Washington Center for Equitable Growth

1500 K Street NW, Suite 850
Washington, DC 20005

Working paper series

Patent Publication and Technology Spillovers

Deepak Hegde
Kyle Herkenhoff
Chenqi Zhu

December 2018

https://equitablegrowth.org/working-papers/patent-publication-and-technology-spillovers/

Patent Publication and Technology Spillovers

Deepak Hegde
New York University

Kyle Herkenhoff
University of Minnesota

Chenqi Zhu
New York University

Abstract

How does publication of new ideas through patents affect technology spillovers? To answer this, we develop a model which predicts that invention disclosure through patents (i) increases technology spillovers at the extensive and intensive margins (ii) increases overlap between distant but related patents and decreases overlap between similar patents (iii) lowers average inventive step, originality, and scope of new patents (iv) decreases patent abandonments and (v) increases patenting. We test these predictions by leveraging the enactment of the American Inventor's Protection Act, which advanced the publication of U.S. patents by about two years. The empirical findings support our model's predictions.

JEL classification: D23, G24, L26, O34

First Version: January 23, 2015; Current Version: September 19, 2018

Author emails: dhegde@stern.nyu.edu, kfh@umn.edu and czhu@stern.nyu.edu respectively.
We are grateful to Wesley Cohen, Alberto Galasso, Bronwyn Hall, Joachim Henkel, Zorina Khan, Hong Luo, Petra Moser, Alan Marco, Rosemarie Ziedonis, and audiences at Stanford, Northwestern, and NYU. This research was supported by National Bureau of Economics Research, Innovation Policy Grant (2014). Hegde acknowledges the support of the United States Patent and Trademark Office's Thomas Alva Edison Visiting Scholars program. All errors are our own.

1. Introduction

Patents divulge inventors' proprietary knowledge to the world. For example, Thomas Edison's light bulb patent (USPTO patent # 223,898—see Figure 1) revealed methods of creating incandescent filaments and paved the way for subsequent innovations in electric lighting by others. The U.S. patent office has published over ten million such inventions, as part of a grand bargain that exchanges invention disclosure for inventors' temporary monopoly rights.[1] The publication requirement seeks to inform the work of follow-on inventors and reduce duplicative research and development (R&D). Thus, the patent system's net effect on technological progress depends critically on the completeness and rate at which patents are disclosed to the public (Scotchmer and Green, 1990; Fromer, 2009).

[Figure 1 here]

In this study, we contribute to the literature on technology spillovers (e.g., Jaffe, 1986; Bloom, Schankerman, and Van Reenan, 2013) by providing causal evidence that patent publication generates substantial knowledge (technology) spillovers and reduces technology overlap. Thus, our study advances what little we know about the interaction of patent disclosure and innovation (Graham and Hegde, 2015; Williams, 2017; Furman et al., 2018).[2] The few existing papers on the topic provide contradictory conclusions: survey-based research suggests that inventors in some industries consider patents the most important source of relevant technical knowledge (Ouellette, 2017) and information on rivals' R&D activities (Cohen et al., 2002); however, some legal scholars argue that "patent disclosures play an insignificant role in promoting R&D spillovers" (Roin, 2005, p. 2027).

We measure the link between patent publication and technology spillovers by leveraging the enactment of the American Inventor's Protection Act of 1999 (P.L.106-113; henceforth, "AIPA") as a natural experiment. AIPA harmonized U.S. patent laws with those of the rest of the world by requiring applications filed on or after November 29, 2000 to be published 18 months from the filing date. Before AIPA, inventors were allowed to keep the existence of their U.S. patent applications secret until the patent was granted, which, in 2000, averaged about 3.5 years. Thus, AIPA reduced the period of secrecy for U.S. patent applications by about two years, on average, allowing us to measure the effects of patent disclosure on knowledge diffusion and follow-on patenting.

[1] The disclosure requirement for patentability in the U.S. states: "the [patent] specification shall contain a written description of the invention . . . in such full, clear, concise and exact terms as to enable any person skilled in the art . . . to make and use the same."(35 U.S.C. § 112).

[2] In contrast, over 100 published and working papers examine the effect of the monopoly rights awarded by patents (see Williams (2017) and Hall et al. (2014a) for surveys of the relevant literature).

To direct our empirical analysis, we develop a theoretical framework that models AIPA as provisioning "news shocks" to inventors about recent technologies. Under AIPA, the technological know-how embedded in patents enters the stock of public knowledge faster. In this new information environment (relative to the pre-AIPA environment), inventors see more related inventions in the patent system earlier, prior to filing their own patents. Better information about competing inventions reduces the likelihood of duplicate patent applications, which, in turn, reduces technological similarity between closely related inventions. A consequence of better information prior to patent filing is fewer duplicate applications and, thus, lower abandonment or rejection of applications. But inventors also draw more heavily on recently disclosed patents and take smaller inventive steps, thus raising the average technology similarity among related, but not substitute, applications. Lastly, overall patent activity and invention may increase, decrease or stay the same, depending on the net effect of the two countervailing forces of lower costs of invention (due to superior information and lower uncertainty) and free riding by follow-on inventors.

Our empirical analysis of the one million patents filed at the U.S. Patent and Trademark Office (USPTO) three years before and three years after AIPA yields the following findings: (i) mean delay to receive 1/3/5/7 follow-on patent citations (our measure of technology spillovers) from application date drops by 12 to 20 percent after AIPA; (ii) follow-on citations to post-AIPA patents within a ten-year window after disclosure date increase by four to 19 percent; (iii) technological overlap increases between distant but related patents and decreases between highly similar patents after AIPA; (iv) patent renewal rates, originality and patent scope decline; (v) post-AIPA patent applications are about six percent less likely to be abandoned or rejected; (vi) post-AIPA patents issue with nearly five percent fewer claims and more words per claim, together indicative of narrower patent scope; and (vi) no evidence of a decrease in overall patenting after AIPA.

The before-and-after analysis may yield unreliable estimates of AIPA's effects if the law also affected variables such as the quality of inventions that select into patenting or pre-grant publication. To address this concern, we adopt a Difference-in-Difference (DID) regression design (we refer to this as the "twin" study design). We build a sample of 316,563 patent applications filed at the USPTO between 1998 and 2003, each of which has an equivalent patent filed at the European Patent Office (EPO).[3] The U.S. patents form our treatment group, while their EP "twins" form our control group. The EPO required 18-month disclosure of applications well before AIPA's enactment—since its establishment in 1977—and we show through a series of tests that the European twins of U.S. patents were not plausibly affected by AIPA's enactment.

[3] We refer to patents filed at the USPTO as U.S. patents and patents filed at the EPO as European patents (EP patents) throughout the paper, regardless of the applicants' country of origin.

This "twin" study design allows us to control for unobserved characteristics of each patent family (comprising the U.S. patent and its EPO twin) and to account for quality-based selection into patenting or early disclosure using family-fixed effects. Thus, we isolate AIPA's effects by analyzing differences in the diffusion patterns of identical twin patents, one filed in the U.S. and the other in Europe, before and after AIPA. This design identifies the effect of the USPTO's 18-month patent disclosure alone since the EPO disclosed the European twins of U.S. patents at 18 months both before and after AIPA (twin applications were published nearly simultaneously after AIPA—18 months from the filing date of the earliest twin, called the "priority date"). By identifying the effect of 18-month disclosure in the U.S. for patents disclosed around the same time in the EPO, this research design provides conservative estimates of the effects of patent disclosure.

The "twin" study design confirms the estimates obtained with the before-and-after analysis. We also find that the increase in citations to U.S. patents are due largely to an uptick in citations to U.S. patents after AIPA, rather than to a decrease in citations to their European twins. The estimates are driven by the fact that applications disclosed by one office (EPO) are not seamlessly disseminated to inventors filing patents in another office (the USPTO) until published by the latter. Indeed, previous studies have noted that inventors and patent examiners are more likely to search for relevant patents originating in countries in which they are filing, or examining, patent applications (Harhoff et al., 2009).

In response to concerns that pre-grant disclosure harms small inventors (as expressed by 26 Nobel laureates in a letter to the Senate: Modigliani et al., 1999), AIPA provided U.S. applicants with a loophole: they could opt out of 18-month disclosure under the condition that they forgo foreign protection. Restricting observations to the post-AIPA period and comparing patents that used the opt-out provision (about eight percent of post-AIPA patents) with those that did not provides evidence consistent with the before-and-after analysis: post-AIPA patents disclosed at 18 months experienced faster diffusion and reduced technological overlap compared to post-AIPA patents published at grant.

The founders of the modern patent system envisioned invention disclosure as the offsetting mechanism against the monopoly rights created by patents. Our study provides large-sample causal evidence that invention disclosures through patents indeed inform follow-on inventors and shape their patenting decisions. Thus, welfare analyses of patent systems that consider only their incentive effect on R&D or their blocking effect on follow-on invention are incomplete.

AIPA affected the timing of additions to, arguably, the world's single largest repository of technical knowledge, and it is considered the most important U.S. patent law enacted in the 20th century. The U.S. Congress's motivation for AIPA's enactment was as follows: *"US researchers and investors are denied the opportunity to learn what their foreign competitors are working on until a US patent issues. This* *causes duplicative research and wasted developmental expenditures, putting U.S. inventors at a serious disadvantage vis-a-vis their foreign counterparts and competitors*."[4] Our evaluation of AIPA confirms that the policy indeed increased technology spillovers, reduced duplicative patenting and reduced patenting costs, without having a negative impact on overall patenting or innovation. Thus, our policy evaluation provides evidence against recent proposed legislation (e.g., H.R. 5980) that seeks to limit pre-grant publication on the assumption that disclosure imposes a net cost on innovation.

The rest of the paper is organized as follows. Section 2 provides institutional background and reviews the literature on patent disclosure. Section 3 develops a theoretical framework that motivates our empirical investigation. Section 4 describes the sample and the results of the before-and-after AIPA analysis. Section 5 reports the findings of the results using the "twin" study design. Section 6 concludes with a discussion of the implications of our findings.

2. Background

2.1 Institutional Details

Prior to AIPA, the disclosure of a U.S. patent application, containing detailed technical descriptions and drawings of the invention, occurred when the patent was issued. Applications that were either rejected by the patent office or withdrawn by their applicants were never published. AIPA required patent applications filed at the USPTO on and after November 29, 2000 to be published by the government 18 months after the application date.[5] Since most foreign countries' patent systems already required 18-month publication of patent applications, AIPA's enactment harmonized the U.S. patent system's disclosure policies with international norms.

However, in response to concerns that pre-grant disclosure harms small inventors (see Modigliani et al. (1999), the Act provided U.S. applicants with a loophole: they could opt out of 18-month disclosure under the condition that they forgo foreign protection. Thus, applicants that opted out of foreign protection post-AIPA (as was the case for applicants that did not pursue foreign protection before the Act) could keep both the presence of their patent application and the application's content secret until patent grant. For patents that take a long time to issue, the additional period of pre-grant secrecy beyond 18 months could be substantial; for example, among U.S. patent applications filed in 2005, 50% took more than 38 months,

[4] See https://www.congress.gov/106/crpt/hrpt287/CRPT-106hrpt287-pt1.pdf

[5] Applications can be published before 18 months if applicants submit an early publication request to the USPTO.

25% more than 51 months, and 10% more than 61 months to issue. Patent applications in these groups could gain at least an additional 20 months, 33 months, and 43 months of secrecy, respectively, by opting out of foreign protection.

2.2 Literature Review

Since the work of Levin et al. (1987) and Cohen et al. (2002), a large literature on the impact of property rights on appropriability has emerged (see Hall et al. (2014a) or Williams (2017) and citations therein). Cohen et al. (2002) find that companies involved in R&D use lead-time, secrecy, and other informal appropriation mechanisms more often than they use patents. More recently, Hall et al. (2014b) show that, from 1998 to 2006, only about 4% of firms in the UK actually patented, and most of them rate lead-time and secrecy as better mechanisms to appropriate return from their innovations. Follow-up work by Hall and Sena (2014) find that firms that use formal IP protection, such as patenting, gain more in terms of productivity relative to other firms, *ceteris paribus*. However, recent work by Graham and Hegde (2015) finds that, of the companies that do patent, only about 7.5% of patent applicants opt out of pre-grant publication. In other related work, Johnson and Popp (2003) find that patents that remain in the patent office longer are cited more often, and, therefore, pre-grant publication diffuses these ideas faster. In a reduced-form simulation, Johnson and Popp (2003) find small positive short-run effects of pre-grant publication through this mechanism, but no long-run impact.

We also build on the work by Bloom, Shankerman, and Van Reenan (2013) who use variation in federal and state R&D tax credits to identify technology spillovers. We empirically advance this literature by measuring the effect of patent publication on technology spillovers. In terms of theory, our work is closely related to Bloom et al (2013); namely, both frameworks flexibly model substitutability and complementarity between own and rival technology, and both model free-riding. However, our frameworks differs along three key dimensions. We incorporate the negative effects of technology disclosure through the outside option of the firm as opposed through product market competition, we explicitly incorporate the stages of patenting, and we enrich the inventors' information space in order to derive testable implications of pre-grant patent publication.

Our work also complements concurrent work by Furman et al. (2018), who show how the introduction of patent libraries in the 1980s increased local patenting, job creation, and citations. Other discussions of AIPA, such as that of Okada and Nagaoka (2015) and concurrent work by Stefano and Simeth (2018), are descriptive in nature and compare patent citation outcomes before and after AIPA. Relative to this literature, the "twin" study methodology allows us to measure the causal effect of AIPA on technology similarity, patent scope, the timing and composition of citations, as well as overall patenting. Furthermore, we put forth the theory of AIPA as a news shock, which allows us to interpret and rationalize our findings.

In terms of theory, some of the earliest work on information disclosure and innovation is by Horstmann et al. (1985), Bhattacharya et al. (1992), Anton and Yao (1994, 2004), and Bhattacharya and Guriev (2006). More-recent work by Aoki and Spiegel (2009), Hopenhayn and Squintani (2011) and Bobtche et al. (2013), among others, endogenizes the timing of patent races. Of particular note, Hopenhayn and Squintani (2011) find that when R&D output is secret, firms take longer to patent inventions, and, thus, invention disclosure slows with R&D secrecy. We contribute to this literature by integrating the patent process into a heterogeneous firm framework with endogenous investment, drawing on elements from Atkeson and Burstein (2007, 2011). We model pre-grant publication as providing advance information (i.e. a "news shock") that can alter the patenting, investment and abandonment decisions of follow-on inventors. At a high level, descriptive data on citations to prior art (disclosed patent applications or issued patents) support this view that pre-grant publications propagate news about recent inventions—both applicants and patent examiners respond to pre-grant patent disclosures (see Appendix Figure A1) with 45% (48%) of applicant (examiner) citations to prior art being to pre-grant publications.[6] The fraction increases to nearly 60 percent in dynamic technological areas such as "Computers and Communications."

