boosting the financing of venture capital-backed firms. But the magnitude of these effects was modest, closing just over one-tenth the patenting gap between India and other leading nations. This modest effect may reflect the implementation of the reform (a back-of-the-envelope calculation suggests that high-quality patenting could have been increased by an additional 9 to 31 percent), as well as the uncertainty surrounding patent litigation. These results underscore the importance of institutional quality in facilitating innovation.

*MIT Sloan School of Management; Harvard University; University of Warwick. Kala and Lerner are affiliates of the National Bureau of Economic Research. We thank Amita Arora, Neeraj Bharadwaj, Rajat Kochhar (discussant), Bo Li (discussant), Franziska Ohnsorge (disucssant), Prof. Unnat P. Pandit, Arul George Scaria, Sacha Wunsch-Vincent, and participants at the 14th AIEA-NBER conference and the India Policy Forum for their guidance regarding the analysis, data sources, and institutional details. We thank the Harvard Business School Division of Research and Doctoral Programs and the MIT Sloan School of Management for financial support. Brooke Jones, Yongjang (Kyle) Seo, Shambhavi Sinha, Zhaoying (Laura) Wang, Zhoutao (Andy) Zhang, and Austin Zhu provided very valuable research assistance in this effort. Lerner has received compensation for advising limited partners in venture funds, venture capital groups, and governments designing policies relevant to venture capital. All errors and omissions are our own.

1 Introduction

Historically, a very small set of high-income countries have undertaken the bulk of investment in innovation, resulting in large global differences in income and productivity (Keller, 2004). One widely acknowledged impediment to knowledge accumulation and innovation in many developing nations has been the absence of strong intellectual property rights, especially patents, to incentivize firms to invest in the development and commercialization of novel products and services. Economists have long argued (e.g., North (1990), Romer (1990), and Landes and Posner (2003)) that intellectual property rights promote growth in emerging countries by preventing free riding on innovative efforts, thus encouraging private sector investment in R&D. Lai (1998) and Branstetter et al. (2006) provide careful empirical corroboration of these claims.

At the same time, patent systems are imperfect mechanisms to promote innovation. Both the design (for instance, Scotchmer (2004) and Budish et al. (2015)) and the implementation of patent systems (e.g., Feng and Jaravel (2020) and Freilich and Kim (2026); for an overview, see Scott Morton and Shapiro (2016)) can lead to distortions. While these issues have been extensively scrutinized in the U.S. context, there has been little prior exploration of the impact of patent system effectiveness on domestic innovation in emerging economies. Importantly, it is not ex-ante obvious whether the impacts of mitigating constraints to patenting would be larger or smaller in developing economies. For instance, since such constraints may be more likely to bind, facilitating patenting could increase innovation. Conversely, if multiple frictions impede innovation, then policies to facilitate patenting by improving the quality of patent granting institutions may not increase aggregate innovation. Yet studying these questions in emerging markets is challenging, given the lack of long-term patenting data and paucity of causal evidence.

This study uses novel data on the universe of patents and a natural experiment in India, one of the world's largest economies and one with a significant policy focus on innovation. The importance of greater innovation has been articulated in numerous speeches by Prime Minister Modi and other leaders. This policy priority has also manifested itself in large public spending on facilities such as the Micron semiconductor plant in Gujarat.[1] Despite this policy focus, India has lagged substantially behind its peers in patenting activity. Figure 1(a) plots the number of patents filed by domestic residents as a share of GDP for a

[1] Examples include https://dst.gov.in/pm-launches-rs-1-lakh-crore-research-development-and-innovation-scheme, https://www.pmindia.gov.in/en/news_updates/pm-addresses-the-emerging-science-technology-and-innovation-conclave-2025/, and https://www.pmindia.gov.in/en/news_updates/india-should-be-among-the-top-three-ai-superpowers-globally-pm-modi-sets-2047-vision/.

range of countries over the 2005 to 2023 period. India's domestic patenting rate is only a tenth that of China and substantially below the rates observed in advanced economies. While India is similar or better than some developing countries when normalizing by GDP, Figure 1(b) normalizes patenting by the number of STEM graduates. On this measure, India is among the very lowest of the major nations; only Mexico ranks comparably among large economies. Figure 1(c) presents a similar picture when normalizing by the number of STEM PhDs, reinforcing that the low patenting rate is particularly striking given the level of human capital that India possesses.

Moreover, when we look at the Indian organizations filing the most patents, the picture is puzzling as well. Table 1 lists the top 20 Indian patent applicants by total applications filed between 2005 and 2022, compiled from our own disambiguated dataset combining granted patents and non-granted applications. The list features on the one hand large firms and public research bodies and on the other, a set of small-to-mid-size private universities. The contrast in *realized grant rates* is stark: established applicants such as the Indian Institutes of Technology (collectively 84.2%), TVS Motors (89.3%), and Samsung's India R&D center (98.7%) convert the large majority of their filings into grants, whereas the five private universities among the top twenty applicants only have a success rate of 10.9% in their patent filings. The inclusion of these private universities among the top applicants is surprising, as none of the five appear on the 2025 Nature Index of the 500 leading global academic research institutions. Moreover, Lovely Professional University and Chandigarh University have been identified as among the four universities worldwide with the most problematic research practices (Meho and Akl, 2025).[2]

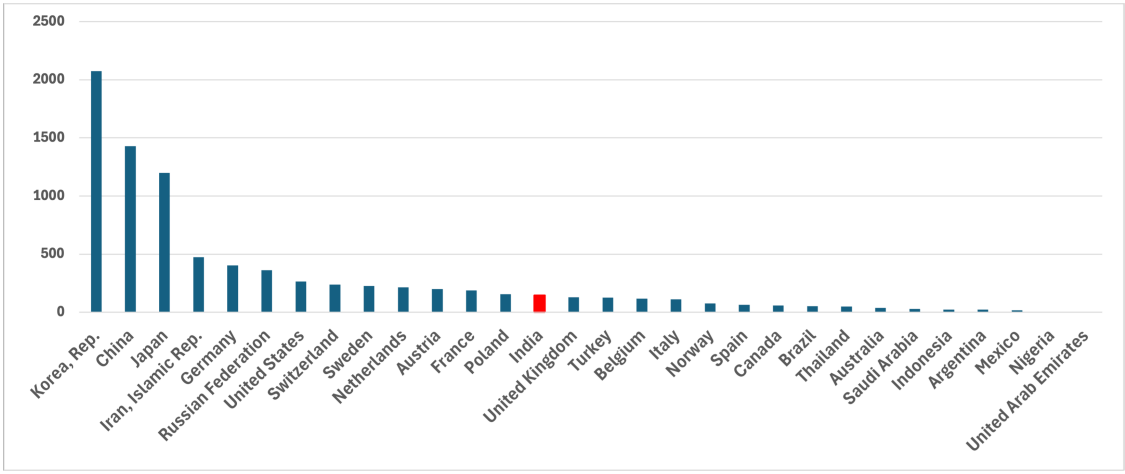
Numerous media accounts highlight the substantial incentives that the National Institutional Ranking Framework and National Assessment and Accreditation Council accreditation provide for private Indian universities to file patents.[3] As noted above, very few of these institutions' applications are ever granted, which may reflect both their quality and the presence of incentive schemes based on applications and not awards. This pattern suggests that a substantial share of recent Indian domestic applications comes from

[2]The Nature Index is at https://www.nature.com/nature-index/research-leaders/2025/institution/academic/all/global and is based on top-tier journal publications. Problematic research universities are identified by Meho and Akl (2025) using indicators such as decreases in the share of first-author listings and increases in the count of "hyper-prolific" authors (those with implausible large number of publications). Another high-filing school, Galgotias University, was the center of a high-profile scandal involving a misattributed robot: see https://www.bbc.com/news/articles/cge8nd5ve00o.

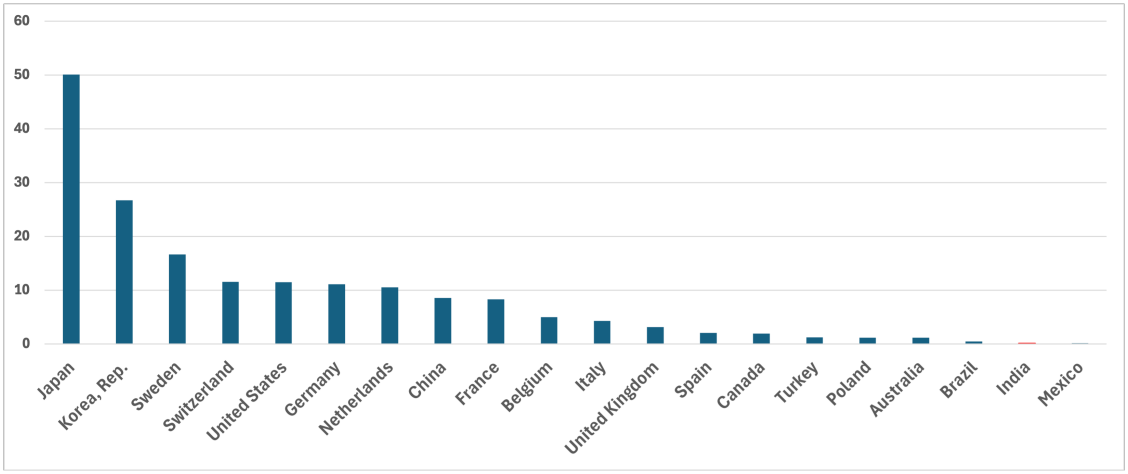
[3]See, for instance, https://www.msn.com/en-in/news/India/the-indian-patent-story-why-private-universities-file-more-patents-than-iits-combined/ar-AA1WQs3k#:~:text=That%20last%20line%20is%20critical,of%20parameters%2C%22%20he%20says.&text=Achal's%20comparison%20of%20the%20NIRF,strong%20allegation:%20gaming%20the%20system.

Figure 1: Indian Patenting Is Low: International Comparisons

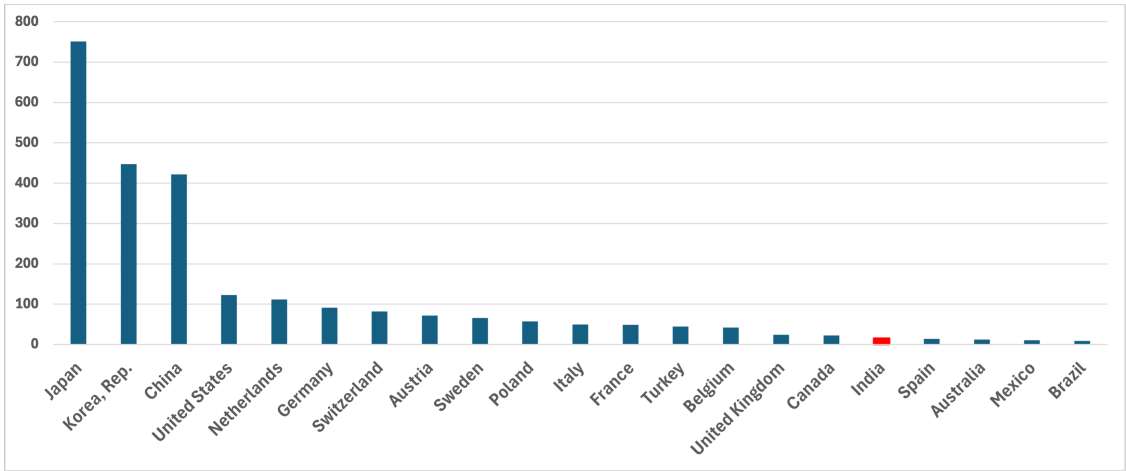
(a) Patents by Domestic Residents as a Share of GDP



(b) Patents per STEM Graduate



(c) Patents per STEM PhD



Note: Panel (a) plots the number of patent filings by domestic residents as a share of GDP (constant 2015 USD in billions) over the 2005–2023 period. Panels (b) and (c) normalize domestic patent filings by, respectively, the number of STEM graduates and STEM PhDs awarded. Patent counts are from WIPO IP Statistics; STEM graduate and PhD data are from the UNESCO Institute for Statistics.

Table 1: Top Indian Patent Applicants and Their Grant Rates

Rank by Apps	Applicant	Apps	Granted	Grant %
1	Indian Institute of Technology (Collectively)	5,158	4,342	**84.2**
2	Council of Scientific and Industrial Research (CSIR)	3,349	2,313	69.1
3	Tata Consultancy Services Limited	2,224	1,647	74.1
4	Bharat Heavy Electricals Limited	2,202	1,833	**83.2**
5	TVS Motor Company Limited	1,926	1,719	**89.3**
6	Chandigarh University	1,902	182	**9.6**
7	Lovely Professional University	1,563	144	**9.2**
8	Bosch Limited (India)	1,483	120	**8.1**
9	Defence Research and Development Organisation (DRDO)	1,396	1,149	**82.3**
10	Amity University	1,361	349	**25.6**
11	Mahindra & Mahindra Limited	1,350	1,110	**82.2**
12	Tata Motors Limited	1,230	1,041	**84.6**
13	Tata Steel Limited	1,222	1,077	**88.1**
14	Samsung R & D Institute India-Bangalore Private Limited	1,153	1,138	**98.7**
15	Wipro Limited	983	725	73.8
16	Larsen & Toubro Limited	938	401	42.8
17	Chandigarh Group Of Colleges	894	65	**7.3**
18	Hindustan Unilever Limited	861	639	74.2
19	Maruti Suzuki India Limited	847	631	74.5
20	Sanskriti University	810	22	**2.7**

Note: The unit of observation is an Indian patent applicant. This table reports the top 20 Indian lead applicants by total applications filed over application years 2005–2022, compiled from our disambiguated dataset combining granted patents and non-granted applications. Grant status as of June 2024. *Apps* is the number of applications, *Granted* the number subsequently granted, and *Grant %* their ratio. Indian lead applicants are identified from the first applicant listed on each application; applicant names are normalized and obvious institutional variants (such as the IITs) are aggregated into a single row. Cells are shaded by grant rate (green high, red low).

institutions for which patenting is plausibly incentive-driven rather than a reflection of underlying R&D. The contrast with realized grants is instructive: the leading domestic recipients of granted patents are dominated by established firms and public research organizations, such as the Council of Scientific and Industrial Research and the Tata group companies (Appendix Table A.1 lists the top twenty domestic grantees by grant year; Appendix Table A.2 reports the same ranking by the application year of granted patents.) Thus, to assess how institutional reforms affect innovation requires understanding both the impacts on average as well as the composition of high vs. low quality patenting.

India's innovative growth at a time when patenting is seemingly of peripheral importance presents a substantial puzzle. Certainly, many of the past periods of rapid acceleration of innovation, such as those in Great Britain, the U.S., Japan, Korea, and China, have been associated with an increasingly important role of patents, even though the precise institutional structures and use of awards have differed. In this paper, we seek to

understand this apparent contradiction. To do so, we draw on an Indian patent dataset that we constructed in Kala, Lerner, and Liu (2025).[4] Section 2 briefly describes how we construct our patent dataset.

In Section 3, we seek to understand the low level of patenting. We examine in turn five explanations.

We first consider the possibility that the paucity of patenting reflects a failure of India to pursue innovation. In contradiction to this claim, we find that India has experienced a rise in innovative spending in recent years, driven both by local firms, government spending, and foreign expenditures. More persuasively, Indian innovation has climbed sharply, as illustrated by metrics as diverse as venture capital investment, scientific publications, and GitHub repository commitments.

A second explanation is industry mix. A growing literature suggests that the mixture of Indian economy is quite different from China and the U.S., with a much greater emphasis on services (e.g., Fan et al. (2023), Peters et al. (2026)). We show that were the mixture of Indian industry to mirror that of the U.S. and China, the volume of patenting would only be 12% and 47% higher, respectively. Thus, the industry mixture does not fully explain this discrepancy seen in Figure 1.

A third explanation may be popular attitudes. As we discuss below, patents in India have had a fraught history, having been used during the British Raj as a mechanism to protect United Kingdom-based exporters from local competition. We test whether negative attitudes toward intellectual property are larger in India than in other countries. In contradiction to this hypothesis, surveys of consumer attitudes suggest that Indian attitudes toward intellectual property are similar to other countries. Moreover, the temporal changes in attitudes towards intellectual property bear little relationship with changes in patenting.

Fourth, we examine the returns to patenting to estimate to what extent patents, once issued, are rewarded by the public market. Were these awards not to affect market valuations, it might be natural to conclude that they are not particularly helpful in creating additional revenues and profits, whether due to low grant quality or an inability to enforce awards. We link the patent records to the Prowess database,[5] and examine the market

[4]We constructed this data using the methodology pursued with the U.S. patent data since the inception of the first patent project at National Bureau of Economic Research (Griliches, 1984). It emulates many of the key steps undertaken in the literature using the U.S. Patent and Trademark Office (USPTO) data, such as the compilation of citations (e.g., Jaffe and Trajtenberg (2002)), the calculation of measures of patent importance (Kelly et al., 2021), the analysis of patent text (Myers and Lanahan (2022); Kalyani et al. (2025)), the construction of international patent families (Putnam, 1997), the disambiguation of assignees (e.g., Jaffe and Trajtenberg (2002); Magerman et al. (2006), and the identification of venture-backed patents (Bernstein et al., 2016).

[5]The linkage combines exact and fuzzy name matching with extensive manual review; every business grantee with more than ten granted patents is individually verified, and the matched firms account for 85.5%

reaction to issued patents following Kogan et al. (2017). The evidence suggests that issued patents are indeed associated with an increase in market value, especially awards to firms who had an established record of patenting. This analysis suggests that the returns to high-quality patents are high.

Finally, we examine the causal role of institutional frictions in the process through which patents are issued, and how they impact the quantity and quality of patenting. We show that Indian patent applications have been historically characterized by extremely long processing times, over eight years in 2016 and falling to 68 months (five and a half years) in 2023, which is several times longer than the US or China. These long delays might have detrimental effects on the value of patent protection for at least two reasons. First, patent protection worldwide is typically only available for twenty years from date of filing. Too short a period of effective protection, a consequence in this case of long delays, may discourage investments in innovation (Budish et al., 2015). Second, to the extent that patent grants facilitate the licensing of technology or the raising of outside funding (as shown, for instance, in Gans et al. (2008) and Farre-Mensa et al. (2020)), delays in grants may constrain technological progress, especially for younger and capital-constrained firms.

To help understand the impact of these delays on the quantity and quality of patenting, we exploit a natural experiment in Indian patent policy in 2016 that we introduce in Section 4. Before that year, the regional patent office responsible for the applicant's city would automatically process the application.[6] After the reform, IP India began reassigning patent awards in this year across the four Indian patent offices. This reform was explicitly motivated in the hopes of reducing the time until issuance.

We use a difference-in-differences (DID) approach to estimate the causal impacts of this reform, where the cross-sectional measure of treatment is how much "outside capacity" existed prior to the reform. For instance, for applicants filing for biotechnology patents in Chennai, the treatment constitutes the number of patent examiners with biotechnology expertise in offices other than Chennai. Since the reform happened in 2016 for everyone, we employ a non-staggered design.

In Section 5, we explore the consequences of this reform. We first show that the reform had a material impact on the length of time it took for a patent to receive its first examination report (FER), a key step in the review process. Applicants in technology field-by-office cells where the reform made available more examiners in the same field in other offices experienced faster application processing times after the reform. A one

of all granted patents whose first grantee is a business. For full documentation of the Prowess matching, see the data appendix of Kala et al. (2025).

[6]The Patents Act of 1970 established four patent offices in regions across the country, in contrast to most nations, which have a single office.

standard deviation additional unit of examiner capacity improved processing time by about 85 days, or 0.22 standard deviations of pre-reform processing time.

We similarly see that the change impacted the rate and quality of innovation. First, we show that the reform shifted who applied subsequently for awards. More specifically, applications in the treated field-by-office cells after the reform were disproportionately likely to be ones from the organizations with historically high success rates in filings. Meanwhile, these reforms do not appear to have incentivized a similar boost in applications by "unserious" applicants.

We also test whether changing patenting institutional quality impacted outcomes for other types of innovation. We focus specifically on venture capital, a form of financing associated with innovation and economic growth. To explore this question and minimize the effects of endogeniety, we explore the consequences of the policy reform on firms that already had filed for (ultimately successful) patents at the time of the awards.

We first show that those venture-backed firms in treated field-by-office cells were more likely to subsequently file patents, but only at statistically insignificant confidence levels. This result may reflect that filing for patent protection as an Indian startup was far rarer than in the U.S., as we show below. But the venture-backed firms that had filed patents prior to the reform, and thus were well-positioned to benefit from the change, were also more likely to receive another round of venture financing. The impacts are economically significant: a one standard deviation increase in outside examiner capacity raises a patent-holding venture-backed firm's annual probability of raising a round by 4.7 percentage points, as compared to a pre-reform mean of 22%. When we augment the sample with a considerably larger set of high-potential startups that were not venture backed at the time of the reform, the effects are similar.

While improving patent institution quality appears to have broader positive effects on Indian innovation, suggesting the relevance of such reforms to emerging economies more generally, the effect on aggregate patenting was relatively modest: we estimate that the reform only closed 11.4% of the gap in patenting to GDP depicted in Panel (a) of Figure 1. This estimate begs the question of whether the impact of the reform could have been greater had it been implemented differently. We explore two dimensions of this question in Section 6.

The first is whether the reform could have been more efficient if the reallocation process were adjusted. We show that while the reform significantly reduced processing delays, offices did not route applications in a manner that minimized queues after the reform. Undertaking a back-of-the-envelope calculation, we show that had patents been allocated in an optimal manner, the number of high-quality applications would have increased

further by between nine and thirty-one percent. This reform may have closed as much as 40% of the gap in the GDP-adjusted patenting rate, as opposed to the 11.4% that we estimate actually occurred. We argue that such a step would be a highly cost-effective way to improve innovation outcomes, since the reform would utilize existing capacity more efficiently.

Second, we show that for at least six years after the reform, patent litigation levels remained relatively modest, with under 100 cases filed each year. The litigation of patents, as we discuss below, was characterized by a high degree of policy uncertainty. It is not until the creation of the Intellectual Property Division at the Delhi High Court (announced in 2021 and implemented in 2022), which was followed by similar courts elsewhere in the nation, that much of the confusion was resolved. Suggestively, in the years since this reform, we have seen an acceleration of domestic patent applications. This correlation suggests the importance of institutional complementarity (e.g., Aoki (2001) and Acemoglu and Robinson (2012)) in this setting: the efficiency of the patent office depends not just on its own operations, but also the presence of an effective judiciary to adjudicate these disputes.

This work relates to several strands of the literature. The first is the substantial literature on institutions, technology adoption, and long-run growth, reviewed in Acemoglu et al. (2005) and Akcigit et al. (2026). Research examining the workings of intellectual property systems specifically, however, has focused on a small number of developed nations, especially the U.S. (a point emphasized by Matcham and Schankerman (2026)). A related literature to this one is the literature on state capacity (see Finan et al. (2017); Okunogbe and Tourek (2024) and Xu et al. (2023) for excellent reviews of this literature). Prior work has focused on court delays to estimate the impact of institutional delays on firm outcomes (Kondylis and Stein, 2023; Boehm and Oberfield, 2020; Amirapu, 2021; Rao, 2021; Chemin, 2020). We study how institutional capacity in patent offices impact innovation in a large developing country, and how a cost-effective reallocation of capacity increases high-quality patenting that more comprehensive reforms like rewarding patent filing find difficult to move (Long and Wang, 2016).

A second strand is the growing body of work on innovation in developing countries. Much work has focused on China (e.g., König et al. (2022), Colonnelli et al. (2024), and Bai et al. (2025)). But innovation in other developing countries is attracting increasing interest. A particular focus has been mechanisms to incentivize domestic innovation that meets local needs (a recent example is Moscona and Sastry (2025)). Finally, a literature examines the impact of the TRIPS accords in the early 2000s on pharmaceuticals in developing nations. These works focus largely on the changes in pricing and availability of drugs, the

considerable majority of which were originally developed in the West. Examples include Chaudhuri et al. (2006), Cockburn et al. (2016), and Duggan et al. (2016).

In addition to shedding light on important academic discussions, this work is relevant to current policy debates on innovation and development. The intense focus of Indian policymakers in recent years on promoting innovation—motivated in part by the success of China in this arena and the desire for more self-determination over emerging technologies such as artificial intelligence—is far from unique: many developing nations have adopted similar objectives. Like India, they are increasingly viewing their intellectual property system as a key element in reaching these goals (Commonwealth Secretariat and United Nations Conference on Trade and Development, 2024; World Intellectual Property Organization, 2024). The successes and challenges associated with the Indian patent reform process, and the implications regarding the complementarity of intellectual property institutions, thus have broader implications for modern development policy.

2 Data

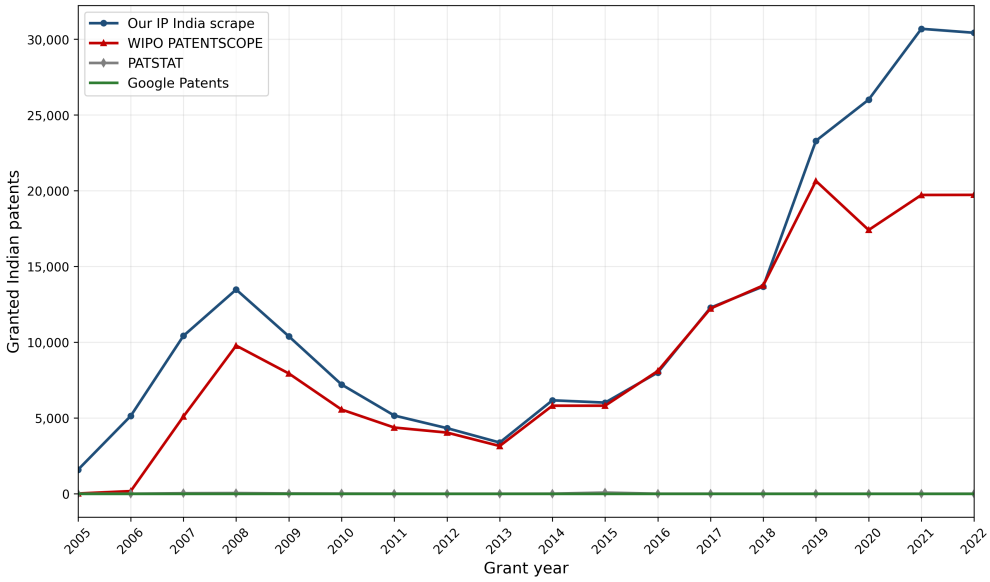
Our analysis draws on a new, patent-level dataset of Indian patenting that we describe in Kala et al. (2025). We summarize its construction here and provide additional detail in Appendix A.

