Patent Visibility and the Diffusion of "Trapped Knowledge"

Randol H. Yao*

Valuable knowledge developed in one part of the world may remain trapped due to frictions in how knowledge is exposed globally. This paper examines how increasing the visibility of foreign innovations—by granting US patents—"untraps" knowledge. Using difference-in-differences with an examiner leniency instrument, I find that US grants of foreign patents significantly increase the intensity and reach of forward citations. Using a novel measure of "trappedness," I show that knowledge from historically more trapped countries and sectors sees larger diffusion benefits after US grants. These findings highlight the central role of the US as a platform of global knowledge diffusion.

*Yao: MIT Sloan School of Management (email: hyyao@mit.edu). I thank Pierre Azoulay, Richard Freeman, Danielle Li, Jacob Moscona, Abhishek Nagaraj, and Scott Stern for advice and comments.

I. Introduction

Not all valuable knowledge finds its way to where it is needed. While the common metaphor of "standing on the shoulders of giants" captures the idea of cumulative innovation, it overlooks a critical challenge: the visible giants may not be the only ones worth standing on. While knowledge originating from certain countries, institutions, or languages circulates widely and becomes foundational, other potentially insightful ideas may remain confined to local or regional contexts. The result is a form of epistemic inertia: valuable innovations developed in one part of the world may remain "trapped" locally—not due to lack of merit, but because of frictions that limit their global exposure and integration. Such frictions can arise from linguistic barriers, differences in institutional norms, and structural biases that favor knowledge from already dominant innovation hubs. As a result, the global innovation system may reward ideas that are more visible, with this visibility gap reinforcing itself over time: well-recognized knowledge continues to accumulate citations and resources, while the "hidden diamonds" may remain buried, waiting to be discovered and integrated into the broader global knowledge network.

This paper explores how increasing the visibility of local knowledge affects global innovation. Specifically, I focus on the United States, whose patent system is globally visible and influential. I examine how granting US patents to foreign innovations that are already established in their home countries can help untrap knowledge that would otherwise remain confined. While foreign inventors often file patents in the US to protect commercial interests in a large market, obtaining a US patent also increases an invention's visibility. Once a foreign-origin invention is granted in the US, it is published in English, indexed in international patent databases, and treated as a vetted piece of prior art, making it more likely to be noticed and cited by researchers around the world.

To study the impact of changes in visibility, I focus on a panel of patent families with granted patents in non-US (home) countries, which subsequently file applications through the USPTO. There are 2,470 patent families observed over 23 calendar years from 2000 to 2022. This sample is especially informative because it allows me to observe the dynamic effects following home-country grants but prior to the issuance of US grants. By comparing foreign patents whose US patent applications are granted with those that are not through a difference-in-differences setting, controlling for patent-specific effects and filing cohort-specific trends, I can isolate the effect of receiving a US grant on the trajectory of knowledge diffusion. Even after difference-in-differences, a key concern remains: the US grant status may be endogenous to observed or unobserved characteristics like patent quality, which may also drive the diffusion of knowledge and therefore bias the estimates. To address this, I instrument the US grant status with examiner leniency, exploiting quasi-random assignment to stricter versus more lenient examiners. This shifts grant outcomes for reasons orthogonal to the application's characteristics like the intrinsic quality. In an ideal world, where all knowledge is treated equally and is equally accessible, this estimated effect would be small, as the home country patents alone would lead to the same exposure with or without the US patents, or may be negative if US protection raises legal costs that deter follow-on use.

The analysis proceeds in several stages. First, I estimate the impact of receiving a US patent grant on forward citations using the panel of foreign-origin patent families. The results show that receiving a US patent grant meaningfully amplifies exposure: US grants increase the likelihood of receiving any forward citation by about 17% and raise the total number of citations by roughly 27%. The dynamic effects plot shows no meaningful difference in citation trajectories in the years leading up to the grants. These findings provide evidence that the US patent makes foreign inventions more visible.

If the US functions as a platform for global exposure, then one would expect the diffusion benefits to appear in regions beyond the United States and the invention's country of origin. Therefore, I examine how the impact spreads across different geographies. In addition to sizeable increases in citations within the US and the home country, citations from third countries rise markedly by about 23%, indicating that US grants help extend knowledge to places that are otherwise more disconnected from its origin. Further, US grants extend the knowledge to more middle- and low-income countries, hinting that these countries may face greater informational and institutional barriers and therefore benefit more from the amplified exposure provided by the US platform. In this sense, US grants not only enhance recognition in the US innovation system, but also act as a bridge that carries knowledge across borders to new, and often more peripheral, regions.

Next, I show that once an inventor gains visibility through a US grant, attention spills over to their broader portfolio of patents. I find that citations to the inventors' other patents increase significantly after the US grants, with the strongest effects for less established inventors with a limited number of patents. This portfolio-level boost suggests that US grants can help elevate an entire line of work, not just a single invention, by signaling credibility and drawing international attention.

Finally, I explore why some inventions gain more citations than others after US grants. A key mechanism is "trappedness": the extent to which knowledge from a particular country or sector is overlooked in the global citation network. I find that the effects are concentrated among patents from countries with historically limited global reach, such as Korea and China, while patents from English-dominant or highly integrated countries show little change. To better quantify the trappedness of each country and sector, I then develop a novel citation-based measure based on whether patents from certain countries and technology sectors are cited by inventors or only surfaced later by examiners. The intuition is that examiner-added citations reveal knowledge that is relevant but not visible or recognized by the inventors. Using this measure, I show that US grants have a larger effect when the originating country is more trapped. This relationship also holds across sectors, reinforcing that US grants play an especially important role in amplifying the global exposure of knowledge that would otherwise remain under-recognized. Taken together, these findings suggest that US patent grants serve as a powerful platform amplifying the exposure of valuable but less utilized knowledge, particularly when those ideas originate from places that are less integrated into the global innovation system.

The idea that innovation builds cumulatively on prior knowledge is foundational. In endogenous economic growth models, past ideas function as non-rivalrous inputs into future invention, enabling sustained economic growth (Romer, 1990; Jones, 1995). A rich empirical literature explores the cumulative nature of innovation and how it is shaped by geographic and institutional frictions (Jaffe et al., 1993; Galasso and Schankerman, 2015; Furman and Stern, 2011). Besides the cumulative production, the accessibility of knowledge is also important. Mokyr (2011) emphasizes that the economic significance of useful knowledge hinges not only on its production but on its accessibility. The ability of a society to stand on the shoulders of giants depends on the quality of mechanisms for storing, certifying, and accessing that knowledge (Furman and Stern, 2011). This paper extends this literature by looking at a scenario when potentially valuable knowledge is locked within local contexts, and asks whether certain institutions can facilitate its integration into the broader cumulative innovation process.

A parallel literature in international economics and innovation studies highlights the barriers to cross-border knowledge diffusion. Classic work emphasizes how technological, geographic, and institutional factors play critical roles in the spread of ideas across national boundaries (Eaton and Kortum, 1996; Keller, 2004; Rogers et al., 2014). Further empirical studies document such frictions in cross-border citation and licensing behavior (Branstetter, 2001; De Rassenfosse and Seliger, 2020; MacGarvie, 2006). Much of the literature on overcoming these barriers focuses on individuals as carriers of knowledge across borders (Borjas and Doran, 2012; Ganguli, 2015; Kerr, 2008; Fry and Furman, 2023). This paper shifts attention from mobile individuals to mobile ideas, emphasizing how institutional signals from USPTO can help the cross-border diffusion of valuable foreign-origin knowledge to reach disconnected areas.

The findings also speak to long-standing debates about the role of intellectual property rights in shaping innovation and diffusion. On one side, intellectual property rights are seen as beneficial for encouraging investment, creating markets for ideas, and bringing more knowledge into the public domain (Gans and Stern, 2000; Hellman, 2007; Merges and Nelson, 1994). On the other side, the anti-commons perspective warns that stronger IP protections may block follow-on research by making knowledge harder to access (Heller and Eisenberg, 1998; David, 2004). Empirical studies also show that patenting can have zero or negative effects on follow-on citations (Murray and Stern, 2007; Sampat and Williams, 2019; Williams, 2013). The evidence in this paper highlights a distinct channel through which patents may facilitate knowledge flows: by increasing the exposure of otherwise overlooked ideas.

Relatedly, the results intersect with a growing literature on the appropriateness of technology in global contexts. This literature highlights that even high-potential innovations may fail to spread if they are not perceived as suitable or relevant across different settings (Moscona and Sastry, 2025; Lerner et al., 2024). This paper raises a complementary concern: even when innovations are appropriate, they may never reach broader audiences due to diffusion barriers. The increase in attention from countries well beyond the origin and the US—particularly from other middle- and low-income countries—hints at the diffusion of appropriate technology facilitated by US grants.

The remainder of this article proceeds as follows. Section II provides background on trapped knowledge and the role of patents in knowledge diffusion. Section III describes the data and empirical strategy. Section IV presents the main results and explores the underlying mechanisms. Section V concludes.

II. Background

A. Trapped Knowledge in Global Innovation

Across sectors and disciplines, many innovations, practices, or discoveries remain less utilized despite their potential to generate immense societal benefits. For instance, Israel's water-saving drip irrigation systems could have transformed agriculture in water-scarce regions globally, but adoption lagged in many countries due to the lack of exposure, training, or adaptation mechanisms. Also, India's high-volume, low-cost cataract surgery model could have reduced blindness in many low-income countries, but knowledge of the approach remained localized for years due to the weak global diffusion of health delivery innovations.

There are many reasons why useful technologies remain less utilized: lack of awareness, institutional barriers, geopolitical fragmentation, high costs, and more. In some cases, lower utilization may be rational—when local conditions make adoption less beneficial or more costly. However, in other cases, valuable knowledge fails to spread not because it lacks relevance or utility, but because it lacks exposure. I refer to this phenomenon as "trapped knowledge": knowledge that remains geographically confined despite its global potential, primarily because of limited exposure. One illustrative case of geographically trapped knowledge comes from China's agricultural sector.

Case Study: China's Hybrid Rice

In the 1960s, Yuan Longping discovered the genetic basis of heterosis in rice, where the offspring of two genetically distinct parents exhibit superior traits. This was a unique discovery because heterosis had been thought impossible in self-pollinating crops like rice. Building on this, Chinese researchers bred high-yield, resilient hybrid varieties in the 1970s that became central to China's food strategy, markedly boosting output and reducing hunger. By 1994, hybrid rice accounted for 57% of national rice production with average yields in some provinces reaching 36–50% above the conventional 5 metric tons per hectare benchmark (Yuan, 1998).

Despite its success in China, this innovation remained largely unknown and less utilized outside the country, especially in Africa, where it is potentially well-suited. African countries like Nigeria, Senegal, and Madagascar have very high rice consumption, but local yields around Sub-Saharan Africa averaged just 2.4 tons per hectare due to poor seed quality, erratic rainfall, and limited inputs, making up only 6% of global total rice production with 26% of total harvested area (FAOSTAT, 2022; Yuan et al., 2024). China's hybrids matched these constraints: they are high-yielding; tolerant to heat, drought, and poor soils; and some lines exhibit resistance to threats such as bacterial blight. Unlike US or European varieties optimized for large-scale, mechanized production under controlled irrigation, Chinese hybrids—often requiring hand-planting of young seedlings—were developed for rain-fed, labor-intensive, smallholder systems common across African regions. Despite this suitability, few efforts were made to introduce it to Africa due to language barriers, institutional disconnects, and a development landscape dominated by Western actors. African governments often looked to institutions like the International Rice Research Institute (IRRI) or the World Bank, but largely overlooked Chinese advances.

Changes began in the 1990s, when Chinese researchers affiliated with US institutions initiated influential hybrid rice research (Yu et al., 1997), which in turn spurred follow-on studies in the United States. In the early 2000s, Chinese institutions began working with African partners on hybrid rice demonstration and training. The Africa Rice Center launched a hybrid rice program in partnership with the Chinese Academy of Agricultural Sciences and IRRI under the Green Super Rice initiative (El-Namaky and Demont, 2013). Countries like Côte d'Ivoire, Liberia, Madagascar, Mozambique, Nigeria, Tanzania, and Uganda began evaluating Chinese hybrids (Zenna et al., 2017). These efforts yielded real gains: for instance, Chinese hybrid rice has allowed Tanzania to become the largest regional producer, showing a fivefold increase in paddy rice exports in 2015 compared with the early 2000s (Siméon et al., 2022).

Appendix Figure A1 Panel (a) shows the rise of hybrid rice research in China, the US, and other middle- and low-income countries based on publication counts. Defining the "emergence year" as the point when a region reaches 10% of its total hybrid rice publications, hybrid rice research emerged in China in 1983, in the US in 1999, and in other mid- or low-income countries around 2005. While many factors may explain this sequence, one plausible mechanism is that the growing body of hybrid rice research in the US during the 90s facilitated greater attention, adoption, and local research in other developing regions. Panel (b) is consistent with this idea: hybrid-rice patents authored by Chinese researchers tend to reach broader audiences when they are filed abroad—primarily in the US. Those filed in the US receive nearly twice as many forward citations on average as those filed in China, and have a much higher share of citations coming from a diverse set of regions.