3. Theoretical Framework

Consider a finite-horizon economy populated by a continuum of potential entrants. Let $t = 0, 1$ denote time. At date 0, firms draw their idea quality from a distribution $F(z) : [0, \infty) \to [0, 1]$. After drawing an idea, firms must decide between entering a competitive market whose payoff at $t = 1$ is independent of the quality of their idea, V^c, or patenting the idea. At t = 1, the payoff to successfully patenting is given by $V^p(\cdot)$. Let Z_0 denote the stock of disclosed knowledge at $t = 0$. The profitability of a firm is determined by three factors: (i) idea quality; (ii) firm investment in the patent, $\Delta \in [0, \infty)$, which we will interpret as the number of claims and the patent scope; and (iii) the complementarity or substitutability of the idea and investments with the existing stock of disclosed knowledge. We assume that the profitability of a patent based on idea z, with scope Δ, and the existing stock of disclosed knowledge $Z_0 > 0$ is given by the following CES functional form:

[6] Note that later-filed applications more frequently cite pre-grant publications. This is because, for example, for patents filed in 2001, only a few cite pre-grant publications, as there are not many patent applications published from which applicants or examiners back then could possibly cite. As the stock of publications accumulates, both applicants and patent examiners cite them much more frequently.

$$V^p(z,\Delta,Z_0) = \left((z+\Delta)^\rho + Z_0^\rho\right)^{\frac{1}{\rho}}$$

The parameter ρ ($0 < \rho < \infty$) determines the substitutability between private patent investments and disclosed public knowledge; $\rho < 1$ corresponds to complements; $\rho = 1$ corresponds to perfect substitutes; and $\rho > 1$ corresponds to substitutes. Based on our empirical evidence, which follows, we will focus on the case in which $\rho > 1$, and, therefore, public knowledge is a substitute for private patent investments.

During the pre-AIPA regime, patenting and investment decisions had to be made without knowing whether a duplicate patent application was pending examination at the patent office. We assume that there is a duplicate patent in the system with probability $1-q$ (we will refer to duplicates as 'close' technologies). If a duplicate patent is present, then the firm abandons the patent and receives the competitive value V^c. The cost of expanding the scope of the patent by Δ is $c(\Delta)$ (we assume that $c(\cdot)$ is increasing and convex, $c(0) = 0$). The value to a firm at $t = 0$ is given by

$$V_0(Z_0) = \int \max\{\max_{\Delta} qV^p(z,\Delta,Z_0) + (1-q)V^c - c(\Delta), V^c\} dF(z)$$

Let $\Delta(z, Z_0)$ denote the optimal investment in patent scope, and define z_p as the minimum idea quality that enters the patent system—i.e., $V^c = V^p(z_p, \Delta(z_p, Z_0), Z_0)$. Under the simple regularity conditions presented in the Appendix, we obtain the following results: (R1) the idea threshold z_p declines as the stock of public knowledge Z_0 increases; and (R2) under the assumptions $c(\Delta) = \frac{\Delta^\gamma}{\gamma}$, $\gamma > \rho$ and $\rho > 1$, patent scope, Δ, is decreasing in the stock of public knowledge Z_0. The second result (R2) means that if public knowledge and private patent investment are substitutes, then the greater the existing stock of disclosed knowledge, the less inventors will invest in patent scope.

Under the post-AIPA regime, duplicate patents are known prior to investing in the patent (i.e., the random event corresponding to duplication, $1-q$, is known in advance); moreover, the current cohort of patents is disclosed and enters public knowledge. Thus, the available stock of public knowledge is given by $Z_1 = Z_0 + \int_{z_p}^{\infty}(z + \Delta(z))dF(z) > Z_0$. The post-AIPA value of a firm at $t = 0$ is given by

$$V_0^{AIPA}(Z_0) = \int \left[q \max\{\max_{\Delta} V^p(z,\Delta,Z_1) - c(\Delta), V^c\} + (1-q)V^c\right] dF(z)$$

Since firms optimize, $Z_1 > Z_0$, and $q < 1$, the value of patenting, conditional on any idea quality z, must increase (consistent with result (R1)). Therefore, we obtain result (R3): the patenting threshold z_p declines with the introduction of AIPA; more patents are filed, and the inventive step size is smaller. We interpret smaller inventive steps as being synonymous with non-duplicative ('distant') technologies becoming more similar, on average.

Lastly, we show in the Appendix that under the hypotheses of (R2) and the assumption of low patent duplication rates—i.e., $q \approx 1$—investment in the scope of patents will decline. In other words, if public knowledge and private patent investments are substitutes, and duplication is a low-probability event, we obtain the final result (R4): conditional on idea quality z, investment in the scope of patents, Δ, declines post-AIPA. Note that, given that patentees had the choice of opting out of pre-grant disclosure, and few patents remained secret after AIPA, we do not model a disclosure penalty.

In summary, modeling AIPA as a news shock yields the following testable implications:

1. The closest technologies increase distance (duplication declines), and abandonments decrease.
2. The furthest technologies decrease distance—i.e., average inventive step size (proportional to z_p) declines. Thus, patentees make smaller inventive steps.
3. Patent scope declines under the assumptions (i) that public knowledge and private patent investment are substitutes; and (ii) that pre-AIPA duplication rates are low.
4. Technology enters the public domain and production sooner.
5. Patent filings increase.
6. If we relax the assumption that the outside option V^c is constant and assume that $\frac{\partial V^c(Z_1)}{\partial Z_1} > 0$—i.e., there is free-ridership—then patent filings may decrease.

4. Before-and-After Analysis

4.1 Sample and Data

We start with the universe of utility patent applications filed at the USPTO during 1998-2003. The applications data are drawn from the agency's Patent Application Information Retrieval (PAIR) files and include both unsuccessful (rejected/abandoned) and successful applications. We track the citations received by these applications until the end of 2016. We supplement these data with the European Patent Office's PATSTAT (2017 Spring version) for information on (i) International Patent Classification (IPC) assignments; (ii) the worldwide patent family table to identify patents with foreign and EPO parallel applications; and (iii) standardized patent applicant names. After excluding applications that do not have information on important variables, have errors, or are reissued patents, our sample has 1,536,346 applications filed at the USPTO. Of these, 675,917 applications (75.4% of which were issued patents) were filed before AIPA's enactment, and 860,429 applications (69.5% of which were issued patents) were filed after AIPA.

Our proxy for the speed of knowledge spillovers (or diffusion) is the amount of time it takes for a patent to receive a certain number of forward citations, and our proxy for the extent of spillovers is the number of patent forward citations (citations received by the focal patent from future patents, after removing self-citations). We also compare technology similarity, inventive step, patent originality, abandonment rates, and claims before and after AIPA to shed light on the effects of pre-grant disclosure on follow-on patenting. In most analyses, we focus on granted patents rather than on applications, except when we compare abandonment rates. [7] We discuss the construction of our main variables in detail below, and Table 1 summarizes all variables used in our analyses.

[Table 1 here]

Citation lag

The first variable of interest, citation lag, proxies for the speed of knowledge diffusion (*intensive margin*). It is measured as the average difference between the application dates of the focal patent and its first 1/3/5/7 forward citations. We use the application date because it is closer to the invention date than to the disclosure or grant date, and we want to measure how rapidly knowledge diffuses from the inception of one invention to the creation of another. Pre-AIPA patents are filed earlier and have a longer time to accumulate citations, which biases the citation lag upward for pre-AIPA relative to post-AIPA patents. Hence, to ensure comparability, we include only citations within a ten-year window from application date in the computation of citation lag. We further exclude self-citations, as they reflect internal cumulative developments rather than knowledge spillover across inventors. Both focal and citing patents are filed at the USPTO.

Forward citations

The number of forward citations, after excluding self citations, measures the extent of knowledge spillovers associated with the focal patent (*extensive margin*). We count citations received by each focal patent in the 3/5/7/10 years after its disclosure date, which is the publication date for patents with pre-grant publications and the grant date for those without. The citation clock starts at the publication date for post-AIPA patents with pre-grant publications because they become visible and, thus, “citable” by follow-on

[7] We exclude abandoned and pending patents for the analyses on citation lags, citation counts, and technology similarity to keep the pre-AIPA and post-AIPA samples relatively comparable. Unsuccessful applications that are published 18 months after application receive 4.3 forward citations in the ten years after application, as shown in Appendix Table A1. Arguably, any knowledge diffusion stemming from abandoned applications is an additional diffusion effect that would not have happened in the pre-AIPA world.

inventors, but pre-AIPA patents (as well as post-AIPA patents that opt out of pre-grant publication) become visible only at grant.[8]

Technology similarity

Technology similarity, which we also refer to as technology overlap, is based on the cosine distance between the focal patent and next-generation patents. Next-generation patents include all patents in the same technology class (IPC 4-digit code) [9] as the focal patent and filed in the 19-36 month window after the focal patent application.[10] We start the window at the 19th month to ensure that the next-generation patents have had the opportunity to use the information in the 18-month disclosures of post-AIPA patents. We stop the window at the 36th month since this is the average grant (and, thus, disclosure) lag for pre-AIPA patents. Presumably, patents filed within this 19-36 month window are the ones most likely to benefit from the technological knowledge revealed by the pre-grant publications after AIPA, although we ensure the robustness of our findings for different windows. To isolate the informational impact of patent disclosure, we exclude next-generation patents that have the same assignee as the focal patent.

Next, we calculate the cosine distance between each focal patent and its next-generation patent as

$$Sim_{ij} = \frac{N_i N_j{}'}{(N_i N_i{}')^{1/2}(N_j N_j{}')^{1/2}},$$

where i represents the focal patent and j represents any patent in the next generation. $N_i = (N_{i1}, N_{i2}, \ldots N_{i7154})$ is a vector with each element representing patent i's fraction of IPC assignments that belong to each of the 7,154 IPC main groups (IPC 7-digit code). The cosine distance is widely used to measure the proximity of two vectors, each representing the location in a pre-defined space (e.g., Jaffe, 1986; Bloom et al., 2013; Kuhn et al., 2017). Thus, for each focal patent i, we have a vector of similarities

[8] Appendix A2 compares cumulative forward citations with their citations clock starting from application date, disclosure date, and grant date, respectively.

[9] Each patent typically receives multiple IPC codes. The IPC is a hierarchical technology classification system used in many patent offices. Each IPC code takes the form of "A01B 1/00." The first four digits ("A01B") are called the subclass, followed by a one-to-three-digit "group" number ("A01B 1" is a group) and a two-to-five-digit "subgroup" number split by "/" ("A01B 1/00" is a subgroup). During our sample period, there are 641 IPC subclasses (IPC 4-digit code), 7,154 IPC main groups (IPC 7-digit code), and 62,654 IPC subgroups (IPC 12-digit code). For patents with multiple IPC codes, we choose the one listed first as the main IPC code for U.S. patents. According to PATSTAT, the USPTO lists the main IPC code first, but other authorities, such as the EPO, list the IPC codes alphabetically. When we compute technology similarity for EP patents, we choose the subclasses (4-digit codes) with the highest frequency as the main subclass.

[10] We exclude next-generation patents in IPC 4-digit codes different from the focal patent to reduce computational burden and restrict attention to the next-generation patents most likely to be related to the focal patents. We also require applications in the next generation to be granted to maintain comparability for patents before and after AIPA, as we do not have information on the IPC assignments for undisclosed abandoned applications.

$\{Sim_{i1}, ..., Sim_{iJ}\}$ between patent *i* and patents in its next generation. We use this vector to compute the distribution of similarity between patent *i* and its next generation, and we record the similarity value at every 5th percentile of this distribution. Figure A3 in the Supplementary Appendix illustrates this procedure with an example. Higher percentiles in the similarity distribution (between a focal patent and its next-generation patents) correspond to technologically distant patents, and lower percentiles in the distribution correspond to close patents. By construction, our technology similarity measure ranges from 0 to 1, with larger values (and higher percentiles) indicating a higher degree of technological overlap. This measurement strategy allows us to compare the differentiation of patents before and after AIPA as a function of the crowdedness of technological areas.[11]

4.2 Before and After Analysis

In this section, we first graphically compare knowledge diffusion and patent characteristics before and after AIPA. We then pursue a parametric regression approach that removes the pre-trends apparent in the data. The regression model is specified as follows:

$$Outcome_{it} = \alpha_0 + \sum_{\tau \in Post} \beta_t * I\{Month_t = \tau\} + f(Month_t) + \alpha_2 EarlyGrant_i + \alpha_3 OptOut_i + TechFE + \epsilon_{it}, \quad (1)$$

where *i* indicates the application filed in calendar month *t*.[12] The dependent variables are the citation lags, citation counts, technology similarity, renewal rates, patent originality, abandonment rates, and claims. Following Gross et al. (2016), we include a function of a continuous variable of application month (*f(Month)*), as well as a set of dummy variables indicating each month in the post-AIPA regime (37 dummies in total, each indicating a month from December 2000 to December 2003). *f(Month)* controls for the pre-trend, while the coefficient on the dummies (β_t) captures the lagged effect of AIPA at different horizons relative to the pre-trend. In the baseline regressions, *f(Month)* is a linear function of the calendar month. In robustness checks, we add second- or third-order polynomials of *Month* to control for potential non-linear pre-trends. To quantify AIPA's immediate and longer-term effects, we take the average of the estimated β_t over different horizons and compute the corresponding standard errors using the delta method.

[11] Table A2 of the Supplementary Appendix reports the results of a validation exercise of our similarity measure.

[12] Since AIPA became effective on November 29, 2000, we allocate 2,536 applications filed on November 29-30, 2000 to December 2000 so that each month is classified as either pre- or post-AIPA throughout the paper.

We control for whether patents are granted before 18 months (*EarlyGrant*) since such patents are *de facto* untreated by AIPA. We also add a dummy variable indicating patents that opt out of pre-grant publication (*OptOut)*. USPC technology fixed effects are added to control for time-invariant industry characteristics.