2.1 Data Collection and Construction.

We scrape the IP India's official application pages and the documents linked to each patent, record by record, covering applications filed between 1995 and 2022. From the raw pages and attached PDFs, we parse applicants, inventors, legal status, first examination reports (FERs), backward citations, controller (examiner) information, and examination timing. We then geolocate applicants and grantees, disambiguate names into consistent entities, classify applicants and grantee types, and link applicants and grantees to firm-level data including Prowess and Pitchbook. Importantly, for examiner information, we carefully disambiguate their names to obtain a uniquely identifiable examiner set. In total we obtain 761,734 applications filed between 1995 and 2022, of which 343,618 have been granted as of June 2024.

Complete Specifications Data A distinctive feature of the data is the full text of each patent's complete specification, that is, the document giving the invention's technical disclosure. Most databases do not provide complete specifications for Indian grants, and

Figure 2: Coverage of Indian Patent Grants Across Data Sources

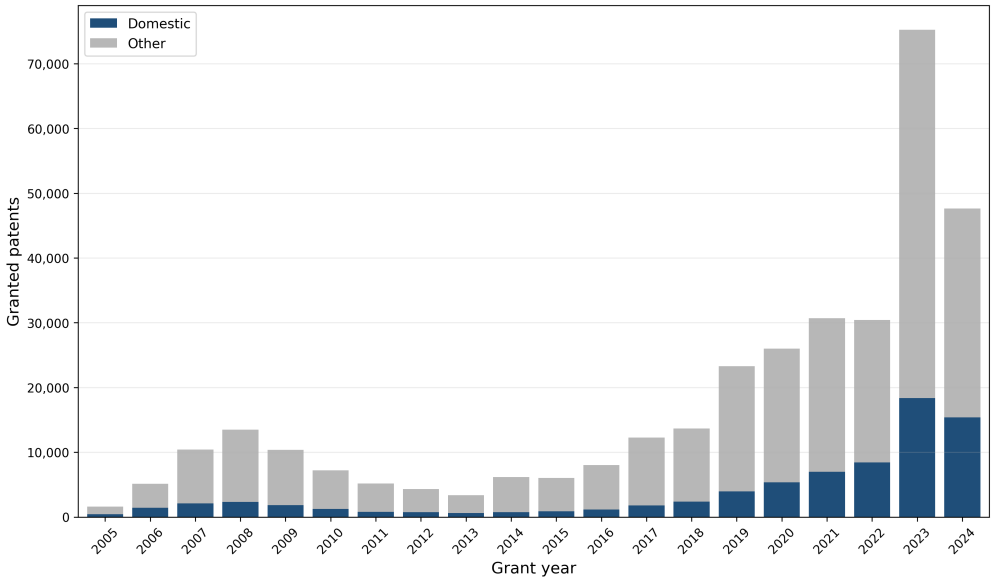


Note: This figure compares coverage of Indian patent grants across major international data sources and our own collection. WIPO reports only limited bibliographic information and lacks coverage at the start and end of the window; PATSTAT misses nearly all Indian grants; and Google Patents does not provide Indian grant documents.

even IP India's own web field is unusable for roughly 16% of grants. We supplement the web-scraped text with the downloaded specification PDFs, using OCR and large language models to reach about 95.5% coverage of granted patents filed between 1995 and 2022.

Why a new dataset is needed. We begin by showing that standard international patent databases cover Indian patenting poorly. Figure 2 compares the coverage of Indian grants across four sources and motivates the need for a purpose-built collection. The World Intellectual Property Organization's PATENTSCOPE database reports only limited bibliographic information and is incomplete at the start and end of our window- for instance, in 2007, it has about half the patent grants as in the IP India data that we collect, and in 2021, about 60% of the IP India data. Note that in addition to aggregate numbers being substantially lower, the gap between the two changes from year to year, which has implications for policy analyses. The European Patent Office's PATSTAT database (also the foundation for many commercially available databases) misses nearly all Indian grants. Finally, Google Patents does not provide Indian patent documents.

Figure 3: Domestic versus Foreign Indian Patent Grants, by Grant Year



Note: Indian patent grants by grant year, split into domestic and foreign grantees. A patent is classified as domestic if at least one grantee or inventor has an address in India. Domestic grants have risen over time but remain a modest share of all Indian grants relative to the levels seen in China or the United States.

2.2 Stylized Facts about Indian Patenting

Trends in Domestic Innovation Domestic patent filings and grants (i.e., where the grantee or inventor has an Indian address) have grown over time, though they remain a modest share of all Indian grants. Figure 3 splits granted patents into domestic and foreign grantees by grant year. The share of successful patent applications that are by domestic entities (22.6% between 2005 and 2024) has stayed low relative to the levels seen in China or the United States, which in recent years have been roughly 90% and 50% domestic in origin respectively.[7]

Citations as Measures of Innovation Quality A second stylized fact worth noting is that unlike in the U.S. or European settings, Indian patents are rarely cited in Indian filings. Only 0.58% of the 675,740 backward-citation entries we extract (3,947 entries) cite an Indian patent document, and only 1.40% of the 233,319 citing applications cite at least one Indian patent. It also underscores the importance of studying patent systems in developing countries to develop and use appropriate quality measures, which may be different from those in developed countries.

[7]The Chinese and U.S. estimates are from https://link.epo.org/ip52024stats and earlier reports in this series.

3 Explaining the Low Rate of Indian Patenting

We now turn to several possible explanations for why Indian patenting has been so low, despite substantial innovative activity in other dimensions.

3.1 India as an Innovator

One possibility is the low level of patenting reflects a dearth of overall innovation in India. There is substantial evidence, however, that innovative activity in India is vibrant and has grown substantially in recent decades, as we discuss here.

Increasing R&D spending. Indian R&D expenditures have climbed sharply in recent years, from 127 thousand crore rupees in 2020-21 to an estimated 313 thousand in 2025-26 (Department of Science and Technology, Ministry of Science and Technology, Government of India, 2026). The private sector share of research spending over the same period rose from 36% to 58%. As a share of GDP, however, India's R&D expenditures remain modest to the major industrialized nations (0.90% in 2025-26).

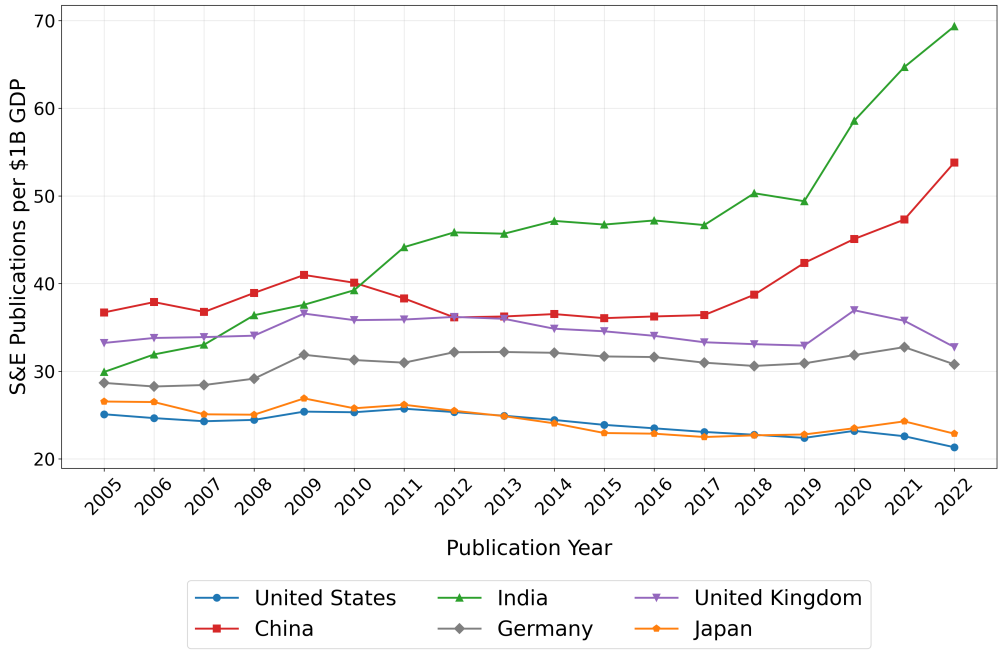
High and rising scientific publication output. Scientific publications are a well-established indicator of innovative capacity. Figure 4(a) shows trends in science and engineering publications for India and a range of comparator countries, normalized by GDP. India's output has grown substantially over the period from 2005 to 2023. Figure 4(b) further shows that the share of Indian publications in the global top 1% by citation impact has also risen, indicating that this growth reflects quality as well as quantity increases.

Venture capital investment is substantial and growing. One of the most direct indicators of private-sector innovation activity is venture capital financing, which targets firms pursuing novel, scalable ideas. Figure 5 plots the aggregate level of venture capital and GDP per capita for China, India, and the U.S., using all years (through 2022) for which data are available. It shows that India has experienced a marked increase in venture capital investment in recent decades. If we compare the GDP per capita when the three nations reached a given level of venture investment, we see India is shifted to the left, i.e., it has achieved a given level of venture activity while considerably poorer than its two peers.

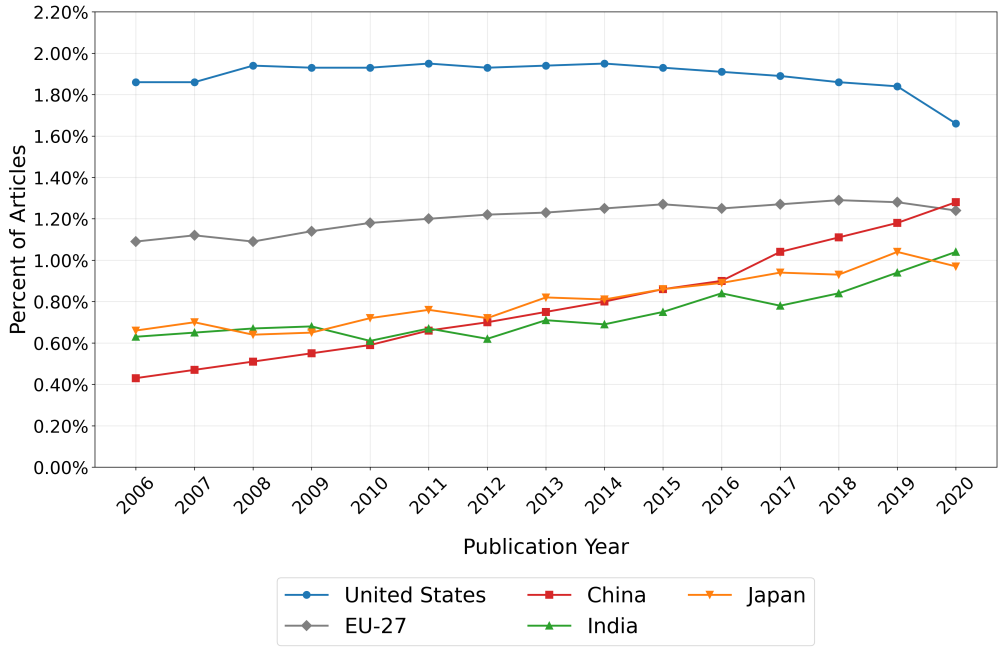
Active participation in global open-source development. Beyond traditional innovation metrics, a novel measure of innovative output comes from open-source software development. Participation in platforms such as GitHub captures technical creativity and

Figure 4: Science and Engineering Publications, Selected Countries

(a) S&E Publications per $1B GDP

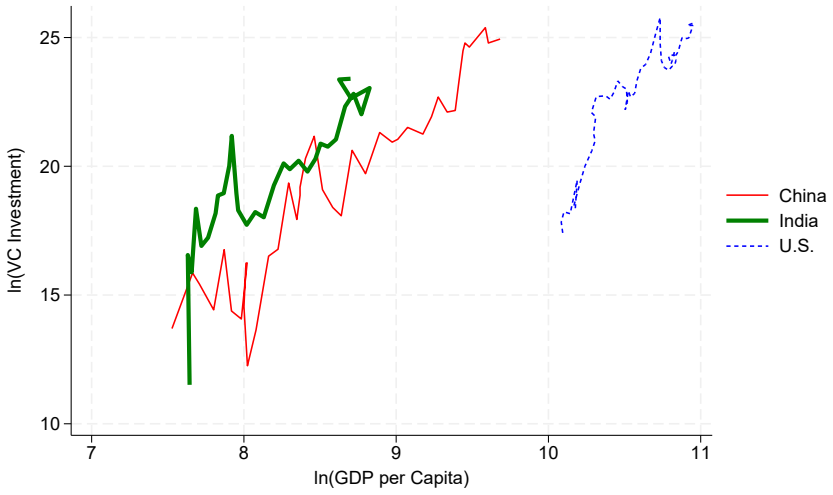


(b) S&E Publications in Top 1% by Citation Impact



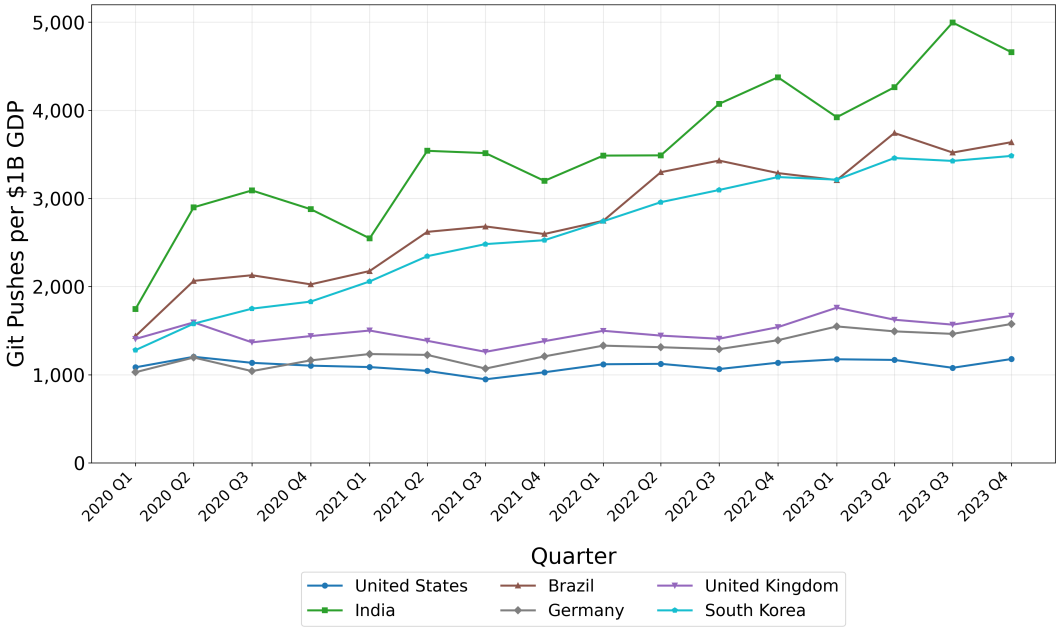
Note: Panel (a) plots science and engineering publications per $1 billion GDP (constant 2015 US dollars) for the United States, China, India, Germany, the United Kingdom, and Japan. Publication counts come from NSF Science and Engineering Indicators Figure PBS-3 (https://ncses.nsf.gov/pubs/nsb202333/figure/PBS-3). GDP values come from World Bank GDP series. Panel (b) plots the share of all S&E journal articles that fall in the global top 1% most-cited articles for the United States, the EU-27, China, India, and Japan, from 2006 to 2020, using NSF Science and Engineering Indicators data on the impact of published research (https://ncses.nsf.gov/pubs/nsb202333/impact-of-published-research).

Figure 5: Venture Capital Investment, Selected Countries



Note: This figure plots trends in venture capital investment for selected countries. VC investment data are sourced from PitchBook.

Figure 6: GitHub Pushes per $1B GDP, Selected Countries



Note: This figure plots quarterly GitHub pushes per $1 billion GDP (constant 2015 US dollars) for the United States, India, Brazil, Germany, the United Kingdom, and South Korea. GitHub activity data come from the GitHub Innovation Graph https://innovationgraph.github.com/global-metrics/git-pushes. Quarterly git pushes are divided by annual GDP for the corresponding calendar year. GDP values come from the World Bank.

output that bypasses the patent system entirely. Figure 6 shows that India is one of the top contributors to GitHub among major economies when normalizing by GDP, placing

Table 2: Probability of Breakthrough Patent by Inventor Nationality

	Pooled regression				Separate sample	
	India	U.S.	China	Other Developing	India vs. Non-U.S.	India vs. U.S.
Panel A: 1-year forward and 5-year backward KPST score						
U.S. patents, 1995–2022	0.0384***	0.0460***	0.0191***	0.0002	0.0557***	0.0121***
	[0.0013]	[0.0003]	[0.0006]	[0.0004]	[0.0016]	[0.0013]
U.S. patents, 2005–2022	0.0382***	0.0415***	0.0088***	0.0039***	0.0571***	0.0137***
	[0.0013]	[0.0003]	[0.0006]	[0.0004]	[0.0016]	[0.0013]
Panel B: 5-year forward and backward KPST score						
U.S. patents, 1995–2016	0.0594***	0.0671***	0.0213***	−0.0039***	0.0785***	0.0197***
	[0.0026]	[0.0004]	[0.0015]	[0.0006]	[0.0029]	[0.0026]
U.S. patents, 2005–2016	0.0669***	0.0642***	0.0266***	−0.0035***	0.0923***	0.0292***
	[0.0027]	[0.0005]	[0.0015]	[0.0007]	[0.0032]	[0.0027]

Note: OLS regressions with robust standard errors clustered at the patent level in brackets. Each patent-inventor pair is used as an observation. The dependent variable in each case is an indicator for whether the patent is in the top 10% of the respective KPST score measure. The first four columns ("Pooled regression") report coefficients from a single specification with all four inventor-nationality dummies; the omitted category is other developed-country inventors. The last two columns ("Separate sample") report the India coefficient from specifications run on restricted subsamples: column (5) excludes U.S.-inventor patents so that the India coefficient compares India with the remaining non-U.S. inventors (China, other developing, and other developed), and column (6) restricts the sample to India- and U.S.-inventor patents so that the India coefficient compares India with the United States. All regressions include controls for CPC section and year of award. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

India's software innovation on a par with, or ahead of, many countries that patent at far higher rates.

Indian inventors invent high quality patents in the U.S. While the level of patenting in India is low, Indian innovators do file patents in the U.S., which allows us to compare how innovative activity from Indian inventors compare with others in the US patent system. To examine this, we use the extended KPST textual novelty scores through 2022 from Kelly et al. (2021). We construct two measures of quality for each patent from these scores, both of which are measures of how much a patent influences future innovation relative to how much it draws on prior art. These are the ratio of the one-year forward textual similarity of the patent to its five-year backward similarity and the ratio of the five-year forward similarity to the five-year backward similarity. (The second measure should be more accurate, but of necessity uses a smaller sample.) Following KPST, we define a "breakthrough" patent as one in the top 10% of this ratio.

Table 2 reports the results of analysis of innovator nationality and high-quality U.S.-

filed patents. Having an Indian inventor is associated with a 3.8 percentage point higher probability of being a breakthrough patent. This is comparable in magnitude to the U.S. inventor coefficient (4.2 pp) and four times the China coefficient (0.9 pp). The last column confirms that this pattern is not driven by the choice of reference group: when the sample is restricted to India- and U.S.-inventor patents only, Indian inventors remain at least as likely as U.S. inventors to produce breakthrough patents, with the India coefficient positive and statistically significant across all specifications (1.2–2.9 pp). This result suggests the Indian inventors undertake innovation that compares favorably in quality with US and Chinese inventors in the US patent system. Thus, lack of innovative capacity does not explain the low patenting rates in India.

3.2 Industry Composition

Economists have increasingly appreciated the unique features of India as a services-based economy. One possibility is that the low level of patenting can be explained by industry mix: if India's economic activity is concentrated in sectors that tend not to patent (e.g., software services and business process outsourcing), then the low aggregate patenting rate may be a mechanical consequence of industrial composition, rather than a reflection of low innovative effort.

To assess the extent to which the low level of Indian patenting reflects the composition of the Indian economy, we calculate the number of patents that would have been filed in India had the distribution of economic activity by industry there mirrored that of the U.S. or Chinese economies.

For the universe of 77,821 Indian domestic patents filed between 1995 and 2022 and granted before 2024, we extracted the primary IPC subclass (e.g., A61K)[8] for each patent. We assigned these to the European Union's NACE Rev. 2 primary industries using the Dorner and Harhoff (2018) probabilistic concordance, which relates IPC subclasses to NACE two-digit industry divisions. We then aggregated these assignments from two-digit divisions to one-letter sections.[9] This allowed us to classify patenting in India by broad

[8]If there are multiple IPC codes assigned to a patent, we determine the primary IPC subclass by truncating a patent's raw IPC codes to their 4-digit levels and selecting the most frequent subclass, resolving ties with the first-listed code.

[9]The industry-technology concordance tables were generated from comprehensive linked inventor-establishment data for Germany, using 235,933 patents registered between 1999 and 2011 at the European Patent Office by inventors residing in Germany. The NACE (Nomenclature statistique des Activités économiques dans la Communauté Européenne) is the European standard classification for economic activities, providing a four-digit framework for statistical data. It is directly derived from the United Nations' ISIC (International Standard Industrial Classification) scheme, and is identical at the highest levels that we use here.

industry.

We then obtained information about value added by NACE (equivalently, ISIC) section in China, India, and the U.S. for the year 2020, to facilitate comparisons (and to reflect the concentration of Indian patents in recent years). The U.S. data were obtained from the OECD STAN database; the Indian and Chinese data from the United Nations National Accounts database. The Chinese data were aggregated in several cases (notably, sectors M through T are reported as a single category), so we conducted the analysis using consolidated categories where necessary.

We computed the number of Indian patents that would have resulted had (a) the ratio of patenting to value added in each Indian sector remained the same but (b) the distribution of value added across industries mirrored that of the U.S. or China. More formally, across all NACE sections i, we compute:

$$\text{Counterfactual Patents} = \sum_i \text{Patents}_i^{\text{India}} \times \frac{\text{VA Share}_i^{\text{Foreign}}}{\text{VA Share}_i^{\text{India}}} \tag{1}$$

Table 3 reports the results. In both cases, the answer is that patenting would have increased only modestly. Had the industry structure mirrored that of the U.S., Indian patenting would have been 12% higher. Had it mirrored that of China, patenting would have been 47% higher. The China result is larger primarily because China's manufacturing sector (NACE Section C) accounts for 25.6% of value added, compared to 15.4% in India—and manufacturing is by far the most patent-intensive sector, accounting for roughly two-thirds of all Indian patents.

These magnitudes are too small to explain the gap in patenting between India and its peers. Recall that India's domestic patenting rate as a share of GDP is roughly one-tenth that of China. Even under the most favorable counterfactual—adopting China's manufacturing-heavy industry structure—Indian patenting would increase by less than 50%. Industry composition thus explains only a modest portion of the shortfall, and the bulk of the gap must reflect other factors.

3.3 Attitudes towards Patenting

In keeping with the literature on the long-lived consequences of institutional rigidities (Acemoglu et al., 2001), we might see a cultural reluctance to patent, even if it is economically attractive. We find limited support for this hypothesis, indicating that it is not the driver for low patenting rates.

Table 3: Counterfactual Indian Patenting Under Alternative Industry Structures

		India		U.S.	China
NACE	Industry	# Patents	VA Share (%)	VA Share (%)	VA Share (%)
A	Agriculture, forestry & fishing	274	20.3	0.9	7.9
B	Mining & quarrying	112	1.7	1.0	
C	Manufacturing	50,064	15.4	11.0	25.6
D	Electricity, gas & steam	679	2.2	1.5	4.7
E	Water supply & waste management	144	0.5	0.3	
F	Construction	658	7.4	4.4	7.0
G	Wholesale & retail trade	3,715	9.3	10.3	9.3
H	Transportation & storage	202	4.2	3.1	3.9
I	Accommodation & food services	52	0.5	2.2	1.5
J	Information & communication	2,600	6.7	7.3	3.7
K	Financial & insurance activities	158	6.0	8.6	8.1
L	Real estate	65	7.1	12.8	7.1
M	Professional & scientific activities	11,985	3.5	8.1	21.3
N	Administrative & support services	450	0.5	4.1	
O	Public administration & defence	174	6.6	8.8	
P	Education	1,810	4.6	5.6	
Q	Human health & social work	1,139	1.8	7.8	
R	Arts, entertainment & recreation	47	0.2	0.8	
S	Other service activities	350	1.3	1.4	
T	Household activities	3	0.2	0.1	
Total (weighted patent count)		74,681	100.0	100.0	100.0
Counterfactual total				**84,010**	**109,432**
Ratio to actual				**1.12**	**1.47**

Note: This table reports the results of a counterfactual exercise assessing the role of industry composition in explaining India's low level of patenting. We assign each Indian patent to NACE Rev. 2 industry sections using its primary IPC subclass and the Dorner and Harhoff (2018) probabilistic concordance. Value added shares are for 2020; Indian and Chinese data are from the UN National Accounts database; U.S. data are from the OECD STAN database. Patent counts span filing years 1995–2022. The counterfactual total computes the number of Indian patents that would have resulted had the distribution of economic activity across industries mirrored that of the comparator country, holding the sector-specific ratio of patenting to value added constant.

3.3.1 A Brief History

The Indian patent system has a long history. The initial legislation was passed in 1856, modeled after the English Act of 1852. Post-independence, there was a considerable degree of skepticism among the new political leadership about the effects of the patent system during the colonial era. (Reddy T. and Chandrashekara (2017) narrates a compelling history of the resistance to embracing intellectual property protection in the first half-century of the Republic of India.)[10] In particular, the British colonial administration was perceived as

[10]See also the legal and political science literature on the Indian patent system, especially Halliburton (2017), Liu and Racherla, eds (2019), and Singh (2024).

having limited India's ability to design a patent policy favorable to local inventors, instead focusing on protecting the interests of exporters based in the United Kingdom.

In response, the process of establishing a patent system in India was a very slow one. The government established the Bakshi Tekchand Committee in 1949 and the Justice Rajagopal Ayyangar Committee in 1957. Both committees expressed deep skepticism about the advisability of any patent system for India.