Taken together, Chinese hybrid rice illustrates a powerful case of "trapped knowledge" being untrapped. Once localized agricultural expertise was finally diffused to the right context, it addressed a long-standing productivity gap, improved food security, and reduced import dependence for countries struggling to meet domestic rice demand. Besides the value of the technology itself, this example highlights the importance of creating mechanisms for knowledge to diffuse to places where it is needed. In this regard, some evidence suggests that the US may play a role in helping locally developed knowledge to diffuse globally. Building on this intuition, this paper focuses on the US as a channel through which geographically confined knowledge may be untrapped and brought into broader circulation.

B. Patents and Knowledge Diffusion

Patents, a major form of intellectual property, serve as an important tool in knowledge protection and dissemination. They are designed to protect inventions by granting exclusive rights to their inventions while making the knowledge public. In principle, the global spread of digital patent databases and translation technologies should make all patents universally searchable. In this sense, patents provide not only a legal shield but also a formal channel through which technical knowledge can, at least in principle, circulate across borders.

Yet in practice, many innovations remain less utilized, especially those originating from less prominent or non-Western countries. People may overlook these patents due to language barriers, unfamiliar legal systems, or skepticism about their quality. As a result, valuable inventions may remain confined within local markets, preventing their broader diffusion to other regions that could benefit from them (Jaffe et al., 1993; Thompson and Fox-Kean, 2005). This highlights a central paradox: while patents are designed to disseminate knowledge, institutional and cultural frictions often limit their reach when the inventions come from outside dominant Western systems.

In reality, the motivation for inventors to patent abroad is primarily commercial: inventors seek to protect their competitive position in foreign markets where their products may be sold, produced, or utilized. Filing patents in key jurisdictions like the United States is often a strategic move to facilitate market entry and commercial exploitation. However, international patenting can inadvertently affect the global diffusion of the underlying knowledge. First, patenting abroad can substantially raise the visibility of an invention. For instance, an innovation patented in the United States can benefit from the prestige associated with the USPTO, which is widely regarded as one of the most influential patent offices globally. Further, a US patent may also serve as a signal of quality, suggesting that the invention has passed rigorous examination and is deemed worthy of legal protection. In this sense, the US spotlights foreign inventions—amplifying their exposure beyond the home market.[1]

III. Data and Methodology

A. Data

I focus on patent families: sets of patent applications and publications for the same technology across different countries. The earliest application in a family is the priority

[1] Admittedly, other mechanisms may also be at play. Patenting can sometimes block follow-on innovation by restricting reuse through exclusive rights, a concern highlighted by the anti-commons hypothesis (Murray and Stern, 2007; Sampat and Williams, 2019; Williams, 2013). However, these dynamics are unlikely to explain any positive diffusion effects of foreign patenting. If anything, they work in the opposite direction: by introducing frictions or legal constraints, they tend to bias against observing increases in knowledge flow. Thus, to the extent I find positive effects of US patent grants on downstream use, these likely represent a lower bound on the true impacts.

application, and the patent office where it is filed is the priority authority (or home country). The main sample of interest consists of patent families with a granted priority patent in a non-US country and subsequently files a US application.[2] Patent families with filings to a third country besides the home country and the US, or with filings to regional offices, are excluded, as the impact of US grants will likely be overestimated if it combines the positive effects from filing in other regions. Furthermore, the home application is granted before the US application is filed, which eliminates the home-country granting process as a confounder of the impacts of the US grants and allows me to observe dynamics as I can track citations in years after home-country grants but before US grants.

I select the patent families where the subsequent US application was made between 2001 and 2019. This window accommodates two institutional features central to the empirical strategy: (1) USPTO patent applications filed on or after November 29, 2000 are published in the public record regardless of whether they are eventually granted as part of the American Inventors Protection Act of 1999; and (2) the random assignment of examiners to patent applications ended in 2019, when the USPTO began implementing changes toward a more targeted routing system (Aneja et al., 2024).

I use patent data from PATSTAT Global, which is published by the European Patent Office, with information on patent applications and granted patents collected from national and regional patent offices worldwide (EPO, 2023). For each application, there is information about the application authority (country), whether the application is being granted, corresponding publication information, and the citation network. I then complement this data with bulk data from USPTO, which includes information for examiners and art units—groups of patent examiners who specialize in a particular area of technology—for all the US applications.

The final balanced panel consists of 2,470 patent families over 23 calendar years from 2000 to 2022, corresponding to 506 art units originally filed in 19 non-US priority

[2]Patent families with applications filed to no longer existing countries are excluded from the sample. This includes Czechoslovakia, the German Democratic Republic, the Soviet Union, and Yugoslavia/Serbia and Montenegro.

authorities.[3] Figure 1 shows the distribution of priority authorities and the average grant rate of the US applications from these authorities. East Asian authorities—South Korea, China, and Japan—lead the distribution, followed by Western countries such as Germany, Australia, France, and the United Kingdom.

B. Empirical Strategy

The goal is to empirically examine the impact of patent grants in the US on follow-on innovation, as a channel through which previously trapped knowledge can become untrapped. The unit of analysis is a patent family-year, so a simple comparison can be done to compare two patent families that were originally filed and granted in the same non-US country, belong to the same sector, and filed a US application in the same year, but where only one of the two patent families' US applications is granted. In an ideal world, where all knowledge is treated equally and is equally accessible, this difference should be negligible because the home patent alone should lead to the same exposure with or without the US patent, or be negative if the US grants lead to more legal barriers that hinder follow-on innovation.

1. Difference-in-Differences

I use a difference-in-differences (DID) framework to estimate the effect of US patent grants. For patent family i at calendar year t with US filing in year y, the specification is

$$\text{Citation}_{it} = \alpha + \beta \text{Post Grant}_{it} + \lambda_i + \gamma_{ty(i)} + \varepsilon_{it}, \tag{1}$$

where Post Grant_{iy} is an indicator equal to 1 in years after the US grant of treated patent families. λ_i are patent family fixed effects, $\gamma_{ty(i)}$ are the filing cohort × calendar year fixed effects. While patent family fixed effects remove time-invariant characteristics, how the citations accumulate is highly correlated with the "age" of the patent, which may also correlate with the US grant status. So I include filing cohort × calendar year fixed effects to allow for cohort-specific time trends. Identification, therefore, comes from within-cohort × year contrasts: among patent families with their US applications filed in the

[3] In addition to the standard 4-digit numeric codes (e.g., 2131, 2873, 1611) that belong to Technology Centers, this also includes administrative offices (e.g., Office of Enrollment and Discipline) and special handling divisions (e.g., LO115, EDAD, CSMD).

same year and observed in the same calendar year, those already granted are compared to those not yet granted. The standard errors are clustered by technology class.

Further, I extend the baseline DID specification to a dynamic framework to examine the effects of US patent grants relative to the year of the US grant decision. Specifically, I estimate the following regression:

$$
\begin{aligned}
\text{Citation}_{it} = & \alpha + \sum_{k=-4, k\neq -1}^{6} \beta_k \cdot \mathbf{1}\{k \text{ years relative to US Grant}\}_{ty(i)} \\
& + \lambda_i + \gamma_{yt(i)} + \varepsilon_{it}. \qquad (2)
\end{aligned}
$$

The coefficients on the lead years (years before US grants), β_{-4} through β_{-2}, allow me to assess how outcomes respond to future changes in patent grant status. The coefficients on the lag years (years after US grants), β_0 through β_6, capture how the treatment effect evolves over time, shedding light on whether the effect is short-lived or persistent. The year immediately before the US grants ($k = -1$) is omitted.

2. Examiner Instrument

OLS may yield biased estimates since the probability of being granted is endogenous to the patent's observed and unobserved characteristics.[4] To address this, I follow recent studies that use examiner leniency as an instrumental variable (IV) for patent grant, exploiting quasi-random assignment of examiners as a source of exogenous variation (Sampat and Williams, 2019; Aneja et al., 2024).

Examiners are assigned to patent applications in a manner that is largely random within art units (i.e., based on the last digit of the sequentially-assigned application number), and examiners retain discretion in evaluating patentability despite the uniform mandate to grant patents to eligible, novel, non-obvious, and useful inventions (Cockburn et al., 2003; Lemley and Sampat, 2012). This examiner discretion leads to substantial variations in outcomes, with lenient examiners being more likely to grant applications compared to a counterfactual scenario where these applications are assigned to harsh

[4]Note that all patent families in the treatment group or the control group file a US application, so the selection bias into application—an applicant's decision to applying for patent in the US is endogenous to both the applicant's and the patent's observed and unobserved characteristics—will not play a role in the empirics and therefore not considered in this context.

examiners. Because of the quasi-random assignment, examiner leniency is plausibly uncorrelated with unobserved characteristics of the application, conditional on covariates such as application year and technology class. The IV approach thus identifies a local average treatment effect (LATE) for marginal applications—those whose outcome depends on the leniency of the assigned examiner.

I implement a two-sample two-stage least squares (TS2SLS) estimator as illustrated in Sampat and Williams (2019) and Angrist and Krueger (1992). Specifically, for each patent examiner, I retrieve the historical record of all patents this examiner has handled. I then estimate the examiner's leniency based on this separate sample as the leave-one-out patent grant rate of an examiner for each application filing year following Aneja et al. (2024). Specifically, for each US application for patent family i filed in year y and examined by patent examiner k in art unit a, the leniency is calculated as

$$\text{Leniency}_{iy(i)k(i)a(i)} = \left(\frac{1}{N_{y(i)k(i)a(i)}}\right)\left(\sum_{l\neq i}^{N_{y(i)k(i)a(i)}} Grant_l\right), \tag{3}$$

where $N_{y(i)k(i)a(i)}$ is the total number of applications examined by examiner k in year y in art unit a, and $Grant_l$ equals one if the (non-focal) application is granted by the examiner.[5]

The instrument satisfies the relevance, exclusion, and monotonicity assumptions required for identifying a local average treatment effect. Cross-sectionally, a 10 percentage point increase in examiner leniency is associated with a 9.16 percentage point increase in the likelihood of being granted (SE = 0.052), supporting instrument relevance. Second, the exclusion restriction is plausible given the quasi-random assignment of examiners within art units, making leniency unlikely to be correlated with unobserved application characteristics. Further, monotonicity is also reasonable in this context, as it is unlikely that any application would become less likely to be granted when assigned to a more lenient examiner; examiner discretion operates in one direction, primarily by approving marginal cases rather than reversing clearly patentable ones.

Appendix Figure A2 shows evidence consistent with the relevance and exclusion as-

[5]The main results are robust to alternative definitions of examiner leniency, i.e., defining the leniency at the examiner-art unit or the examiner-year level. In practice, over 85% of examiners in the sample are only associated with one art unit in a given year.

sumptions. Examiner leniency is plotted on the horizontal axis, and the vertical axis shows the actual patent grant rates and the predicted grant rates using observable characteristics that may be correlated with patent value and are fixed at the time of application, including the technology class, year, the priority authority, the size of the patent family, and the number of inventors (Sampat and Williams, 2019; Aneja et al., 2024). The strong association with actual grant rates supports instrument relevance, while the weak relationship with predicted grant rates supports exogeneity.

The endogenous variable $Post\ US\ Grant_{it}$ can be viewed as the interaction of a post-period indicator $Post_{it}$ and an indicator for granted families $Treated_i$, and as a result, is instrumented by $Post_{it} \times Leniency_{t(i)k(i)}$. The assumption is that the endogenous part of the treatment is whether a patent will ever be granted or not, instead of the timing of treatment. Namely, conditional on being granted a patent, the exact time when it will be granted is fairly random. As a result, the instrument $Post_{it} \times Leniency_{t(i)k(i)}$ will be relatively exogenous. To empirically show this relationship, I use the predicted grant rates based on *ex-ante* characteristics, as discussed above, to capture patent value (Sampat and Williams, 2019; Kline et al., 2019). Appendix Figure A3 shows that the predicted grant rate is not associated with when a patent is granted, indicating no evidence that applications that appear more likely to be patented based on covariates that are fixed at the time of filing are differentially more likely to be granted earlier. To further alleviate the concerns about endogenous timing, I also constructed a predicted post-grant indicator $\widehat{Post}_{it} = \mathbf{1}\{t > \widehat{Grant\,Year}_i\}$. Specifically, I predict a grant date for all the patents based on their filing date, the art unit being assigned, and the leniency of the examiner. All these characteristics are pre-determined and are therefore orthogonal to observed and unobserved characteristics like patent quality. The alternative instrument is then $\widehat{Post}_{it} \times Leniency_{t(i)k(i)}$. The robustness of results to this alternative method further validates my empirical strategy.