4.3 Empirical Results

Table 2 provides summary statistics for the sample of U.S. applications filed during 1998-2003. As discussed in Section 2, after AIPA's effective date, patents with foreign parallel applications were published by the USPTO at 18 months from the first filing date, while those without foreign applications could opt out of the pre-grant publication requirement. 8.6% (8.1%) of the post-AIPA successful applications opted out of pre-grant publication. 18.7% of patent applications filed in the pre-AIPA period and 15.6% of applications filed in the post-AIPA regime were granted before 18 months.[13] On average, pre-AIPA patents receive four citations within a three-year window of disclosure (grant date), while post-AIPA patents receive 3.7 in the same window (the disclosure date is the earlier of the pre-grant publication date and the grant date). As we show later, this lower number of citations to post-AIPA patents is due to a citations truncation bias, which increased for later patents. The time it takes for a patent to receive one citation dropped from 38 months to 35 months after AIPA. Patents that were the least similar were closer after AIPA (i.e., the 5th percentile of the technology similarity distribution *increased*), whereas the closest patents were more differentiated (i.e., the 95th percentile of the technology similarity distribution *decreased*). Lastly, the average number of claims increased from 17.5 to 18.8 after AIPA (our regression analyses reverse this finding), and the average number of words per independent claim appears to have remained unchanged (160.0 pre-AIPA and 160.5 post-AIPA).

[Table 2 here]

We next directly examine and quantify changes in the speed of knowledge diffusion. Figure 2 plots monthly average citation lags for all U.S. patents filed from 1998 to 2003. The average time to receive the first (first seven) citation(s) ranged from 30 to 33 (42 to 45) months before AIPA, and fell to 27 to 29 (39 to 41) after. The citation lag increased modestly when we compute the time lag from the first citation to the first ten citations, presumably because patents with at least ten forward citations are more valuable, and their knowledge diffuses faster than that of patents with fewer forward citations. The takeaway from Figure 2 is that citation lags dropped sharply and substantially after AIPA, suggesting that patent disclosure speeds up knowledge diffusion.

[Figure 2 here]

[13] The lower percentage of within-18-month grants is consistent with the growing backlog of unexamined applications and increased application-grant lag during our study period at the USPTO (Hegde, 2012).

We next estimate the before and after regressions described in Section 4.2 to compare citation lags around the effective date of AIPA. The dependent variable is the natural logarithm of one plus the citation lag. To be consistent with the analyses on forward citations and technology similarity discussed later, here, we report the regression results with a linear pre-trend in Panel A of Table 3. Our variable of interest is the series of dummies indicating each month in the post-AIPA period. We report the average coefficient on these dummies in the first row (*Pre-Post-Dif*). The averages suggest that the delay to receive one to seven forward citations decreased by 11.6%-20.1%, relative to the predicted delay based on the pre-trend. The coefficient on *OptOut* is significantly positive, ranging from 17.5%-27.6%, which more than offsets the average AIPA effect (*Pre-Post-Dif*). With the caveat that quality-driven selection into opt-outs may be driving the finding, this lends additional support to the idea that pre-grant patent disclosure, rather than macro trends, hastened knowledge diffusion.[14]

[Table 3 here]

In addition to the speed of knowledge diffusion (*intensive margin*), we are also interested in the overall scope of knowledge spillovers (*extensive margin*). We measure knowledge spillovers by the number of citations received in 3/5/7/10 years after the disclosure date (recall that the disclosure date is the publication date for patents with pre-grant publications and the grant date for those without). Figure 3 plots the monthly average of forward citations for all granted applications filed at the USPTO from 1998 to 2003. There are clear downward trends for both pre- and post- patents, which are caused mainly by data truncation (a greater fraction of potential citing patents had not yet been granted as one approached the end of our observation period—December 31, 2016—increasing undercounting with time).

Figure 3 shows that post-AIPA patents received more forward citations than pre-AIPA patents, although the increase in magnitude appears negligible when citations are counted three years after disclosure.[15] Note that the way that we count forward citations favors pre-AIPA patents since they are granted, and patents generally experience an increase in citations upon grant. As we increase the duration of the forward citation windows, this effect of patent grant is attenuated, and the number of forward citations received by post-AIPA patents clearly exceeds that of pre-AIPA patents five, seven or ten years after disclosure. Relatedly, the larger jump in ten-year than in five- or 7-year forward citations also reflects the cumulative

[14] As Figure 2 shows, a linear pre-trend does not fit the pre-AIPA patents very well; hence, we add the second- or third-order polynomial of the calendar month and report the results in Appendix Table A3. The coefficients on both *Month^2* and *Month^3* are statistically significant, suggesting that the inclusion of non-linear pre-trends is warranted. Nevertheless, we still observe a statistically significant drop in citation lags in the post-AIPA period, although the economic magnitude varies with the degree of polynomials included. The estimated drop in citation delays is 30.0%-41.4% (11.0%-16.3%) when we model the pre-trend using a quadratic (cubic) function.

[15] If we start to accumulate citations after the application date or the grant date, we observe that post-AIPA patents receive more citations than pre-AIPA patents across all four citation windows.

effect of pre-grant patent disclosure on follow-on invention. The jump occurs immediately after AIPA's enactment, making other factors, such as concurrent economic conditions, less likely to be behind this increase. To examine the effect of confounding factors that fall within our study period, such as the dot.com bubble and burst, we exclude computer and software patents and find the pattern almost unchanged.

[Figure 3 here]

The corresponding regression results are reported in Panel B of Table 3. The dependent variable is the natural logarithm of one plus the number of forward citations, counted in windows ranging from three to ten years after disclosure. The average pre- and post-AIPA difference (*Pre-Post-Dif*) is significantly positive across all four citation measures, and post-AIPA patents receive an estimated 3.8-19% more forward citations, on average, than the predicted citations based on the pre-trend. The pre-trend is captured by the coefficient on *Month*, which is significantly negative, consistent with the downward trend evident in Figure 2.[16] To graphically illustrate AIPA's estimated impact, in Figure 2, we add a line that fits pre-trends and extrapolate it to the post-AIPA period to indicate the expected citations for post-AIPA patents (if AIPA had not been enacted) *ceteris paribus*. Almost all monthly average citation counts in the post-AIPA period stay above this line, suggesting that AIPA increased knowledge spillovers. The estimated coefficient on *OptOut* is significantly negative and of similar magnitude as the average AIPA effect (*Pre-Post-Dif*), suggesting that pre-AIPA patents and opt-out patents experienced similar knowledge diffusion patterns.

Our theory predicts that, as a result of faster and more-complete knowledge diffusion after AIPA, technology similarity between post-AIPA patents disclosed at 18 months and their follow-on patents would have increased. At the same time, early patent disclosure could also force out close rivals and reduce duplicative patent applications, resulting in lower technology similarity between the disclosed patents and their closest rivals. To empirically test these predicted effects, we compute the 5th, 10th, 15th, 25th, 50th, 75th, 85th, 90th, and 95th percentiles of technological similarity between each focal patent in our sample and its next-generation patents. As described in Section 4.1, similarity at lower percentiles proxies for similarity between technologically remote patents, whereas similarity at higher percentiles proxies for similarity between technologically close patents.

Figure 4 plots the monthly average similarity for all U.S. patents filed from 1998 to 2003. We observe a large increase in similarity between technologically remote patents (5th-15th percentiles of similarity) and

[16] In robustness checks, we control for non-linear pre-trends by including second- or third-order polynomials and find that our inferences are not changed appreciably; hence, for simplicity, we do not report these results.

technologically moderate patents (25^{th}- 85^{th} percentiles of similarity) and a sharp drop in similarity focusing on the top 5% or 10% closest patents (90^{th} -95^{th} percentiles of similarity). These findings are confirmed by the corresponding regression estimates reported in Panel C of Table 3.

[Figure 4 here]

One may be concerned that post-AIPA U.S. patents received a higher number of citations not due to higher knowledge diffusion, but because follow-on patentees shifted their backward citations from foreign patents to their equivalent U.S. pre-grant publications. To address this concern, we redraw Figures 1 through 3 by plotting the monthly averages for patents with and without foreign parallel applications separately in Appendix Figure A4. We see similar patterns of increased forward citations, shortened citation delays, and reduced technology similarity in the post-AIPA regime for U.S. patents with foreign parallel applications, as well as for U.S. patents that did not have foreign parallel applications. The increase in citations was even more pronounced for U.S. patents that did not have foreign parallel applications, alleviating the concern that our results are driven by a migration of citations from foreign patents to their U.S. equivalents, rather than to a true increase in knowledge spillovers.

Our model predicts that the patenting threshold would decrease due to reduced uncertainty in the patenting process and the positive externality from recent patents becoming public knowledge faster. Thus, patentees would make smaller inventive steps. While we cannot directly observe the patenting threshold, we investigate whether inventors pursued patenting for less valuable or less original inventions after AIPA, measuring value through 3.5-year patent renewal rates and originality through the originality index developed by Trajtenberg et al. (1997) Patents that renew after grant are considered more valuable than those that do not; patents that refer to a broader class of prior art are considered more original.

Panel A of Figure 5 plots the monthly average 3.5-year renewal rate for patents filed from 1998 and 2003 and granted by mid-2014.[17] We find that the renewal rate went up before AIPA but gradually decreased in the post-AIPA period. The regression analysis with technology-class fixed effects, reported in Column 1 of Table 4, confirms this graphic evidence suggesting that, indeed, inventors pursued patenting for less-valuable ideas post AIPA. Panels B and C of Figure 5 plot the monthly average originality. Originality is measured as the Herfindahl dispersion of backward citations in the focal patent across different technology classes of granted patents in Panel A (all patent applications, regardless of the grant status in Panel B). Similar to renewal rates, we find that patent originality rose steadily before AIPA and gradually fell

[17]The renewal data were downloaded from the USPTO (https://bulkdata.uspto.gov/data/patent/maintenancefee/) on April 23, 2018. As the sample patents were granted by mid-2014, four years before the record date of renewals, there is no truncation errors in the computation of renewal rates.

afterwards. The regression estimates, reported in Columns 2-3 of Table 4, confirm that inventors patent less-original ideas in the post-AIPA period.

[Figure 5 here]

[Table 4 here]

Abandonments and claims

Our model predicts that as more information on pending applications becomes available, inventors make a more informed decision on whether or not to patent their inventions, leading to fewer unsuccessful applications. Panels D and E of Figure 5 show application abandonment rates before and after AIPA. The two panels respectively examine all abandonments and abandonments that are not followed by continuation filings.[18] Once we account for the increasing trend of abandonments before AIPA, we find that terminal abandonment rates declined after AIPA. The regression estimates in Columns 4 and 5 of Table 4 confirm the graphical evidence and suggest up to a 9.2% (=0.019/0.206) decrease in abandonments relative to the pre-AIPA period.

Next, we examine patent scope around AIPA. Following the prior literature, such as Kuhn and Thompson (2017), we use three measures related to patent scope: the total number of allowed claims; the number of independent claims; and the average number of words in the independent claims. A larger number of claims indicates broader scope, and a greater number of words indicates that claims defined with greater precision and clarity. Panels F through H of Figure 5 plot the monthly average patent scope and word count per independent claim for patents filed from 1998 to 2003. We find that patent scope decreased in the post-AIPA period and precision increased. The regression estimates in Columns 6-8 of Table 4 confirm these results.[19]

While the sharpness of the jumps that coincide with AIPA's enactment suggest that these differences are due to AIPA, the magnitude of differences may be contaminated by other coincident changes, such as the dot.com bubble and burst or other macroeconomic cycles that altered the quality of patents filed in the two periods. One could also argue that the greater number of citations to post-AIPA patents reflects the selection of higher-quality patents into the pre-grant disclosure regime after AIPA, rather than enhanced

[18] Not all abandoned patents can be considered "dead and buried" since applicants frequently abandon applications only to file continuation applications with some modifications, which claim lineage with the abandoned application (Hegde et al., 2009).

[19] Our theory model predicts a decrease in patent scope through lower investments in patenting post-AIPA. While we do not directly observe investments in patenting, we find that patent renewal rates went down by up to 1.4 percentage points after AIPA, indicating lower investments in patenting.

knowledge spillovers. To allay these concerns, we test AIPA's effects using a DID approach in the next section.

5. Difference-in-Differences Analysis

5.1 Empirical Design

To identify AIPA's causal impact on technology spillovers, we use a difference-in-differences framework to compare patents filed in the USPTO and their corresponding EP applications before and after AIPA's enactment. This approach allows us to purge the effects of plausible quality-based selection into patenting and pre-grant publication after AIPA. As discussed in Section 2, U.S. applications with foreign parallel applications were mandated to be published 18 months after application if they were filed on or after November 29, 2000. In contrast, patents filed at the EPO were always published 18 months after application (by the EPO), thus providing us with an ideal control group. These parallel, or "twin," applications protect the same underlying invention and, thus, have the same technological value and technology cycles. This twin design of using USPTO and EPO parallel applications has been validated in previous studies, such as Graham et al. (2003).

We focus on EP equivalents rather than on equivalent applications filed at other foreign locations since the EPO is a large patent office with prosecution standards that are relatively similar to the USPTO's. The EPO is also the most favored foreign location for U.S. patent applicants, which allows us to construct a sizeable sample of twin applications. The comparison of patents in two relatively comparable jurisdictions that cover the same technology sharpens our identification of AIPA's effects.

One may question whether disclosure in the U.S. really matters for applications that are filed at both the USPTO and EPO since the EPO publishes virtually the same application 18 months after filing (or nearly simultaneously at the USPTO and EPO after AIPA). Yet the prior literature suggests that applicants and examiners are more likely to search locally for prior art, and our identification rests critically on this assumption. We believe that this identification strategy provides conservative estimates of the effect of patent disclosure since it captures only the marginal effect of disclosure by the USPTO for identical applications that are also disclosed simultaneously by the EPO.