While the Patents Act of 1970 did establish a patent system, it limited the awards in many important respects. Among other provisions, Chapter II, Article 5 stated that patents granted for "substances intended for use, or capable of being used, as food or as medicine or drug" or "relating to substances prepared or produced by chemical processes (including alloys, optical glass, semiconductors and inter-metallic compounds)" would be limited to the protection of processes (rather than product patents). The food and medicine process patents would only last for the lesser of five years from the patent issue or seven years from the application date (Chapter VIII, Article 53). Finally, Chapter XVI gave the government sweeping rights to revoke or compulsorily license patents for a variety of grounds, including not being "available at reasonably affordable prices to the public."[11]

The promulgation of the World Trade Organization's Trade-Related Aspects of Intellectual Property Rights (TRIPS) agreement in the mid-1990s coincided with a period of economic liberalization in India. While India signed the TRIPS agreement in 1994, it was not until 2005 that a patent law fully compliant with global standards (e.g., full patent protection for food and medicine, and patent terms of 20 years from the application date) was enacted.[12]

3.3.2 Evidence about Attitudes

Attitudes toward intellectual property. We draw on survey data from the World Economic Forum's Executive Opinion Survey (EOS) to examine perceptions of intellectual property protection. Figure 7a plots the average annual patent applications against the average EOS IP protection score for each country over 2008–2022. There is a strong positive cross-country relationship: countries with higher perceived IP protection tend to file more patents. India sits well above the fitted line, with an EOS score around four, suggesting that weak perceptions of IP protection may not explain India's low patenting rate. Figure 7b

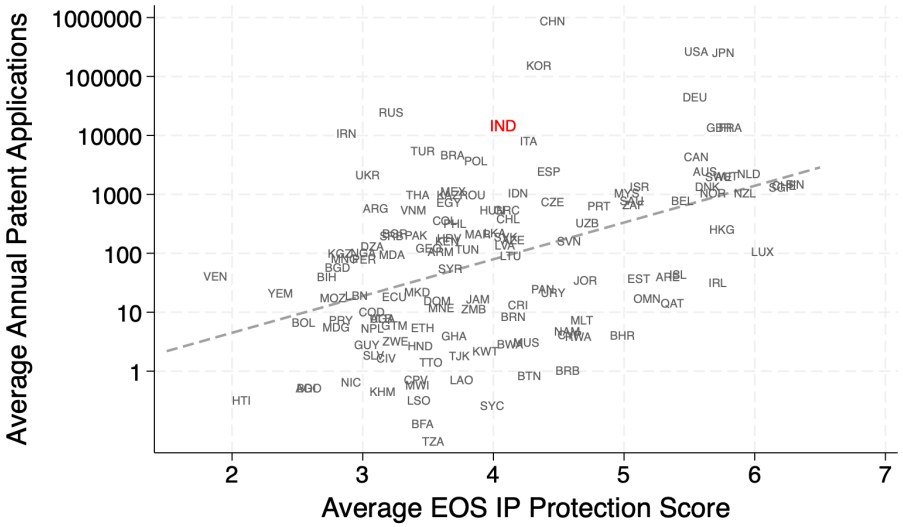
[11] https://ipindia.gov.in/writereaddata/portal/ipoact/1_31_1_patent-act-1970-11march2015.pdf.

[12] In April 1995, the Indian government agreed to receive patent applications for pharmaceutical and agricultural products and to store them for processing once the Indian patent law to make such subjects patentable took effect. See for details, https://www.wto.org/english/tratop_e/dispu_e/79r.pdf

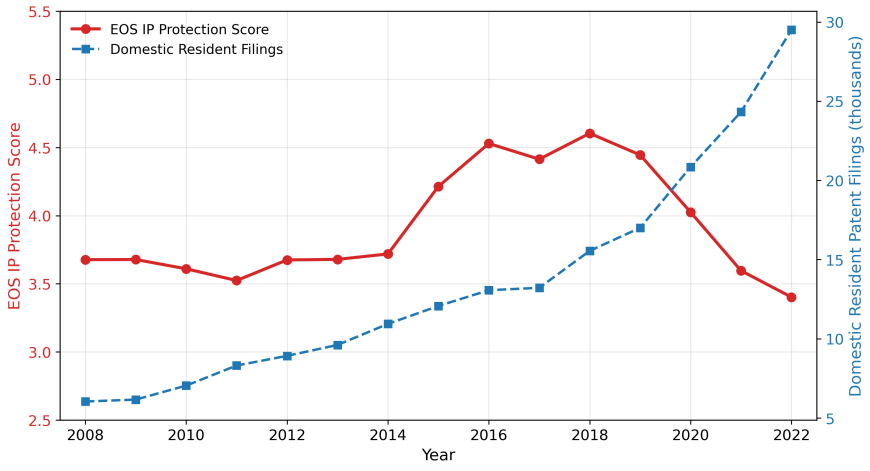
shows that India's EOS IP score peaked around 2018 and has since declined, even as patent applications have risen steadily—indicating that perceptions of IP protection have not kept pace with the growth in patenting, though.

Figure 7: Patent Applications and WEF IP Protection Score

(a) Cross-Country: Patents vs. EOS Score



(b) India's EOS IP Score Over Time



Note: In Panel (a), each point represents a country, labeled by ISO3 code. The horizontal axis is the country's average WEF Executive Opinion Survey score for intellectual property protection (1 = not at all; 7 = to a great extent), averaged over 2008–2022. The vertical axis is the country's average annual resident patent applications over the same period (log scale). India is highlighted in red. The dashed line is the OLS fit in log-level space. In Panel (b), India's annual EOS IP protection score from 2008 to 2022. Patent data are from WIPO; EOS data from the World Economic Forum Global Competitiveness Reports.

We additionally examine responses from the World Values Survey (WVS). Figure A.1a plots Indian responses to question E016, which asks whether "more emphasis on technology" is a good or bad thing. While a majority of respondents consistently view technology positively, the share saying "good thing" declined from about 85% in 1990 to roughly 50–65% in subsequent waves, indicating some ambivalence. Figure A.1b plots responses to WVS question F114B, which asks whether stealing property is justifiable. The overwhelming majority of Indian respondents (over 90%) view property theft as "never justifiable," and this share has remained stable across waves—suggesting that respect for property rights per se is a dominant cultural feature.

These patterns suggest that cultural attitudes toward technology and property are broadly supportive of innovation. This result suggests that impediments to patenting lie more in net economic returns to patenting and the institutional environment than in cultural resistance.

3.4 Lack of Patent Value

Another possibility is that the returns to high-quality patents are low: in other words, it is too difficult to realize value from patents in India once they are issued. We focus on one manifestation of value realization: the extent to which patents are rewarded by the equity markets. If Indian capital markets do not reward patenting, firms' incentives to seek patent protection may be substantially reduced. Such an absence of a reaction might reflect that the private returns to patenting are too low to justify the costs of filing, for instance, due to an inability to license, litigate, or otherwise monetize awards effectively.

To test this hypothesis, we match Indian patents grantees to firms listed on the National Stock Exchange (NSE) and the Bombay Stock Exchange (BSE) through the Prowess database. We estimate abnormal returns and share turnover patterns for firms around the publishing of granted patents.[13]

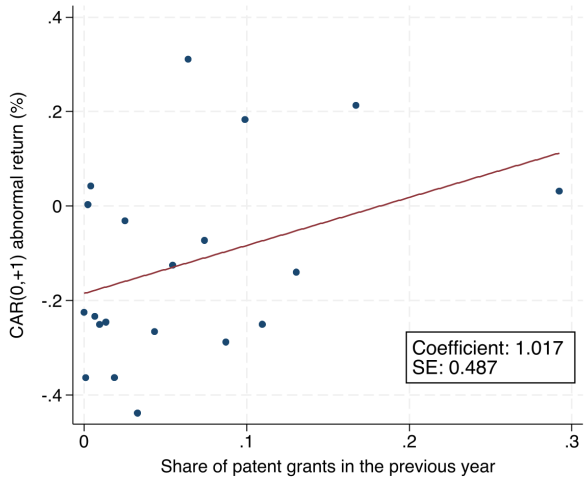
Figure 8 plots binscatters of stock-market reactions against the firm's patenting footprint upon the publication of new patent grants for the firm,[14] weighted by the number of grants published for the firm on that Monday. Figure 8a shows the day-of-publication cumulative abnormal return (or CAR, computed using the (0,+1) window and expressed

[13]While the patent applicants know the result of a grant on the grant day, the public will only know new granted patents every Friday around 6 PM from IP India's journal, and thus only able to respond the next Monday.

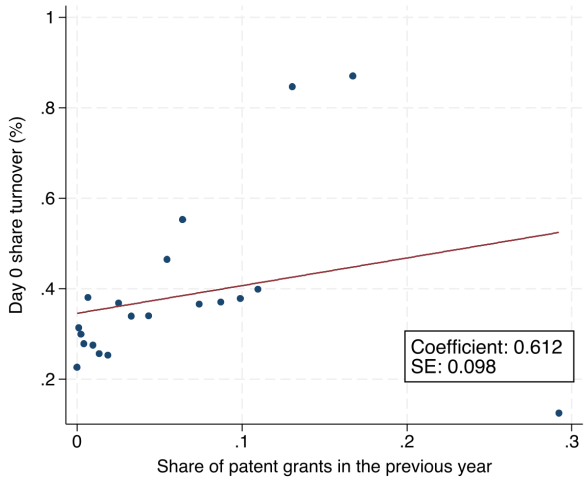
[14]It is reasonable to assume that stock markets will have higher reactions to patents from already established patent applicants. Therefore, we use the share of total patent-grant publications in the previous year, a rolling sum over trading days -246 to -6, divided by the corresponding total across all firms, to measure how established a given firm is in patenting.

Figure 8: Stock-Market Reaction to Patent Grant Publication

(a) Day-of-Publication CAR(0,+1)



(b) Day-0 Raw Share Turnover



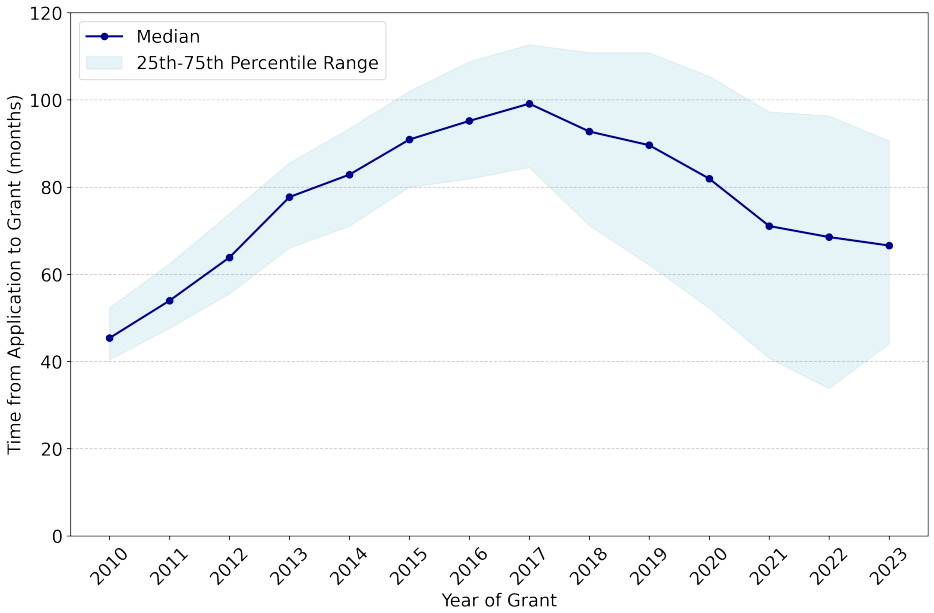
Note: This figure plots binscatters of stock-market reactions on patent grant publication days (the Monday following Friday gazette publication by the Indian Patent Office) against the firm's share of total patent-grant publications in the previous year, weighted by the number of grants published on that Monday for the firm, conditional on the firm having any patents being granted on that event. Figure 8a plots the day-of-publication cumulative abnormal return CAR(0,+1), where abnormal returns are computed as raw returns net of the value-weighted market return. Figure 8b plots share turnover on day 0. The sample includes all publicly listed Indian firms in the Prowess database matched to Indian patent grants over 2005–2022, including both National Stock Exchange (NSE) and Bombay Stock Exchange (BSE) listings. For dual-listed firms, we use NSE as the preferred measure given its higher liquidity.

in percent). The slope is positive and significant at the 5% level, with a coefficient of 1.017 and a standard error of 0.487, implying positive responses from the stock market. Figure 8b shows raw day-0 share turnover (in percent). The relationship is also positive and significant (coefficient 0.612, SE 0.098), with more patent-active firms experiencing greater trading activity around grant publication. For U.S. patent grants, Kogan et al. (2017) document a mean market-adjusted return of 0.07% over the three-day issuance window. In India, the mean CAR(0,+1) across grant-publication events is -0.12% overall, but 0.12% for firms in the top decile of the previous year's share of patent grants, suggesting a positive response to patent filings for the most established patent filers. Thus, the market seems to recognize and reward established patentees.

3.5 Institutional Frictions

A final set of explanations focuses on a state capacity explanation, namely, institutional frictions in the issuing process that may impede the quantity and quality of patenting. In this section, we examine an important feature of the Indian patent issuance process, namely,

Figure 9: Time from Application to Grant, Indian Patents



Note: This figure plots the median time from application to grant (in months) for Indian patents, by year of grant. The shaded band shows the 25th–75th percentile range.

the long delay involved in obtaining an awards, and the causal impacts of mitigating this constraint.

3.5.1 Long Delays in Patent Processing

Figure 9 shows the time from application to grant for Indian patents, by grant year. The median wait time peaked at nearly 100 months (over eight years) for patents granted around 2016–2017, before declining to approximately 68 months for patents granted in 2023. Table 4 shows that in 2016, the average pendency of 84 months in India was over three times higher compared to the US or China. Even at its recent lows, the wait time in India far exceeds that of peer countries.[15]

A natural conjecture is that India's patent office is simply understaffed. Table 4 compares patent office statistics across India, the United States, and China. Panel A shows that circa 2016, India had approximately 78 applications per examiner—comparable to the U.S. ratio of 74 per examiner and well below China's 115 per examiner. Yet India's average pendency (84 months) was roughly four times that of the U.S. (25 months) and China (22 months). By the most recent available data (Panel B), India's pendency has improved to

[15] Schultz and Madigan (2016) compared patent processing times in 2015 across a wide variety of developed and emerging economies, and found that only Brazil and Thailand consistently exceeded India on this metric.

38 months, but it remains well above the U.S. (26 months) and China (16 months). These comparisons suggest that the delay problem may indicate lack of institutional capacity rather than a simple lack of manpower.[16]

These concerns have been echoed by the Indian government's own analyses of its patent system. For instance, a 2022 report by the Economic Advisory Panel to the Prime Minister (Sanyal and Arora, 2022) highlighted comparative statistics similar to those we reported above. They highlighted both the desirability of increasing patent office staffing and the need to increase productivity. Among their recommendations were the hiring of more controllers (or senior examiners) responsible for the final approval of patent grants, the imposition of deadlines for each stage of the review process, and the reduction of cumbersome compliance requirements.

4 The 2016 Auto-Allocation Reform

4.1 Institutional Background

India has had four regional patent offices since 1970, located in Chennai, Delhi, Kolkata, and Mumbai. Prior to 2016, the IP India operated under a system in which applicants filed in the regional office associated with their location. Patents filed in a given regional office were ultimately examined by examiners within that office. Within each office, there were distinct examination groups. While nominally there were four broad groups, IP India has recruited individuals with 14 distinct skill clusters, such as Chemistry and Electrical Engineering.[17] This approach created substantial heterogeneity in pendency across offices and technology groups: some offices-by-technology groups accumulated large backlogs, while examiners elsewhere had shorter queues.

4.2 2016 Auto-Allocation Reform

In 2016, IP India changed how patents were processed, with the objective of reducing processing delays. It introduced an auto-allocation system that reassigned patents *across* offices. As the 2016-17 *Annual Report* of IP India explained, the reform "has brought into a single group-wise queue of Requests for Examination to have uniform examination at all

[16] One indication of potential productivity issues is, of the 580 patent examiners reported by IP India, only 97 completed FERs in 2015, a number that we utilize in the calculations below.

[17] While the patent data does not record examiner's skill group directly, a related variable from patent data, "field of invention," closely mirrors the skill clusters, and is used in the empirical analysis.

Table 4: Patent Office Comparison: India, United States, and China

	India	**United States**	**China**
Panel A: Circa 2016			
Number of patent examiners	580	8,160	11,600
Patents applied	45,444	607,753	1,339,000
Patents processed	30,271	645,907	675,000
Applications per examiner	78	74	115
Patents processed per examiner	50	79	58
Average pendency (months)	84	25.3	22
Panel B: Most Recent Available (Circa 2024)			
Number of patent examiners	626	8,599	16,000
Patents applied	110,375	598,982	1,678,000
Patents processed	40,606	577,135	921,000
Applications per examiner	176	70	105
Patents processed per examiner	65	67	58
Average pendency (months)	37.8	26.3	16
Patent office budget	$38M	$4.20B	$933M
Budget per patent processed	$936	$7,244	$1,013

Note: Panel A covers data around 2016 (India FY 2016-17, U.S. FY 2016, China 2016) and Panel B covers the most recent year available at compilation (India FY 2024-25, U.S. FY 2024, China 2023, the last available Chinese patent examiners information). India reports on April–March fiscal years, the U.S. on October–September fiscal years, and China on the calendar year. *Patent examiners*: For the U.S., the number of examiners covers utility, plant and reissue (UPR) examiners only. For China, the figure aggregates headquarters and regional patent examiners. *Patents applied*: U.S. — utility applications only. China — invention-patent applications (comparable to U.S. utility patents) only. *Patents processed*: India — applications disposed of. U.S. — UPR applications allowed plus abandoned. China — examined-and-concluded invention applications in 2016; granted invention patents only in 2023 due to reporting limitation. *Average pendency (months)*: average time from filing (or request for examination, where applicable) to final disposition. Each office's own reported pendency is used where available; for India the WIPO figure is used. Patent office budget: total agency appropriation. India covers the Intellectual Property sub-head in Demand No. 30 (DPIIT). U.S. covers total USPTO obligations. China covers the CNIPA final-accounts total. USD equivalents use the annual-average exchange rates (83.682 INR/USD in 2024; 7.077 CNY/USD in 2023). For India, Number of examiners, patents applied, and patents processed come from the Office of the Controller General of Patents, Designs and Trade Marks (CGPDTM) Annual Reports 2016-17 and 2024-25 respectively;for the U.S., from the USPTO FY 2016 Performance and Accountability Report and the USPTO FY 2024 Workload Tables respectively; for China, from the SIPO 2016 Annual Report and the CNIPA 2023 Annual Report respectively, with the 2023 China examiner count taken from a CNIPA press release dated 13 June 2024. Average pendency for India is from the WIPO World IP Indicators 2017 (Panel A) and the WIPO IP Statistical Country Profile 2024 (Panel B); pendency for the U.S. and China is from the same USPTO and CNIPA Annual Report sources cited above. Patent office budget is from the Union Budget of India 2024-25; from the USPTO FY 2024 Agency Financial Report for the U.S.; and from the CNIPA 2023 final-accounts disclosure for China.

Patent Office locations with respect to date of filing of Request of examination. The system has removed anomaly of different time of examination at different locations."[18]

4.3 Post-Reform Rate of Reallocation

We first show trends in the likelihood that a patent is assessed outside of the office where it is filed, i.e., the rate of reallocation. Figure 10 illustrates the resulting change in the geographic patterns of patent examination. The first panel shows that essentially no applications were examined outside their origin office before 2016, after which the cross-office share rises sharply. On average, about half of applications were reallocated to a different office from the application office. The second panel summarizes the resulting flows of applications across offices. Chennai and Mumbai reallocated a larger share of their applications, whereas Delhi processed most applications at home.

Descriptive Determinants of Reallocation. As a descriptive exercise, we examine how patents are allocated post-reform with a conditional logit model. It predicts the likelihood a patent is assigned to an office using three attributes: the office is the origin office (where the patent was filed), the office's field examiner capacity, and the office's per-examiner caseload for an incoming patent. In Section 6.1, we further analyze the mechanism that best matches actual allocation using "optimized" allocation simulations and compare office and examiner field match distribution in the simulation to the actual data. Our logit model predicts the likelihood of office assignment with the following equation:

$$P_{ijt} = \frac{\exp\left(\alpha_j + \beta_1 \text{IsHome}_{ij} + \beta_2 \text{FieldSupply}_{ijt} + \beta_3 \text{Queue}_{ijt}\right)}{\sum_{k=1}^{4} \exp(\alpha_k + \beta_1 \text{IsHome}_{ik} + \beta_2 \text{FieldSupply}_{ikt} + \beta_3 \text{Queue}_{ikt})} \tag{2}$$

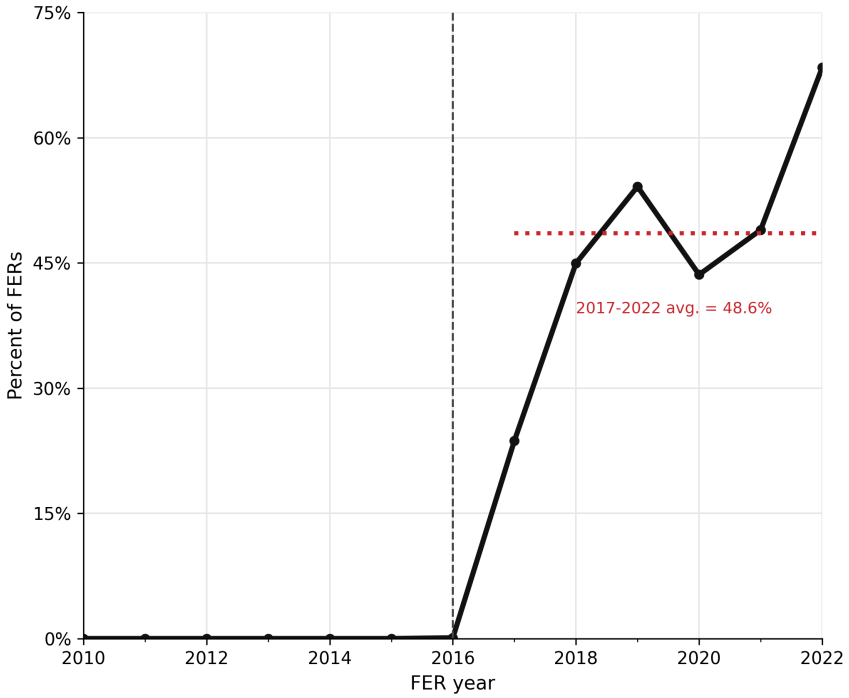
where i indexes an incoming patent application, j indexes the candidate office, t indexes RFE year, and k indexes the four candidate offices in the denominator. IsHome_{ij} is a dummy which takes a value of 1 if patent i is filed in office j. FieldSupply_{ijt} is office j's pre-reform capacity in the field of patent i, measured in examiner-equivalents.[19] Queue_{ijt} is the standardized average queue per examiner in office j at patent i's RFE date. α_j is an

[18] https://ipindia.gov.in/storage/uploads/docs-operator/d87014ba-d805-4b5c-bb4e-dd0e9061a318.pdf, page 18.

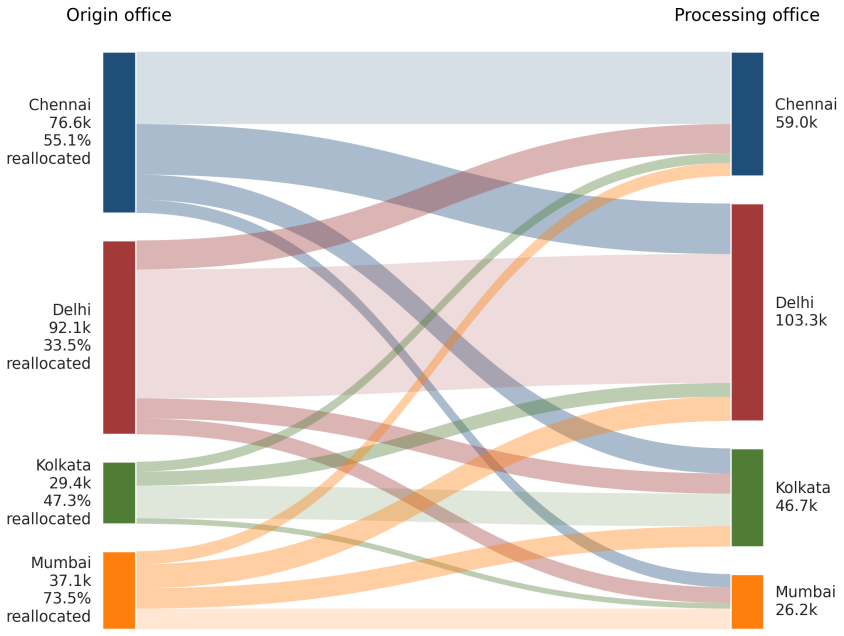
[19] Each pool examiner's unit of capacity is split across fields by their pre-reform caseload shares, and those shares are summed over the office's examiners sitting in that office in year t. For example, if Delhi has three examiners whose pre-reform shares of Chemistry patents were 20%, 0%, and 10%, Delhi's Chemistry FieldSupply is 0.3 examiner-equivalents.

Figure 10: Patterns of Patent Reallocation Across Offices

(a) Share examined outside origin office, by FER year



(b) Aggregate cross-office flows of applications, FER years 2017-2022



Note: Panel (a) plots the share of applications examined outside their origin office by year of first examination report (FER); essentially no applications were examined outside the origin office before 2016, after which the cross-office share rises sharply. Panel (b) summarizes the resulting flows of applications across the four patent offices.

alternative-specific constant for candidate office j, the conditional-logit analogue of an office fixed effect, absorbing any persistent office-level pull not captured by field supply or queue. Only differences $\alpha_j - \alpha_k$ are identified, so we normalize $\alpha_{\text{Delhi}} = 0$. Each patent contributes four alternatives with exactly one chosen office; the logit model estimates the binary outcome of whether j is chosen $(= 1)$.

We run the logit model over the benchmark–or actual–data to understand how the real allocation was decided. We consider the 134,012 patents in the logit analysis, focusing on those whose RFE date is on or after the reform date and before the end of 2021. We then compare the benchmark's results to the logit coefficients of several allocation rules. The first rule is purely random allocation at the examiner level. The second rule is queue based, where we simulate an optimized allocation based on examiner workload. At patent i's RFE date, we calculate the per-examiner caseload at each office $(\frac{\text{\# active patents}}{\text{\# examiners in that office and RFE year}})$, then allocate randomly among the two least congested offices.[20] The third rule is field based, where we weight assignment to examiners based on their pre-reform field share. The last rule combines the queue and field rules: we weight examiner assignment by pre-reform field share across examiners at the two least-congested offices.