IV. Empirical Results

The sample of interest includes annual observations from 2000 to 2022 for 2,470 patent families, where the US applications for 1,582 of the families are ultimately granted. Appendix Table A1 shows the distribution of CPC classes (3-digit) in the sample.

Table 1 presents the summary statistics. The sample includes 17,052 observations in the pre-treatment period (never-treated families and pre-grant years for treated families), and 13,192 observations in the post-treatment period for treated families. Citation outcomes rise notably following the US grant. The average number of forward citations nearly doubles, from 0.72 to 1.40, and the probability of being cited increases from 31% to 53%. Citations from the US exhibit the sharpest shift, with both citation likelihood and counts roughly doubling post-grant. Third-country citations also increase meaningfully, from 0.16 to 0.28 citations per year on average, while home-country citations show a more modest rise. Metrics of citation breadth also show that the number of distinct third countries citing the patent doubles post-grant (from 0.07 to 0.14), with gains observed across both high-income and middle/low-income countries. Finally, measures of average "nonfocal" citations—citations to other patents held by the same inventors—also rise post-grant (from 1.05 to 1.70), especially those coming from US and third-country patents.

Taken together, these patterns suggest that receiving a US patent grant is associated with an amplified global exposure of non-US inventions, boosting direct citations to the focal patent family, broadening diffusion across international borders—including into jurisdictions with weaker prior ties—and elevating the visibility of the inventors' broader portfolio. The consistent post-grant increases across citation sources reinforce the idea that US grants serve as a lever for expanding the global exposure of foreign-origin knowledge. While these differences do not on their own establish causality, they provide strong motivation for the following empirical analyses that estimate the effects of US grants on knowledge diffusion.

The empirical analysis proceeds in several stages. First, I estimate the effect of receiving a US patent grant on forward citations and examine the dynamic effects to assess pre-trends and temporal shifts following the grant. Next, I explore the diffusion of impact by disaggregating citations by geographic origin, measuring changes in the number and composition of citing countries, and analyzing how attention spills over to other patents by the same inventors. Finally, I investigate a potential mechanism by studying whether the effects of US grants are larger for countries and sectors with more "trapped" knowledge, using both group-level comparisons and a novel citation-based measure of trappedness.

A. Main Results

Figure 2 plots the average normalized number of citations received in each lead and lag year relative to the US *filing* year, separately for patent families whose US applications are ultimately granted versus those that are never granted.[6] The citations are normalized to mitigate the truncation problem: more recent cohorts of patents will be mechanically less cited, even though they are not less innovative, and the truncation problem varies across technology classes (Lerner and Seru, 2022). Therefore, I normalized the citations by dividing the raw number of citations by the average number of citations for the same patent cohort in each patent class.

On and before the filing year, there is no noticeable difference in citation levels between the two groups, suggesting that pre-existing trends do not drive the citation gap. Also, there is no clear gap in the first year after filing, corresponding well with the fact that the majority of the patents take two years or longer to be granted. Starting from the second year after US filing, however, a clear divergence emerges: granted families begin to accumulate more citations. The difference is relatively stable over time, indicating that, on average, families with granted US applications receive 0.2–0.3 more normalized citations per year. This correspondence between grant timing and citation divergence provides suggestive evidence that it is the US grant itself—rather than other unobserved factors—that drives the observed citation advantage.

Table 2 presents the baseline IV estimates using the difference-in-differences specification described in equation (1). The results show that receiving a US patent grant leads to a statistically significant increase in forward citations: the probability of receiving any citation increases by 7.2 percentage points (a 23% increase over the pre-grant mean), and the average number of citations rises by 0.27 (a 38% increase). The normalized outcomes, which account for variation across filing year and technology class, show similarly strong effects: a 17% increase in citation likelihood and a 27% increase in citation counts.

Appendix Table A2 reports OLS estimates using the same specification as the baseline

[6] Applications are never categorically rejected, but abandoned by applicants following non-final rejections issued by patent examiners (Lemley and Sampat, 2008). Therefore, there is no clear "rejection data" for failed US applications.

IV model. The results are very similar in magnitude and direction, suggesting that any selection bias is likely driven by time-invariant characteristics that are already absorbed by the fixed effects. Appendix Table A3 and A4 demonstrate that the estimated effects remain stable across alternative fixed effects, outcome transformations, and instrument constructions, reinforcing the robustness of the findings. Additionally, the first-stage F-statistic is very large, whose magnitude is consistent with recent literature using examiner-based instruments (Aneja et al., 2024). At the same time, to address concerns that the constructed IV may mechanically overstate the first-stage F-statistics (Hull, 2017), I report Anderson–Rubin Wald statistics to provide weak-instrument robust inference for Table 2. The significant test results indicate that the estimated effects are not sensitive to weak instrument concerns.

Figure 3 presents the dynamic effects based on equation (2), which includes three lead years and six lag years relative to the year immediately before US grants. The coefficients on the lead terms show little evidence of any significant difference in citations prior to the US grants, which is reassuring and inconsistent with the concern that higher-quality patents are both more likely to be cited early and more likely to be granted. Following the grant, citations increase almost immediately, with the effect peaking around two to three years after the grant before gradually tapering off. This pattern suggests that the grant triggers a short-run surge in exposure—bringing the invention more prominently into view within the global innovation system. The effects begin to attenuate after year 4 and become insignificant after year 6, pointing to a largely transitory boost rather than a permanent shift in citation dynamics.

To address potential concerns about heterogeneous treatment timing and dynamic effects that can bias conventional two-way fixed effects models, I also estimate dynamic group-time average treatment effects using the *csdid* estimator developed by Callaway and Sant'Anna (2021). These results, shown in Appendix Figure A4, display similar post-grant increases and flat pre-trends, further reinforcing the main findings and the credibility of the identification strategy.

Overall, these results provide compelling evidence that US patent grants meaningfully amplify the exposure of foreign-origin inventions. The effects are consistent across raw and normalized outcomes and robust to a range of alternative specifications. The identification strategy helps address concerns about endogeneity, and the event-study plots confirm that the citation increase coincides closely with the timing of the grant. Taken together, the findings point to a clear shift in how granted inventions are integrated into the global innovation landscape.

B. The Diffusion of Knowledge

An important dimension of global exposure is not only whether an invention receives more citations, but where those citations come from and how far the knowledge travels. To understand how US patent grants shape the geographic diffusion of foreign-origin inventions, this subsection disaggregates the citation effects by the location of citing inventors, the diversity of countries reached, and the broader spillovers to the inventors' other patents.

Table 3 disaggregates the citation effects of US patent grants by the geographic location of the citing patents. Citations from the US show the strongest response: receiving US grants increase the probability of being cited by US patents by 27% and raises the number of citations by 31%. This sizable shift highlights how the grant meaningfully boosts the invention's exposure within the US innovation system. Citations from the home country also increase, though the effect is more modest—around 20%. This smaller response is consistent with the idea that home-country inventors are already familiar with the invention through local databases, networks, or parallel filings. For them, the US grants likely enhance perceived credibility rather than visibility, while the effect within the US likely reflects both. Interestingly, citations from third countries—those outside both the US and the home country—also rise significantly by 20–23%. These gains suggest that US grants help foreign-origin patents reach broader international audiences that may otherwise have limited access or awareness. This pattern reinforces the idea that the USPTO serves as a global amplifier of exposure, extending the reach of non-US inventions well beyond their origins.

Motivated by this sizable increase in citations from third countries, Table 4 examines whether US grants expand not only the volume of follow-on citations but also their geographic breadth. US grants increase the overall number of citing countries by 0.12, a 35% rise relative to the pre-grant mean. Looking at the third countries, US grants lead to a

37% gain in the number of citing countries—driven by the significant 39% increase in the number of citing mid- and low-income countries. This suggests that the boost in the geographic reach is not merely driven by increased attention from the US or the home country, but reflects meaningful global diffusion. Although the magnitude of the effect remains relatively modest in absolute terms, the increase is both statistically significant and practically meaningful. More importantly, it hints that many home-country inventions may be relevant and appropriate for use in third-country settings, but face structural diffusion barriers that a US grant helps overcome. Appendix Figure A5 illustrates that the citing and cited country-classes tend to be closer in technological composition—measured by the Euclidean distance between their CPC subclass distributions—when the corresponding US application is granted. This difference is highly statistically significant (two-sample t-test of difference in means: $P < 0.0001$) and provides suggestive evidence that US patent grants facilitate the diffusion of innovations to regions where the underlying knowledge is particularly well-suited.

Table 5 approaches the question of diffusion from a different angle: rather than focusing on citations to the focal patent family, it examines whether the grant increases downstream attention to other patents from the same inventors. This allows us to assess whether US grants raise the broader scientific profile of the inventors and attract more attention to their surrounding body of work. Unlike focal patent outcomes that are closely connected to the cohort-specific trends, non-focal outcomes aggregate citations to the inventors' other patents, which span a wide range of filing years and ages. As a result, the differential trends in non-focal citations are shaped more by the size and maturity of the inventor's broader patent portfolio. Intuitively, given that there exists any non-focal patent, the less established the scientists are, the more they stand to benefit from the added exposure following the US grants. Therefore, I focus on patent families associated with below-median inventor prominence, as measured by the number of weighted non-focal patents at the time of US filing.[7] I also include age cohort $\times$ year fixed effects, where the age cohort is based on the average age of the non-focal patents at the time of focal US filing, to account for broader portfolio citation dynamics.

[7] Specifically, for each focal patent, I get a list of inventors and count the number of associated non-focal patents before the US filing of the focal patent that ever-cited to approximate how established the inventors are. The median value is about 15—equivalent to a focal patent with 3 inventors, where each inventor has, on average, 5 patents receiving citations during the years before the US filing of the focal patent.

The results show a statistically significant 15.5% increase in citations to non-focal patents following the grant. This effect is primarily driven by increased attention from the US and from third countries, with no meaningful change in citations from the inventors' home country. The fact that third-country citations rise by over 23% shows that US patent grants expand the global exposure of the inventors—not just of the focal invention—within the global innovation network. The grant serves as a gateway, pulling a broader portfolio of work into the global spotlight that might otherwise remain peripheral to international audiences. Appendix Figure A6 further decomposes the impact on non-focal citations by quartiles of how established the focal scientists are. The largest gains appear in the bottom quartile (Q1), followed by smaller but still significant effects in Q2; and the effects are insignificant for Q3 and Q4.[8] These results align with the intuition that less established scientists benefit more from the added global exposure provided by US patent grants, as their broader body of work may otherwise remain under the radar.

Taken together, the results underscore that US patent grants play a meaningful role in diffusing knowledge. They facilitate access to third countries, broaden the geographic reach of citations, and draw attention to the inventors' broader portfolio—especially for less established inventors. These patterns reinforce that the US patent system functions not just as a legal mechanism, but as a platform for global exposure.

C. The Trappedness Mechanism

Not all foreign knowledge benefits equally from being granted in the US. Innovations from countries that are already deeply embedded in global innovation networks may need little additional exposure, while those from more peripheral or linguistically isolated countries face greater barriers. Therefore, US grants may play a disproportionately important role in enabling diffusion for countries that are more "trapped". To better understand this mechanism, I first examine how the effect of US patent grants on forward citations varies by the linguistic and regional background of the priority (home) country.

Table 6 shows that the increase in citations following a US patent grant is primarily driven by Korean and Chinese patents: US grants meaningfully amplify the global ex-

[8] Inventor prominence is highly skewed: the 25^{th}, 50^{th}, and 75^{th} percentiles correspond to approximately 4, 15, and 92 average scientist-weighted non-focal patents, respectively. The 90^{th} percentile exceeds 1,300, likely reflecting large firms with very different underlying dynamics; thus, Q4 is capped at the 90th percentile.

posure of these inventions—both within the US system and across third countries that may otherwise have limited ties to the origin. In contrast, there is little evidence of increased citations for Japanese, English-dominant, or non-English-speaking European patents. This likely reflects their already high levels of visibility due to long-standing integration into global innovation networks. For instance, Japan was a global technological leader during the 1990s, while countries like Germany, France, the UK, and Canada have long-standing institutional, legal, and linguistic alignment with the US and broader international patenting systems. Interestingly, citations from the home country increase only for Chinese patents. This likely reflects a unique domestic mechanism: in a setting where scientific output has grown rapidly but concerns about patent quality remain, a US grant provides a credible external signal of quality, which not only boosts international attention but also enhances local uptake.

Taken together, the results suggest that US patent grants help elevate innovations from more peripheral or linguistically isolated countries. For inventors from places like Korea and China—where language barriers, institutional distance, or concerns about domestic patent quality may limit international engagement—the US grants provide a powerful boost in exposure. This pattern highlights the asymmetric value of US grants in reducing diffusion barriers for otherwise overlooked knowledge.