Our main regression specification is summarized below:

$$Outcome_{ijt} = \alpha_1 + \alpha_2 US_j + \alpha_3 US_j * PostAIPA_t + \delta W_j + Family_i + Month_t + \epsilon_{ijt}, \quad (2)$$

where j designates a patent in family i filed in month t. *US* indicates whether the patent application is at the USPTO, and *PostAIPA* indicates whether the patent application is submitted after the effective date of AIPA, which itself is not identified due to the month fixed effects. The variable of interest is the interaction between *US* and *PostAIPA*, which captures the impact of AIPA on the outcome of interest. W_j represents control variables, which include whether the patent is granted (*Granted*) and whether it is granted before 18 months (*EarlyGrant*). By sample construction, which we will discuss in detail in the next section, all U.S. patents in the "twin" sample are granted, while EP equivalents can be granted, abandoned, or pending. The patent prosecution process is, on average, longer in the EPO than in the USPTO; hence, we interact *US* with *EarlyGrant* to allow them to have different coefficients. We add application month fixed effects to control for global trends and business cycles. Most importantly, the "twin" study design allows us to control for unobservable characteristics, such as patent quality and sector-specific time-variant shocks, by adding patent family fixed effects. Thus, our strategy isolates the effect of AIPA, as well as pre-grant disclosure by the USPTO, using patents with identical twins that are disclosed simultaneously in Europe.

One could argue that that the propensity to file for EP parallel applications may have changed after AIPA because of the mandated disclosure requirement in the U.S. We check and do not find any noticeable change in the proportion of U.S. patents with EP parallel applications (or other foreign applications) during 1998-2003.[20] Another concern is that future patentees might simply shift their backward citations from EP patent applications to their U.S. counterparts after AIPA, so we want to check how citations to EP patents changed after AIPA. To test for this possibility, we estimate the following regression specification:

$$Outcome_{ijt} = \alpha_0 + f(Month_t) + \alpha_2 US_j + \delta W_j + \sum_{\tau \in Post} \beta_t I\{Month_t = \tau\} + \sum_{\tau \in Post} \gamma_t US_j * I\{Month_t = \tau\} + \epsilon_{ijt}. \quad (3)$$

Equation (3) controls for a reduced-form common pre-trend in a DID setting. The common pre-trend is captured by $f(Month_t)$, and the level difference between U.S. and EP equivalents is captured by the coefficient on the *US* dummy. We also include the interactions of *US* with each dummy of the month in the post-AIPA period. Therefore, each β_t captures the average changes for both U.S. and EP patents relative to the pre-trend due to the overall economic or technological changes, while each γ_t captures the additional effect of AIPA on U.S. patents. We compute the average changes for EP patents (denoted as *CommonDif*) by taking the mean of the estimated β_t, and the *AIPA-Effect* is captured by the average estimated γ_t.

One disadvantage of equation (3), compared to equation (2), is that we put a restrictive functional form on the common pre-trend. In equation (2), we control for the common trend implicitly through the patent

[20] The propensity to file foreign or EP parallel applications from 1998 to 2003 is depicted in Appendix Figure A5.

family fixed effects. Whether the data exhibit a macro trend or a technology-specific trend, the advantage of the "twin" study design with family-fixed effects is that it controls for trends without specifying its explicit functional form. In equation (3), we specify the common pre-trend as a function of the application month, by which we invoke the parallel trend assumption and impose an additional restriction by requiring it to be linear, quartic, or cubic. We discuss this assumption in detail below as we graphically examine the data.

5.2 Sample Selection

As in our before-and-after analysis, data on U.S. applications are drawn from the USPTO's PAIR files. We supplement these data with information on EPO twin applications available from the European Patent Office's PATSTAT. The comparisons we highlight here are between the U.S. and EP parallel applications.

We identify our treatment group from the universe of utility patent applications filed in the USPTO from January 1, 1998 to December 31, 2003. We identify the corresponding foreign applications filed in the EPO as the control group, based on the patent family table from PATSTAT, which records the complete set of equivalent patent applications filed across different national patent offices.

Appendix Table A4 describes the sample selection process for our treatment group. Our initial sample selection procedure yields 403,292 U.S. patents, which account for 36.4% of all granted patents filed during the same period. Presumably, these patents are relatively more important than the average U.S. patents that do not seek foreign protection. The number of U.S. patents included in the sample further reduces to 316,117 after we impose the following requirements: (i) EP parallel applications are filed within 18 months of the application of their associated U.S. applications according to the Patent Cooperation Treaty (1970); and (ii) exclude(d) EP applications are the PCT filings in the international phase with the EPO designated as the receiving office.[21] We report the corresponding summary statistics in Appendix Table A5. We observe significant increases in the forward citations received in the five, seven or ten years after disclosure (grant date before AIPA and pre-grant publication date after AIPA) for U.S. patents following AIPA's enactment, relative to EP equivalents. Post-AIPA U.S. patents do not receive more citations than pre-AIPA patents when we count citations within three years after disclosure. This is likely the case because pre-AIPA patents are granted at the time of disclosure and, thus, are likely to receive more citations. We also see a sharp drop in the time lag of receiving forward citations following the enforcement of AIPA, relative to EP equivalents. The technological similarity with remote or average patents increases, while that with close patents decreases (*90th percentile* and *95th percentile*), which

[21] PCT filings in the international phase with EPO designated as the receiving office are identified as those filed in the EPO with kind code of "W" in PATSTAT.

suggests the combined effect of enhanced knowledge spillovers and reduced duplicative research. Changes in similarity for EP patents surrounding AIPA appear negligible.

5.3 Difference-in-Differences Results

Citation lag

Our first variable of interest is the speed of knowledge diffusion. To track forward citations to U.S. patents made by patents filed in other patent offices, we supplement the USPTO's database with information from PATSTAT. [22] We require the citing U.S. patent to be granted to avoid mechanical inflations in the forward citations made by U.S. pre-grant publications. For EPO applications, citation data are from PATSTAT, and both citing and cited patents could be granted patents or pre-grant publications. Following Harhoff, Hoisl, and Webb (2009), we adjust for patent equivalents when counting forward citations for EPO patents. Specifically, if a future EPO patent cites a U.S. patent but not its EPO equivalent, it is counted as one forward citation for the EPO equivalent. As we point out in the variable construction section, only patents with at least 1/3/5/7 forward citations are included in the graphs and analyses. This requirement singles out relatively important patents. The more forward citations required, the more selective this requirement is. The selection is more prominent for EP patents, given that EP patents receive, on average, fewer forward citations due to institutional differences related to citing prior art across the two patent offices. On average, only 53% (9%) of the EP applications in our sample receive one (seven) non-self forward citations within ten years after application. But this selection effect does not necessarily bias our DID estimates.

Figure 6 plots the monthly average citation lags for U.S. and EP parallel applications, respectively. It shows a consistent and compelling drop in the citation lags for U.S. patents in the post-AIPA period across the four different citation lag measures. The drop was concentrated in a short window right after AIPA's enactment. By the second quarter of 2002, the time lag for U.S. patents had roughly stabilized. Meanwhile, there is no noticeable change in the speed of knowledge diffusion for EP applications around AIPA. More importantly, the U.S. and EP applications shared a similar trend before AIPA, which alleviates the concern about violating the parallel trend assumption for valid DID analysis.

[Figure 6 here]

In Table 5, we estimate equation (2) to test the impact of AIPA on the pace of knowledge diffusion. The dependent variable is the natural logarithm of one plus the citation lags. Our variable of interest is the

[22] We combine the USPTO data and PATSTAT data to construct forward citations for U.S. patents because the citation data in PATSTAT are limited to 99 citations per patent for application citations and examiner citations, respectively. As U.S. patents tend to cite more often, they are likely affected by the data limitations, while EP patents are likely unaffected, according to the data manual of PATSTAT.

interaction term, *PostAIPA*US*, which is statistically negative across the four regressions. The economic magnitude is also significant. The point estimates indicate that it takes U.S. patents 25% to 29% less time to receive 1/3/5/7 forward citations after AIPA, relative to EP equivalents. In unreported results, we also find a significant reduction in citation lags across different technology classes, and the effect is strongest for "Computers and Communications" patents, consistent with the rapid innovation cycles in this industry. Overall, the evidence provides strong support that pre-grant disclosure speeds up knowledge diffusion.

[Table 5 here]

Forward citations

The analysis of citation lags helps uncover AIPA's impact on the speed of knowledge diffusion (*intensive margin*). Here, we examine AIPA's effect on the magnitude of knowledge spillovers (*extensive margin*). Figure 7 plots the quarterly average forward citations for equivalent U.S. (left axis) and EP (right axis) applications, respectively. First, we see an obvious level difference between U.S. and EP patents, both before and after AIPA. In the DID design, the level difference itself does not bias our estimation. The key is whether they follow a parallel trend prior to AIPA's enactment. To examine this, we plot the natural logarithm of one plus the adjusted forward citations in Appendix Figure A6. We adjust the forward citations by the average citations received by patents filed in 1998 in the same NBER 2-digit code at the same patent office, so that U.S. and EP patents can be aligned on the same scale. We use a linear function to fit the pre-AIPA U.S. and EP patents separately to ease the visual examination of the pre-trend. The two fitted lines appear roughly parallel, providing us with confidence in our DID design. The overall downward trend in EP patents is due mainly to data truncation. The downward trend is shifted upward for U.S. patents immediately after AIPA's enactment but remains largely the same for EP patents in the post-AIPA period. The pattern is largely similar across the different windows to count forward citations.

[Figure 7 here]

Panel A of Table 6 reports the DID estimates for forward citations using equation (2). The dependent variables are the natural logarithm of one plus the number of forward citations. Consistent with the graphic evidence, we observe a significant positive coefficient on *PostAIPA*US* for the 5/7/10-year forward citations, although it is significantly negative for three-year forward citations. Economically, U.S. patents receive 5.7% (14.7%) more five-year (ten-year) forward citations in the post-AIPA period, relative to their EP equivalents. The economic magnitude increases as we extend the horizon of citation counts, which is probably due to the cumulative effect of knowledge in pre-grant publications being transferred to subsequent generation patents, which, in turn, stimulate further follow-on innovation.

[Table 6 here]

One may be concerned that forward citations received by U.S. patents increased because before AIPA, only EP equivalents were published; hence, future patents would have no choice but to cite the visible EP equivalents. After AIPA, since both U.S. and EP pre-grant publications were public, future patents could cite either U.S. or EP publications, thus boosting forward citations received by U.S. patents after AIPA through a substitution effect. If this was the case, we should observe an increase in forward citations for U.S. applications and a decrease for EP equivalents of roughly the same magnitude. To check this, we estimate regressions specified in equation (3) and report the corresponding results in Panel B of Appendix Table A6. Overall, we find that rather than a drop in citations, EP equivalents also received more forward citations after AIPA, although the increase was economically small and only statistically significant at the 10% level when we count the citations in a ten-year window after application. Similar to the estimates of equation (2), estimates of equation (3) show that post-AIPA U.S. patents received 3.0% and 11.3% more citations in the five- and ten-year windows after application, respectively. Therefore, the increase in forward citations received by post-AIPA U.S. patents could not be due to a substitution of citations.

Since detailed information about U.S. inventions with EPO equivalents was already publicly available before AIPA through EP pre-grant publications, knowledge spillovers associated with U.S. disclosure suggests the existence of search frictions across patent offices. Such frictions may have arisen from search costs, language barriers, or a lack of other institutional conditions that facilitate knowledge diffusion across national patent office jurisdictions. Given these search frictions, we expect that AIPA caused a larger increase in knowledge diffusion in the U.S. than in the EP. To test this, we count the forward citations made by future U.S. patents and EP patents separately and run the same regressions as before. The estimation results of the key coefficient (*PostAIPA*US*) are reported in Panel B of Table 6. Again, the coefficients on *PostAIPA*US* are significantly positive for 5/7/10-year forward citations made by future U.S. patents, as well as by future EP patents. More importantly, the magnitude is much larger for forward citations made by future U.S. patents, and its difference is statistically significant at conventional levels. Taken together, our evidence suggests that pre-grant publications increase knowledge spillovers by reducing search costs across patent offices and countries. In unreported analyses, we repeat the same regressions for each of the NBER one-digit technology classes. We find AIPA's effect on the extent of knowledge diffusion to be strongest for "Computers and Communications" patents, characterized by rapid innovation cycles.

Technology similarity

In this section, we study the impact of pre-grant publication on technology similarity. On the one hand, adequate and timely knowledge diffusion can spur follow-on innovation, which could decrease the technology distance between the focal patent and subsequent patent applications. On the other hand, the prompt availability of information on competing inventions can reduce duplicative research in the subsequent period, which would increase technology distance among the closest patents.

To shed light on the potential impact of AIPA on technology similarity, in Figure 8, we plot the monthly average technology similarity with technologically remote or close patents filed in the future. Overall, we see an increase in similarity among technologically remote patents after AIPA (*5th percentile* to *50th percentile*). The pre-post difference becomes smaller in magnitude at the 75^{th} percentile. At the highest percentiles of similarity (90*th percentile* and *95th percentile*), we observe a drop in technology overlap among U.S. patents after AIPA.

[Figure 8 here]

The regression results using equation (2) in Table 7 confirm this graphic finding, revealing an increase in similarity at the 50^{th} or 75^{th} percentile to be about 12.0% (=0.013/0.108) and 3.4% (=0.010/0.289), respectively. Note that the estimated reduction of duplicative patenting is likely to be underestimated, as the *90th percentile* and *95th percentile* are pushed upwards by increased knowledge spillovers after AIPA. Therefore, the 2.1% (=0.013/0.611) reduction in the *95th percentile* should be interpreted as the lower bound for the reduction in substitutive patenting that can be attributed to AIPA.

[Table 7 here]

Figure 9 plots the coefficients on the interaction term in the regressions, with the similarity measure at every 5th percentile as the dependent variables. The coefficient initially increases, reaches a plateau from the 60^{th} to the 70^{th} percentile, and then decreases quite sharply. Collectively, both the graphic evidence and regression results provide robust evidence that, after AIPA, technologically distant U.S. patents become more similar, and similar patents become more differentiated.

[Figure 9 here]

Patenting Intensity

Finally, we empirically analyze the effect of AIPA on patenting intensity. Our theory predicts that patenting activities increase after AIPA due to a richer information environment that decreases the cost of patenting. This prediction is derived under the assumption that the value of the outside option (entering the competitive market without patenting) remains unchanged by AIPA. If this value increases as public knowledge accumulates faster, the patenting rate could decrease after AIPA.