The results of these separate logit models are presented in the rows of Table 5, where we further differentiate by early $(2016 - 2018)$ and late $(2019 - 2021)$ post-reform periods. Table 5 suggests home bias is strong in the benchmark data; in the early post-reform period, the 1.05 log-odds coefficient is statistically significant and suggests that being the origin office increases the odds of getting a patent by 2.86x. This bias decreases by 11.5% in the late post-reform period, to about a 2.53x advantage.

The field supply coefficient rises from 0.07 to 0.36 across the two periods, indicating patents are allocated toward offices with more field capacity. The coefficient is z-scored, so the late value implies that a one standard deviation increase in an office's field supply raises its odds of receiving the patent by 1.44x.

The queue column has a natural scale: random and field-based allocation, which ignore congestion, both sit at essentially zero, while the two queue-based rules run between -0.69 and -1.78. The benchmark more closely matches the random and field allocations, with a null (0) and modestly positive (0.07) coefficient across both periods, so the actual allocation shows no tendency to divert patents away from busier offices. Conversely to queue optimization, the actual data is slightly more likely to send patents to busier offices in the later period. Removing the office constants only strengthens this: in the

[20] We consider patents to be "active" at date x if the RFE date $< x \leq$ min(FER date, grant date). Note that this considers pre-reform patents that are folded in as the "backlog" at the start of the reform for all allocation rules.

Table 5: Determinants of Indian Patent Allocation

Allocation rule	N patents	Home bias	Field supply	Queue
2016–2018				
Benchmark (Actual)	59,733	1.05^{***}	0.07^{***}	0.00
		(0.01)	(0.01)	(0.01)
Random across office	59,733	-0.01	0.00	-0.01
		(0.01)	(0.01)	(0.01)
Field-based across office	59,733	0.00	0.35^{***}	-0.01^{*}
		(0.01)	(0.01)	(0.01)
Queue-based across office	59,733	0.00	0.03	-0.93^{***}
		(0.01)	(0.01)	(0.01)
Field + queue across office	59,733	0.03^{**}	0.26^{***}	-0.78^{***}
		(0.01)	(0.01)	(0.01)
2019–2021				
Benchmark (Actual)	74,279	0.93^{***}	0.36^{***}	0.07^{***}
		(0.01)	(0.01)	(0.02)
Random across office	74,279	0.01	0.02	0.03
		(0.01)	(0.01)	(0.02)
Field-based across office	74,279	-0.02^{*}	0.37^{***}	0.01
		(0.01)	(0.01)	(0.02)
Queue-based across office	74,279	-0.03^{***}	-0.02^{*}	-0.69^{***}
		(0.01)	(0.01)	(0.02)
Field + queue across office	74,279	-0.02^{*}	0.20^{***}	-1.78^{***}
		(0.01)	(0.01)	(0.05)

Note: Coefficients from separate conditional logit models (McFadden 1974), standard errors in parentheses. Each patent contributes four observations (one per candidate office). *Home bias* $= 1$ if candidate office equals patent's origin office; *Field supply* $=$ examiner-equivalents in the patent's field at the candidate office (pre-reform caseload shares summed over its pool examiners), standardized within RFE year. *Queue* $=$ the candidate office's caseload per examiner at the patent's own RFE date, counting cases that arrived before it and had not yet left: the common pre-reform backlog plus that allocation's own post-reform assignments, divided by the number of examiners at that office in that year, and standardized within RFE year. All models include alternative-specific constants for the candidate office (Delhi omitted). Periods are RFE years, i.e. when the allocation decision is made, not FER years. The sample comprises the 134,012 patents whose RFE falls on or after the reform date and before the end of 2021, with the FER observed in any year; about 12% of 2021 arrivals were still awaiting examination or grant at the data vintage, so the latest cohort excludes the slowest cases.

by-office models of Appendix Table A.3 and Appendix Figure A.2,[21] the benchmark's queue coefficient is positive for every office.

Field supply, by contrast, is positive and significant in every specification with and without the constants. The optimized queue-based rule confirms the model would detect workload routing if it were present, as its queue coefficient is large and negative in both periods. Consequently, our logit model results suggest allocation is driven by where a patent was filed and by field capacity, rather than by congestion. Nonetheless, since the reform increased available capacity in certain fields depending on capacity in other offices, it may have reduced delays, which we examine next.

5 The Causal Impacts of the Auto-Allocation Reform

We now turn to examining the effects of this reallocation policy, looking at the effects on examination lags, the quality of applicants, and venture capital financing.

5.1 Empirical Strategy

We use a DID design to estimate the impacts of the 2016 reform. Since patent officers are specialized in particular fields, the unit of analysis is a field-of-invention by origin-office by year cell. Our main cross-sectional exposure measure is the amount of same-field examiner capacity outside the origin office in 2015, with examiner-years assigned fractionally across fields.

We measure examiner capacity from the First Examination Reports themselves: each FER names the issuing office and the controller (examiner) who signed it, which we extract from the scanned PDFs using a combination of OCR and large language models and then disambiguate into consistent examiner identities (Appendix Figure A.3 shows an example). A cell's 2015 outside-office same-field capacity is the examiner-equivalent capacity working in that field in offices other than the origin office. The capacity is similarly defined as in section 4.3, where each examiner's unit of capacity is split across the fields she examined in proportion to her caseload.

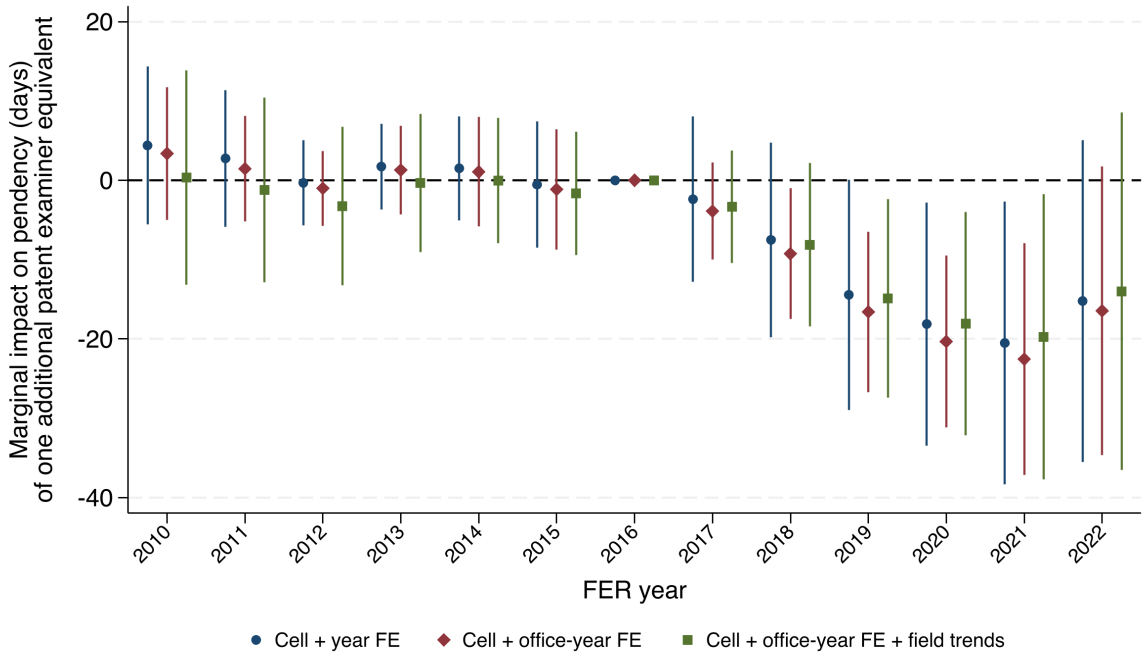
Specifically, we estimate:

$$Y_{fot} = \sum_{\tau \neq 2016} \beta_\tau \left(\text{OutsideCapacity}_{fo,2015} \times 1[t = \tau] \right) + \alpha_{fo} + \delta_{ot} + \epsilon_{fot}, \tag{3}$$

where f indexes field, o indexes origin office, and t indexes year. The fixed effects absorb time-invariant field-office differences and office-specific annual shocks. Our first main

[21] Both omit the office constants, which are collinear with IsHome within a single origin office.

Figure 11: Examiner Capacity Reallocation and Pendency



Note: This figure reports event-study coefficients from field-of-invention by origin-office by year cells. Treatment is 2015 outside-office same-field examiner capacity, with examiner-years assigned fractionally across fields. The dependent variable is mean RFE (request for examination)-to-FER (first examination report) pendency in days. The omitted year is 2016; standard errors are clustered by field-office cell and 95% confidence intervals are shown.

outcome variable Y_{fot} measures the days between the applicant requested an examination report (RFE) to the date her application received the first examination report (FER).[22] We describe the other outcome variables as we introduce them below.

5.2 Results

Figure 11 shows that cells with more newly accessible same-field examiner capacity experienced larger post-2016 reductions in RFE-to-FER pendency. The pre-reform coefficients are small relative to the post-reform decline. Taking the 2019 estimate, one additional outside-office same-field examiner-equivalent of capacity is associated with about 16.6 fewer RFE-to-FER days. Because a one-standard-deviation increase in outside-office capacity across cells is 5.1 examiner-equivalents, this corresponds to roughly 85 fewer days, or about 0.22 standard deviations of the pre-reform pendency distribution—an economically

[22]We choose this measure because it is what the reform actually targets, i.e., reducing the time under examination. The full application-grant cycle also involves how quickly the applicant responds and whether a hearing is needed, among other considerations.

meaningful reduction of nearly three months in examination delay.

Robustness. The pendency result is robust to several identification concerns. First, the results are robust to alternative measurements. The post-reform decline is unchanged under alternative capacity definitions, including threshold-count assignment of examiners to fields and a 2005–2015 average of active examiner-years (Figure A.4); under alternative field-share thresholds and minimum-activity requirements (Figure A.5); when examiners are weighted by 2015 examination throughput rather than counted as individuals (Figure A.6); and when capacity is split across fields by pre-2015 portfolio shares (Figure A.7). Nor does it hinge on the linear specification: cells above the median capacity increase display the same widening post-reform pendency gap (Figure A.8). The second concern is that a single office or the sample choice drives the result: the negative post-reform pattern appears within each origin office (Figure A.9) and in both the full set of cells and the main-field subsample (Figure A.10). The pattern survives field-specific linear trends. The final concern is inference with a modest number of field-office cells: wild cluster bootstrap confidence intervals are similar to the main results (Figure A.11).

5.3 Quality of Applicants

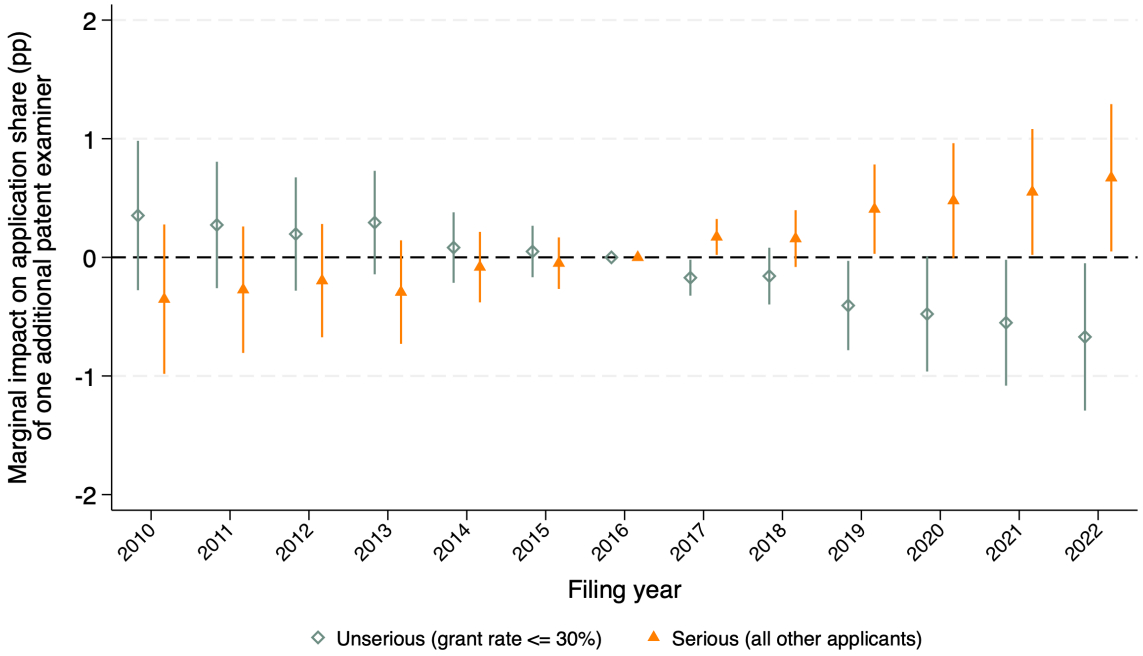
If the delays in patent processing were an important barrier to patenting, the 2016 reforms—which reduced delays for field-offices with greater available processing capacity—should have had consequences on final patenting outcomes. We first focus on the changing mixture of applicants and then on total filings.

In particular, as highlighted in Table 1, domestic applicants vary significantly in success rates. The low success rate of certain applicants, particularly private universities, suggest that they were motivated primarily by a desire to fulfill quotas or rating schemes based on the number of patent filings, rather than desire to obtain formal property rights. We might anticipate that these applicants would be far less likely to respond to a change in patent processing times than those applicants who have traditionally had much more success (since their incentives appear to be to simply file, without regard to whether they are granted a patent).

To do this analysis, we divide the applicants by their success rate in applications prior to the 2016 reform. We identify a subset of applicants that we term "unserious": those with at least three patent applications that were filed between 2010 and 2015 and a success rate of applicants 30% or below.[23] We compare these applicants to all others, whom we

[23]Success rates are computed from eventual outcomes: for each application filed between 2010 and 2015,

Figure 12: Examiner Capacity Reallocation and Applicant Composition



Note: This figure reports event-study coefficients for application shares by applicant quality category, measured as shares of all filings. Applicant quality is defined using pre-reform disposal histories: "unserious" applicants have at least three pre-reform disposals and a realized grant rate of 30% or below, while "serious" applicants are all other applicants; the two groups are complements. Treatment is 2015 outside-office same-field examiner capacity, with examiner-years assigned fractionally across fields, and coefficients are percentage-point changes per one additional examiner-equivalent. Estimates include cell and office-year fixed effects plus field-specific linear trends. The omitted year is 2016; standard errors are clustered by field-office cell and 95% confidence intervals are shown. Recent filing cohorts, especially 2020–2022, are still maturing.

refer to as "serious."

Figure 12 shows how applications from these two sets of parties change after the reform, again looking at location-by-technology cells that were more or less affected by the new policy. The share of applications from "unserious" applicants falls in more exposed cells after the reform, while the share from the "serious" applicants rises. Taking the 2019 estimate, one additional examiner-equivalent of newly accessible same-field capacity reduces the unserious share of filings by 0.41 percentage points; a one-standard-deviation increase in exposure (5.1 examiner-equivalents) thus implies a decline of about 2.1 percentage points, or roughly 8% of the pre-period unserious share of 25%.

we measure its final legal status as of 2026. An application is disposed once it has been granted, refused, abandoned (including being "deemed to be abandoned" by the IP India, if the applicant does not respond to the examiner report by the deadline), or withdrawn; the success rate is the number granted divided by the number disposed, with withdrawn and abandoned applications counting as failures.

Table 6: Examiner Capacity Reallocation: Difference-in-Differences Estimates

Dep. var.:	RFE-to-FER pendency (days)		Share of good applications	Total applications
	(1)	(2)	(3)	(4)
Outside capacity × post-2016	-12.745***	-13.540***	0.0023**	11.734**
	(4.328)	(3.214)	(0.0011)	(5.792)
Cell FE	Yes	Yes	Yes	Yes
Year FE	Yes	–	–	–
Office × year FE	No	Yes	Yes	Yes
Number of Obs	980	980	895	912
Mean of Dep. Var	1,340.149	1,340.149	0.201	541.851
SD of Dep. Var	569.708	569.708	0.124	739.134

Note: The unit of observation is a technology-field by patent-office cell by year (76 cells; FER years 2010–2022 for pendency, filing years 2010–2021 for all applications). Each column reports the static difference-in-differences coefficient on 2015 outside-office same-field examiner capacity interacted with a post-2016 indicator, with each examiner's unit of capacity split fractionally across fields. The dependent variables are mean RFE-to-FER pendency in days (columns (1)–(2)), the share of applications filed by good applicants (column (3)), and the total number of applications filed (column (4)); an applicant is good if it had at least three disposed applications pre-reform and a grant rate of at least 70%. Standard errors clustered by field-office cell (76 clusters) are in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Table 6 reports the difference-in-differences counterparts of these event studies: the coefficient on outside-office capacity interacted with a post-2016 indicator, for pendency under both fixed-effects structures and for the two application outcomes under the preferred cell plus office-by-year specification. One additional examiner-equivalent of newly accessible same-field capacity reduces RFE-to-FER pendency by about 13.5 days (column (2)) and raises the share of applications from *good* applicants—those with at least three disposed applications pre-reform and a grant rate of at least 70%—by 0.23 percentage points (column (3)). Total filings also rise, by about 12 applications per cell-year (column (4)), though this estimate should be read with caution: its event study fails the pre-trends test. The corresponding event studies are in Appendix Figures A.12 and A.13. Thus, institutional reforms that reduce long processing delays change the composition of innovation, significantly increasing the proportion of high-quality patent filings.

5.4 Effects on Startups and Venture Investment

Does improving institutional capacity for patenting also lead to changes in other dimensions of innovation? One crucial area, a focus of academics and policymakers alike, is the venture capital sector. The ability of young firms to raise external financing, frequently from venture capitalists who provide governance and connections in addition to capital,

has been associated with increased innovation and growth in both developed and emerging economies (e.g., Bernstein et al. (2016) and Lerner et al. (2024)). Prior work on the U.S. have shown a relationship between patent grants (as a result of randomly assigned examiners) and a greater ability of start-ups to raise finance (Farre-Mensa et al. (2020) and Gaulé (2018)).

Here, we proceed in two parts. First, we examine whether the patent allocation reform boosted patenting by venture-backed firms. We then examine whether the reform boosted the probability of firms raising subsequent venture financing rounds, focusing on the subset of venture-backed firms that had sought patent protection prior to the 2016 changes.

We estimate the impacts on venture capital outcomes using a DID specification similar to that above. To examine this question, we match the PitchBook universe of Indian startups to patent assignees by name and corporate registry. We examine the results using two samples. The first is the set of firms with at least one completed VC round by 2015, yielding a panel of 1,187 firms. The second, broader set of firms also includes founded by 2015 that may have had the potential to raise venture capital, including those financed by angel investors, accelerators, and incubators (an additional 5,001 companies).[24]

For firm i, located through its headquarters and industry in field-office cell $c(i)$ with office $o(i)$, we estimate

$$\text{AnyFiling}_{it} = \beta \left(\text{OutsideCapacity}_{c(i),2015} \times \text{Post}_t \right) + \alpha_i + \delta_{o(i)t} + \epsilon_{it}, \tag{4}$$

where AnyFiling_{it} is an indicator for filing any (eventually granted) patent application in year t ($\times 100$), $\text{OutsideCapacity}_{c(i),2015}$ is the same outside-office same-field capacity measure as in the pendency analysis, $\text{Post}_t = 1[t \geq 2016]$, and α_i and $\delta_{o(i)t}$ are firm and office-by-year fixed effects.

The estimate in Table 7 is positive but insignificant in both specifications: e.g., in Panel A, 0.04 percentage points per examiner-equivalent (wild-cluster bootstrap $p = 0.62$; column (1)). Appendix Figure A.14 shows the event-study version, which replaces Post_t with year indicators.

It is worth noting the paucity of patenting by venture-backed firms in India versus global benchmarks, which may explain the lack of significance of the above results. We can compare, for instance, patenting among Indian venture-backed firms to their counterparts in the U.S. Among the 35,472 U.S. firms with a first venture financing round between 1976 and the end of 2012, at least 22.3% filed a successful patent application within two

[24] In this analysis, we require that the founding year, location, and industry of the firm be known as well. We further require that any completed venture round carry a date. The Startup India Program studied by Chhina (2026) did not begin accepting registrations until 2016 and only offered tax credits to firms incorporated after April 2016.

Table 7: Patenting and Venture Financing Responses among Indian Startups

	Patenting first stage (post ≥ 2016)	Financing triple-difference (post ≥ 2018)	
Dep. var.:	P(any patent filing)	P(any deal)	No. of deals
	(1)	(2)	(3)
Panel A: Firms VC-backed by 2015			
Outside capacity × post	0.043	-0.001	0.000
	(0.080)	(0.120)	(0.001)
	[0.622]	[0.995]	[0.761]
Outside capacity × post × pre-2016 patent holder		1.140**	0.014**
		(0.512)	(0.007)
		[0.018]	[0.027]
Firm FE	Yes	Yes	Yes
Office × year FE	Yes	Yes	Yes
Patent-holder × year FE	No	Yes	Yes
Number of Obs	14,450	14,450	14,450
Mean of Dep. Var	1.260	17.135	0.183
SD of Dep. Var	11.152	37.683	0.418
Panel B: All frame firms			
Outside capacity × post	0.047	-0.053	-0.000
	(0.028)	(0.062)	(0.001)
	[0.191]	[0.408]	[0.502]
Outside capacity × post × pre-2016 patent holder		0.687**	0.008**
		(0.308)	(0.003)
		[0.064]	[0.081]
Firm FE	Yes	Yes	Yes
Office × year FE	Yes	Yes	Yes
Patent-holder × year FE	No	Yes	Yes
Number of Obs	65,961	65,961	65,961
Mean of Dep. Var	0.834	6.856	0.073
SD of Dep. Var	9.093	25.270	0.279

Note: The frame is PitchBook India companies founded by 2015 with a known founding year, an office-field cell, and dated venture rounds, each observed from its founding year (or 2010) through 2022. Panel A is the 1,187 frame firms with a completed Early or Later Stage VC round by 2015 (53 pre-2016 patent holders). Panel B is all 6,188 frame firms (126 holders): Panel A plus the 5,001 firms not VC-backed by 2015, funded later or never. Column (1) reports the difference-in-differences estimate of an indicator for filing any patent application in the year (×100) on outside-office same-field examiner capacity interacted with a post-2016 indicator, with firm and office-by-year fixed effects. Columns (2) and (3) report triple-difference estimates with the post period dating from 2018; the dependent variable is an indicator for raising any completed venture round in the year (×100) in column (2) and the number of completed rounds in column (3). The first row is the response of firms without pre-2016 patent filings, and the second the differential response of pre-2016 patent holders, adding patent-holder-by-year fixed effects. Standard errors clustered by office-field cell (48 clusters in Panel A, 57 in Panel B) are in parentheses and wild-cluster bootstrap p-values in brackets; *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

years of its first financing.[25] A similar calculation using Indian patent data, even after

[25]We use patent data through 2017 to make this calculation. Thus, in the most extreme case——a firm financed at the end of 2012 who did not file its first patent until two years after the first round date——the observation would be included if the patent issues within three years of the filing date. (The long run mean pendency is the US is about 26-27 months (https://www.uspto.gov/dashboard/patents/).) We

adjusting for the longer review times at IP India, suggests the probability of a patent filing rate within two years of a first venture financing was only 2%. The limited patenting by venture-backed firms is consistent with the statistics on the low rates of Indian patenting overall documented in Figure 1.

We then examine the consequences of these grants. If faster examination raises the value of patent protection, the firms best positioned to benefit on the financing margin are those that filed for awards at the time of the reform in the targeted areas. The greater likelihood that their pending awards will be granted in a timely manner and be available for exploitation may be a positive signal for venture investors.

We therefore estimate the triple-difference specification

$$\begin{aligned}\text{AnyDeal}_{it} = \beta_1 \left(\text{OutsideCapacity}_{c(i),2015} \times \text{Post}_t\right) \\ + \beta_2 \left(\text{OutsideCapacity}_{c(i),2015} \times \text{Post}_t \times \text{PreHolder}_i\right) \\ + \alpha_i + \delta_{o(i)t} + \lambda_t \,\text{PreHolder}_i + \epsilon_{it},\end{aligned} \tag{5}$$

where AnyDeal_{it} is an indicator for raising any venture round in year t ($\times 100$) and PreHolder_i marks the firms whose first patent filing predates 2016. The fixed effects α_i, $\delta_{o(i)t}$, and $\lambda_t \,\text{PreHolder}_i$ absorb, respectively, persistent firm differences, office-wide funding shocks, and the common time path of patent-holding firms, so β_2 compares patent holders in more versus less exposed cells, relative to how other venture-backed firms diverged across the same cells. Because the median lag from first examination report to grant is roughly two years in our data, accelerated examination translates into granted rights from about 2018 onward, and we date $\text{Post}_t = 1[t \geq 2018]$ accordingly.

Columns (2) and (3) of Table 7 report the estimates of Equation 5. Appendix Figure A.15 shows its event-study version, which replaces Post_t with year indicators (2017, the last pre-treatment year, omitted) so that each point is $\beta_{2,\tau}$: the difference between patent holders' and other firms' financing response to a unit of outside capacity in year τ. The pre-treatment differentials are flat (a joint wild-bootstrap test of the 2010–2016 leads has $p = 0.72$), and turn positive from 2018, when accelerated grants arrive. The static estimate is $+1.14$ percentage points per examiner-equivalent. A one-standard-deviation increase in exposure (4.1 examiner-equivalents) raises a patent holder's annual probability of raising a venture round by about 4.7 percentage points, as compared to a pre-reform base rate of 22 percent. Deal counts and amounts (not reported) show the same pattern.

The aggregate counterpart is telling: for venture-backed firms without pre-reform patents, the corresponding coefficient is slightly negative and statistically indistinguishable from zero, so total venture activity in more exposed cells does not rise. The evidence is

thank Rick Townsend for this analysis.

thus suggestive of a composition effect: the reform shifted venture funding toward firms holding accelerated patent rights rather than raising its level—consistent with patents serving as a financing asset whose value rose when examination delays fell, but operating only within the minority of venture-backed firms that participate in the patent system.

5.5 Estimating the Aggregate Impact

While the analyses above suggest that the reallocation reform was associated with an increase in the quality of Indian patent filings and a greater ability of entrepreneurial firms with patents to raise venture financing, the question remains of the overall effect of the policy reform. This question is salient given the striking historical disparity between India and other leading countries in patenting noted in Figure 1.

We undertake a simple back-of-the-envelope calculation. We employ the regression in column (4) of Table 6: the impact of an additional unit of outside capacity in the post-reform period on the number of applications filed, computed at the examination field-by-office-by year level. The coefficient implies that a cell with one more unit of outside capacity filed about 11.7 more applications per year after 2016 than before.