1. Citation-based Trappedness

To further explore the mechanism, I quantify the extent to which knowledge from a given country and technology class is overlooked by global inventors with a novel citation-based measure of "trappedness". This approach is motivated by the distinction between two types of backward citations in patent applications: citations added by applicants and citations added by examiners in search reports.[9] The key assumption is that examiner-added citations serve as a proxy for knowledge that was not readily known or visible to the inventors at the time of filing. Such knowledge, while formally relevant, likely did not contribute to the invention itself (Alcácer et al., 2009). Therefore, a high share of such examiner-added citations suggests that applicants may be systematically unaware of relevant prior art from that source—a signal that the knowledge is "trapped" in its

[9]There are other types of backward citations, like the citations introduced when filed for appeal by applicant/proprietor or the citations introduced during opposition. All these other types constitute a minimal share and are excluded from the analysis.

country of origin.

It is important to emphasize that this framework does not assume that examiners are omniscient. Rather, the assumption is that examiners, by virtue of their role to check the originality of the applications and the professional search resources, are systematically more capable than inventors and their patent attorneys in identifying relevant prior art. In this sense, the space of potential citations is less "trapped" with respect to examiners than it is with respect to inventors, making examiner-added citations a useful signal of invisible or under-recognized knowledge. This assumption is consistent with prior findings that citations from inventors are much more likely to be from near the inventor's location than examiner-added citations (Thompson and Fox-Kean, 2005).

To address concerns that the measure may be confounded by general home bias—i.e., the tendency of both applicants and examiners to cite patents from their own country due to familiarity or institutional norms—I compute the ratio of examiner-added to applicant-added citations. This relative measure helps cancel out factors that similarly affect both parties, under the assumption that applicants and examiners from the same country share a comparable baseline familiarity with foreign knowledge and face the same institutional environment. Specifically, I calculate this measure using all backward citations made by international patent filings submitted through the Patent Cooperation Treaty (PCT) to WIPO, which provides a relatively consistent basis for comparison across countries.[10] For each country c being cited, I calculate:

$$TrappedRatio_c = \log\left(\frac{\sum_i \mathbf{1}\{ExaminerAdded_i = 1\} \times \mathbf{1}\{CitedCountry_i = c\}}{\sum_i \mathbf{1}\{ExaminerAdded_i = 0\} \times \mathbf{1}\{CitedCountry_i = c\}}\right), \quad (4)$$

where i indicates each single citation, and $ExaminerAdded_i$ is an indicator if this citation is added by the patent examiner. A trapped ratio equal to 2 means that for 1 citation added by applicants, $e^2 \approx 7.38$ citations are added by examiners.

Recognizing that trappedness may vary not only across countries but also across technology fields, I extend the measure to the country-sector level ($TrappedRatio_{cj}$), where j denotes the technology class (CPC) of the cited patent. For instance, knowledge

[10]National route filings are excluded because citation tagging practices vary widely across patent offices. For example, the China National Intellectual Property Administration (CNIPA) does not indicate whether a citation was added by the applicant or the examiner.

in China's engineering and chemistry may be less trapped than knowledge in its social science or health-related fields.

Figure 4 plots the $TrappedRatio_c$ for all top-cited countries. A higher value indicates greater reliance on examiner intervention relative to applicant awareness. South Korea exhibits the highest trappedness, suggesting that the majority of citations to Korean prior art are added by examiners rather than applicants. Russia, the former Soviet Union, and China also rank high on this measure, while countries such as Germany, Spain, the UK, and Canada are among the least trapped—indicating higher baseline awareness of their knowledge by global inventors.

Table 7 Panel A investigates the trappedness mechanism by interacting the grant indicator with this measure of country-level trappedness. The idea is straightforward: the more a country is trapped, the more likely its innovations are to be overlooked by the broader global audience. If the US serves as a platform to validate and surface such knowledge, then the benefit of US grants should be greater for these more trapped countries. The estimates confirm this mechanism. The interaction term between grant status and country-level trappedness is positive and statistically significant, indicating that US patent grants are especially effective for countries that start from a disadvantaged position. For example, patents from the most trapped countries (with $TrappedRatio_c \simeq 4$) receive 0.182 more normalized citations post-grant, on top of the insignificant baseline effect of 0.0687.

There are important dynamics when looking across citing geographies. Citations from the US show a significantly positive baseline response to grant status, likely reflecting the direct institutional role of the USPTO dominating the effect regardless of country origin—making it harder to detect differential effects based on trappedness. Further, citations from the home country remain largely unaffected by trappedness, consistent with the intuition that most home inventors are already aware of local knowledge (with China being a notable exception discussed earlier). Most strikingly, third-country citations exhibit a strong interaction effect that drives the overall effect. This pattern reinforces the idea that the US grants serve as a global exposure mechanism, helping innovations break into distant or disconnected regions of the innovation network that would otherwise remain out of reach.

One potential concern is that the results in Panel A may be driven by a few large countries that exhibit high levels of trappedness and account for a sizable share of the sample. To address this, I turn to the more granular country–sector level measure of trappedness: $TrappedRatio_{cj}$. As shown in Appendix Figure A7, this measure introduces greater variation across country–technology combinations and leads to a more balanced distribution of observations, reducing the risk that the effects are dominated by a few dominant countries. Panel B of Table 7 presents consistent results using this refined measure. These patterns further strengthen the interpretation that US patent grants are especially effective at elevating under-recognized knowledge—not just from specific countries, but also across a broader set of technologically specialized domains where diffusion barriers are highest.

Together, these findings suggest that US grants help reduce frictions in the global innovation system, particularly for knowledge that is otherwise less visible. The patterns are consistent at both the national and sectoral levels, confirming that untrapping is especially impactful when the knowledge is the most trapped.

V. Conclusion

This paper investigates how granting US patents to foreign-origin inventions affects their global diffusion. While the primary function of the US patent system is legal protection, I show that it also acts as a powerful informational platform that surfaces knowledge previously confined within regional or linguistic boundaries. US grants significantly increase the breadth of follow-on citations and elevate the role of the knowledge in global innovation networks.

Notably, US grants help the global diffusion of foreign patents by increasing citations from third countries and reaching more mid- or low-income countries. They also elevate attention to the inventors' broader portfolio, with the strongest gains observed among less established scientists. Moreover, the diffusion benefits are largest for knowledge originating from countries and sectors that are *ex ante* more trapped—whose innovations are less visible in global citation networks. This suggests that US grants play a disproportionately important role in elevating knowledge from more disconnected parts of the global innovation system, helping surface ideas that might otherwise remain under-recognized.

These findings carry two central implications. First, they underscore the dual role of the US patent system—not only as a legal institution but as a central platform in global knowledge diffusion. Granting a patent in the US elevates an invention's visibility and credibility, and connects it to wider scientific and commercial audiences. Second, they highlight the unequal landscape of global innovation diffusion, in which some countries' knowledge requires external validation to gain recognition. While the US grant does help overcome key frictions, it cannot fully substitute for more systemic solutions. Reducing informational, legal, and linguistic frictions in the global innovation ecosystem—through better citation tools, interoperable metadata, institutional translation services, etc.—is critical for ensuring that valuable ideas can travel as freely as their merit deserves.

Finally, this study calls for greater attention to science and technology originating outside the United States. There may be many "hidden diamonds" that either never pursued US patent protection or failed to get a patent due to procedural, linguistic, or institutional barriers, rather than any lack of merit. These ideas may remain invisible not because they lack value, but because they were never integrated into the global systems that confer enough exposure. Much as Bell et al. (2019) highlight the role of exposure in revealing "lost Einsteins," the findings here point to exposure as the force that can reveal "lost ideas." Expanding efforts to identify, validate, and support high-potential knowledge from underrepresented regions is important not only for equity but also for accelerating global innovation and solving global challenges.

REFERENCES

Alcácer, J., Gittelman, M., and Sampat, B. (2009). Applicant and examiner citations in us patents: An overview and analysis. *Research Policy*, 38(2):415–427.

Aneja, A., Reshef, O., and Subramani, G. (2024). Attrition and the gender patenting gap. *Review of Economics and Statistics*, pages 1–31.

Angrist, J. D. and Krueger, A. B. (1992). The effect of age at school entry on educational attainment: an application of instrumental variables with moments from two samples. *Journal of the American statistical Association*, 87(418):328–336.

Bell, A., Chetty, R., Jaravel, X., Petkova, N., and Van Reenen, J. (2019). Who becomes an inventor in america? the importance of exposure to innovation. *The Quarterly Journal of Economics*, 134(2):647–713.

Borjas, G. J. and Doran, K. B. (2012). The collapse of the soviet union and the productivity of american mathematicians. *The Quarterly Journal of Economics*, 127(3):1143–1203.

Branstetter, L. G. (2001). Are knowledge spillovers international or intranational in scope?: Microeconometric evidence from the us and japan. *Journal of international Economics*, 53(1):53–79.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.

Cockburn, I., Kortum, S., and Stern, S. (2003). Are all patent examiners equal? examiners, patent characteristics, and litigation outcomes. *Patents in the knowledge-based economy*, 19:52–53.

David, P. A. (2004). Can" open science" be protected from the evolving regime of ipr protections? *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft*, pages 9–34.

De Rassenfosse, G. and Seliger, F. (2020). Sources of knowledge flow between developed and developing nations. *Science and Public Policy*, 47(1):16–30.

Eaton, J. and Kortum, S. (1996). Trade in ideas patenting and productivity in the oecd. *Journal of international Economics*, 40(3-4):251–278.

El-Namaky, R. A. and Demont, M. (2013). Hybrid rice in africa: challenges and prospects. *Realizing Africa's rice promise*, pages 173–178.

EPO (2023). PATSTAT Global. https://doi.org/10.7910/DVN/FX3DIE.

FAOSTAT (2022). Statistical database. https://www.fao.org/faostat/en/#data/QCL. Accessed April 10, 2022.

Fry, C. and Furman, J. L. (2023). Migration and global network formation: Evidence from female scientists in developing countries. *Organization Science*.

Furman, J. L. and Stern, S. (2011). Climbing atop the shoulders of giants: The impact of institutions on cumulative research. *American Economic Review*, 101(5):1933–1963.

Galasso, A. and Schankerman, M. (2015). Patents and cumulative innovation: Causal evidence from the courts. *The Quarterly Journal of Economics*, 130(1):317–369.

Ganguli, I. (2015). Immigration and ideas: what did russian scientists "bring" to the united states? *Journal of Labor Economics*, 33(S1):S257–S288.

Gans, J. S. and Stern, S. (2000). Incumbency and r&d incentives: Licensing the gale of creative destruction. *Journal of Economics & Management Strategy*, 9(4):485–511.

Heller, M. A. and Eisenberg, R. S. (1998). Can patents deter innovation? the anticommons in biomedical research. *Science*, 280(5364):698–701.

Hellman, T. (2007). Bridging the science to market gap: the role of patents. *Journal of Economic Behavior and organization*, 63:624–647.

Hull, P. (2017). Examiner designs and first-stage f statistics: A caution. *Unpublished Working Paper*, 5.

Jaffe, A. B., Trajtenberg, M., and Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *the Quarterly journal of Economics*, 108(3):577–598.

Jones, C. I. (1995). R & d-based models of economic growth. *Journal of political Economy*, 103(4):759–784.

Keller, W. (2004). International technology diffusion. *Journal of economic literature*, 42(3):752–782.

Kerr, W. R. (2008). Ethnic scientific communities and international technology diffusion. *The Review of Economics and Statistics*, 90(3):518–537.

Kline, P., Petkova, N., Williams, H., and Zidar, O. (2019). Who profits from patents? rent-sharing at innovative firms. *The quarterly journal of economics*, 134(3):1343–1404.

Lemley, M. A. and Sampat, B. (2008). Is the patent office a rubber stamp. *Emory LJ*, 58:181.

Lemley, M. A. and Sampat, B. (2012). Examiner characteristics and patent office outcomes. *Review of economics and statistics*, 94(3):817–827.

Lerner, J., Liu, J., Moscona, J., and Yang, D. Y. (2024). Appropriate entrepreneurship? the rise of china and the developing world. Technical report, National Bureau of Economic Research.

Lerner, J. and Seru, A. (2022). The use and misuse of patent data: Issues for finance and beyond. *The Review of Financial Studies*, 35(6):2667–2704.

MacGarvie, M. (2006). Do firms learn from international trade? *Review of Economics and Statistics*, 88(1):46–60.

Merges, R. P. and Nelson, R. R. (1994). On limiting or encouraging rivalry in technical progress: The effect of patent scope decisions. *Journal of Economic Behavior & Organization*, 25(1):1–24.

Mokyr, J. (2011). The gifts of athena: Historical origins of the knowledge economy. In *The gifts of Athena*. princeton university press.

Moscona, J. and Sastry, K. A. (2025). Inappropriate technology: Evidence from global agriculture. Technical report, National Bureau of Economic Research.