[Figure 10 here]

We measure patenting intensity with the number of patent applications. Figure 10 shows that the monthly count of patent applications filed at the USPTO grew steadily before AIPA and that the growth rate slowed slightly after AIPA while remaining positive. The monthly volume of EP patent applications peaked around AIPA's effective date and marginally slowed down after AIPA, making it hard to graphically infer the impact of AIPA on patent intensity.[23] However, regression analysis suggests a rise in the number of applications or eventually granted applications filed in the USPTO and EPO after AIPA. The corresponding results are reported in Table 8. In Column 1 (2), we define observations at the patent-office *X* month level and count the total number of applications (eventually granted applications) filed in each month from 1998 to 2003 at the USPTO and EPO, respectively. DID estimates indicate that the number of US patent applications (eventually granted applications) increased by 2,304 (1,159) per month after AIPA, relative to EP applications. This increase is large in magnitude—for example, Column 1 indicates that U.S. patent applications increased by 12.3% (=2304/18775) compared to the pre-AIPA average (18,775). As a robustness check, we also conduct the analysis at the patent-office *X* technology-class *X* month level and obtain qualitatively and quantitatively similar results as reported in Columns 3-6. For example, Column 3 indicates that the number of U.S. patent applications in a typical technology class increased by 11.7% (=3.686/31.390) after AIPA, relative to its pre-AIPA average (31.39). Taken together, these findings suggest that patenting intensity increased after AIPA due to the reduced cost and uncertainty associated with patenting.

[Table 8 here]

6. Concluding thoughts

In this study, we develop a theoretical framework to understand the effects of disclosing inventions through patents on technology spillovers and follow-on patenting. The framework is motivated by the enactment of AIPA, which expedited the disclosure of U.S. patent applications by nearly two years, on average. Consistent with the framework's predictions, we find that AIPA had the following effects: (i) increased the rate and magnitude of knowledge diffusion associated with U.S. patents; (ii) increased overlap between technologically distant patents; decreased overlap between similar patents; and lowered inventive steps; (iii) decreased patent abandonments; (iv) decreased patent scope; and (v) increased patenting. These aftereffects of AIPA are absent in U.S. applications' equivalent "twin" applications, filed

[23] In Appendix Figure A7, we analyze the patenting activity in each of the eight single-digit IPC classes. We observe a greater patent propensity in several large IPC technology classes, except for "human necessities" patents (agriculture, food, tobacco, personal or domestic articles, health, life-saving, amusement patents).[23] The pattern does not change appreciably if we exclude EP patents filed by US firms. We also observe a similar pattern when we focus on patent applications that are eventually granted.

in Europe. Evidence of enhanced knowledge diffusion is also absent for the subsample of U.S. patent applications that opt out of pre-grant publication after AIPA.

Overall, we provide causal evidence that patent disclosure matters and that rules governing the timing of disclosure can have a profound impact on knowledge diffusion and follow-on patenting. Early disclosure appears to promote knowledge diffusion, to lower patenting costs and to reduce patenting duplication without decreasing patenting. We are cautious not to interpret these findings as evidence that patent disclosure accelerates the pace of invention, reduces R&D duplication, and decreases the cost of innovation, although such a conclusion would be reasonable given the commonly held assumption that patents are proxies for innovation. Nevertheless, the finding that AIPA had a profound and direct impact on patenting is, in itself, an important contribution to the scarce literature on patent disclosure. Our findings imply that welfare analyses of patents should incorporate their disclosure effects, and we leave it for future empirical work to provide further evidence linking patent disclosure to innovation.

.

7. References

Anton, James J. and Yao, Dennis A. (1994). Expropriation and inventions: Appropriable rents in the absence of property rights. *The American Economic Review*, pages 190-209, 1994.

Anton, James J. and Yao, Dennis A. (2004). Little patents and big secrets: managing intellectual property. *RAND Journal of Economics*, pages 1-22, 2004.

Aoki, Reiko and Spiegel, Yossi. (2009). Pre-grant patent publication and cumulative innovation. *International Journal of Industrial Organization*, 27(3):333-345, 2009.

Atkeson, Andrew and Burstein, Ariel. (2007). Innovation, firm dynamics, and international trade. Technical report, *National Bureau of Economic Research*, 2007.

Atkeson, Andrew and Burstein, Ariel. (2011). Aggregate implications of innovation policy. Technical report, *National Bureau of Economic Research*, 2011.

Bhattacharya, Sudipto and Guriev, Sergei. (2006). Patents vs. trade secrets: Knowledge licensing and spillover. *Journal of the European Economic Association*, 4(6):1112-1147.

Bhattacharya, Sudipto, Glazer, Jacob and Sappington, David EM. (1992). Licensing and the sharing of knowledge in research joint ventures. *Journal of Economic Theory*, 56(1):43-69.

Bobtche, Catherine, Bolte, Jerome, and Mariotti, Thomas. (2013). Researchers dilemma. Technical report, *Institut d'Economie Industrielle* (IDEI), Toulouse.

Bloom, N., Schankerman, M., & Van Reenen, J. (2013). Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4), 1347-1393.

Cohen, W. M., Goto, A., Nagata, A., Nelson, R. R., & Walsh, J. P. (2002). R&D spillovers, patents and the incentives to innovate in Japan and the United States. *Research policy*, 31(8), 1349-1367.

Fromer, J. C. (2009). The Compatibility of Patent Law and the Internet. *Fordham Law Review*, 78, 2783.

Furman, J. L., Nagler, M., & Watzinger, M. (2018). Disclosure and Subsequent Innovation: Evidence from the Patent Depository Library Program (No. w24660). National Bureau of Economic Research.

Graham, S., Hall, B., Harhoff, D., & Mowery, D. (2003). Patent quality control: A comparison of US patent re-examinations and European Patent oppositions. *Patents in the knowledge-based economy*, 74, 84-85.

Graham, S., & Hegde, D. (2015). Disclosing patents' secrets. *Science*, 347(6219), 236-237.

Gross, T., Notowidigdo, M. J., & Wang, J. (2016). The marginal propensity to consume over the business cycle (No. w22518). *National Bureau of Economic Research*.

Hall, Bronwyn, Christian Helmers, Mark Rogers and Vania Sena. (2014). The Choice between Formal and Informal Intellectual Property: A Review. *Journal of Economic Literature*, 52(2): 375-423.

Hall, Bronwyn, and Vania Sena. (2014). Appropriability Mechanisms, Innovation and Productivity: Evidence from the UK. NBER Working Paper No. 20514.

Harhoff, D., Hoisl, K., & Webb, C. (2009). European Patent Citations–How to Count and How to Interpret Them, unpublished manuscript, University of Munich.

Hopenhayn, Hugo and Squintani, Francesco. (2011). Preemption games with private information. The *Review of Economic Studies*, 78(2):667-692, 2011.

Hegde, D. (2012). Funding and performance at the US Patent and Trademark Office. *Nature Biotechnology*, 30(2), 148.

Hegde, D., Mowery, D. C., & Graham, S. J. (2009). Pioneering inventors or thicket builders: Which US firms use continuations in patenting?. *Management Science*, 55(7), 1214-1226.

Horstmann, Ignatius, MacDonald, Glenn M and Slivinski, Alan. (1985). Patents as information transfer mechanisms: To patent or (maybe) not to patent. *The Journal of Political Economy*, pages 837-858.

Jaffe, A. B. (1986). Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits and Market Value. *American Economic Review* 76(5): 984-1001.

Johnson, D. K., & Popp, D. (2003). Forced out of the Closet: The Impact of the American Inventors Protection Act on the Timing of Patent Disclosure. *RAND Journal of Economics*, 96-112.

Kuhn, J., & Thompson, N. (2017). How to Measure and Draw Causal Inferences with Patent Scope. *International Journal of the Economics of Business*.

Levin, Richard C., Alvin K. Klevorick, Richard R. Nelson, and Sidney G. Winter. (1987) Appropriating the Returns from Industrial Research and Development. *Brookings Papers on Economic Activity* 3:783–831.

Modigliani, F., et al. An open letter to the U.S. Senate (1999). www.eagleforum.org/patent/nobel_letter.html.

Okada, Yoshimi & Nagaoka, Sadao, (2015). Effects of early patent disclosure on knowledge dissemination: evidence from the pre-grant publication system introduced in the United States," *IIR Working Paper* 15-12, Institute of Innovation Research, Hitotsubashi University.

Roin, B. N. (2005). The disclosure function of the patent system (or lack thereof). *Harvard Law Review* 118, 2007-2028.

Scotchmer, S., & Green, J. (1990). Novelty and disclosure in patent law. *The RAND Journal of Economics*, 131-146.

Stefano, B., & Simeth, M. (2018). Patents and Knowledge Diffusion: The Effect of Early Disclosure. Available at SSRN: https://ssrn.com/abstract=3160533.

Trajtenberg, M., Jaffe, A. and R. Henderson. (1997). University versus Corporate Patents: A Window on the Basicness of Invention. *Economics of Innovation and New Technology* 5 (1), pp. 19-50

Ouellette, L. L. (2017). Who reads patents? *Nature Biotechnology* 35, p 421–424.

Williams, H. L. (2017). How Do Patents Affect Research Investments?. *Annual Review of Economics*, 9, 441-469.

8. Figures and Tables

Figure 1. Thomas Edison's "Light Bulb" patent

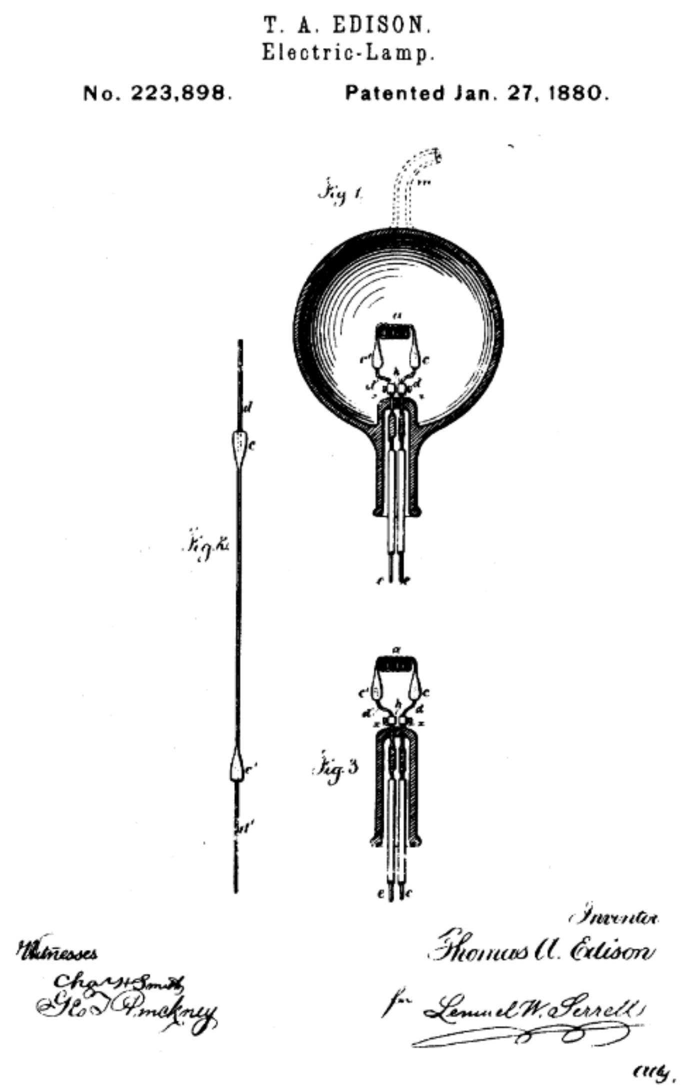

UNITED STATES PATENT OFFICE.

THOMAS A. EDISON, OF MENLO PARK, NEW JERSEY.

ELECTRIC LAMP.

SPECIFICATION forming part of Letters Patent No. 223,898, dated January 27, 1880.

Application filed November 4, 1879.

To all whom it may concern:

Be it known that I, THOMAS ALVA EDISON, of Menlo Park, in the State of New Jersey, United States of America, have invented an Improvement in Electric Lamps, and in the method of manufacturing the same, (Case No. 186,) of which the following is a specification.

The object of this invention is to produce electric lamps giving light by incandescence, which lamps shall have high resistance, so as to allow of the practical subdivision of the electric light.

The invention consists in a light-giving body of carbon wire or sheets coiled or arranged in such a manner as to offer great resistance to the passage of the electric current, and at the same time present but a slight surface from which radiation can take place.

The invention further consists in placing such burner of great resistance in a nearly-perfect vacuum, to prevent oxidation and injury to the conductor by the atmosphere. The current is conducted into the vacuum-bulb through platina wires sealed into the glass.

The invention further consists in the method of manufacturing carbon conductors of high resistance, so as to be suitable for giving light by incandescence, and in the manner of securing perfect contact between the metallic conductors or leading-wires and the carbon conductor.

Heretofore light by incandescence has been obtained from rods of carbon of one to four ohms resistance, placed in closed vessels, in which the atmospheric air has been replaced by gases that do not combine chemically with the carbon. The vessel holding the burner has been composed of glass cemented to a metallic base. The connection between the leading-wires and the carbon has been obtained by clamping the carbon to the metal. The leading-wires have always been large, so that their resistance shall be many times less than the burner, and, in general, the attempts of previous persons have been to reduce the resistance of the carbon rod. The disadvantages of following this practice are, that a lamp having but one to four ohms resistance cannot be worked in great numbers in multiple arc without the employment of main conductors of enormous dimensions; that, owing to the low resistance of the lamp, the leading-wires must be of large dimensions and good conductors, and a glass globe cannot be kept tight at the place where the wires pass in and are cemented; hence the carbon is consumed, because there must be almost a perfect vacuum to render the carbon stable, especially when such carbon is small in mass and high in electrical resistance.

The use of a gas in the receiver at the atmospheric pressure, although not attacking the carbon, serves to destroy it in time by "air-washing," or the attrition produced by the rapid passage of the air over the slightly-coherent highly-heated surface of the carbon. I have reversed this practice. I have discovered that even a cotton thread properly carbonized and placed in a sealed glass bulb exhausted to one-millionth of an atmosphere offers from one hundred to five hundred ohms resistance to the passage of the current, and that it is absolutely stable at very high temperatures; that if the thread be coiled as a spiral and carbonized, or if any fibrous vegetable substance which will leave a carbon residue after heating in a closed chamber be so coiled, as much as two thousand ohms resistance may be obtained without presenting a radiating-surface greater than three-sixteenths of an inch; that if such fibrous material be rubbed with a plastic composed of lamp-black and tar, its resistance may be made high or low, according to the amount of lamp-black placed upon it; that carbon filaments may be made by a combination of tar and lamp-black, the latter being previously ignited in a closed crucible for several hours and afterward moistened and kneaded until it assumes the consistency of thick putty. Small pieces of this material may be rolled out in the form of wire as small as seven one-thousandths of a inch in diameter and over a foot in length, and the same may be coated with a non-conducting non-carbonizing substance and wound on a bobbin, or as a spiral, and the tar carbonized in a closed chamber by subjecting it to high heat, the spiral after carbonization retaining its form.