We multiply the coefficient by the cumulative outside capacity. Since each cell is a field-office in a given year, to compute outside capacity we add up the examiners in other offices in the same field, and then add those numbers up across fields. Suppose there are N examiners in biotech, with n_a, n_b, n_c, and n_d in each of the four offices, and N=n_a+n_b+n_c+n_d. The four biotech cells' outside capacity added together will be (n_b+n_c+n_d) + (n_a+n_c+n_d) + (n_a+n_b+n_d)+(n_a+n_b+n_c) = 3*(n_a+n_b+n_c+n_d) = 3*N. Conceptually, after the reform, each examiner can serve each of the three other offices when needed. Our regression was not measuring work *per se*, but rather measuring how many new examiners you now have access to. We then add up those outside capacity measures across fields.

Using the total number of active examiners in 2015, 97, the total outside capacity is three times this amount, or 291. We then take the product of the coefficient (11.7) and the sum of outside capacity summed over cells (291), roughly 3,415 additional applications per year. Because this estimate is of the total applications, domestic and foreign, we adjust for the proportion of domestic applications between 2016 and 2021 (36.4%), so attribute to the reform an 1,244 additional Indian domestic applications per year (out of a total of 95,063 for the six-year period). This calculation suggests that the reform led to a 8.5% boost the domestic application in the 2016-21 period (6*1,244)/(95,063-(6*1,244)).

We can then quantify how much of the gap between India and other leading nations the reform addressed. To do this, we return to the ratio of annual filings to GDP in Panel (a)

of Figure1. We focus on the years after the reform, from 2016 to 2023. The ratio of annual filings per billion dollars of mean GDP (again in 2015 dollars) was 9.36 for India, while a simple average of the other nations was 15.04. Without the 8.5% boost in patenting from the reform, India's ratio would have been instead 8.62. This implies that the reform closed 11.4% of the gap between India and the mean leading economy (9.36-8.62)/(15.04-8.62).

6 Why Weren't the Reform Effects Bigger?

The reader may find the 11.4% differential surprisingly low: why did the alleviation of these patent delays not have more dramatic effects on innovation? In this section, we examine two potential explanations.

6.1 Could Existing Examiner Capacity Have Been Better Utilized?

One possibility is that the reallocation process was not optimal, as Table 5 already suggested above. In this analysis, we quantify the consequences of the allocation process.

6.1.1 How are Patents Allocated?

We supplement the logit results of Section 4.3 by expanding the set of simulations and comparing the distribution of patents in the simulated versus actual data. We examine the distribution of patents along two dimensions: (1) at the office level (did a patent reviewed in Delhi get routed to Delhi in the simulation?), (2) at the field level (did a chemistry patent get reviewed by an examiner who had handled chemistry before the reform?). We then calculate the mean absolute distance of the office-correct and field-match percentages to observe which simulation most closely matches the actual data. There are six simulated allocation rules we test, and the results are presented in Table 8.

Panel A presents allocation rules which route patents within the origin office either randomly, based on the field of specialization, or based on the shortest queue. Mechanically, these three simulations match office distribution in the actual data 49.5% of the time. Put another way, 49.5% of the patents are reviewed at the origin office. On the field level, we find that patents are overwhelmingly allocated to an examiner who had seen the field before in the actual data (98.1%). Consequently, the overall average gap of the allocation rules to the benchmark is minimized for *Field-based within office* with a difference value of 26.2 percentage points. To visualize the comparison, we plot the percentage of in-field assignments for the benchmark and the random and field-based within-office rules in Appendix Figure A.16.

Table 8: Simulation Outcomes

Allocation rule	Home rate (%)	Office-correct (%)	Field-match (%)	**Avg. gap vs. actual (pp)**
Benchmark (Actual)	49.5	100.0	98.1	**—**
Panel A: Within origin office				
Random within office	100.0	49.5	68.9	**39.9*****
Field-based within office	100.0	49.5	100.0	**26.2*****
Queue-based within office	100.0	49.5	67.0	**40.8*****
Panel B: Across offices (including origin)				
Random across office	28.5	29.0	69.3	**49.9*****
Field-based across office	28.6	29.8	100.0	**36.0*****
Queue-based across office	28.8	28.7	68.8	**50.3*****

Note: $N = 134{,}012$ patents whose RFE falls on or after the 2016 reform and before the end of 2021, so the allocation decision was made under auto-allocation. Cases already queued at the reform are carried as the starting backlog and are never re-routed. *Home rate*: share assigned to their origin office. *Office-correct*: share whose assigned office matches the actual processing office. *Field-match*: share whose assigned examiner had reviewed the patent's field before the reform (scored for all examiners, pool and non-pool). *Avg. gap vs. actual*: mean absolute difference from the benchmark on office-correct and field-match (pp); smaller is closer. Panel A rules route each patent to its origin office (home rate 100% by construction; office-correct = the share never reallocated); Panel B rules route to any of the four offices, the origin office included. Stars: paired t-test of the per-patent composite score (mean of the office-correct and field-match indicators) against the benchmark, testing whether the gap could arise from randomness in the simulated assignment draws.

Panel B presents simulation rules that route patents across all four offices, including the origin office. These rules match the respective simulated rules in the logit model from Table 5. The overall gap between the allocations and the benchmark is again minimized for the *Field-based* rule, further suggesting that there is strong field matching for patent allocation post-reform.

Comparing the two panels, routing every patent home reproduces the actual office 49.5% of the time, against just under 30% for rules that route across offices – consistent with the strong home-office bias in the logit model. Interestingly, the queue-based allocations perform about as well as random allocation in both panels. Since our outcomes are at the office level, this result suggests workload balancing does not describe how patents were actually assigned. Table 8 ultimately reinforces the logit model results, in which the real allocation's queue coefficient is near-zero, home bias is high, and there is some field consideration.

6.1.2 A Back-of-the-Envelope Calculation of the Impacts of Minimizing Delays

Our allocation analysis revealed that offices were not routing patents in a queue-optimized manner post-reform, resulting in an inefficient distribution of examiner workload. In this section, we quantify the pure-efficiency gains, had IP India allocated patents based purely on minimizing examiner queue, sans new hiring. Ultimately, we find that had patents been allocated optimally, the number of good applications would have increased by between 9% and 31% and total applications would have increased by 20%.

We begin by constructing two panels at the examiner-day level, and consider only the set of examiners who were active pre-reform.[26] The first panel is based on the actual allocation data, and contains information about the number of incoming patents,[27] the number of outgoing patents,[28] the count of patents in the examiner's backlog at the reform date, and the queue. The queue is calculated as the backlog plus the number of incoming patents minus the number of outgoing patents. The second panel holds the same information, but the number of patents incoming is simulated based on the within-field examiner with the shortest queue.

On average, the queue facing an incoming patent is 373.2 patents in the actual data and 212.9 patents in the simulated data, with a difference of 160.3. To understand how queue length impacts pendency time, we run the following OLS model:

$$\text{DaysToFER}_i = \alpha + \beta\, \text{QueueLength}_{io} + \epsilon \tag{6}$$

where QueueLength_{io} indexes the size of examiner o's queue when patent i is assigned, DaysToFER_i indexes the number of days from RFE to FER for patent i, and β represents the additional days of pendency associated with one additional patent in an examiner's queue. Results are presented in column (1) of Panel A of Table 9 and suggest that one additional patent in an assigned examiner's queue is associated with a 0.81-day increase in pendency. Column (2) of Panel A shows the estimate is robust to including a field fixed effect.

Scaling the column (1) estimate by the 160.3 patent queue difference between the actual and simulated data, optimal allocation reduces an incoming patent's pendency by $0.8078 \times 160.3 = 129.5$ days. This scaling presumes that a patent's position in its examiner's queue is what determines when it is examined; were examiners working in an arbitrary order, a shorter queue would not translate into a shorter wait. Ordering is in fact close to first-in-first-out: for 88% of patent pairs within an examiner's caseload, the patent

[26] Our data show only about 1% of post-reform patents were assigned to an examiner who entered the data post-reform.

[27] Based on the patent RFE date.

[28] Based on the patent FER date.

that arrived first also received its FER first. Appendix Figure A.17 shows the results of the OLS relative to a FIFO calibration; the data track the OLS, so we treat the FIFO slope as an upper bound.

The remaining inputs are the reform's reduced-form effects on pendency, the share of good applications, and total applications, which are the difference-in-differences estimates reported in Table 6.

We then estimate how optimal allocation might have changed patenting post-reform using the following "back-of-the-envelope" (BOTE) calculation. Let $D = \beta_b \times \Delta\text{Queue}$ denote the days of pendency saved per application (129.5), and $\gamma_o = \beta_o / \beta_{pendency}$ the effect of one pendency day on outcome o:

$$\text{Additional Total Patents/Year} = -\gamma_{total} \times D \times 76 \tag{7}$$

$$\text{Additional Good Patents/Year} = -\gamma_{share} \times D \times \bar{T} \tag{8}$$

where $\Delta\text{Queue} = 160.3$ and $\beta_b = 0.8078$ come from the simulation and the OLS in column (1) of Panel A, β_{share} and β_{total} are the DiD estimates for the share of good applications and total applications (cell + office-year FE; columns (3) and (4) of Table 6), $\beta_{pendency} = -13.54$ is the pendency DiD (column (2)), and $\bar{T} = 43{,}485$ is average observed total filings per year over 2016-2021.

The results are in Panel B of Table 9. Because the total-applications DiD fails its pre-trends test (Appendix Figure A.13), Equation 8 holds total filings at their observed post-reform level $\bar{T}$ and credits only the shift in the *share* of good applications. The "full counterfactual" row instead also credits the estimated volume response and its interaction with the share, $\Delta\text{Good} = \Delta s \times \bar{T} + s_{post}\Delta T + \Delta s \Delta T$, raising the gain from roughly 945 to 3,127 good applications per year. We conclude that an optimal allocation would have increased good applications by at least 9% and by as much as 31%.

We can again do a back-of-the envelope calculation of what would have been the impact on the gap in the patenting-to-GDP ratio, had the ideal reallocation been implemented. The 8,529 additional annual filings would have translated, using the same ratio as before to 3,107 additional Indian applications each year. The hypothetical reform's contribution would have been a 29.8% boost in domestic patenting (1,244+3,107)*6 / ((95,063+3,107*6)-(1,244+3,107)*6), as opposed to the estimated 8.5% increase.

We again look at the years the reform was in force (2016–2023). This increase would have brought India's ratio of annual filings per billion dollars of GDP up from 8.6 to 11.2 (8.6*1.298). Thus, the reform would have narrowed the gap between India and the average major nation by 40.0%, rather than the 11.4% estimated above (11.2-8.6)/(15.0-8.6).

One final question that we examine here relates to the heterogeneity of these effects.

Because examiner capacity varies across fields, the simulated queue reduction differs by field, ranging from about 7.8 in Food to 282.5 in General Engineering. The baseline calculation in the BOTE equation applies the economy-wide average reduction to all 76 cells equally. As a robustness check, we instead assign each cell the queue reduction simulated for its own field and sum across cells:

$$\text{Additional Total Patents/Year} = -\gamma_{total} \times \beta_b \sum_{c=1}^{76} \Delta\text{Queue}_c \tag{9}$$

$$\text{Additional Good Patents/Year} = -\gamma_{share} \times \beta_b \sum_{c=1}^{76} \Delta\text{Queue}_c \times \bar{T}_c \tag{10}$$

where c indexes a cell, ΔQueue_c is the queue reduction simulated for cell c's field, and $\bar{T}_c$ is cell c's average observed annual filings. Results are presented in Panel C of Table 9, and show that the original BOTE results in Panel B are broadly robust to the field-specific queue specification.

6.2 The Impact of Patent Litigation

Courts to litigate patents—whether to adjudicate claims of patent infringement or to appeal decisions of patent examiners—are an essential complementary institution to patent offices. (Schankerman and Schuett (2022) models the interactions between courts and patent offices.) To the extent that the ability to litigate Indian patents has lagged that the reforms of the patent examination, it may have rendered these changes less effective.

Indeed, Indian patent litigation has been characterized by substantial changes and policy uncertainty. There are three kinds of cases in the Indian system. The first class of proceeding is an infringement suit, brought by a patentee against a rival alleged to be infringing a patent under Section 104 of the Patents Act 1970. The other two proceedings are the revocation petitions, brought by a challenger against a patentee under Section 64 of Patents Act to cancel a granted patent. The third is the appeal under Section 117A, brought by a party aggrieved by a decision of the Controller (typically, a refusal to grant a patent).

Initially, all proceedings were heard in one of the several hundred District Courts and a select number of High Courts. In 2007, India moved patent revocation petitions and appeals out of the courts to a specialist tribunal called the Intellectual Property Appellate Board. The specific stipulation of board membership, however, meant that when the technical member for patent reviews stepped down in 2016, the review process came to a halt. (The vacancy persisted because the appointment became part of a wider dispute between the government and the Supreme Court over the control of tribunal appointments.) After a brief attempt by the Board to address the backlog in 2020, the

Table 9: Back-of-the-Envelope Gains from Optimal Queue Allocation

Panel A: Pendency and Queue				
	(1)	(2)		
Queue at assignment (β_b)	0.8078***	0.7980***		
	(0.0470)	(0.0498)		
Field fixed effects	No	Yes		
Observations	133,218	133,218		
R^2	0.279	0.352		
Panel B: BOTE Gains				
	γ	Additional filings/year	Additional filings (2016–2021)	% of baseline
Good applications	−0.00017	944.6	5,668	9.27
Full counterfactual	–	3,127.4	18,764	30.71
Total applications	−0.8666	8,528.5	51,171	19.61
Panel C: Robustness – Field-specific queue				
	γ	Additional filings/year	Additional filings (2016–2021)	% of baseline
Good applications	−0.00017	927.0	5,562	9.10
Total applications	−0.8666	6,951.7	41,710	15.99

Note: Panel A: patent-level regressions of days from request for examination to first examination report on the assigned examiner's queue at assignment; standard errors clustered by examiner. Panel B: γ is the ratio of each outcome's preferred DiD estimate (cell + office-year FE; Table 6) to the pendency estimate; the chain uses β_b from Panel A column (1), under which the simulated queue reduction (160.3 patents) saves 129.5 days per patent. Good applications are filings by applicants with at least three pre-2016 disposed applications and a pre-2016 grant rate of at least 70%; baseline 10,185 of 43,485 total filings/year. For scale, pre-reform (2010–2015) cell-year means are 1,482 days of RFE-to-FER pendency, a good-application share of 0.190, and 512 filings. The full-counterfactual row adds the volume response and its interaction with the composition margin. Panel C substitutes each cell's own queue reduction and filings for the economy-wide averages. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Tribunals Reforms Ordinance abolished the tribunal in 2021. All cases where transferred to the High Courts: Delhi alone receiving nearly three thousand pending cases. Motivated by the need to respond to this massive influx, as well as the desire to improve the quality of the adjudication of these cases, the Delhi High Court in 2022 began hearing these cases in a new Intellectual Property Division. This step that was subsequently emulated by the other High Courts located in cities with offices of IP India.

To assess the consequences of these changes, we build as comprehensive as possible a dataset on patent litigation in India, as described in Appendix Appendix D. Figure 13 illustrates the dramatic shifts associated with these changes: the left panel tabulates cases filed by year and the right panel the cases by year of decision. (We only observe in some cases resolved cases, so it is likely the last few years of the left panel understate the level of activity.) The table make clear that there were an explosion of cases in 2022, with an

increase of decisions in subsequent years.

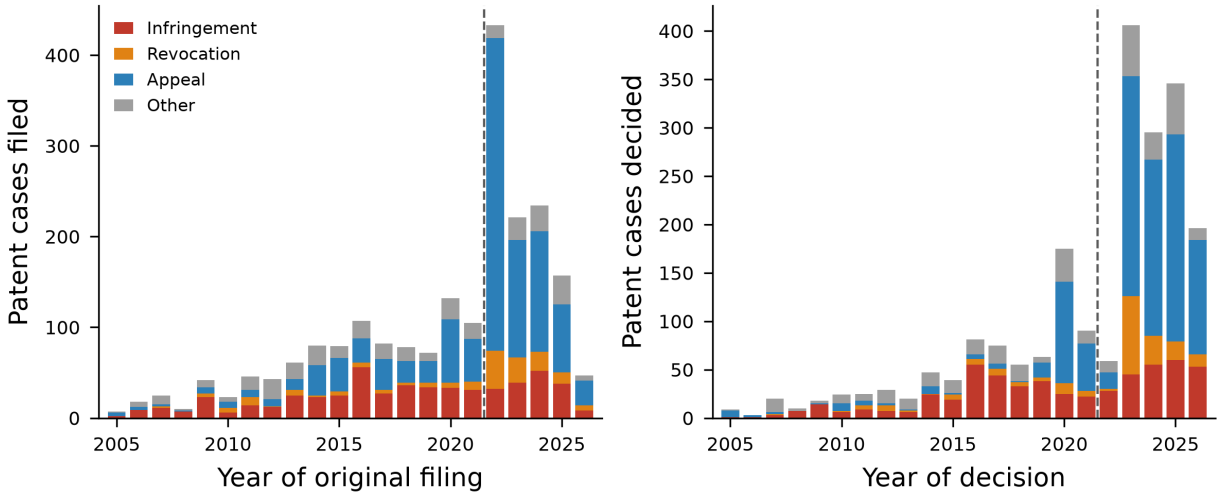


Figure 13: Patent litigation over time by filing and decision date. The left panel reports cases by year of original filing; the right panel reports cases by year of decision.

One hypothesis is that the substantial uncertainty surrounding the adjudication of patents prior to 2022 might have retarded the willingness of inventors to file for patent awards. The consequences of the litigation reform cannot be fully analyzed, since many of the recent applications are still being processed (and Table 1 suggests the need for caution when interpreting counts of Indian patent filings)./footnoteAssessing the impact of the creation of the specialized patent courts is an interesting question, but the recency of the shift and the limited scope of the natural experiment that we study make it challenging task. We leave the assessment of this shift, as well as other patent policies that India could have enacted–from seeking to accelerate the second half of the examination process from FER to issue (a more complex process, due to the back-and-forth interactions between examiners and applicants) to a shift in technologies covered by patent protection—to future work. But IP India's *2024-25 Annual Report* indicates the number of domestic filings has increased from 29,508 in fiscal year 2021-22 to 68,201 in 2024-25. This increase suggests that the reform of a complementary institution may have needed to realize the full effects of the 2016 reform to patent examination.

7 Conclusion

India's innovation puzzle is that low patenting coexists with substantial innovative activity outside the domestic patent system. As indicated by metrics ranging from scientific

publications and venture capital to open-source software contributions and the quality of Indian inventors' U.S. patents, India is highly innovative. Yet its domestic patenting remains far below the levels that its national income, and especially its human capital, would predict.

We have examined several explanations. Industry composition, popular attitudes toward intellectual property, and the difficulty of realizing value from issued patents each appear to account for, at most, only a portion of the gap. The evidence is most consistent with institutional frictions in the patent office, in particular, very long examination delays. Exploiting the 2016 auto-allocation reform, we find that field-office cells that gained access to more same-field examiner capacity saw substantially faster RFE-to-FER processing. In addition, the composition of applications shifted away from low-grant-rate ("unserious") filers toward more credible applicants and boosted venture capital activity.

Nonetheless, the results suggest that the reforms had a relatively modest effect in the years after the reform. We explore two possible explanations: first, the way in which the reform was implemented and second, the fact that major changes to a complementary institution—the system to litigate patent awards—did not happen until eight years later.

More broadly, these results suggest that the value of patents in a developing economy depends not only on the formal strength of the law but also on state capacity of the patent system to issue useful rights in a timely manner. This margin that has received little attention in the literature, despite the fact that it is conducive to policy reform.

References

Acemoglu, Daron, Simon Johnson, and James A. Robinson, "The Colonial Origins of Comparative Development: An Empirical Investigation," *American Economic Review*, 2001, *91* (5), 1369–1401.

_ , _ , and _ , "Institutions as a Fundamental Cause of Long-Run Growth," in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth, Volume 1A*, New York: Elsevier, 2005, pp. 386–472.

Acemoglu, Darron and James A. Robinson, *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*, New York: Crown Business, 2012.

Akcigit, Ufuk, Daron Acemoglu, and Simon Johnson, "Technology and Economic Development," in Pascaline Dupas, Pinelopi K. Goldberg, and Rohini Pande, eds., *Handbook of Development Economics, Volume 6*, New York: Elsevier, 2026.

Amirapu, Amrit, "Justice Delayed Is Growth Denied: The Effect of Slow Courts on Relationship-Specific Industries in India," *Economic Development and Cultural Change*, 2021, *70* (1), 415–451.

Aoki, Masahiko, *Toward a Comparative Institutional Analysis*, MIT Press, 2001.

Ash, Elliott, Sam Asher, Aditi Bhowmick, Sandeep Bhupatiraju, Daniel Chen, Tanaya Devi, Christoph Goessmann, Paul Novosad, and Bilal Siddiqi, "In-Group Bias in the Indian Judiciary: Evidence from 5 Million Criminal Cases," *Review of Economics and Statistics*, 2026.

Bai, Jie, Panle Jia Barwick, Shengmao Cao, and Shanjun Li, "Quid Pro Quo, Knowledge Spillover, and Industrial Quality Upgrading: Evidence from the Chinese Auto Industry," *American Economic Review*, 2025, *115* (11), 3825–3852.

Bernstein, Shai, Xavier Giroud, and Richard R. Townsend, "The Impact of Venture Capital Monitoring," *Journal of Finance*, 2016, *71* (4), 1591–1622.

Boehm, Johannes and Ezra Oberfield, "Misallocation in the Market for Inputs: Enforcement and the Organization of Production," *Quarterly Journal of Economics*, November 2020, *135* (4), 2007–2058.

Branstetter, Lee G., Raymond Fisman, and C. Fritz Foley, "Do Stronger Intellectual Property Rights Increase International Technology Transfer? Empirical Evidence from U.S. Firm-Level Panel Data," *Quarterly Journal of Economics*, 2006, *121* (1), 321–349.

Budish, Eric, Benjamin N. Roin, and Heidi Williams, "Do Firms Underinvest in Long-Term Research? Evidence from Cancer Clinical Trials," *American Economic Review*, 2015, *105* (7), 2044–85.

Chaudhuri, Shubham, Pinelopi K. Goldberg, and Panle Jia Jia, "Estimating the Effects of Global Patent Protection in Pharmaceuticals: A Case Study of Quinolones in India," *American Economic Review*, 2006, *96*, 1477–1514.

Chemin, Matthieu, "Judicial Efficiency and Firm Productivity: Evidence from a World Database of Judicial Reforms," *Review of Economics and Statistics*, 2020, *102* (1), 49–64.

Chhina, Raman S., "Startups and the State," Working Paper, University of Chicago 2026.

Cockburn, Iain M., Jean O. Lanjouw, and Mark Schankerman, "Patents and the Global Diffusion of New Drugs," *American Economic Review*, 2016, *106* (1), 136–164.

Colonnelli, Emanuele, Bo Li, and Ernest Liu, "Investing with the Government: A Field Experiment in China," *Journal of Political Economy*, 2024, *132* (1), 248–294.

Commonwealth Secretariat and United Nations Conference on Trade and Development, *Harnessing Intellectual Property Rights for Innovation, Development and Economic Transformation in Least Developed Countries*, New York: United Nations, 2024.

Department of Science and Technology, Ministry of Science and Technology, Government of India, *Research and Development Statistics, 2025-26*, Government of India, 2026.

Dorner, Matthias and Dietmar Harhoff, "A Novel Technology-Industry Concordance Table Based on Linked Inventor-Establishment Data," *Research Policy*, 2018, *47* (4), 768–781.

Duggan, Mark, Craig Garthwaite, and Aparajita Goyal, "The Market Impacts of Pharmaceutical Product Patents in Developing Countries: Evidence from India," *American Economic Review*, 2016, *106* (1), 99–135.

Fan, Tianyu, Michael Peters, and Fabrizio Zilibotti, "Growing Like India – The Unequal Effects of Service-Led Growth," *Econometrica*, 2023, *91* (4), 1457–1494.

Farre-Mensa, Joan, Deepak Hegde, and Alexander Ljungqvist, "What Is a Patent Worth? Evidence from the U.S. Patent 'Lottery'," *Journal of Finance*, 2020, *75* (2), 639–682.

Feng, Josh and Xavier Jaravel, "Crafting Intellectual Property Rights: Implications for Patent Assertion Entities, Litigation, and Innovation," *American Economic Journal: Applied Economics*, 2020, *12* (1), 140–81.

Finan, Frederico, Benjamin A Olken, and Rohini Pande, "The personnel economics of the developing state," *Handbook of economic field experiments*, 2017, *2*, 467–514.

Freilich, Janet and Soomi Kim, "Is the Patent System Sensitive to Incorrect Information?," *Review of Economics and Statistics*, 2026, *108* (2), 1–38.

Gans, Joshua S., David H. Hsu, and Scott Stern, "The Impact of Uncertain Intellectual Property Rights on the Market for Ideas: Evidence from Patent Grant Delays," *Management Science*, 2008, *54* (5), 982–997.

Gaulé, Patrick, "Patents and the Success of Venture-Capital Backed Startups: Using Examiner Assignment to Estimate Causal Effects," *Journal of Industrial Economics*, 2018, *66* (2), 350–376.

Griliches, Zvi, *R&D, Patents, and Productivity*, University of Chicago Press, 1984.

Halliburton, Murphy, *India and the Patent Wars: Pharmaceuticals in the New Intellectual Property Regime*, Cornell University Press, 2017.

Jaffe, Adam B. and Manuel Trajtenberg, *Patents, Citations, and Innovations: A Window on the Knowledge Economy*, MIT Press, 2002.

Kala, Namrata, Josh Lerner, and Junxi Liu, "Constructing the Indian Patent Dataset," Unpublished Working Paper, Massachusetts Institute of Technology 2025.

Kalyani, Aakash, Nicholas Bloom, Marcella Carvalho, Tarek A. Hassan, Josh Lerner, and Ahmed Tahoun, "Diffusion of New Technologies," *Quarterly Journal of Economics*, 2025, *140*, 1299–1365.

Keller, Wolfgang, "International Technology Diffusion," *Journal of Economic Literature*, 2004, *42* (4), 752–782.

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

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

Kondylis, Florence and Mattea Stein, "The Speed of Justice," *Review of Economics and Statistics*, May 2023, *105* (3), 596–613.

König, Michael, Kjetil Storesletten, Zheng Song, and Fabrizio Zilibotti, "From Imitation to Innovation: Where is All that Chinese R&D Going?," *Econometrica*, 2022, *90* (4), 1615–1654.