Murray, F. and Stern, S. (2007). Do formal intellectual property rights hinder the free flow of scientific knowledge?: An empirical test of the anti-commons hypothesis. *Journal of Economic Behavior & Organization*, 63(4):648–687.

Rogers, E. M., Singhal, A., and Quinlan, M. M. (2014). Diffusion of innovations. In *An integrated approach to communication theory and research*, pages 432–448. Routledge.

Romer, P. M. (1990). Endogenous technological change. *Journal of political Economy*, 98(5, Part 2):S71–S102.

Sampat, B. and Williams, H. L. (2019). How do patents affect follow-on innovation? evidence from the human genome. *American Economic Review*, 109(1):203–236.

Siméon, N., Li, X., and Xiao, S. (2022). China's agricultural assistance efficiency to africa: Two decades of forum for china-africa cooperation creation. *Journal of Agriculture and Food Research*, 9:100329.

Thompson, P. and Fox-Kean, M. (2005). Patent citations and the geography of knowledge spillovers: A reassessment. *American Economic Review*, 95(1):450–460.

Williams, H. L. (2013). Intellectual property rights and innovation: Evidence from the human genome. *Journal of Political Economy*, 121(1):1–27.

Yu, S., Li, J., Xu, C., Tan, Y., Gao, Y., Li, X., Zhang, Q., and Maroof, M. S. (1997). Importance of epistasis as the genetic basis of heterosis in an elite rice hybrid. *Proceedings of the National Academy of Sciences*, 94(17):9226–9231.

Yuan, L. P. (1998). Hybrid rice breeding in china. *Advances in Hybrid Rice Technology. Philippines: International Rice Research Institute*, pages 27–33.

Yuan, S., Saito, K., van Oort, P. A., van Ittersum, M. K., Peng, S., and Grassini, P. (2024). Intensifying rice production to reduce imports and land conversion in africa. *Nature communications*, 15(1):835.

Zenna, N., Senthilkumar, K., and Sie, M. (2017). Rice production in africa. *Rice production worldwide*, pages 117–135.

Main Figures

Figure 1: Distribution of Priority Authorities

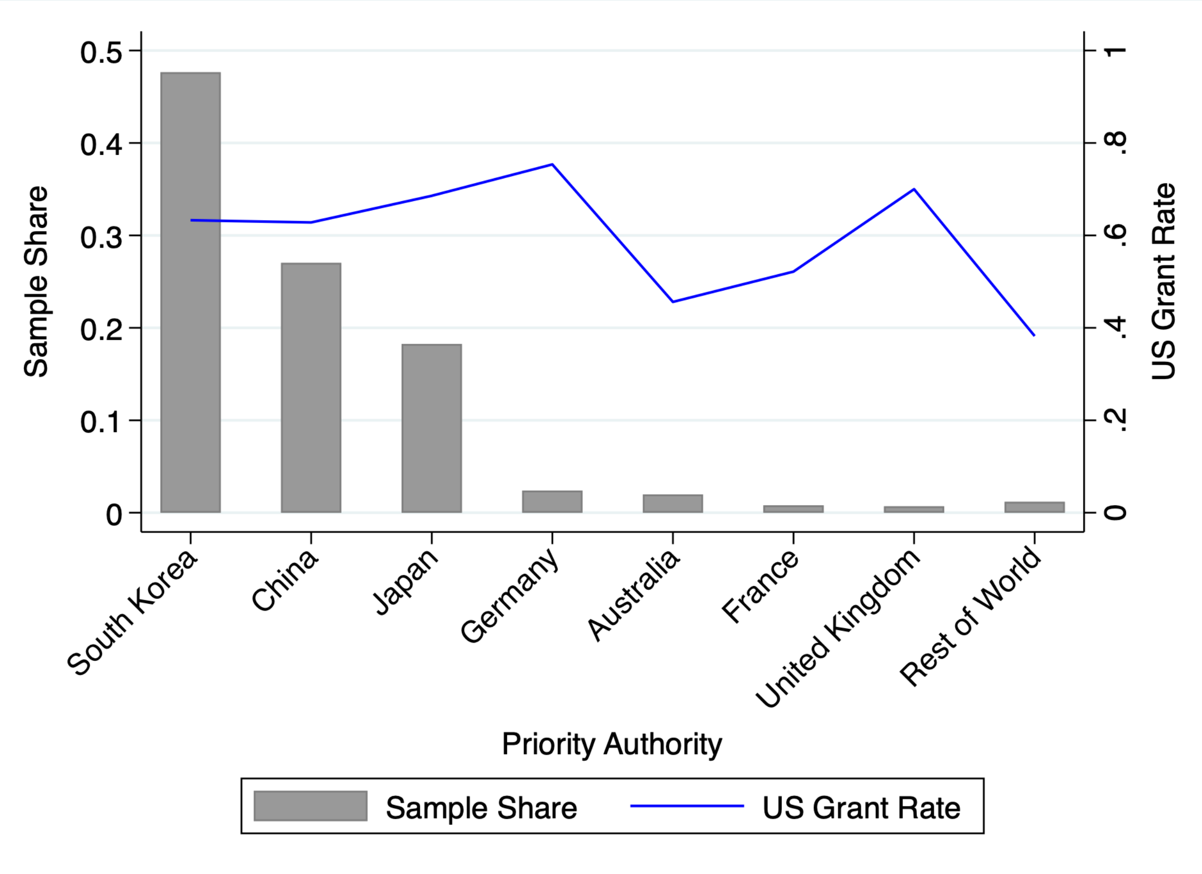

Note: This figure shows the share of patent families from different priority authorities (left vertical axis) and the share of patent families whose US filings were granted (right vertical axis) in the sample of interest.

Figure 2: Citations Trends by US Grant Status

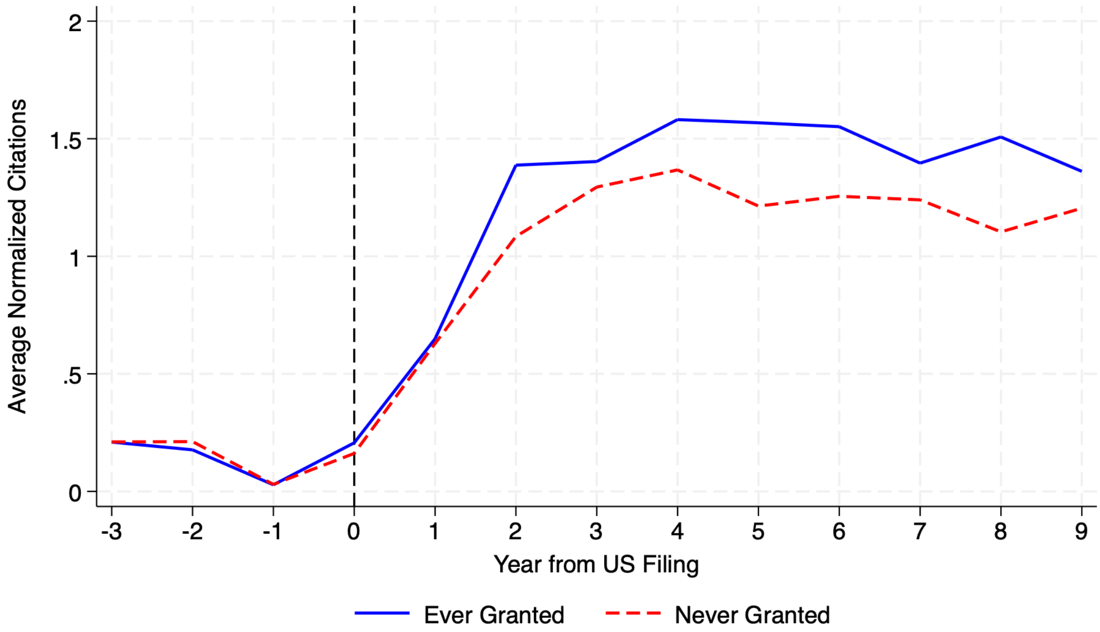

Note: This figure shows the average number of normalized citations relative to the US filing year (left vertical axis), with 3 lead years and 9 lag years, separately for patent families whose US application is granted (blue line) and patent families whose US application is not granted (red line).

Figure 3: Dynamic Effect of US Grants

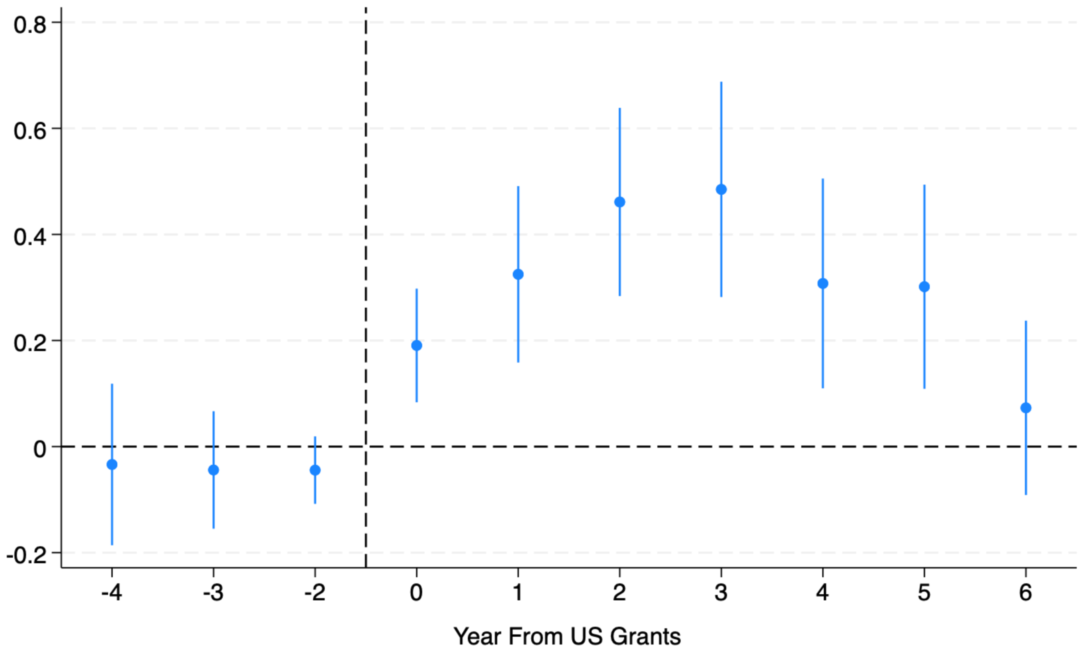

Note: This figure shows an event study plot from an IV regression with 3 lead years and 6 lag years. The year immediately before the US grants is omitted (vertical dashed line). Year -4 includes all years before, and year 6 includes all years after. The unit of observation is a patent family-year. Patent family fixed effects and cohort × year fixed effects are included. Standard errors are clustered by CPC subclass, and the 95% confidence interval is plotted.

Figure 4: Citation-Based Trappedness Measure

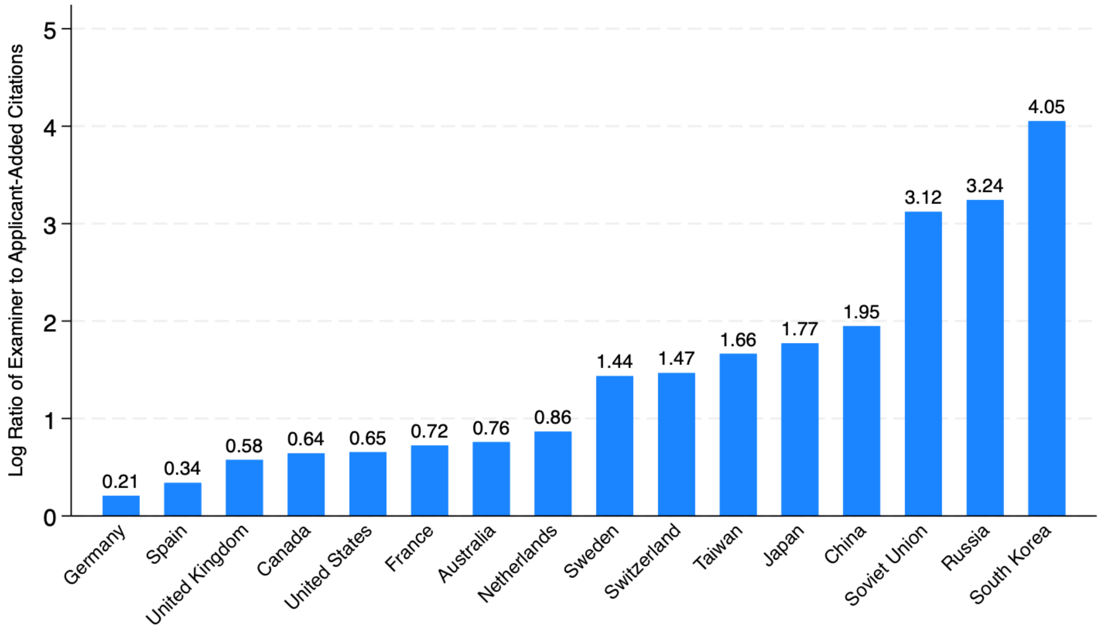

Note: This figure shows the trappedness ratio of top-cited countries/regions, where trappedness is calculated as the log ratio of the number of backward citations added by examiners to the number of backward citations added by applicants. The higher the log ratio, the more "trapped" the country is. Note that the figures display this measure for the top-cited countries, although some, such as the Soviet Union, are not present in the sample for analysis.