All these forms are fragile and cannot be clamped to the leading wires with sufficient force to insure good contact and prevent heating. I have discovered that if platinum wires are used and the plastic lamp-black and tar material be molded around it in the act of carbonization there is an intimate union by com-

Figure 2. Citations lag for U.S. patents before and after AIPA

The figures plot the monthly average citation lags for U.S. patents filed during 1998-2003. Citation lag is measured as the number of months between the application date of a focal patent and the application dates of its 1st, 3rd, 5th, or 7th non-self forward citations. Only patents that have accumulated the required number of forward citations within ten years after application are included. The vertical dashed line represents AIPA's effective date (November 29, 2000).

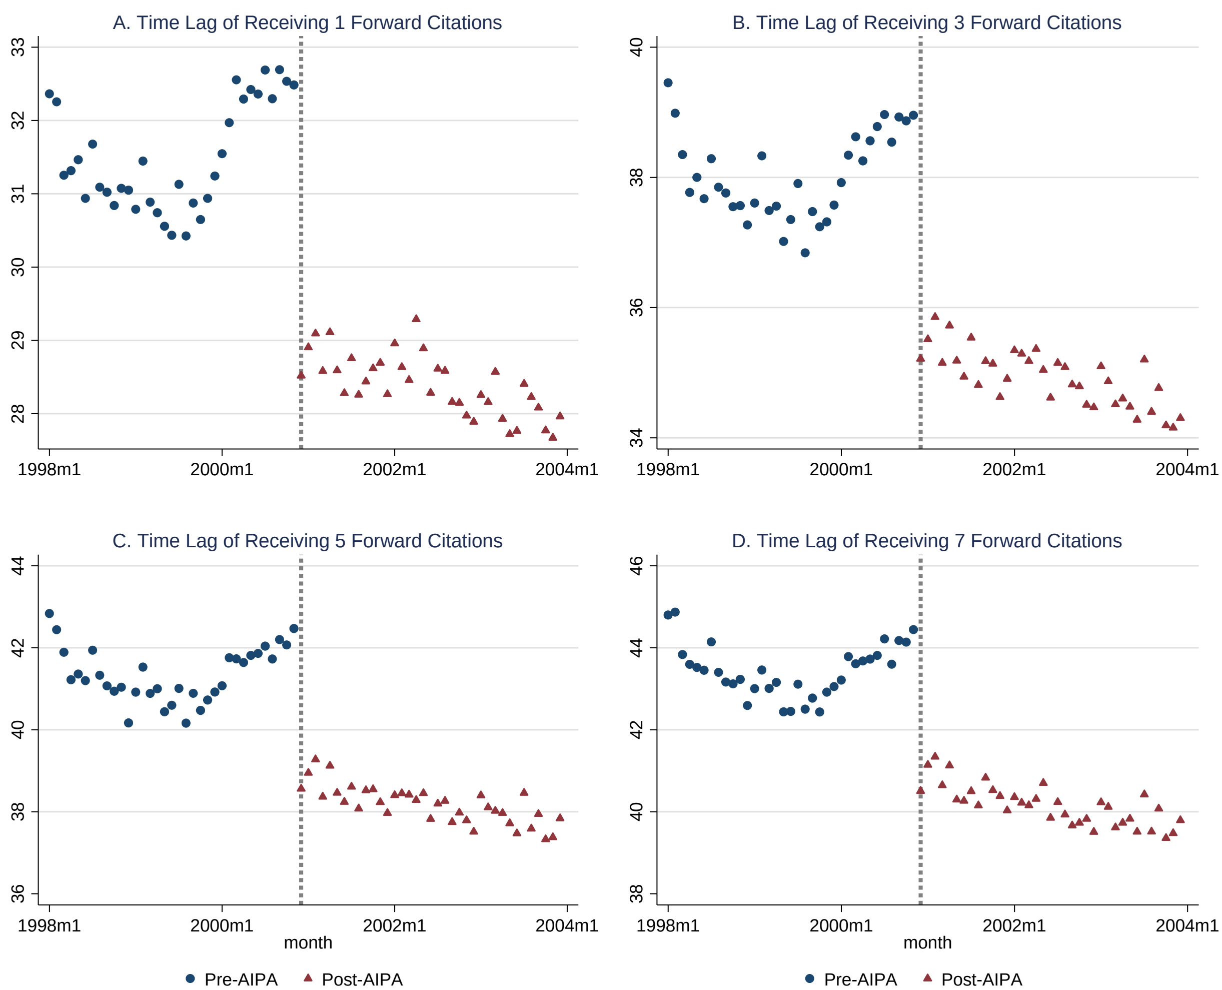

Figure 3. Citations to U.S. patents before and after AIPA

The figures plot the monthly average number of forward citations (excluding self-citations) to U.S. patents filed during 1998-2003. Forward citations are counted cumulatively 3/5/7/10 years after patent disclosure (i.e., publication date for patents with pre-grant publications and grant date for those without). The solid line in each graph represents an OLS regression line fit using only pre-AIPA patents. The vertical dashed line represents AIPA's effective date (November 29, 2000).

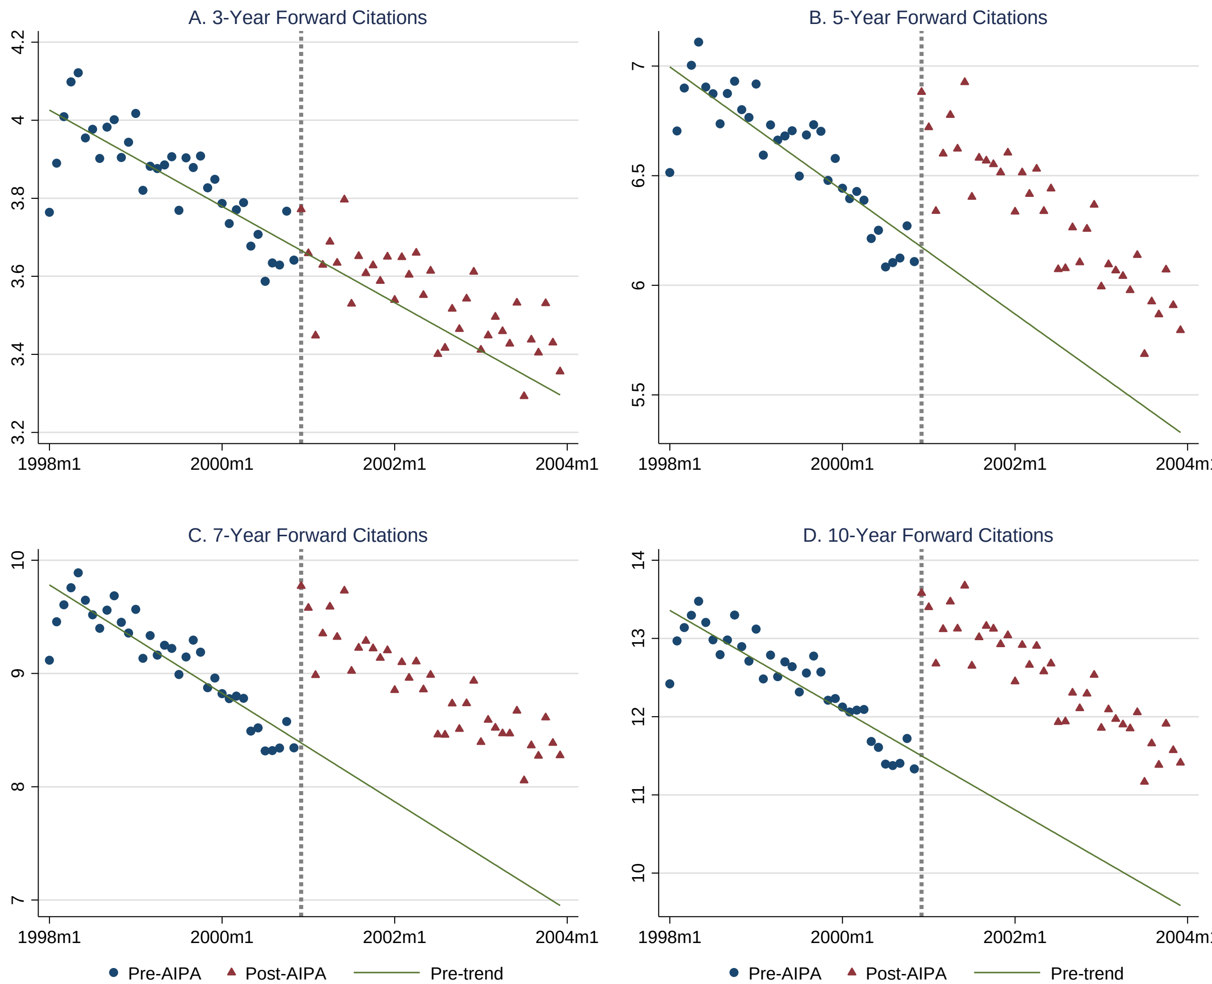

Figure 4. Technology similarity of U.S. patents before and after AIPA

The figures plot the monthly average technology similarity between U.S. patents filed during 1998-2003 and "next-generation" patents. Similarity is measured as the pair-level cosine distance, based on the distribution of IPC main groups (IPC 7-digit code), between the focal patent and patents in its next generation. "Next-generation" patents are those that were filed in the same IPC technology subclass (IPC 4-digit code) within the window of 19-36 months after the focal patent's filing. We then take the 5th, 10th, 15th, 25th, 50th, 75th, 85th, 90th, and 95th percentile values across all "next-generation" patents to construct a patent-level similarity for each focal patent. The solid line in each graph represents an OLS regression line fit using only pre-AIPA patents. The vertical dashed line represents AIPA's effective date (November 29, 2000).

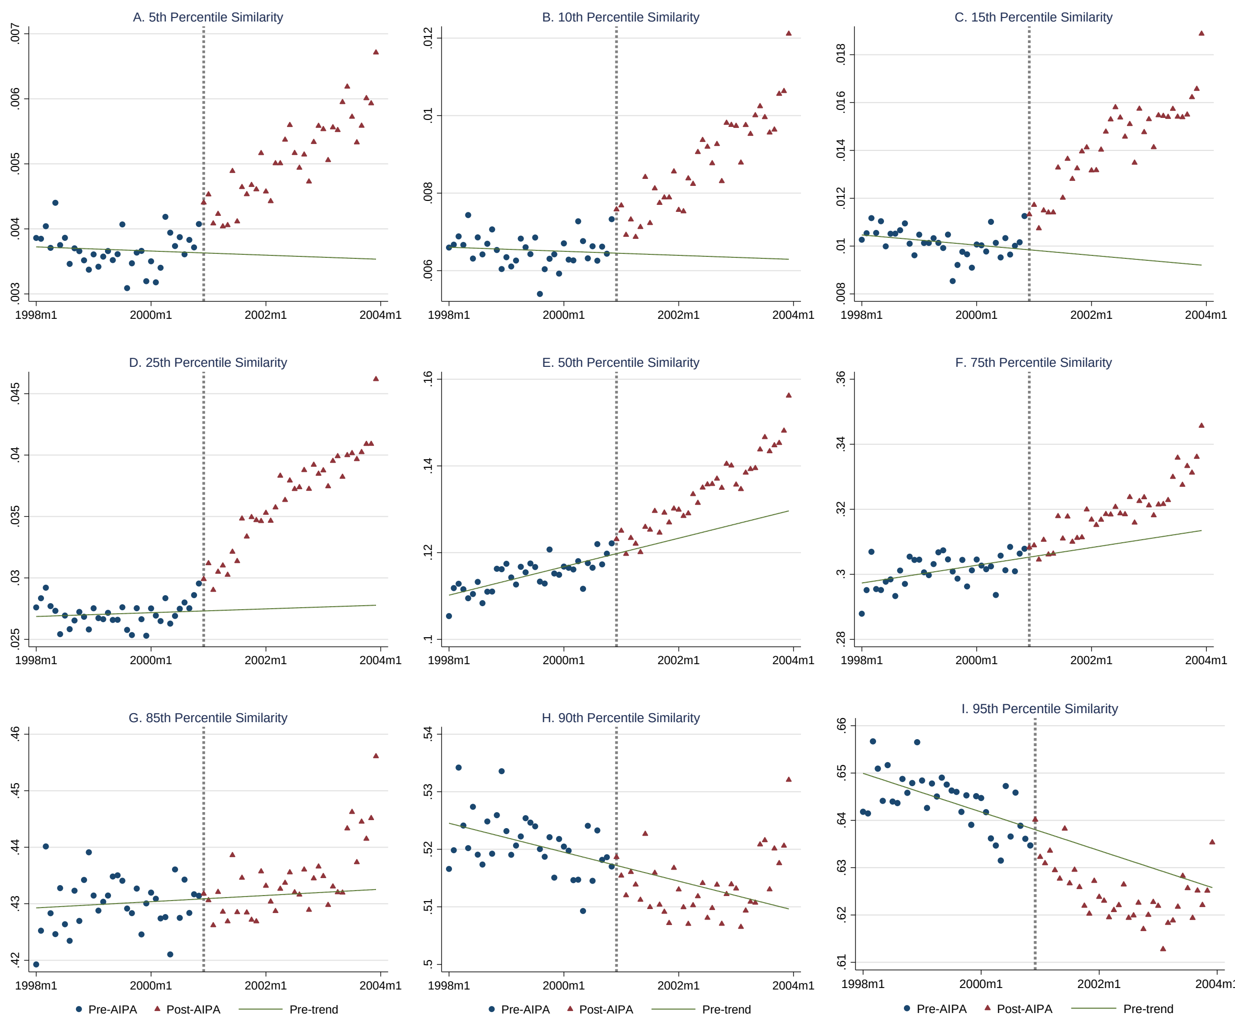

Figure 5. U.S. patent renewal, originality, abandonment, and claims

The figures plot the average renewal rates (Panel A) and originality index (Panels B and C) of eventually granted applications filed during each month during 1998-2003. Panels D and E plot abandonment rates and abandonments without subsequent continuation filings for all applications filed from 1998 to 2003. Panels F-H plot the average number of total allowed claims, independent claims, and average words per independent claim for issued patents. The solid line in each graph represents an OLS regression line fit using only pre-AIPA patents. The vertical dashed line represents AIPA's effective date (November 29, 2000).