Lai, Edwin L.-C., "International Intellectual Property Rights Protection and the Rate of Product Innovation," *Journal of Development Economics*, 1998, *55* (1), 133–153.

Landes, William M. and Richard A. Posner, *The Economic Structure of Intellectual Property Law*, Cambridge: Harvard University Press, 2003.

Lerner, Josh, Junxi Liu, Jacob Moscona, and David Yang, "Appropriate Entrepreneurship? The Rise of China and the Developing World," Working Paper 32193, National Bureau of Economic Research 2024.

Liu, Kung-Chung and Uday S. Racherla, eds, *Innovation, Economic Development, and Intellectual Property in India and China: Comparing Six Economic Sectors*, SpringerOpen, 2019.

Long, Cheryl Xiaoning and Jun Wang, "Evaluating patent promotion policies in China: Consequences for patent quantity and quality," in "Economic impacts of intellectual property-conditioned government incentives," Springer, 2016, pp. 235–257.

Magerman, Tom, Bart Van Looy, and Xiaoyan Song, "Data Production Methods for Harmonized Patent Statistics: Patentee Name Harmonization," SSRN Scholarly Paper 944470, Social Science Research Network 2006.

Matcham, William and Mark Schankerman, "Screening Property Rights for Innovation," *Econometrica*, 2026, *forthcoming*.

Meho, Lokman I. and Elie A. Akl, "Using Bibliometrics to Detect Questionable Authorship and Affiliation Practices and their Impact on Global Research Metrics: A Case Study of 14 Universities," *Quantitative Science Studies*, 2025, *6*, 63–98.

Moscona, Jacob and Karthik Sastry, "Inappropriate Technology: Evidence from Global Agriculture," Working Paper 33500, National Bureau of Economic Research 2025.

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

North, Douglass C., *Institutions, Institutional Change and Economic Performance*, Cambridge: Cambridge University Press, 1990.

Okunogbe, Oyebola and Gabriel Tourek, "How can lower-income countries collect more taxes? The role of technology, tax agents, and politics," *Journal of Economic Perspectives*, 2024, *38* (1), 81–106.

Peters, Michael, Youdan Zhang, and Fabrizio Zilibotti, "Skipping the Factory: Service-Led Growth and Structural Transformation in the Developing World," Working Paper 34692, National Bureau of Economic Research 2026.

Putnam, Jonathan, "The Value of International Patent Rights," Unpublished Working Paper, Yale University 1997.

Rao, Manaswini, "Courts Redux: Micro-Evidence from India," October 2021. Working paper. Previously circulated as "Judges, Lenders, and the Bottom Line: Court-ing Firm Growth in India" and "Judicial Capacity Increases Firm Growth Through Credit Access: Evidence from Clogged Courts of India".

Reddy T., Prashant and Aparajita Lath, "143 Patent Infringement Lawsuits between 2005 and 2015," *SpicyIP*, 2017.

_ **and Sumathi Chandrashekara**, *Create, Copy, Disrupt: India's Intellectual Property Dilemmas*, Oxford University Press, 2017.

Romer, Paul M., "Endogenous Technological Change," *Journal of Political Economy*, 1990, *98* (5), S71–S102.

Sanyal, Sanjeev and Aakanksha Arora, "Why India Needs to Urgently Invest in its Patent Ecosystem," EAC-PM Paper Series EAC-PM/WP/1/2022, Economic Advisory Council to the Prime Minister 2022.

Schankerman, Mark and Florian Schuett, "Patent Screening, Innovation, and Welfare," *Review of Economic Studies*, 2022, *89* (4), 2101–2148.

Schultz, Mark and Kevin Madigan, *The Long Wait for Innovation: The Global Patent Pendency Problem*, Arlington, VA: Center for the Protection of Intellectual Property, Antonin Scalia Law School, George Mason University, 2016.

Scotchmer, Suzanne, *Innovation and Incentives*, Cambridge: MIT Press, 2004.

Scott Morton, Fiona and Carl Shapiro, "Patent Assertions: Are We Any Closer to Aligning Reward to Contribution?," *Innovation Policy and the Economy*, 2016, *16*, 89–133.

Singh, Pratiba M., *Patent Law*, Bangalore: Thomson Reuters, 2024.

World Intellectual Property Organization, *World Intellectual Property Report: Making Innovation Policy Work for Development*, World Intellectual Property Organization, 2024.

Xu, Guo, Erika Deserranno, Diana Moreira, and Edoardo Teso, "Bureaucracy," *VoxDevLit*, 2023, *8*, 3.

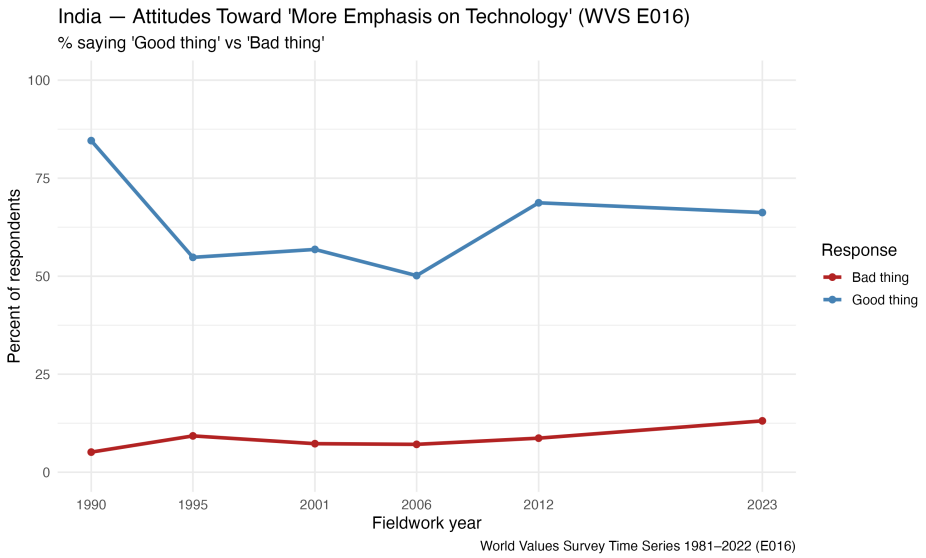
Online Appendix for: The Indian Innovation Puzzle

by Namrata Kala, Josh Lerner, and Junxi Liu

Appendix A Additional Figures and Tables

Figure A.1: Indian Attitudes Toward Technology and Property Rights

(a) Attitudes Toward Technology (WVS E016)



(b) Is Stealing Property Justifiable? (WVS F114B)



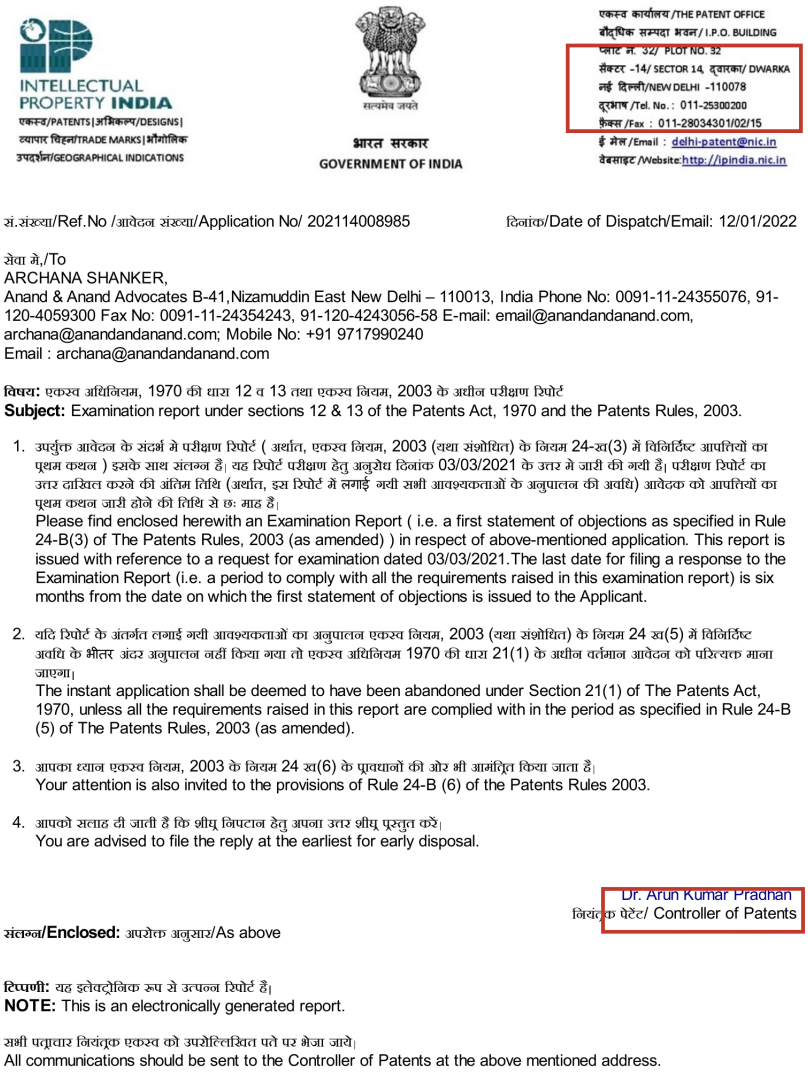
Note: Panel (a) plots responses to WVS question E016 ("More emphasis on technology development": good thing vs. bad thing) for India, 1990–2023. Panel (b) plots responses to WVS question F114B ("Is stealing property justifiable?") for India across available waves. Data from the World Values Survey Time Series 1981–2022.

Figure A.2: Determinants of Indian Patent Allocation: Logit Model Results



Note: These figures display the results of equation 2 for the variable indicated in the graph sub-header, estimated separately by origin office and RFE year. Home office = 1 if the candidate office equals the patent's origin office; Field supply = examiner-equivalents in the patent's field at the candidate office, summing each examiner's pre-reform caseload share over the examiners sitting there that year; Queue = the candidate office's caseload per examiner at the patent's own RFE date. Field supply and queue are standardized within RFE year. These models omit the alternative-specific office constants, which are collinear with Home office within a single origin office. Bands are 95% confidence intervals. Off-scale points are marked with a triangle and labeled. The queue-based rule nearly equalizes office queues by 2020, inflating its queue coefficient and making it incomparable to earlier years.

Figure A.3: Example First Examination Report (FER)



INTELLECTUAL PROPERTY INDIA
एकस्व/PATENTS|अभिकल्प/DESIGNS|
व्यापार चिह्न/TRADE MARKS|भौगोलिक
उपदर्शन/GEOGRAPHICAL INDICATIONS

सत्यमेव जयते
भारत सरकार
GOVERNMENT OF INDIA

एकस्व कार्यालय /THE PATENT OFFICE
बौद्धिक सम्पदा भवन/ I.P.O. BUILDING
प्लाट नं. 32/ PLOT NO. 32
सैक्टर -14/ SECTOR 14, द्वारका/ DWARKA
नई दिल्ली/NEW DELHI -110078
दूरभाष /Tel. No.: 011-25300200
फैक्स /Fax : 011-28034301/02/15
ई मेल /Email : delhi-patent@nic.in
वेबसाइट /Website:http://ipindia.nic.in

सं.संख्या/Ref.No /आवेदन संख्या/Application No/ 202114008985 दिनांक/Date of Dispatch/Email: 12/01/2022

सेवा मे,/To
ARCHANA SHANKER,
Anand & Anand Advocates B-41,Nizamuddin East New Delhi – 110013, India Phone No: 0091-11-24355076, 91-120-4059300 Fax No: 0091-11-24354243, 91-120-4243056-58 E-mail: email@anandandanand.com, archana@anandandanand.com; Mobile No: +91 9717990240
Email : archana@anandandanand.com

विषय: एकस्व अधिनियम, 1970 की धारा 12 व 13 तथा एकस्व नियम, 2003 के अधीन परीक्षण रिपोर्ट
Subject: Examination report under sections 12 & 13 of the Patents Act, 1970 and the Patents Rules, 2003.

1. उपर्युक्त आवेदन के संदर्भ मे परीक्षण रिपोर्ट (अर्थात, एकस्व नियम, 2003 (यथा संशोधित) के नियम 24-ख(3) में विनिर्दिष्ट आपत्तियों का प्रथम कथन) इसके साथ संलग्न है। यह रिपोर्ट परीक्षण हेतु अनुरोध दिनांक 03/03/2021 के उत्तर मे जारी की गयी है। परीक्षण रिपोर्ट का उत्तर दाखिल करने की अंतिम तिथि (अर्थात, इस रिपोर्ट में लगाई गयी सभी आवश्यकताओं के अनुपालन की अवधि) आवेदक को आपत्तियों का प्रथम कथन जारी होने की तिथि से छः माह है।
Please find enclosed herewith an Examination Report (i.e. a first statement of objections as specified in Rule 24-B(3) of The Patents Rules, 2003 (as amended)) in respect of above-mentioned application. This report is issued with reference to a request for examination dated 03/03/2021.The last date for filing a response to the Examination Report (i.e. a period to comply with all the requirements raised in this examination report) is six months from the date on which the first statement of objections is issued to the Applicant.

2. यदि रिपोर्ट के अंतर्गत लगाई गयी आवश्यकताओं का अनुपालन एकस्व नियम, 2003 (यथा संशोधित) के नियम 24 ख(5) में विनिर्दिष्ट अवधि के भीतर अंदर अनुपालन नहीं किया गया तो एकस्व अधिनियम 1970 की धारा 21(1) के अधीन वर्तमान आवेदन को परित्यक्त माना जाएगा।
The instant application shall be deemed to have been abandoned under Section 21(1) of The Patents Act, 1970, unless all the requirements raised in this report are complied with in the period as specified in Rule 24-B (5) of The Patents Rules, 2003 (as amended).

3. आपका ध्यान एकस्व नियम, 2003 के नियम 24 ख(6) के प्रावधानों की ओर भी आमंत्रित किया जाता है।
Your attention is also invited to the provisions of Rule 24-B (6) of the Patents Rules 2003.

4. आपको सलाह दी जाती है कि शीघ्र निपटान हेतु अपना उत्तर शीघ्र पूस्तुत करें।
You are advised to file the reply at the earliest for early disposal.

Dr. Arun Kumar Pradhan
नियंत्रक पेटेंट/ Controller of Patents

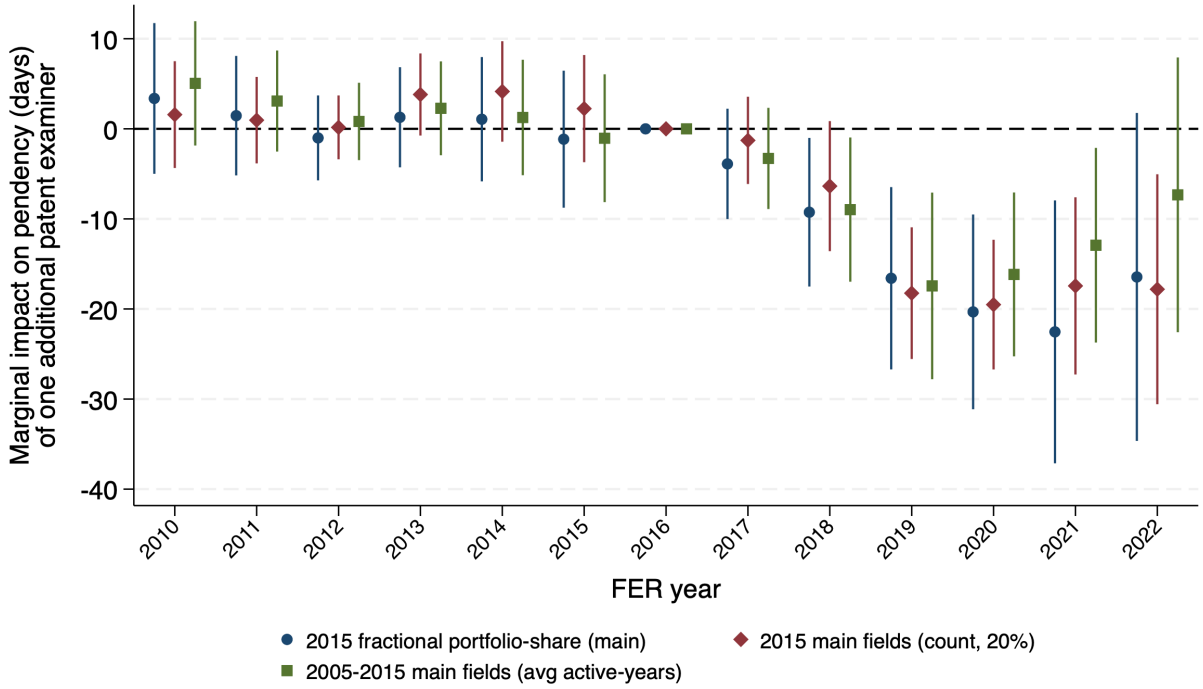
संलग्न/Enclosed: अपरोक्त अनुसार/As above

टिप्पणी: यह इलेक्ट्रोनिक रूप से उत्पन्न रिपोर्ट है।
NOTE: This is an electronically generated report.

सभी पत्राचार नियंत्रक एकस्व को उपरोल्लिखित पते पर भेजा जाये।
All communications should be sent to the Controller of Patents at the above mentioned address.

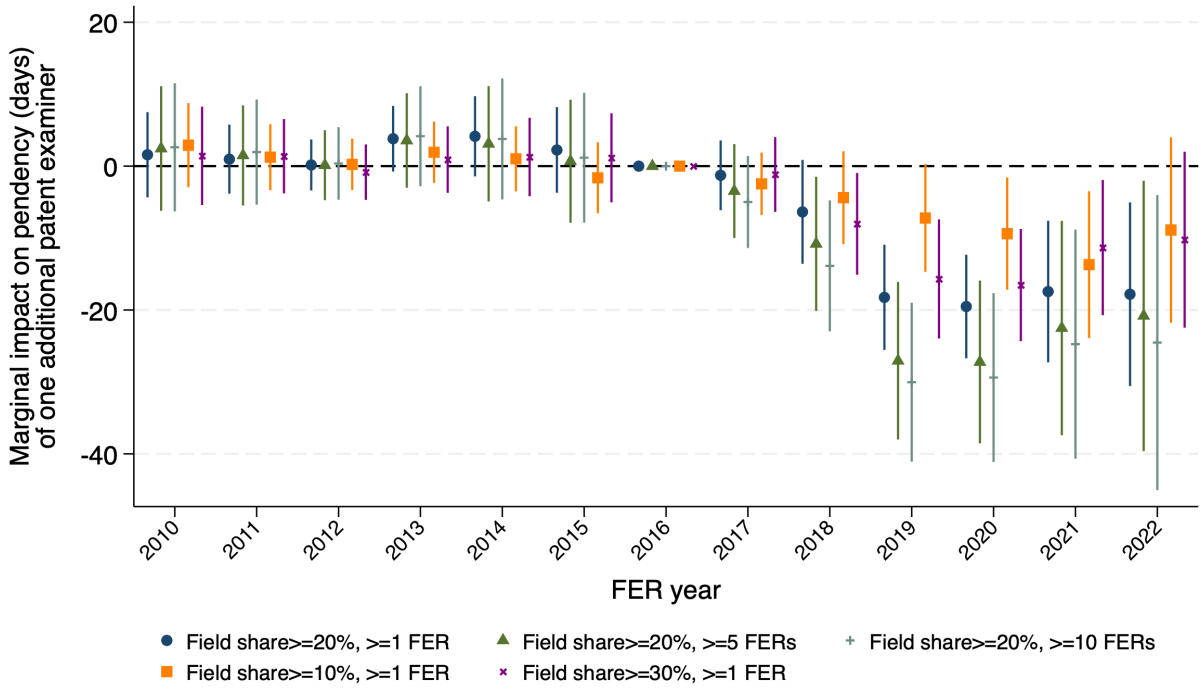
Note: Each application's First Examination Report is a scanned PDF that names the issuing office and the controller (examiner) who signed it, together with the dispatch date. We extract these fields using a combination of OCR (for printed documents) and large language models (for scanned or handwritten documents), and disambiguate controller names into consistent examiner identities.

Figure A.4: Robustness of the Pendency Effect: Capacity Definition



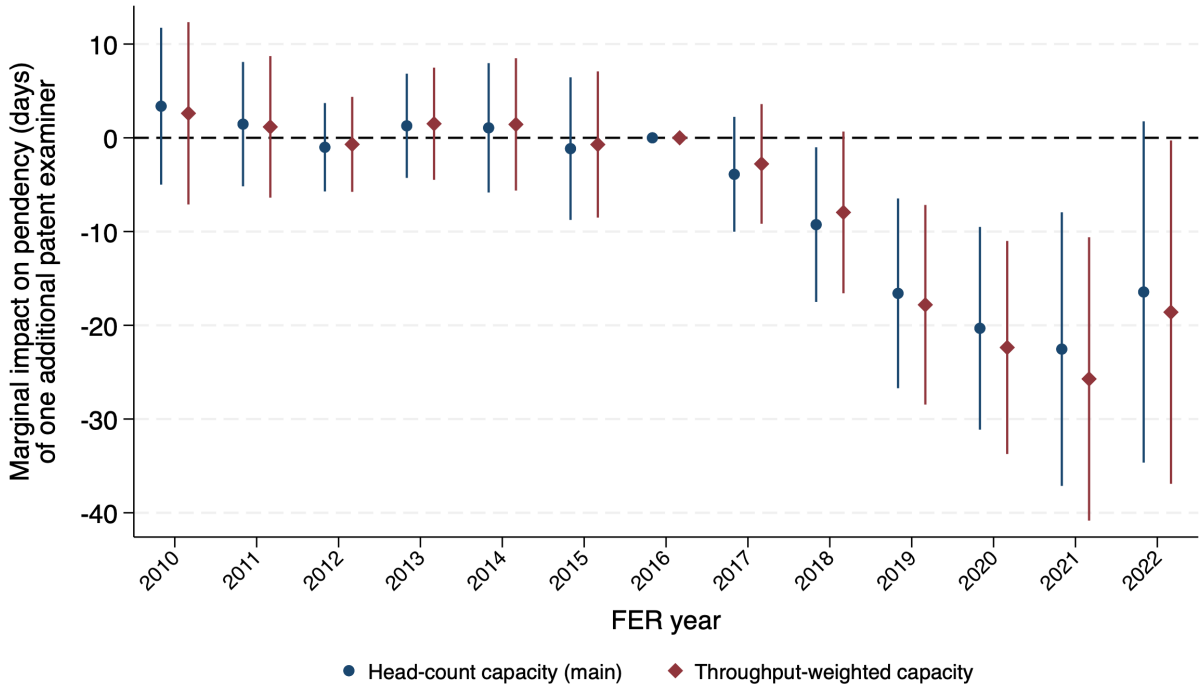
Note: This figure reproduces the main RFE-to-FER pendency event study varying the capacity definition. In each case, a cell's exposure is the capacity in its field at offices other than the origin office. The baseline *2015 fractional portfolio-share* measure splits each 2015 examiner's single unit of capacity across fields in proportion to her 2015 caseload. The *2015 main-field count* instead counts an examiner as one full unit in every field that accounts for more than 20% of her 2015 caseload. The *2005–2015 main-field* measure extends the window: examiners with at least ten FERs over 2005–2015 are assigned to the fields exceeding 20% of their pre-period caseload and contribute the number of years in which they issued at least one FER, averaged over the eleven-year window, with each examiner's home office defined as the office issuing the most of her pre-window reports. The omitted year is 2016, standard errors are clustered by field-office cell, and 95% confidence intervals are shown.

Figure A.5: Robustness of the Pendency Effect: Treatment Measure



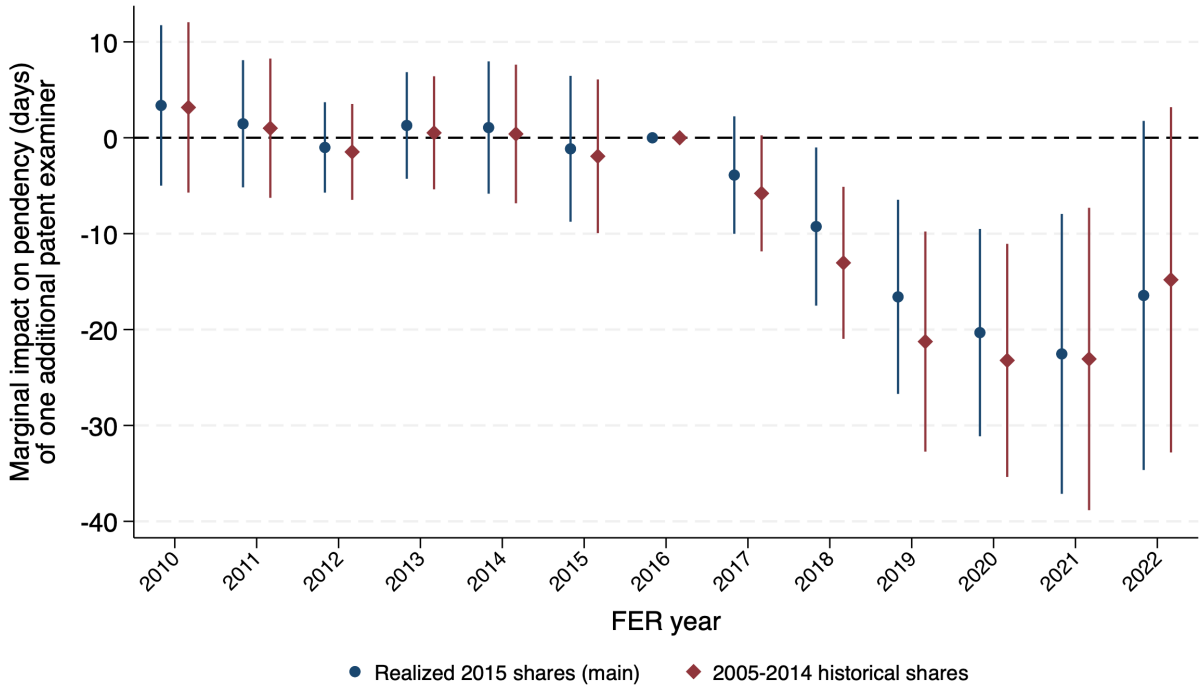
Note: This figure reproduces the main RFE-to-FER pendency event study under alternative versions of the threshold-count treatment measure, in which a cell's exposure is the head-count of outside-office examiners for whom the cell's field is a main field. The series vary the main-field threshold (the field must exceed 10%, 20%, or 30% of the examiner's 2015 caseload) and the minimum activity requirement (at least 1, 5, or 10 FERs issued in 2015); the 20% threshold with no activity minimum corresponds to the count measure of Figure A.4. The omitted year is 2016, standard errors are clustered by field-office cell, and 95% confidence intervals are shown.