Main Tables

Table 1: Summary Statistics

	Pre-Grant ($N = 17,052$)				*Post-Grant* ($N = 13,192$)			
Variable	Mean	SD	Min	Max	Mean	SD	Min	Max
Any Cite (0/1)	0.31	0.46	0	1	0.53	0.50	0	1
Number of Cites	0.72	1.87	0	47	1.40	2.45	0	70
Any Cite from US (0/1)	0.21	0.41	0	1	0.40	0.49	0	1
Number of Cites from US	0.45	1.57	0	47	0.95	2.05	0	69
Any Cite from Home (0/1)	0.07	0.26	0	1	0.11	0.31	0	1
Number of Cites from Home	0.11	0.54	0	18	0.17	0.66	0	16
Any Cite from Third Countries (0/1)	0.11	0.31	0	1	0.19	0.39	0	1
Number of Cites from Third Countries	0.16	0.55	0	10	0.28	0.72	0	10
Number of Countries	0.35	0.60	0	4	0.64	0.75	0	5
Number of Third Countries	0.07	0.27	0	3	0.14	0.38	0	3
Number of Third High-Income Countries	0.02	0.15	0	2	0.04	0.21	0	3
Number of Third Mid/Low-Income Countries	0.05	0.22	0	2	0.09	0.29	0	2
Average Nonfocal Cites	1.05	1.16	0	20	1.70	0.87	0	12.51
Average Nonfocal Cites from US	0.53	1.01	0	20	0.92	0.95	0	12.29
Average Nonfocal Cites from Home	0.29	0.47	0	10	0.37	0.34	0	3.33
Average Nonfocal Cites from Third Countries	0.23	0.40	0	15	0.41	0.31	0	4

Note: This table presents summary statistics for the panel, which includes 2,470 patent families, where 888 families are never-treated and 1,582 are treated at some point. “Pre-Grant” includes all never-treated families and the pre-treatment periods of treated families. “Any Cite” is an indicator variable if the patent is ever cited. “Average Nonfocal Cites” is the average (weighted) number of citations received by other patents of the focal inventors. For all “Average Nonfocal Cites” measures, summary statistics are presented for patent families with a below-median number of (weighted) nonfocal patents before US filing of the focal patent.

Table 2: US Grant and Forward Citations

	Raw Outcomes		*Normalized Outcomes*	
DV:	*Exist Cite (0/1)*	*Num Cite*	*Exist Cite (0/1)*	*Num Cite*
	(1)	(2)	(3)	(4)
Post Grant	0.0715***	0.272***	0.134***	0.203***
	(0.012)	(0.061)	(0.031)	(0.046)
Patent Family FE	Yes	Yes	Yes	Yes
Cohort × Year FE	Yes	Yes	Yes	Yes
Mean (Pre-Grant)	0.31	0.72	0.80	0.76
Percent Δ	23.1%	37.8%	16.8%	26.7%
F-Stat	17795	17795	17795	17795
Weak-IV Robust Wald χ^2	37.70	20.32	19.82	20.94
N	30141	30141	30141	30141

Note: This table estimates the impact of US patent grants on an indicator of the existence of any forward citation and the number of forward citations. The unit of observation is a patent family-year. Columns 1 and 2 present the raw outcomes, and columns 3 and 4 present the normalized outcomes, where each outcome is divided by the average value of that outcome within the corresponding US filing year and technology class. All estimates are IV estimates. Patent family fixed effects and cohort × year fixed effects are included. Kleibergen-Paap weak identification F-statistic and Anderson–Rubin (AR) Wald χ^2-statistic are reported. Standard errors are clustered by CPC subclass, and *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Table 3: US Grant and Forward Citations by Citing Location

	DV: Normalized Number of Citations					
from:	US		Home		Third Countries	
	Exist (0/1) (1)	*Num* (2)	*Exist (0/1)* (3)	*Num* (4)	*Exist (0/1)* (5)	*Num* (6)
Post Grant	0.204***	0.226***	0.163**	0.197**	0.152***	0.174***
	(0.043)	(0.059)	(0.078)	(0.087)	(0.053)	(0.065)
Patent Family FE	Yes	Yes	Yes	Yes	Yes	Yes
Cohort × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Norm. Mean (Pre-Grant)	0.76	0.73	0.81	0.81	0.75	0.75
Percent Δ	26.8%	31.0%	20.1%	24.3%	20.3%	23.2%
Raw Mean (Pre-Grant)	0.21	0.45	0.07	0.11	0.11	0.16
F-Stat	16357	16357	16425	16425	16148	16148
N	28218	28218	23080	23080	27167	27167

Note: This table estimates the impact of US patent grants by the authority of the citing patents. The unit of observation is a patent family-year. All outcomes are normalized within the corresponding US filing year and technology class, and all estimates are IV estimates. Columns (1) and (2) focus on citations from the US; columns (3) and (4) focus on the home country; and columns (5) and (6) focus on citations from countries other than the US and the home country. Patent family fixed effects and cohort × year fixed effects are included. Kleibergen-Paap weak identification F-statistic is reported. Standard errors are clustered by CPC subclass, and *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Table 4: US Grant and Geographic Reach of Innovation

DV:	Num Countries (1)	Num Third Countries (2)	Num High Income Third Countries (3)	Num Mid/Low Income Third Countries (4)
Post Grant	0.123***	0.0261***	0.00632	0.0197***
	(0.017)	(0.007)	(0.004)	(0.006)
Patent Family FE	Yes	Yes	Yes	Yes
Cohort × Year FE	Yes	Yes	Yes	Yes
Mean (Pre-Grant)	0.35	0.07	0.02	0.05
Percent Δ	35.1%***	37.3%***	31.6%	39.4%***
F-Stat	17795	17795	17795	17795
N	30141	30141	30141	30141

Note: This table estimates the impact of US patent grants on the geographic reach of the patents. The unit of observation is a patent family-year. High-income or mid/low-income are classified based on the World Bank Income Classification. Patent family fixed effects and cohort × year fixed effects are included. Kleibergen-Paap weak identification F-statistic is reported. Standard errors are clustered by CPC subclass, and *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Table 5: US Grant and Increased Scientist Attention

	DV: Average Number of Citations to Non-Focal Patents			
	Total (1)	*From US* (2)	*From Home* (3)	*From Third Countries* (4)
Post Grant	0.163***	0.0872**	0.0222	0.0539***
	(0.036)	(0.039)	(0.014)	(0.013)
Patent Family FE	Yes	Yes	Yes	Yes
Age Cohort × Year FE	Yes	Yes	Yes	Yes
Mean (Pre-Grant)	1.05	0.53	0.29	0.23
Percent Δ	15.5%***	16.5%**	7.7%	23.4%***
F-Stat	12233	12233	12233	12233
N	23713	23713	23713	23713

Note: This table estimates the impact of US patent grants on the attention the home country scientists get. The unit of observation is a patent family-year. The outcomes are the average (weighted) number of citations received by other patents of the focal inventors. The regression includes the set of patent families with a below-median number of (weighted) nonfocal patents before US filing of the focal patent. Patent family fixed effects and age cohort × year fixed effects are included, where age cohort is the quintiles of the average age of all non-focal patents at the time of the focal patent's US filing. Kleibergen-Paap weak identification F-statistic is reported. Standard errors are clustered by CPC subclass, and *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Table 6: US Grant and Follow-on Innovation by Home Country Languages

Home Country Language:	Japanese (1)	Korean (2)	Chinese (3)	Eng.-Dominant (4)	Non-Eng. (EU) (5)
DV: Normalized Number of Citations					
Post Grant	0.231	0.227**	0.296***	0.00680	0.327
	(0.210)	(0.095)	(0.102)	(0.250)	(0.211)
DV: Normalized Number of Citations from US					
Post Grant	0.317	0.225**	0.204**	0.0325	0.0838
	(0.194)	(0.095)	(0.093)	(0.424)	(0.233)
DV: Normalized Number of Citations from Home					
Post Grant	0.204	0.132	0.437***	0.281	0.523
	(0.251)	(0.109)	(0.133)	(0.655)	(0.603)
DV: Normalized Number of Citations from Third Countries					
Post Grant	0.0341	0.236**	0.174*	-0.325	-0.0255
	(0.202)	(0.100)	(0.098)	(0.572)	(0.284)
Patent Family FE	Yes	Yes	Yes	Yes	Yes
Cohort $\times$ Year FE	Yes	Yes	Yes	Yes	Yes

Note: This table estimates the effect of US patent grants on the number of forward citations separately for 5 country groups. The country groups are categorized based on language and regional factors. The "Japanese" group includes Japan; the "Korean" group includes South Korea; and the "Chinese" group encompasses China, including Taiwan province and Hong Kong SAR, where Chinese is the main language. These three languages are separated out because they have a large representation in the sample. The "English-dominant" group consists of countries where English is the predominant language, including the United Kingdom, Canada, Australia, Ireland, Singapore, and New Zealand. The "Non-Eng. (EU)" group includes countries in the European Union, such as Germany, France, Italy, Spain, Netherlands, Sweden, Finland, Norway, Denmark, etc. where languages other than English are primarily spoken. The unit of observation is a patent family-year. All outcomes are normalized within the corresponding US filing year and technology class. All estimates are IV estimates. Patent family fixed effects and cohort $\times$ year fixed effects are included. Standard errors are clustered by CPC subclass, and *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Table 7: US Grant and Follow-on Innovation by Home Trappedness

(a) Country-Level Trappedness

	DV: Normalized Number of Citations			
	Total (1)	*From US* (2)	*From Home* (3)	*From Third Countries* (4)
Post Grant	0.0687	0.146^{**}	0.0497	-0.0783
	(0.041)	(0.068)	(0.171)	(0.076)
Post Grant$\times$Trapped Ratio$_c$	0.0454^{***}	0.0282	0.0426	0.0812^{***}
	(0.011)	(0.018)	(0.058)	(0.018)
Patent Family FE	Yes	Yes	Yes	Yes
Cohort $\times$ Year FE	Yes	Yes	Yes	Yes
N	30212	28289	23145	27238

(b) Country-Sector-Level Trappedness

	DV: Normalized Number of Citations			
	Total (1)	*From US* (2)	*From Home* (3)	*From Third Countries* (4)
Post Grant	0.0776	0.139^{*}	0.0767	0.0502
	(0.054)	(0.065)	(0.216)	(0.095)
Post Grant$\times$Trapped Ratio$_{c,j}$	0.0391^{**}	0.0311^{*}	0.0309	0.0391^{*}
	(0.016)	(0.016)	(0.069)	(0.020)
Patent Family FE	Yes	Yes	Yes	Yes
Cohort $\times$ Year FE	Yes	Yes	Yes	Yes
N	25805	24244	21141	23730

Note: This table estimates the differential effects of US patent grants by the measure of trappedness at the country level and the country-sector level. The unit of observation is a patent family-year. All outcomes are normalized within the corresponding US filing year and technology class. All estimates are OLS estimates. Patent family fixed effects and cohort $\times$ year fixed effects are included. Standard errors are clustered by priority authority, and *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Appendix Figures and Tables