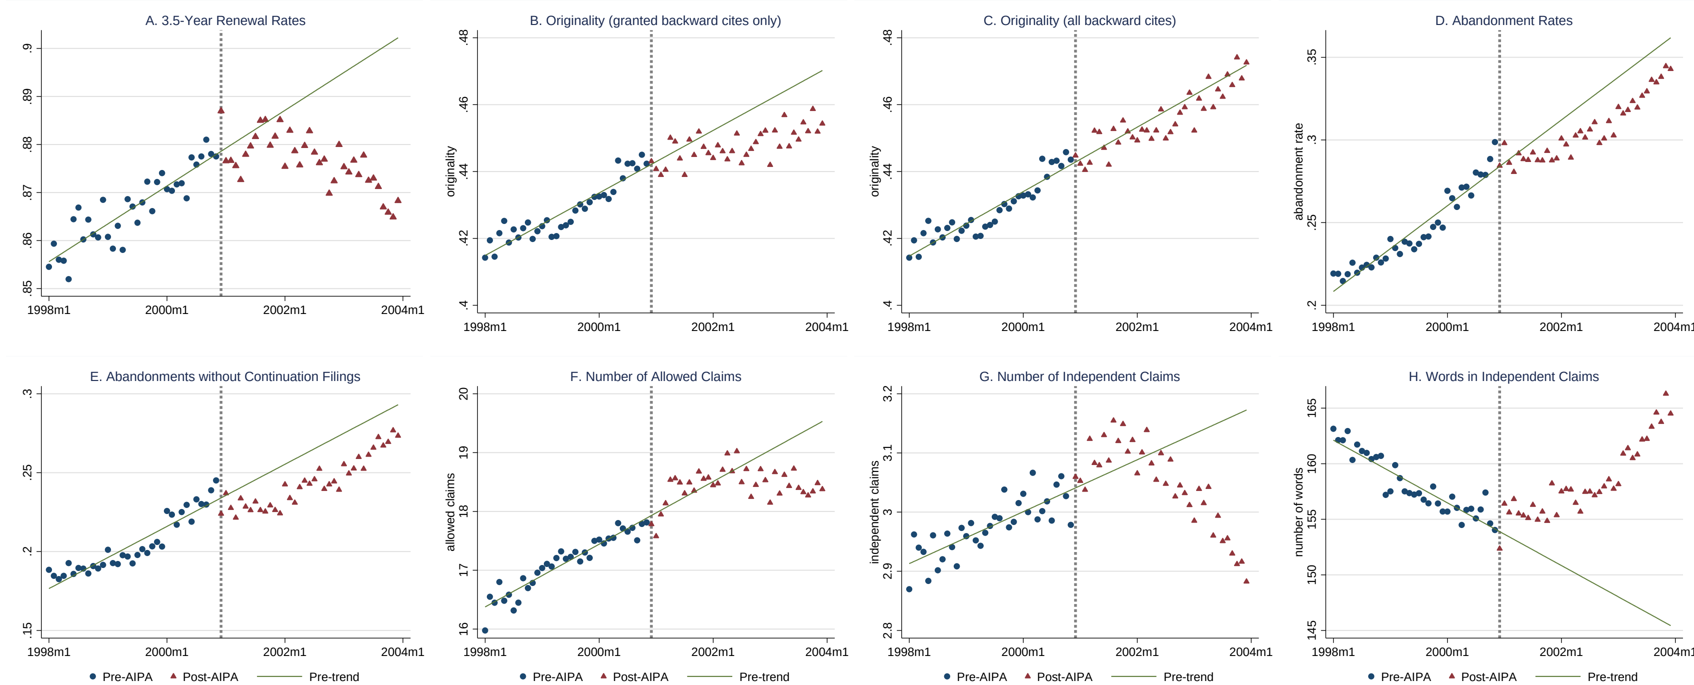

Figure 6. Citation lags to U.S. patents and EPO “twins” before and after AIPA

The figures plot the monthly average citation lags of U.S. patents and their equivalent “twins” at the European Patent Office (EPO) filed during 1998-2003. Time lag is measured as the number of months between the application date of a focal patent and the application dates of its 1st/3rd /5th/7th forward citations. Only patents that have accumulated the required number of forward citations within ten years after application are included. The vertical line indicates AIPA’s effective date (November 29, 2000).

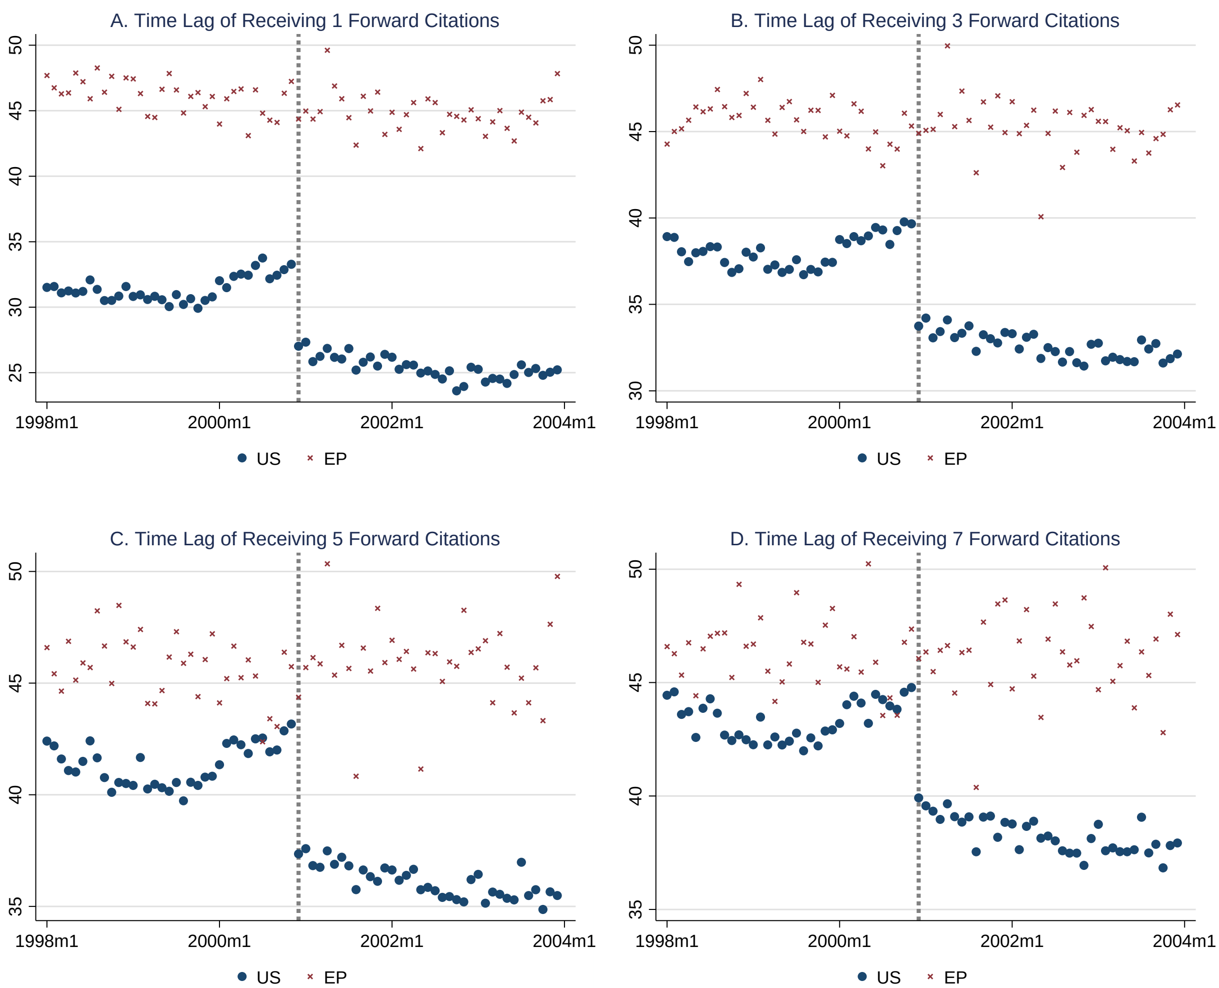

Figure 7. Citations to U.S. patents and EPO "twins" before and after AIPA

The figures plot the monthly average number of forward citations (excluding self-citations) to U.S. patents and their equivalent "twins" at the European Patent Office (EPO) filed during 1998-2003. Forward citations are counted cumulatively 3/5/7/10 years after patent disclosure (i.e., publication date for patents with pre-grant publications or grant date for those without). Citations data are obtained from the USPTO and PATSTAT. The vertical line represents AIPA's effective date (November 29, 2000).

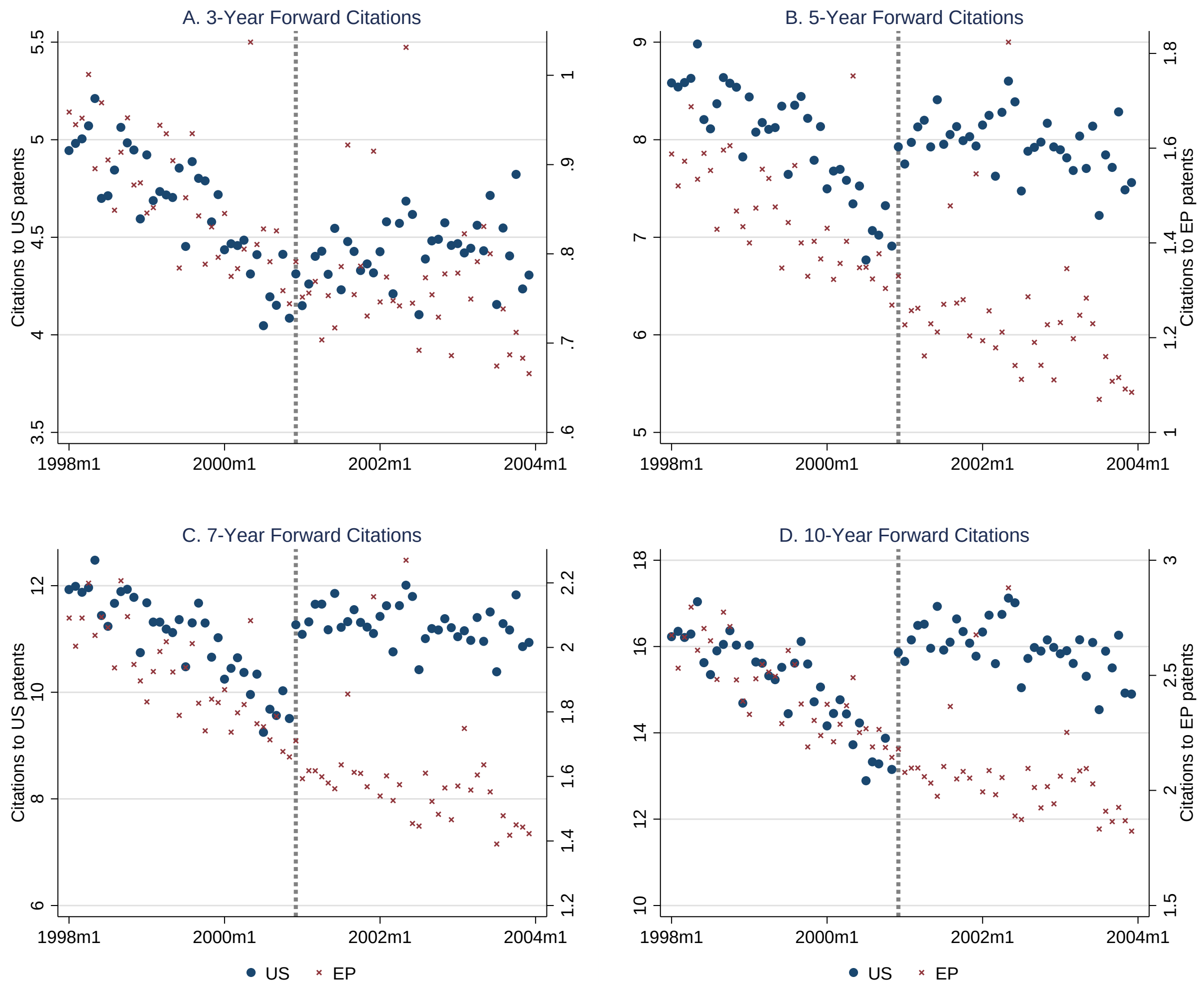

Figure 8. Similarity of U.S. patents and EPO “twins” before and after AIPA

The figures plot the monthly average technology similarity between patents filed during 1998-2003 and “next-generation” patents for U.S. patents and their equivalent “twins” filed at the EPO. Similarity is measured as the cosine distance, based on the distribution of IPC main groups (IPC 7-digit code), between the focal patent and patents in its next generation. “Next-generation” patents are those that were filed in the same IPC technology subclass (IPC 4-digit code) within the window of 19-36 months after the focal patent’s filing. We then take the 5th, 10th, 15th, 25th, 50th, 75th, 85th, 90th, and 95th percentile values across all “next-generation” patents to construct a patent-level similarity for each focal patent. The vertical line indicates AIPA’s effective date (November 29, 2000).