Figure A.6: Robustness of the Pendency Effect: Throughput-Weighted Capacity



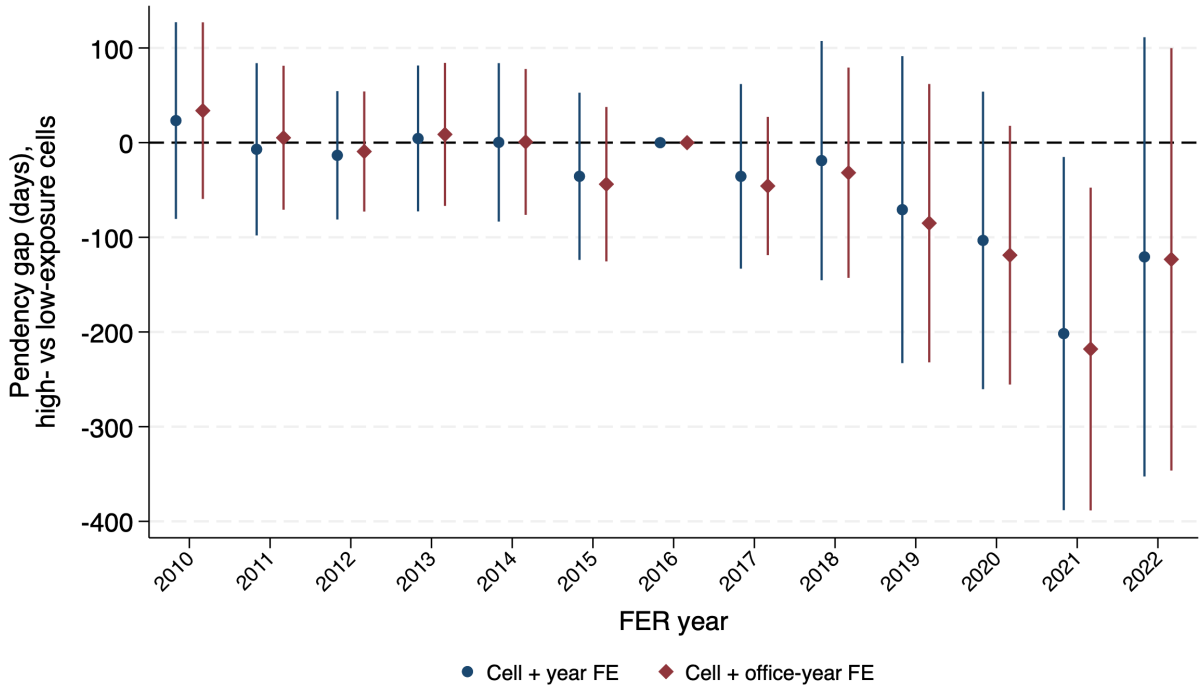
Note: This figure reproduces the main RFE-to-FER pendency event study weighting each examiner by the number of first examination reports issued in 2015 rather than counting heads, scaled so that an average-throughput examiner contributes one unit. The omitted year is 2016, standard errors are clustered by field-office cell, and 95% confidence intervals are shown.

Figure A.7: Robustness of the Pendency Effect: Historical Field-Share Weights



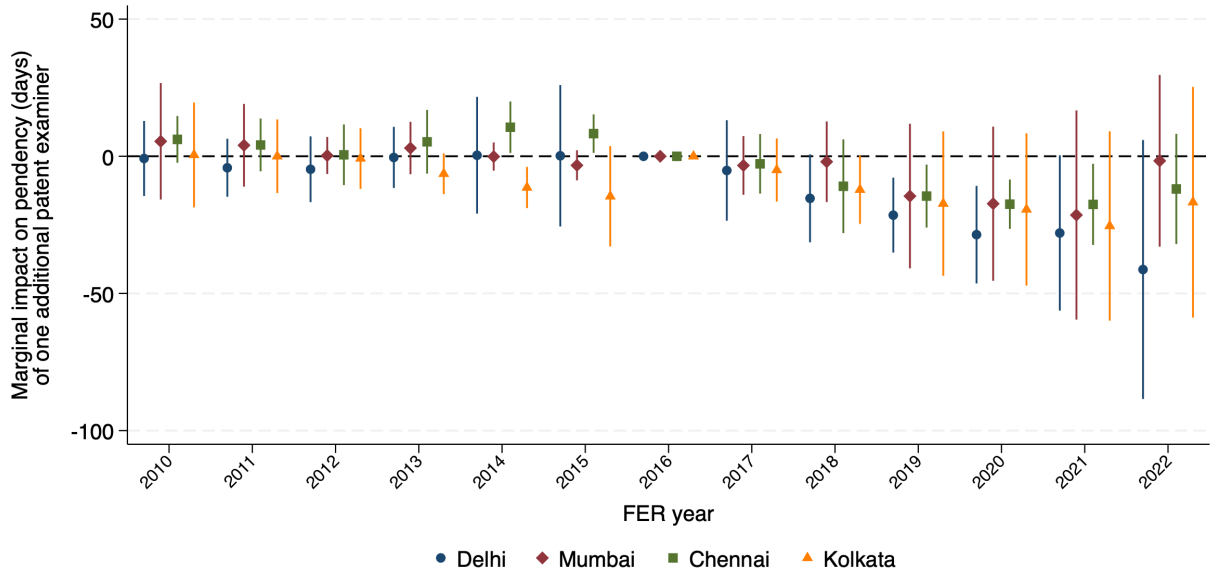
Note: This figure reproduces the main RFE-to-FER pendency event study splitting each 2015 examiner's unit across fields by their 2005–2014 portfolio shares rather than their realized 2015 caseloads; the 8 of 97 examiners with no pre-2015 examination history keep their 2015 shares. The omitted year is 2016, standard errors are clustered by field-office cell, and 95% confidence intervals are shown.

Figure A.8: Robustness of the Pendency Effect: Binarized Treatment



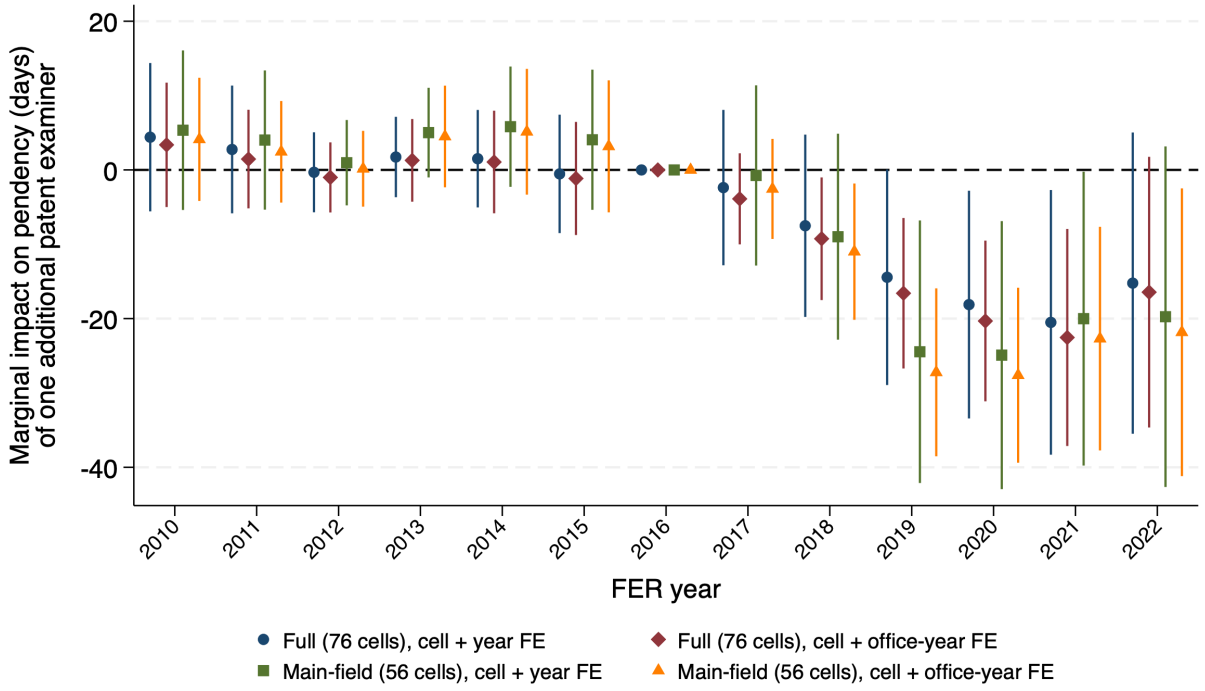
Note: This figure reproduces the main RFE-to-FER pendency event study replacing the continuous capacity measure with an indicator for cells above the median capacity increase, so the coefficients are the pendency gap between high- and low-exposure cells. High-exposure cells are the 38 of 76 cells above the median outside capacity of 1.75 examiner-equivalents; the corresponding static difference-in-differences gap is -87.8 days (cluster $p = 0.046$). The omitted year is 2016, standard errors are clustered by field-office cell, and 95% confidence intervals are shown.

Figure A.9: Robustness of the Pendency Effect: By Origin Office



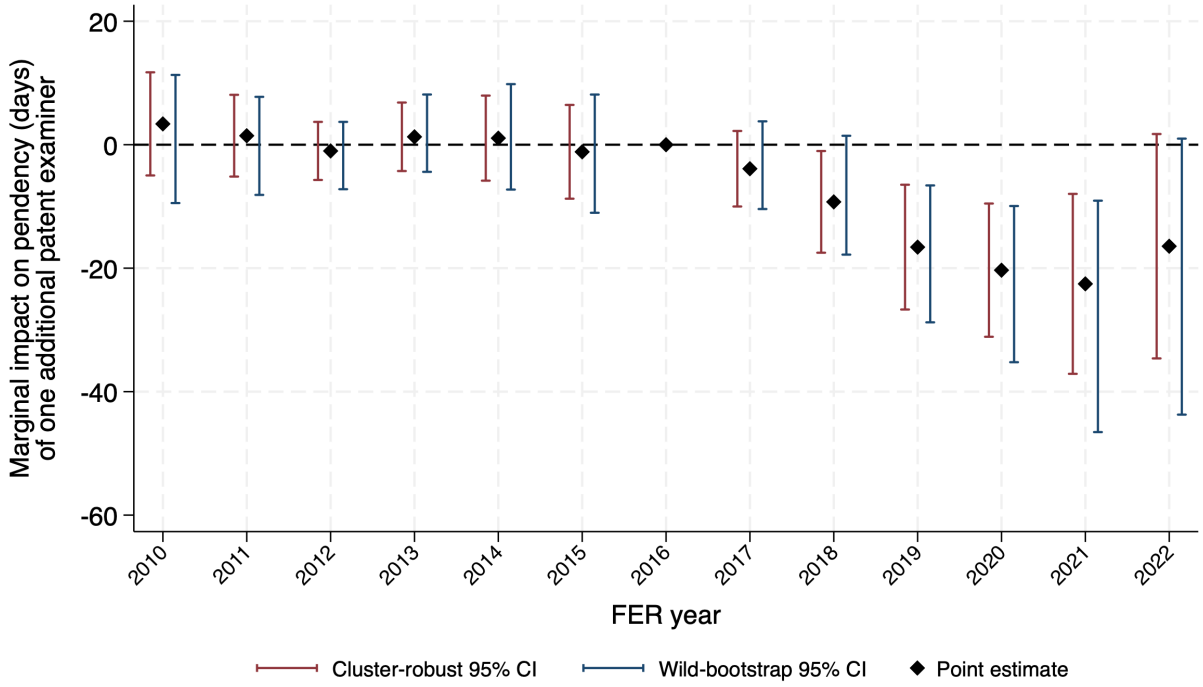
Note: This figure reproduces the main RFE-to-FER pendency event study estimated separately for each origin office, using the baseline 2015 fractional portfolio-share capacity; within a single office, the cell and office-by-year fixed effects reduce to field and year fixed effects. The omitted year is 2016, standard errors are clustered by field-office cell, and 95% confidence intervals are shown.

Figure A.10: Robustness of the Pendency Effect: All Cells versus Main Fields



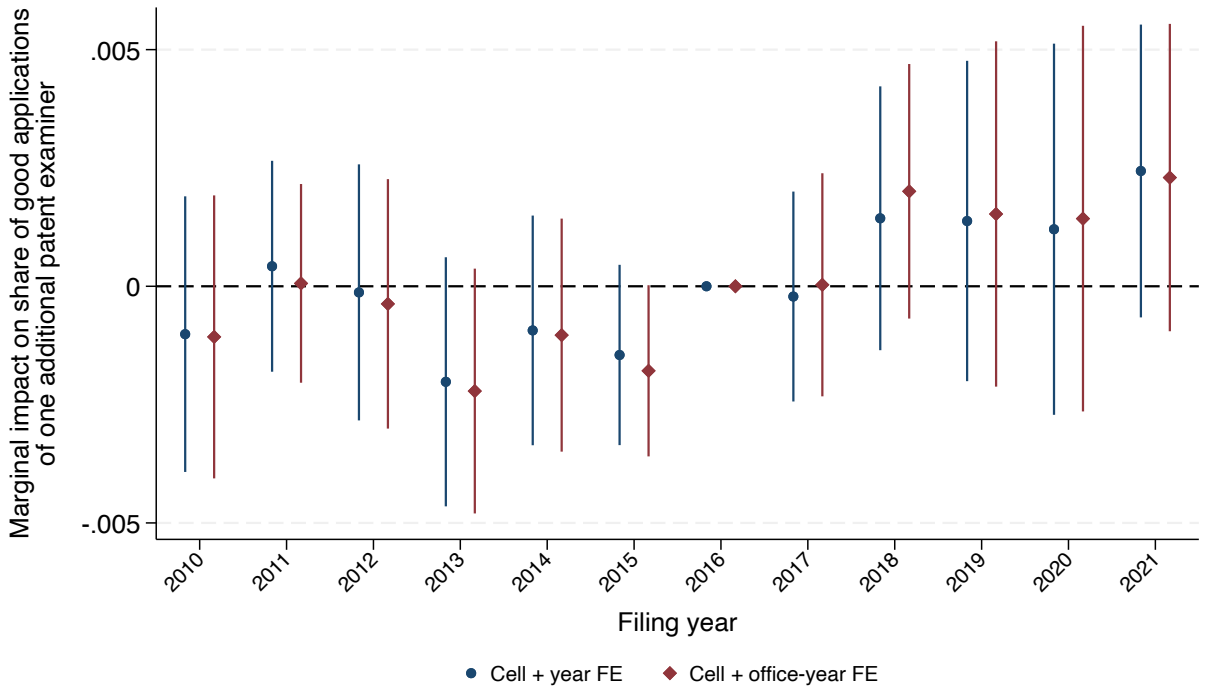
Note: This figure reproduces the main RFE-to-FER pendency event study comparing the full set of field-office cells to the main-field sample, under both cell plus year and cell plus office-by-year fixed effects. Main fields are those accounting for more than 20% of at least one 2015 examiner's caseload, so the main-field sample keeps exactly the cells for which the threshold-count measure of Figure A.4 is defined. The omitted year is 2016, standard errors are clustered by field-office cell, and 95% confidence intervals are shown.

Figure A.11: Robustness of the Pendency Effect: Wild Cluster Bootstrap



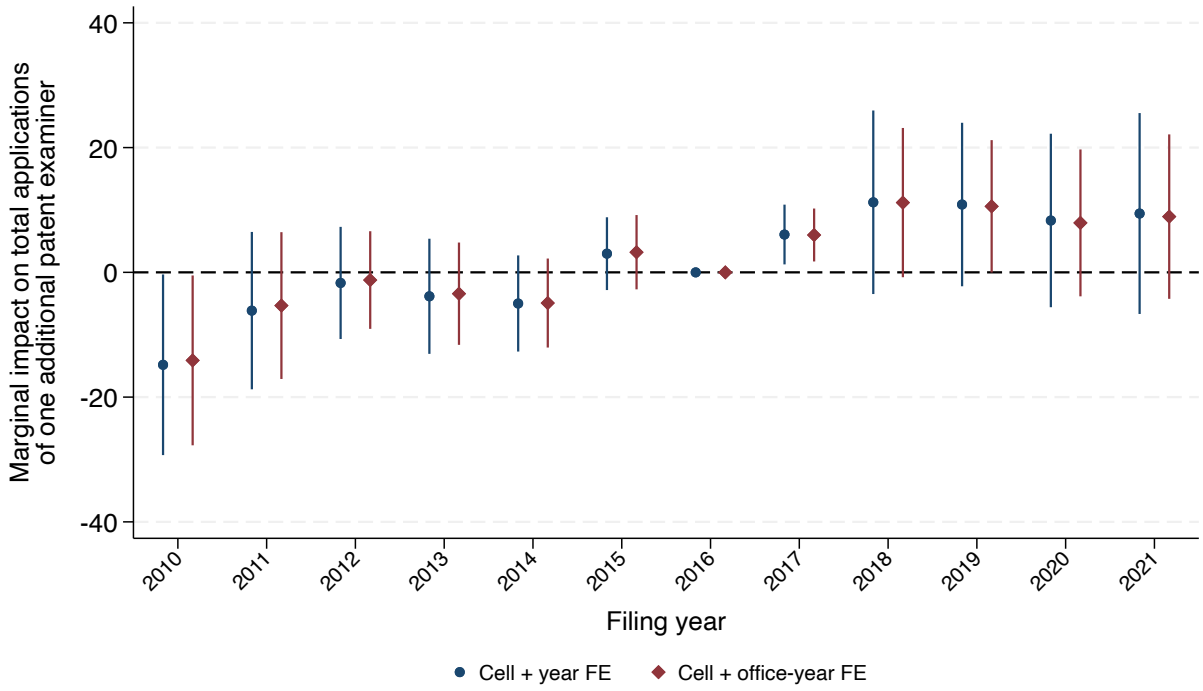
Note: This figure reproduces the main RFE-to-FER pendency event study comparing analytical cluster-robust confidence intervals to wild cluster bootstrap intervals (Rademacher weights, 9,999 replications), given the modest number of field-office cells. The corresponding static difference-in-differences estimate (-13.54 days per examiner-equivalent) has analytical and wild-bootstrap p-values below 0.001. The omitted year is 2016, standard errors are clustered by field-office cell, and 95% confidence intervals are shown.

Figure A.12: Examiner Capacity Reallocation and Share of Good Applications Filed



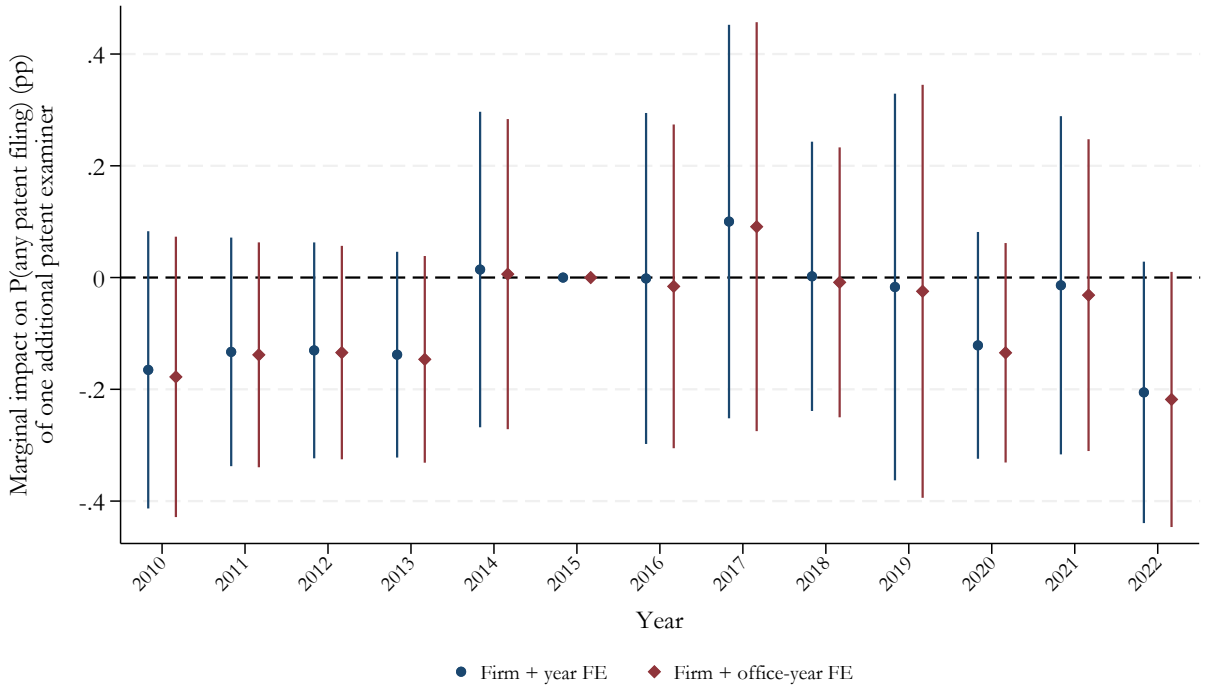
Note: This figure reports event-study coefficients from field-of-invention by origin-office by year cells over filing years 2010–2021. Treatment is 2015 outside-office same-field examiner capacity, with examiner-years assigned fractionally across fields. The dependent variable is the share of good applications filed; an applicant is good if it had at least three disposed applications pre-reform and a grant rate of at least 70%. The omitted year is 2016; standard errors are clustered by field-office cell and 95% confidence intervals are shown.

Figure A.13: Examiner Capacity Reallocation and Total Applications Filed



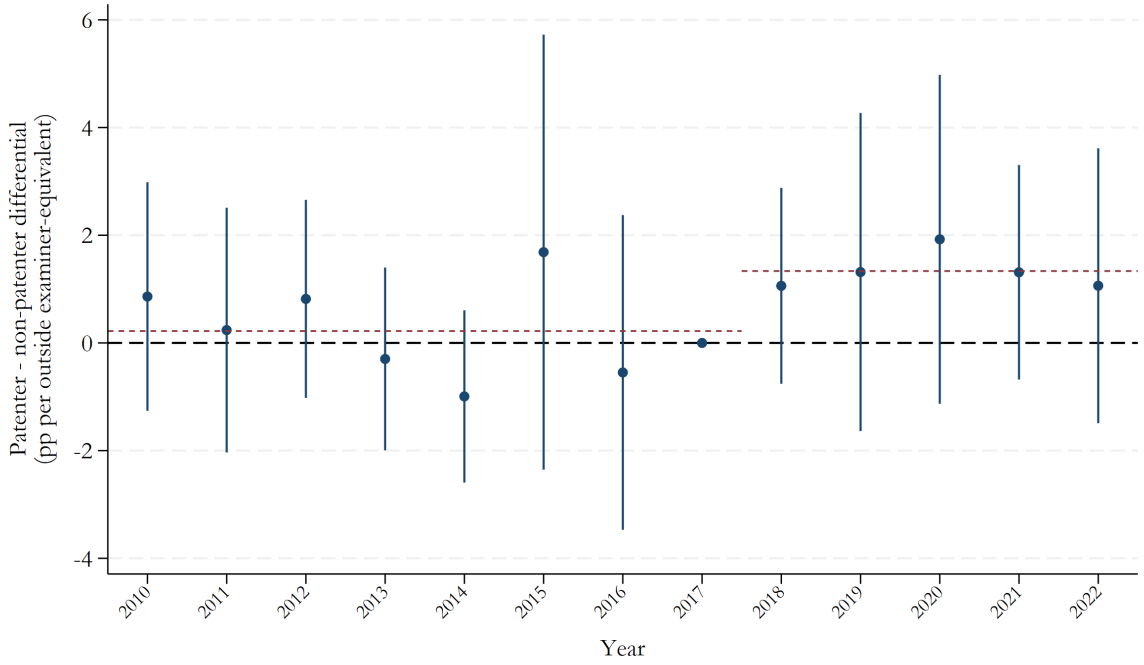
Note: This figure reports event-study coefficients from field-of-invention by origin-office by year cells over filing years 2010–2021. Treatment is 2015 outside-office same-field examiner capacity, with examiner-years assigned fractionally across fields. The dependent variable is the total number of applications filed. The omitted year is 2016; standard errors are clustered by field-office cell and 95% confidence intervals are shown.

Figure A.14: Patenting by Venture-Backed Firms and Examiner Capacity



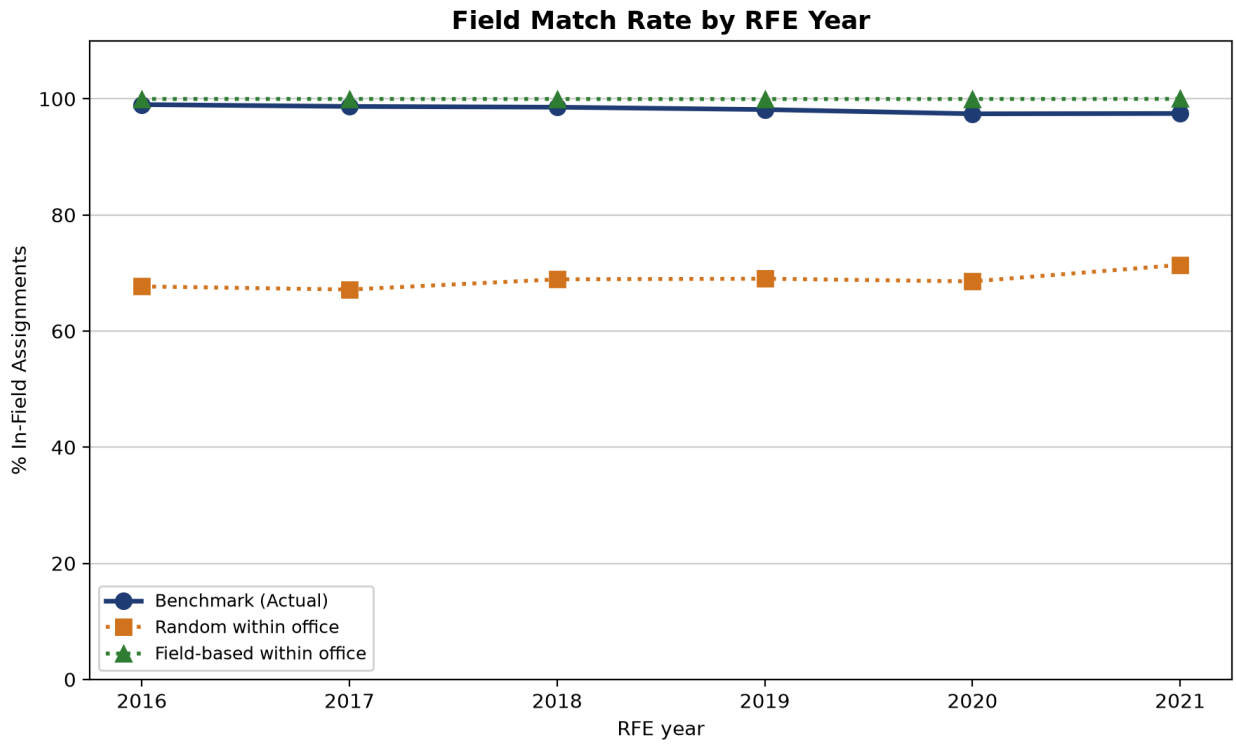
Note: The unit of observation is a firm-year in a panel of 1,187 Indian companies (PitchBook) with a completed venture financing round by 2015, 2010–2022, matched by name to patent assignees. Each point is an event-study coefficient of an indicator for filing any (eventually granted) patent application in the year ($\times$100) on outside-office same-field examiner capacity interacted with the year, with firm and year fixed effects (circles) or firm and office-by-year fixed effects (diamonds). The omitted year is 2015, standard errors are clustered by field-office cell, and 95% confidence intervals are shown.