Figure A1: Hybrid Rice Publications and Patents

(a) Cumulative Distribution of Publications by Region

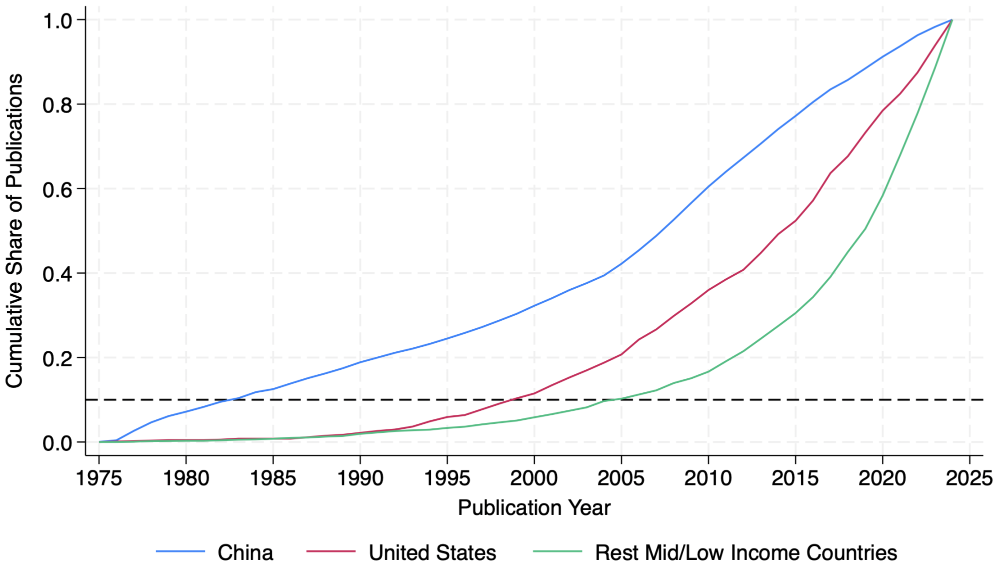

(b) Forward Citation Distribution of Chinese Rice Patents

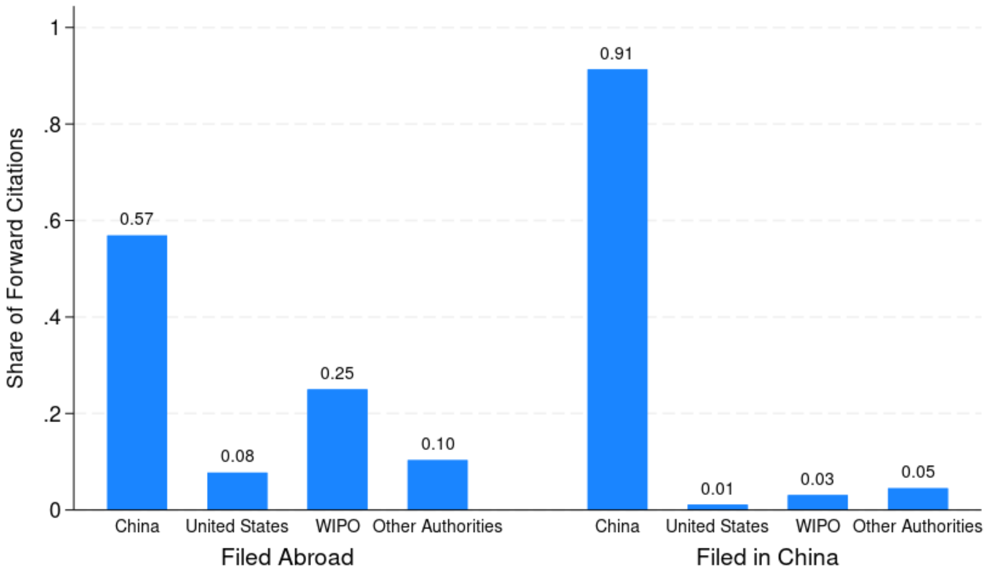

Note: This figure focus on hybrid rice publications and patents. Panel (a) shows the cumulative distribution of the number of hybrid rice publications over time for China, the US, and other mid- or low-income countries. The publications are retrieved from China National Knowledge Infrastructure (CNKI) for China and from Dimensions.org for US and other mid- or low-income countries. Panel (b) shows the distribution of forward citations received by all hybrid rice patents authored by Chinese researchers, separately by whether the patent is filed in China or abroad, and by the jurisdictions where the forward citations are made (China, US, WIPO, and others). China is a global leader in hybrid rice: out of a total of 3,737 hybrid rice patents, 85.4% are from China. Among these, 122 patents are filed in a foreign country, primarily in the US. Conditional on ever being cited, the patents filed domestically receive, on average, 4.8 forward citations, while those filed in the US patents that received 8.3 forward citations on average.

Figure A2: Probability of Patent Grant by Examiner Leniency

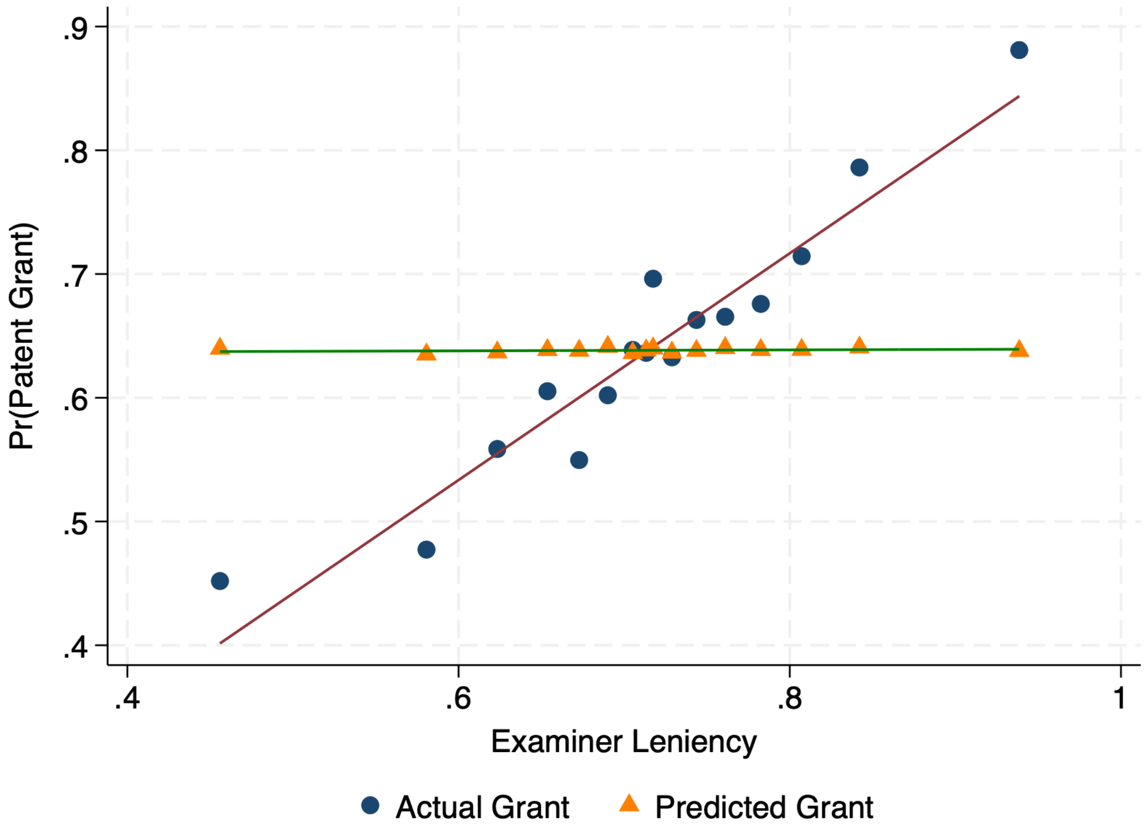

Note: This figure relates the examiner leniency measure, residualized by priority authority, to two variables: (1) the patent grant rate, shown in blue, and (2) the predicted patent grant rate, shown in orange. The predicted patent grant probability is based on observable characteristics that may be correlated with patent value and are fixed at the time of application, including the sector, year, the priority authority, the number of inventors, and the size of the patent family.

Figure A3: Probability of Patent Grant and Timing of Patent Grant

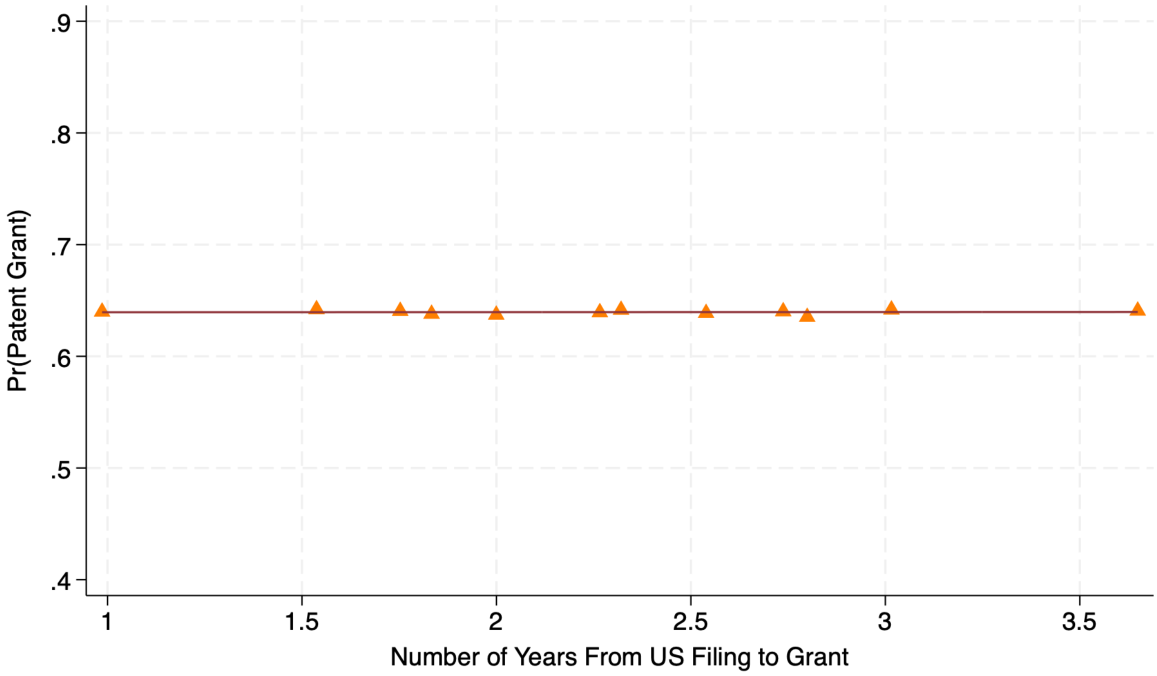

Note: This figure relates the lag between US patent grant dates and patent filing dates, residualized by priority authority, to the predicted patent grant rate. The predicted patent grant probability is based on observable characteristics that may be correlated with patent value and are fixed at the time of application, including the sector, year, the priority authority, the number of inventors, and the size of the patent family.

Figure A4: Additional Panel Event Study (Callaway and Sant'Anna, 2021)

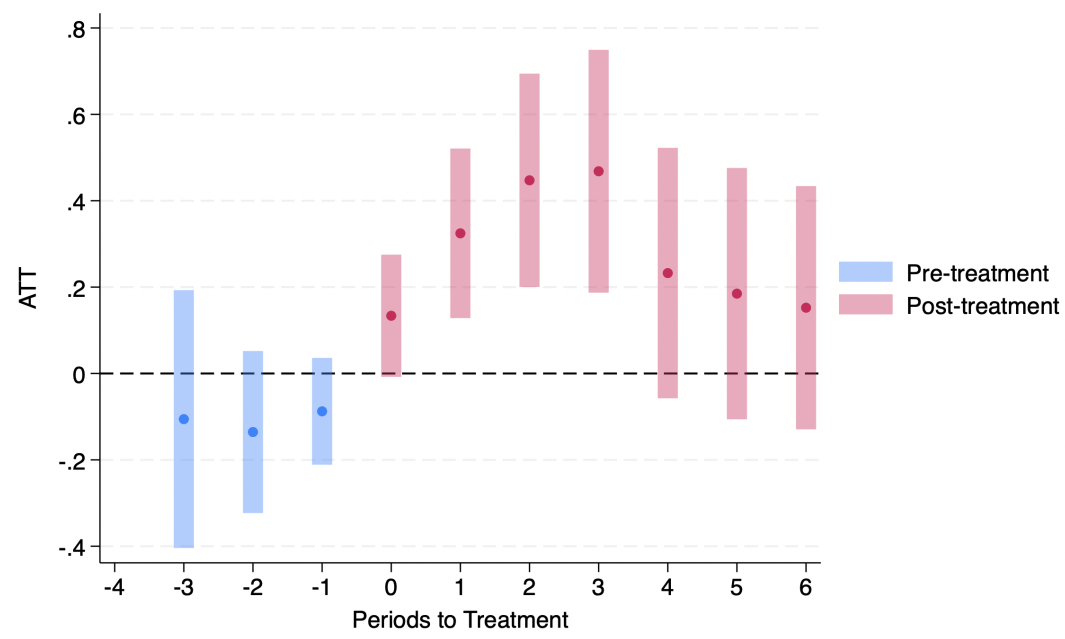

Note: This figure shows the event study plot following Callaway and Sant'Anna (2021) with 3 lead years and 6 lag years. The years before the US patent grant are marked blue, and the years after the US grant are marked red. The unit of observation is a patent family-year. 95% confidence interval is plotted.

Figure A5: Distance between Citing and Cited Sectors by Grant Status

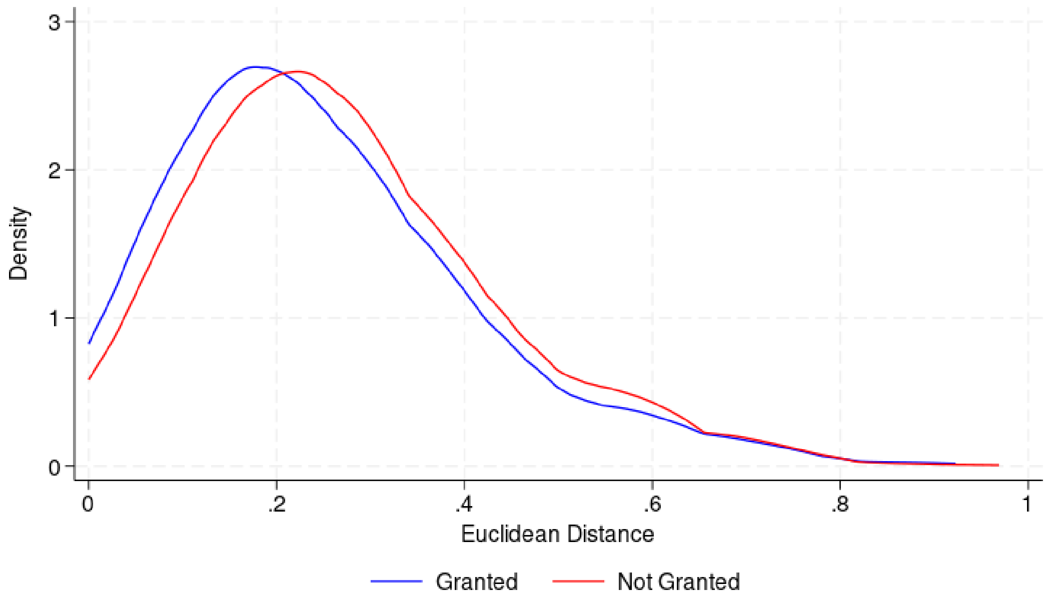

Note: This figure plots the distribution of the Euclidean distance between the citing and cited country-class' technological profiles, separately for patent families whose US application is granted (in blue) and those that are not granted (in red). For each citation, the distance is calculated at the country-pair-class-pair level. For citing country A with CPC class X and cited country B with CPC class Y, I first represent each country–class combination as an N-dimensional vector of CPC subclass shares (where N is the total number of 4-digit CPC subclasses under the 3-digit CPC classes X and Y), where the elements corresponding to subclasses under that class sum to one and all other elements are zero. To compare both sectors between the two countries, I combine information from both classes within each country by summing the two relevant vectors, yielding one aggregated vector for the citing side $(A, X+Y)$ and one for the cited side $(B, X+Y)$. The Euclidean distance between these two vectors is then computed as:

$$d = \sqrt{\sum_{i=1}^{N} \left(v_i^A - v_i^B\right)^2},$$

where v_i^A and v_i^B denote the shares for subclass i in the aggregated citing and cited vectors, respectively. A smaller distance indicates that the citing and cited countries are concentrated in similar subclasses, suggesting technological proximity, while a larger distance reflects specialization in different subclasses.

Figure A6: US Grant and Attention Diffusion – Subsample by Pre-filing Stock

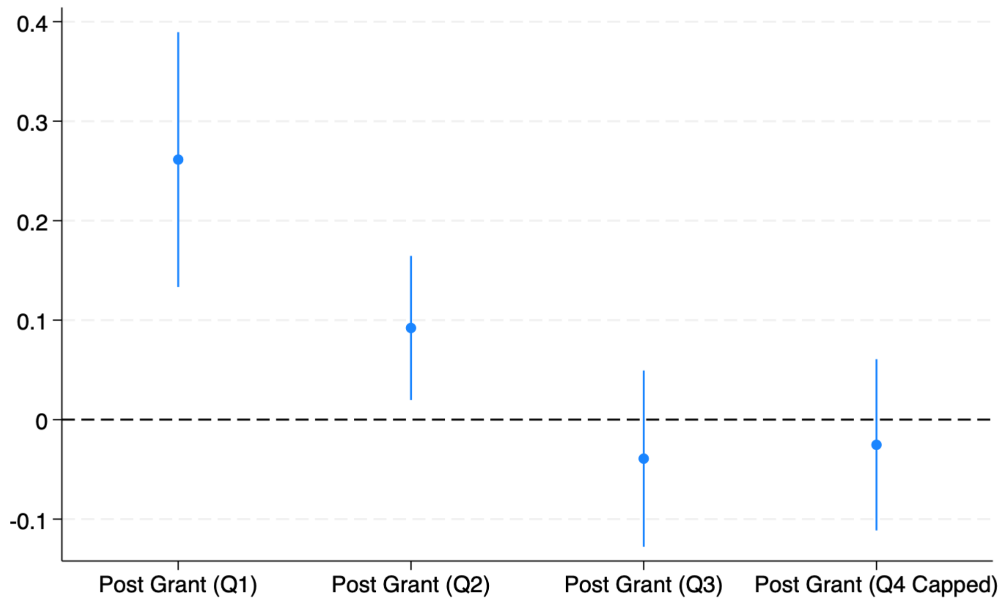

Note: This figure shows the impact of US patent grants on the attention the home country scientists get, separately for 4 groups of focal patents varying in how established the scientists are. The unit of observation is a patent family-year. The outcome is the average (weighted) number of citations received by other patents of the focal inventors. Patent family fixed effects and age cohort × year fixed effects are included, where age cohort is the quintiles of the average age of all non-focal patents at the time of the focal patent's US filing. Inventor prominence is highly skewed: the 25^{th}, 50^{th}, and 75^{th} percentiles correspond to approximately 4, 15, and 92 average scientist-weighted non-focal patents, respectively. The 90th percentile exceeds 1,300, likely reflecting large firms with very different underlying dynamics; thus, Q4 is capped at the 90th percentile. Standard errors are clustered by CPC subclass, and the 95% confidence interval is plotted.

Figure A7: Trappedness Measures at Country and Country-Sector Level

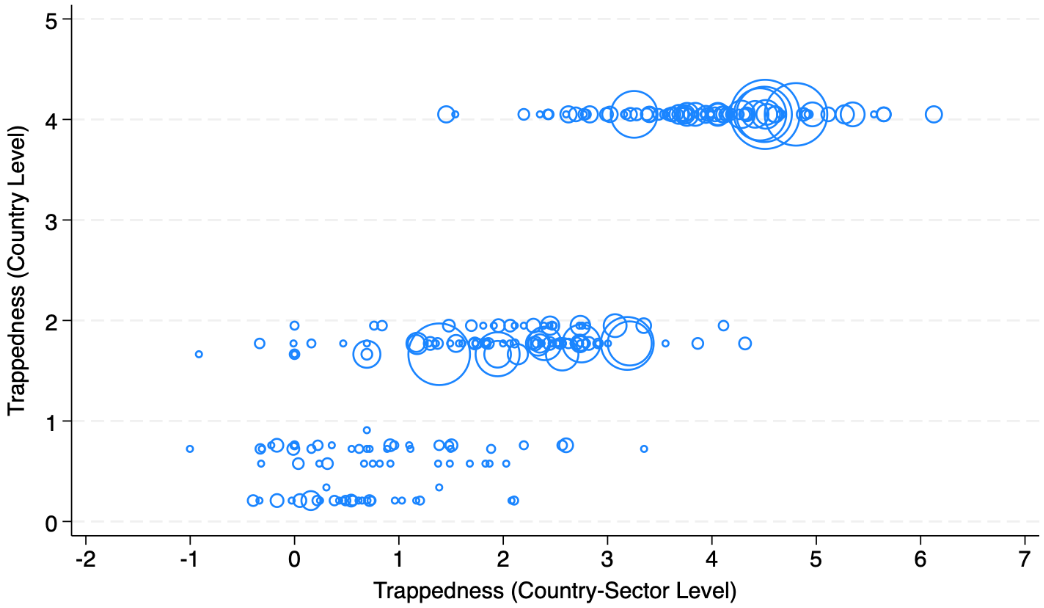

Note: This figure shows a scatter plot between $Trappedness_c$ and $Trappedness_{c,j}$. The size of the circles is proportional to the number of patent families in the specific country-sector in the sample.

Table A1: Distribution of CPC Class (3-digits)

CPC Class	Frequency	Percent	Short Description
H01	440	15.31%	Basic electric elements
G06	408	14.20%	Computing; Calculators
H04	258	8.98%	Electric communication
G01	213	7.41%	Measuring; Testing
A61	153	5.32%	Medical or veterinary science
A63	126	4.38%	Sports; Games
G02	70	2.44%	Optical devices
H03	57	1.98%	Basic electronic circuits
B60	55	1.91%	Vehicles in general
H02	54	1.88%	Generation of electricity
G03	42	1.46%	Photographic technology
G11	42	1.46%	Information storage
B01	39	1.36%	Physical or chemical processes
B41	38	1.32%	Printing
G09	38	1.32%	Education; Display devices
F16	37	1.29%	Engineering elements
A47	33	1.15%	Furniture; Domestic articles
H05	31	1.08%	Other electric techniques

Note: This table shows the frequency of top CPC classes in the panel sample.

Table A2: US Grant and Forward Citations – Ordinary Least Square

	Raw Outcomes		*Normalized Outcomes*	
DV:	*Exist Cite (0/1)*	*Num Cite*	*Exist Cite (0/1)*	*Num Cite*
	(1)	(2)	(3)	(4)
Post Grant	0.0682***	0.260***	0.122***	0.196***
	(0.011)	(0.056)	(0.031)	(0.048)
Patent Family FE	Yes	Yes	Yes	Yes
Cohort × Year FE	Yes	Yes	Yes	Yes
Mean (Pre-Grant)	0.31	0.72	0.80	0.76
Percent Δ	22.0%	36.1%	15.3%	25.8%
N	30235	30235	30235	30235

Note: This table estimates the impact of US patent grants on an indicator of the existence of any forward citation and the number of forward citations. The unit of observation is a patent family-year. Columns 1 and 2 present the raw outcomes, and columns 3 and 4 present the normalized outcomes, where each outcome is divided by the average value of that outcome within the corresponding US filing year and technology class. All estimates are OLS estimates. Patent family fixed effects and cohort × year fixed effects are included. Kleibergen-Paap weak identification F-statistic is reported. Standard errors are clustered by CPC subclass, and *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Table A3: US Grant and Forward Citations – Baseline Robustness Checks

	Alternative FE		Alternative DV		Alternative *Leniency*	
			Log(y+ϵ)	*Asinh(y)*	*Unadjusted*	*Examiner-Year*
	(1)	(2)	(3)	(4)	(5)	(6)
Post Grant	0.333***	1.023***	1.083***	0.146***	0.200***	0.205***
	(0.071)	(0.037)	(0.174)	(0.022)	(0.045)	(0.046)
Patent Family FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE		Yes				
Cohort × Year FE	Yes		Yes	Yes	Yes	Yes
Country × Tech Class × Year FE	Yes					
F-Stat	11630	23648	17795	17795	18849	26689
N	21589	30150	30141	30141	30235	30210

Note: This table estimates the impact of the US patent with alternative specifications. The unit of observation is a patent family-year. Columns 1 and 2 present alternative fixed effects. Columns 3 and 4 present alternative dependent variable transformations. Columns 5 and 6 present alternative instrumental variable adjustments. Kleibergen-Paap weak identification F-statistic is reported. Standard errors are clustered by CPC subclass, and *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Table A4: US Grant and Forward Citations - Alternative IV

	Raw Outcomes		*Normalized Outcomes*	
DV:	*Exist Cite (0/1)* (1)	*Num Cite* (2)	*Exist Cite (0/1)* (3)	*Num Cite* (4)
Post Grant	0.137***	0.294**	0.331***	0.405***
	(0.026)	(0.147)	(0.078)	(0.133)
Patent Family FE	Yes	Yes	Yes	Yes
Cohort × Year FE	Yes	Yes	Yes	Yes
Mean (Pre-Grant)	0.31	0.72	0.80	0.76
Percent Δ	44.19%	40.83%	41.38%	53.29%
F-Stat	1172	1172	1172	1172
Weak-IV Robust Wald χ^2	30.86	4.33	19.58	10.19
N	30235	30235	30235	30235

Note: This table estimates the impact of US patent grants on an indicator of the existence of any forward citation and the number of forward citations. The unit of observation is a patent family-year. Columns 1 and 2 present the raw outcomes, and columns 3 and 4 present the normalized outcomes, where each outcome is divided by the average value of that outcome within the corresponding US filing year and technology class. All estimates are IV estimates based on $\widehat{Post}_{it} \times Leniency_{t(i)k(i)}$. Patent family fixed effects and cohort × year fixed effects are included. Kleibergen-Paap weak identification F-statistic and Anderson–Rubin (AR) Wald χ^2-statistic are reported. Standard errors are clustered by CPC subclass, and *, **, and *** indicate significance at the 10%, 5%, and 1% levels.