Figure A.15: Venture Financing and Examiner Capacity: Pre-Reform Patent Holders versus Other VC-Backed Firms



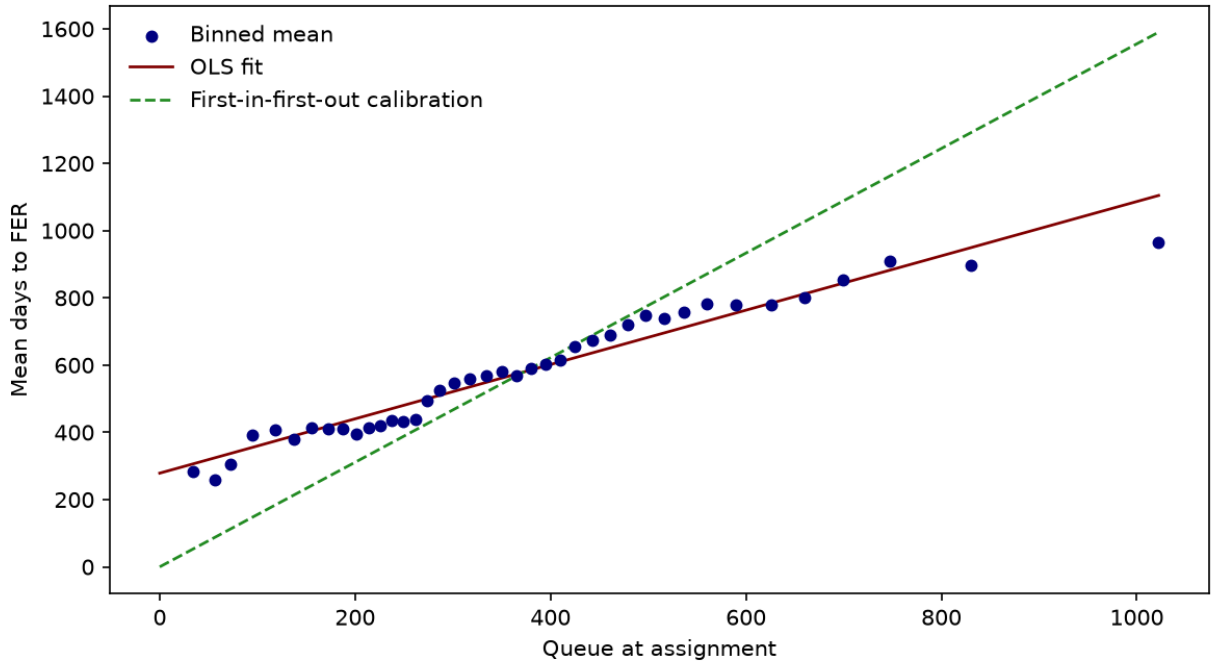
Note: The unit of observation is a firm-year in a panel of 1,187 Indian companies (PitchBook) with a completed venture financing round by 2015, 2010–2022, matched by name to patent assignees. The dependent variable is an indicator for raising any venture round in the year ($\times 100$); each point is the coefficient on outside-office same-field examiner capacity interacted with the year and with an indicator for the 53 pre-2016 patent holders, so it measures the patent-holder–non-holder financing differential per examiner-equivalent. The specification includes firm, office-by-year, and patent-holder-by-year fixed effects. The omitted year is 2017, the last year before the 2018-dated post period, standard errors are clustered by field-office cell, and 95% confidence intervals are shown. The dashed horizontal segments mark the average of the plotted coefficients over the pre-period (2010–2017) and the post-period (2018–2022).

Figure A.16: Field Match Rate



Note: This figure displays the proportion of post-reform patents (N=134,012) which were reviewed by an examiner who had seen that patent's field at least once pre-reform. The blue line indicates the percentages for the actual data; orange represents the allocation rule of random assignment within origin offices; green is the field-based allocation within origin offices. Both simulations do not allocate patents outside the origin office and correspond to the second and third rows of Table 8. The benchmark line corresponds to the first row of the same table.

Figure A.17: Days to FER Against Queue at Assignment



Note: Each dot is the mean days from RFE to FER among patents facing similar queue lengths at assignment (40 equal-sized bins). The solid line is the OLS fit; the dashed line is the wait implied by strict first-in-first-out service at the average pace.

Table A.1: Top Domestic Patent Grantees, Grant Year 2015-2022

Rank	Grantee	Total Patents
1	Council of Scientific and Industrial Research	1,503
2	Bharat Heavy Electricals Limited	1,038
3	Samsung R & D Institute India-Bangalore Private Limited	785
4	TVS Motor Company Limited	705
5	Director General, Defence Research & Development Organisation	698
6	Tata Consultancy Services Limited	696
7	Tata Motors Limited	567
8	Tata Steel Limited	485
9	Indian Institute of Technology Madras	448
10	Indian Institute of Technology Bombay	442
11	Larsen & Toubro Limited	356
12	Mahindra & Mahindra Limited	347
13	Reliance Industries Limited	327
14	Indian Council of Agricultural Research	265
15	Steel Authority of India Limited	259
16	Hindustan Unilever Limited	253
17	Indian Institute of Technology Delhi	236
18	Indian Institute of Technology Kanpur	231
19	Indian Institute of Technology Kharagpur	221
20	Schneider Electric India Private Limited	210

Note: This table reports the top 20 domestic patent grantees in India by year of grant for 2015–2022. Indian patents are defined as those for which either one of the inventors' or one of the grantees' addresses is in India.

Table A.2: Top Domestic Patent Grantees, Application Year 2015-2022

Rank	Grantee	Total Patents
1	TVS Motor Company Limited	936
2	Tata Consultancy Services Limited	911
3	Council of Scientific and Industrial Research	855
4	Bharat Heavy Electricals Limited	818
5	Mahindra & Mahindra Limited	814
6	Indian Institute of Technology Madras	715
7	Indian Institute of Technology Bombay	660
8	Wipro Limited	606
9	Tata Steel Limited	562
10	Director General, Defence Research & Development Organisation	510
11	Maruti Suzuki India Limited	498
12	Indian Institute of Technology Delhi	433
13	SRM Institute of Science and Technology	406
14	Tata Motors Limited	402
15	Indian Institute of Technology Kanpur	336
16	Schneider Electric India Private Limited	328
17	Indian Institute of Science	294
18	Samsung R & D Institute India-Bangalore Private Limited	260
19	Indian Institute of Technology Kharagpur	219
20	Hero MotoCorp Limited	213

Note: This table reports the top 20 domestic patent grantees in India by year of application between 2015–2022 for all the successful applications (granted patents). Indian patents are defined as those for which either one of the inventors' or one of the grantees' addresses is in India.

Table A.3: Determinants of Indian Patent Allocation by Origin Office

Allocation rule	N patents	Home bias	Field supply	Queue
Delhi				
Benchmark (Actual)	55,981	1.33^{***}	0.43^{***}	0.05^{***}
		(0.02)	(0.01)	(0.01)
Random across office	55,981	0.40^{***}	0.25^{***}	0.00
		(0.01)	(0.01)	(0.01)
Field-based across office	55,981	0.21^{***}	0.52^{***}	0.03^{***}
		(0.01)	(0.01)	(0.01)
Queue-based across office	55,981	0.20^{***}	0.25^{***}	-0.88^{***}
		(0.01)	(0.01)	(0.01)
Field + queue across office	55,981	0.31^{***}	0.53^{***}	-0.72^{***}
		(0.01)	(0.01)	(0.01)
Mumbai				
Benchmark (Actual)	22,679	0.25^{***}	0.16^{***}	0.18^{***}
		(0.02)	(0.02)	(0.01)
Random across office	22,679	-0.53^{***}	0.34^{***}	-0.02
		(0.02)	(0.02)	(0.01)
Field-based across office	22,679	-0.44^{***}	0.55^{***}	-0.01
		(0.02)	(0.02)	(0.01)
Queue-based across office	22,679	-0.58^{***}	0.19^{***}	-0.90^{***}
		(0.02)	(0.02)	(0.02)
Field + queue across office	22,679	-0.64^{***}	0.51^{***}	-0.74^{***}
		(0.02)	(0.02)	(0.01)
Chennai				
Benchmark (Actual)	41,420	1.18^{***}	0.35^{***}	0.43^{***}
		(0.01)	(0.01)	(0.01)
Random across office	41,420	0.10^{***}	0.50^{***}	0.06^{***}
		(0.01)	(0.01)	(0.01)
Field-based across office	41,420	0.13^{***}	0.72^{***}	0.07^{***}
		(0.01)	(0.01)	(0.01)
Queue-based across office	41,420	0.19^{***}	0.44^{***}	-0.86^{***}
		(0.01)	(0.01)	(0.01)
Field + queue across office	41,420	0.27^{***}	0.84^{***}	-0.62^{***}
		(0.01)	(0.01)	(0.01)
Kolkata				
Benchmark (Actual)	13,932	1.32^{***}	0.26^{***}	0.15^{***}
		(0.02)	(0.02)	(0.01)
Random across office	13,932	-0.10^{***}	0.47^{***}	0.06^{***}
		(0.02)	(0.02)	(0.01)
Field-based across office	13,932	-0.05^{*}	0.71^{***}	0.04^{***}
		(0.02)	(0.02)	(0.01)
Queue-based across office	13,932	0.04	0.40^{***}	-0.93^{***}
		(0.02)	(0.02)	(0.02)
Field + queue across office	13,932	-0.11^{***}	0.77^{***}	-0.63^{***}
		(0.02)	(0.02)	(0.02)

Note: Coefficients from separate conditional logit models (McFadden 1974), standard errors in parentheses. Each patent contributes four observations (one per candidate office). *Home bias* = 1 if candidate office equals patent's origin office; *Field supply* = examiner-equivalents in the patent's field at the candidate office (pre-reform caseload shares summed over its pool examiners), standardized within RFE year. *Queue* = the candidate office's caseload per examiner at the patent's own RFE date, counting cases that arrived before it and had not yet left: the common pre-reform backlog plus that allocation's own post-reform assignments, divided by the number of examiners at that office in that year, and standardized within RFE year. Alternative-specific office constants are omitted here: within a single origin office they are collinear with *Home bias*.

Appendix B Identifying and Disambiguating Examiners

Unit of observation and source fields. The examiner crosswalk is constructed for the 343,618 granted applications in our data, with each observation corresponding to a granted application indexed by its application identifier (`appID`). For each grant, we seek to identify the controller (examiner) who issued the First Examination Report (FER) and assign that examiner a consistent identifier that can be tracked across applications and years. We obtain the examiner information directly from the FERs, whose signature block reports both the controller's name and the issuing office. These fields are extracted from the scanned documents using optical character recognition (OCR) and large language models. Appendix Figure A.3 illustrates the relevant portion of a typical FER. This process yields a nonempty examiner-name string for 299,100 grants, while the remaining 44,518 grants have no examiner name in the available source documents. Across the 299,100 extracted records, we observe 1,492 distinct raw name strings. This figure substantially overstates the number of unique examiners because the same individual may appear under multiple name variants, including full names, initials, alternative spellings or transliterations, and strings that retain salutations or professional titles. OCR errors introduce additional variation, and in a small number of cases the extraction captures unrelated text, such as an inventor's name or the name of a foreign patent officer appearing in a PCT or other foreign patent-office document associated with the Indian application. We therefore clean the extracted strings and reconcile these variants into consistent examiner identities before constructing examiner-level measures of capacity.

Cleaning examiner names. We begin by converting each raw examiner string into a standardized name key. Punctuation is replaced with spaces, leading salutations and titles are removed, repeated spaces are collapsed, and the remaining text is converted to lowercase. The removed terms include *Dr*, *Shri*, *Sh*, *Sri*, *Mr*, *Mrs*, *Ms*, *Smt*, *Md*, *Mohd*, *Prof*, *Er*, *Late*, *Shree*, and *Miss*. For each standardized key, we also record the years in which it appears and its home office, defined as the office issuing the plurality of the associated FERs. These additional fields help distinguish genuine name variation from distinct individuals. In particular, before the 2016 auto-allocation reform, applications were examined within their regional office, making agreement in office assignment highly informative about examiner identity.

Assigning preliminary examiner identities. We next group standardized keys using similarities in their names. Specifically, we index each key by surname and first token,

and allow two keys to be combined when their given names are identical or when an abbreviated form is consistent with a full form. An automated match is accepted only if it is unique and both keys share the same home office. Keys with multiple candidate matches or conflicting offices remain separate at this stage and are deferred to manual review. We then apply a small hand-coded dictionary of known spelling and transliteration equivalences. Because accepted pairwise matches may link several name variants, we take the transitive closure of the resulting network and assign a preliminary examiner identifier to each connected component. This conservative automated procedure reduces the 1,492 raw name strings to 957 preliminary examiner identities.

Reconciliation with official records. To validate the automated matches, we compare the preliminary examiner identities against the 2023 and 2024 IP India staff rosters and the signature blocks in the underlying FERs. The staff rosters report each officer's name, office, designation, and official email address, providing an independent benchmark for validation. These sources allow us to correct both types of matching error. We separate distinct examiners that were incorrectly merged, most clearly when similar names correspond to different official email addresses, and we combine records for the same examiner when spelling or transliteration differences prevented an automated match. The review also removes names extracted from the wrong part of a document and repairs OCR fragments created when text from the printed title or surrounding signature block was attached to the examiner's name. After these corrections, the data contain 925 examiner identities.

Adjudicating the remaining cases. The remaining ambiguous cases are adjudicated by combining several signals: the FER signature, the pattern of offices and years, the technical field of the applications, the rank history, the roster entry, and the official email address. We give priority to the FER signature when the underlying document is available and legible. Each of the other signals speaks to whether two name variants belong to one person or to two, and they are most decisive when they point the same way. Offices and timing come first. Before 2016, applications were examined within their regional office, so the same name issuing reports from two offices in the same year almost always means two different examiners. We treat a change of office as a transfer only when activity in the first office ends at about the time it begins in the second, with no more than one year of overlap; a genuine transfer empties the old office as it fills the new one, whereas two people keep both offices busy at once. Because the reform was introduced during 2016 and 2017, we allow more uncertainty in those two years, and from 2017 onward we read the office from the FER itself rather than from the application's origin. The technical

field provides a second signal, because examiners specialize. A block of reports in a field far from an examiner's usual work, such as chemistry records attached to a name that otherwise handles mechanical applications, points to a shared name rather than a shared person. Rank history provides a third, because the progression from Examiner to Assistant Controller, Deputy Controller, Joint Controller, and Senior Joint Controller moves in one direction. A later report at a lower rank, or a span of activity that cannot be reconciled with the observed ranks, indicates that a candidate group holds more than one person. Rank also resolves a pattern that would otherwise inflate the count. A large group of examiners issues reports in 2007 through 2009, falls silent through 2015, and resumes in 2016. Read mechanically, each such record looks like two people. It is not: the silent years are the interval in which the examiner was promoted to a controller rank, whose duties center on hearings and grants and generate few first examination reports. The signature confirms the progression, since the same person signs as "Examiner" in 2008 and as "Deputy Controller" in the later period, and the 2023–24 roster lists the senior rank. We verified every member of this group individually and kept each as one examiner, so a gap in activity triggers review but never forces a split on its own. A shared name can also hide two people in the other direction. In four cases a block of older reports carried a field and office different from the modern examiner's work; an early Kolkata chemical examiner signing "A. K. Pradhan," for example, had been merged with a later mechanical examiner, "Arun Kumar Pradhan," and we separated the two. We combine these signals with an asymmetric rule. A hard contradiction, such as a rank reversal, sustained activity in two offices at the same time, or two distinct official email addresses, is enough to keep two identities apart. A merge, by contrast, requires at least two independent signals to agree, and when the evidence is inconclusive we keep the identities separate rather than force a match. As a final pass, we review any remaining identities that share a surname, office, and overlapping years using the same evidence. This process produces 634 examiner identities.

Reconstructing the issuing office. The regional structure of examination before 2016 implies that an application's office of origin also identifies the examiner's office. The auto-allocation reform broke this link: after 2016, the origin office records where an application was filed rather than where the examiner was located. A separate field identifying the examiner's own office appears only from 2017 onward, covering about 55% of reports in that year and nearly all reports from 2018 onward. As a result, examiner location is not directly observed for 2016 and part of 2017. We therefore reconstruct the examiner's office using a small set of hierarchical rules designed to avoid imputing transfers. When

the examiner's own office is recorded, we use it directly. For reports through 2015, we use the application's origin office, which coincides with the examiner's office under the pre-reform assignment system. For reports from 2016 or 2017 without a recorded office, we assign the examiner's office when the offices observed before 2016 and after 2018 coincide, since agreement on both sides of the transition implies that no transfer occurred during the intervening period. For examiners who first appear after the reform and therefore have no earlier record, we use their stable post-reform office. Reports that satisfy none of these rules—primarily those involving examiners who transferred during the transition—retain their recorded origin office and are explicitly flagged rather than imputed. These rules resolve approximately 95% of reports from 2016–17 with high confidence, and the final crosswalk reports both the raw origin office and the reconstructed examiner office. The FER letterhead cannot resolve the remaining cases because it records the application's origin office rather than the examiner's physical location.

Final accounting. The resulting crosswalk has the following coverage:

- The crosswalk contains 343,618 rows, one for each granted application's `appID`.
- Of these grants, 299,100 have a nonempty examiner-name extraction and 44,518 have no examiner name in the available source documents.
- Among the grants with a nonempty extraction, 299,076 are assigned to one of the 634 examiner identities.
- The remaining 24 grants are left unassigned because the extracted text identifies a foreign officer or cannot be read reliably.

Appendix C Linking Indian Patent Assignees to PitchBook

We link 6,455 Indian companies with at least one Early or Late Stage VC deal in PitchBook (2000–2026) to 6,592 disambiguated applicant clusters from the Indian patent corpus, where each cluster represents one firm identity resolved across spelling variants of the applicant name. Since the sources share no common identifier, we match on company names. Names on both sides are lower-cased, stripped of punctuation and parenthetical disambiguators, and singularized; a second "core" version additionally drops legal and generic tokens (pvt, ltd, technologies, solutions, india, and similar) so that matching turns on the distinctive part of the name. We then accept matches in three tiers: exact agreement of the normalized full name (126 links); agreement of the core name where that core is distinctive, defined as at least two tokens or six characters (68 links); and fuzzy agreement requiring a token-set ratio of at least 90 together with a full-string ratio of at least 75 and a core-name ratio of at least 80 (26 links). The latter two thresholds block matches driven only by a shared generic suffix. Where several companies match the same cluster we keep the best pair by tier and score. This yields 220 automatic links and 902 remaining candidate pairs.

The candidate pairs are adjudicated by a large language model (Claude Opus 4.8) prompted to weigh the distinctive portion of the name against corroborating evidence on geography, industry–technology coherence, and the timing of filings relative to financing, returning accept, reject, or unsure. Twelve hand-labeled pairs are used as worked examples and twenty-one are held out for validation. Of 890 pairs classified, the model accepted 79 and rejected 808; three unsure cases were resolved by a second pass with web search and all rejected. On the held-out pairs the model agreed with the human label in all twenty-one cases. The final linkage contains 301 firm–cluster pairs — 220 automatic, 68 from the model, and 13 confirmed by hand — of which 219 fall inside the 2010–2022 estimation panel and 212 record a filing within the window.

Assigning firms to office $\times$ field cells. Treatment varies at the patent office $\times$ technology field level, so each company is assigned to a cell using time-invariant attributes rather than its realized patenting. For the office, headquarters state strings are resolved to a canonical state through an alias dictionary and fuzzy matching, then mapped to one of the four patent offices using the IP India's official territorial jurisdiction (Delhi for the north and north-west, Mumbai for the west and center, Chennai for the south, Kolkata for the east and north-east); headquarters city serves as a fallback where the state is missing. This assigns an office to 99.8 percent of deals. For the field, PitchBook's 208 industry

codes are mapped to the patent "Field of Invention" vocabulary, each code inheriting its parent industry group's mapping subject to code-level overrides — necessary because PitchBook's single "Pharmaceuticals and Biotechnology" group would otherwise leave the biotechnology field empty. Industries with no patent counterpart (finance, retail, media, generic services) are routed to a residual category and drop out, accounting for 9.7 percent of firms. Fifteen fields have venture-backed activity, which crossed with four offices gives the 53 cells that identify the design.

Appendix D The Indian Patent Litigation Database

We assemble a case-level census of Indian patent litigation from five sources. In this section, we first review these sources and how we identified patent cases therein:

- The first and largest is the corpus of High Court orders released through the eCourts system and mirrored on the AWS Open Data Registry. We screen all 19,128,733 orders published between 2001 and 2026 for five signals of a case involving a patent: a reference to a section of the Patents Act, a patent or its application number, the Patent Office or Controller of Patents named as a party, the invocation of a patent-specific doctrine, or the word patent in the cause title. An order enters the candidate pool if it contains at least one of these signals. This broad screen admits 16,359 candidate orders. We then use a language-model reader to review the full text of each candidate and determine whether the order substantively concerns a patent or patent application. This review excludes false signals, including Letters Patent Appeals and matters involving other forms of intellectual property administered by the same Controller General. It confirms 2,041 patent orders, corresponding to 1,819 distinct cases before the subsequent merge and audit.

- The second source is the case registry of the Delhi High Court's Intellectual Property Division, which publishes cause lists and case-status pages for matters filed since 2021. These records identify patent matters through their proceeding codes, including cases that have not produced a published order. For example, the case number `C.A.(COMM.IPD-PAT)-2/2021` identifies a patent appeal filed in 2021, although the registry records no decision date and the High Court corpus contains no corresponding order. We match the registry entries to the High Court corpus using standardized case numbers and retain cases not already represented in that corpus. This procedure contributes 291 cases, of which 285 have no recorded decision date.

- The third is a community backup of the Intellectual Property Appellate Board's website, made after the Board's abolition in April 2021. The archive contains 250 order documents from its patent collection, covering decisions issued from July 2019 through March 2021. We use optical character recognition and a language-model reader to recover the case number, parties, bench, order date, and available disposition text. After consolidating multiple orders from the same proceeding and matching records across sources, the archive contributes 181 cases.

- The fourth and fifth are the Development Data Lab's release of 80,935,944 district-court cases filed between 2010 and 2018 (derived from Ash et al. (2026)) and a similar

collection of 6,960,842 more recent district records. Neither source provides a patent-specific case identifier, but both record the statute governing the case and, where available, the section invoked. We first identify candidate cases that reference the Patents Act in the statute field, then validate them using the recorded statutory provision and case type to ensure that they involve patent litigation. For example, we retain a civil suit filed under section 104 as an infringement suit but exclude criminal cases that carry a Patents Act label in error. After consolidating overlaps between the two sources, this procedure contributes 76 unique district-court cases.

We then combine the five sources into a database with one entry for each case. This requires identifying both multiple orders issued in the same case and records of that case appearing in more than one source. We use court and tribunal identifiers to bring these records together. Distinct cases involving the same underlying dispute remain separate entries and are linked for the dispute-level analysis. Transfers from the IPAB to the High Courts require additional matching because a continuing case may receive a new registration number. For matters transferred following the Board's abolition in 2021, we use the complete tribunal case number quoted in the High Court's cause title, including its bench suffix, to establish the link. The matched records enter as one case, retaining the original tribunal filing year, where recoverable, together with the High Court re-registration date.

Harmonizing the five sources yields a final database of 2,419 unique cases, with recorded filing years ranging from 1995 to 2026. Of these, 1,871 are attributed to the High Court corpus, 291 to the Delhi registry, 181 to the tribunal archive, and 76 to the two district-court collections. By type of proceeding, appeals against decisions of the Controller of Patents form the largest category, comprising 1,223 cases, followed by 569 infringement suits and 262 revocation or validity proceedings. The remaining 365 comprise other patent matters, such as writ petitions and interim applications.

We further examine how courts handle patent cases through a detailed review of their orders, with results available for 1,860 High Court cases. Orders granting or denying relief account for 756 cases: relief was granted in 283, partly granted in 75, and denied in 398. The other 1,104 cases record settlements, withdrawals, remands, or procedural outcomes. Settlements were recorded in 131 cases and withdrawals in 252. We treat settlements and withdrawals together in the analysis, since a withdrawal may reflect an unrecorded settlement. Courts remanded another 350 cases for reconsideration and disposed of 172 as infructuous. The available orders in 149 cases concern only an adjournment or interim decision, while orders in another 50 record a transfer or state no outcome. These last categories require caution when assessing success, since an interim order or an unstated

outcome does not establish how the case ultimately ended.

Our database is undoubtedly incomplete. Some High Courts appear not to publish every resolved case, and their caseloads may therefore be understated in our records. The tribunal archive begins in July 2019 and provides no direct coverage of the Board's earlier output. But two external benchmarks put the coverage in perspective. The first, for infringement suits, is the hand count by Reddy T. and Lath (2017) of 143 patent infringement suits filed in selected High Courts between 2005 and 2015. Our census holds 135 High Court infringement suits with a recorded filing year in the same window, although our count covers a broader set of courts. The second is the Delhi High Court's monthly statistics as a benchmark for patent appeals. Across seven published monthly reports for 2022, the court records 456 patent appeals instituted, compared with 289 appeals in our database with a recorded filing year of 2022. The court's total exceeds our count for the full year, although the different reporting periods prevent a direct estimate of coverage. This shortfall is doubtless due to the fact that many cases are likely to still be pending matters that have yet to produce a published order.