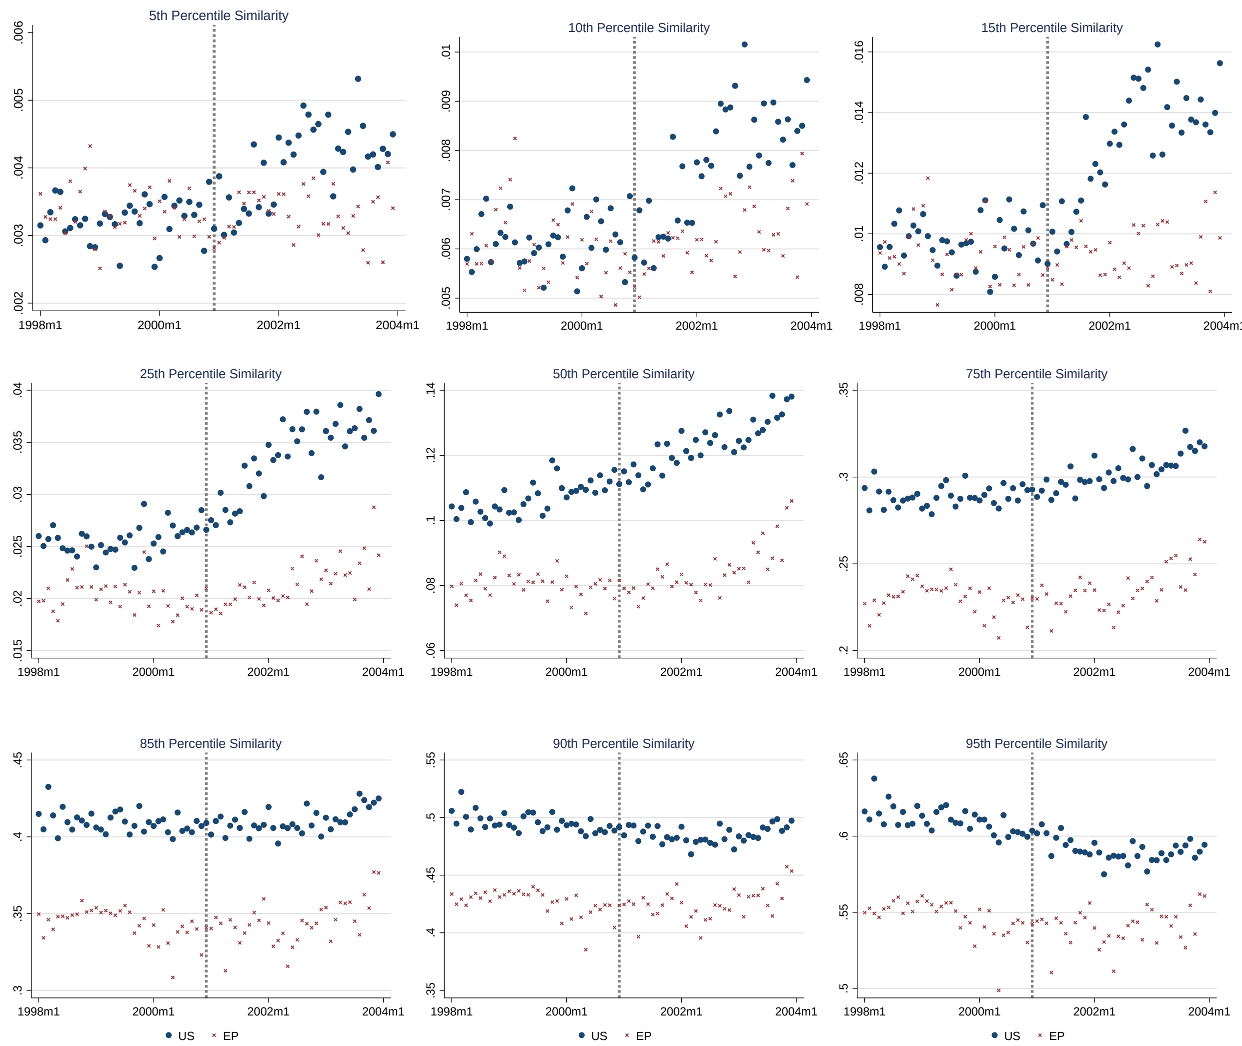

Figure 9. AIPA's effect on similarity at different levels of technology overlap

This figure plots the estimated AIPA effect on technology similarity measured at different percentiles (between focal patents and next-generation patents). The estimated AIPA effect is the coefficient on the interaction term *US #PostAIPA*. Refer to Table 7 notes for a description of the regression specifications.

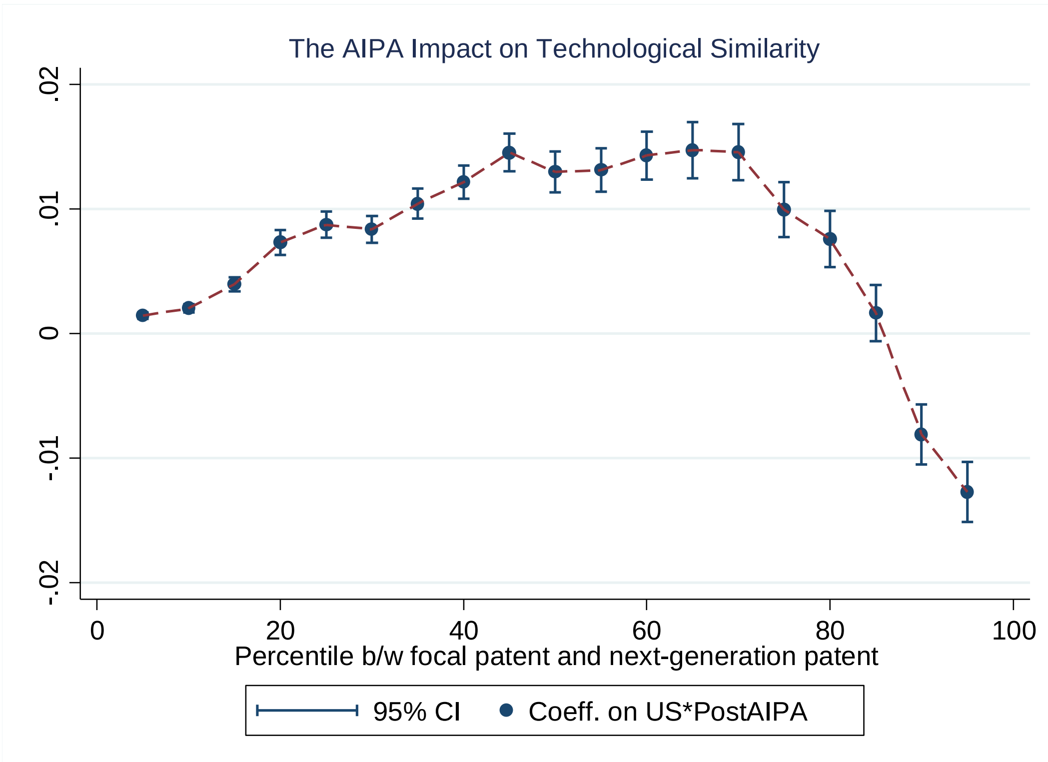

Figure 10. Patenting Intensity (U.S. v. EP)

This figures plot the number of patent applications (the left graph) and eventually granted patent applications (the right graph) filed during each month during 1998-2003 at the USPTO and EPO, respectively. The vertical dashed line represents AIPA's effective date (November 29, 2000).

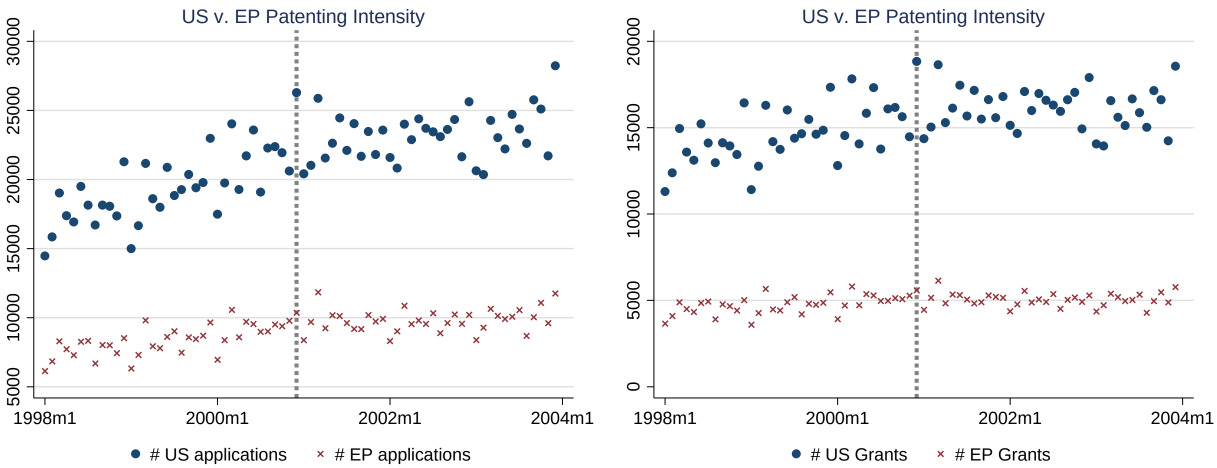

Table 1. Variable definitions

Variable	Definition
Optout	Dummy variable, equal to one if the patent application is filed after the enactment of AIPA and opts out of the pre-grant publication requirement.
EarlyGrant	Dummy variable, equal to one if the patent application is granted 18 months after application.
Fcite3Y	The number of forward citations received within three years after disclosure (publication date for patents with pre-grant publications and grant date for patents without). When followed by suffix 'US' ('EP'), the forward citing patents included in the computation are restricted to those applied in the USPTO (EPO).
Fcite5Y	The number of forward citations received within five years after disclosure. When followed by suffix 'US' ('EP'), the forward citing patents included in the computation are restricted to those applied in the USPTO (EPO).
Fcite7Y	The number of forward citations received within seven years after disclosure. When followed by suffix 'US' ('EP'), the forward citing patents included in the computation are restricted to those applied in the USPTO (EPO).
Fcite10Y	The number of forward citations received within ten years after disclosure. When followed by suffix 'US' ('EP'), the forward citing patents included in the computation are restricted to those applied in the USPTO (EPO).
Lag1Fcite	The average time lag to receive the first forward citations conditional on having at least one forward citation within ten years of application (unit: month).
Lag3Fcite	The average time lag to receive the first three forward citations conditional on having at least three forward citations within ten years of application (unit: month).
Lag5Fcite	The average time to receive the first five forward citations conditional on having at least five forward citations within ten years of application (unit: month).
Lag7Fcite	The average time lag to receive the first seven forward citations conditional on having at least seven forward citations within ten years of application (unit: month).
Xth-IPC7Sim	The X^{th} percentile of the pair-wise cosine similarity based on the shares of IPC main group assignments (IPC 7-digit codes) of the focal patent and the next cohort patents (patents that are applied in the same IPC subclass [IPC 4-digit codes] within the window of 19-36 months after the application date of the focal patent). X ranges from 5 to 95.
3.5-Year Renewal	Dummy variable, equal to one if payment of renewal fees due in 3.5 years from grant date is made.
Originality	One minus the Herfindahl index of the patent’s backward citations in each U.S. patent classification system (USPC) technology class. Only backward citations of patents that are granted when the citations are made are included.
Originality2	One minus the Herfindahl index of the patent’s backward citations in each U.S. patent classification system (USPC) technology class. Both backward citations of pre-grant publications and granted patents are included.
Claims	Total number of claims allowed at grant.
IndClaims	The number of independent claims allowed at grant.
IndClaim_Wrd	The average number of words per independent claim.
Abandon	Dummy variable, equal to one if the application is abandoned.
Abandon2	Dummy variable, equal to one if the application is abandoned and does not file continuation applications that claim priority from it.

Table 2. Summary statistics for U.S. patent applications, 1998-2003

This table reports the summary statistics of key variables of interest of patent applications filed in the USPTO from 1998 to 2003. For more details on the variable definitions, please refer to Table 1.

	Pre-AIPA Grants #=509,924			Post-AIPA Grants #=597,786		
VARIABLES	Mean	S.D.	Median	Mean	S.D.	Median
Optout				0.081	0.273	0.000
EarlyGrant	0.187	0.390	0.000	0.156	0.363	0.000
Fcite3Y	4.070	7.904	2.000	3.727	7.326	2.000
Fcite5Y	6.961	13.518	3.000	6.624	12.690	3.000
Fcite7Y	9.633	19.109	4.000	9.335	17.997	4.000
Fcite10Y	13.205	26.967	5.000	13.179	26.501	5.000
Lag1Fcite	38.080	21.455	34.590	34.976	21.547	30.842
Lag3Fcite	41.396	20.342	39.252	38.214	20.427	35.430
Lag5Fcite	43.477	19.757	41.981	40.197	19.849	38.108
Lag7Fcite	45.478	19.158	44.630	42.105	19.317	40.623
5th-IPC7Sim	0.004	0.030	0.000	0.006	0.037	0.000
10th-IPC7Sim	0.007	0.043	0.000	0.010	0.052	0.000
15th-IPC7Sim	0.011	0.055	0.000	0.016	0.068	0.000
25th-IPC7Sim	0.028	0.089	0.000	0.038	0.110	0.000
50th-IPC7Sim	0.116	0.208	0.000	0.135	0.219	0.000
75th-IPC7Sim	0.301	0.323	0.218	0.319	0.316	0.258
85th-IPC7Sim	0.430	0.341	0.408	0.434	0.328	0.408
90th-IPC7Sim	0.521	0.335	0.530	0.514	0.319	0.507
95th-IPC7Sim	0.644	0.305	0.707	0.625	0.293	0.667
3.5-Year Renewal	0.867	0.339	1.000	0.877	0.329	1.000
Originality	0.429	0.276	0.490	0.448	0.275	0.500
Originality2	0.429	0.276	0.490	0.456	0.273	0.500
Claims	17.530	17.357	15.000	18.845	17.042	16.000
IndClaims	3.023	2.566	2.000	3.091	2.572	2.000
IndClaims_Wrd	159.958	103.039	141.000	160.505	104.907	141.167
	Pre-AIPA Applications #=675,917			Post-AIPA Applications #=860,429		
Optout				0.086	0.280	0.000
HasChild	0.230	0.421	0.000	0.258	0.438	0.000
Abandon	0.247	0.431	0.000	0.306	0.461	0.000
Abandon2	0.206	0.404	0.000	0.244	0.429	0.000

Table 3. AIPA's effect on knowledge diffusion and patent similarity: before-and-after analysis

This table reports Ordinary Least Squares (OLS) regression estimates of AIPA's effect on the citation lags (Panel A), number of forward citations (Panel B), and technology similarity (Panel C). The regressions are estimated using the following specification:

$$Outcome_{it} = \alpha_0 + \sum_{\tau \in Post} \beta_t * I\{Month_t = \tau\} + \alpha_1 Month_t + \alpha_2 EarlyGrant_i + \alpha_3 OptOut_i + TechFE + \epsilon_{it},$$

where i indicates the patent application filed in month t. We include *Month* (a continuous variable indicating the calendar month when the patent is filed) and a set of dummy variables indicating each month in the post-AIPA regime ($I\{Month_t = \tau\}$); hence, the continuous variable *Month* captures a linear pre-trend. Control variables include a dummy variable (*EarlyGrant*) indicating patents that are granted before 18 months after application and another (*OptOut*) indicating patents that opt out of the 18-month disclosure requirement in the post-AIPA period. To control for technology heterogeneity, we also include technology class fixed effects (3-digit USPC code). Standard errors are clustered by the application month. We obtain the pre-post difference by taking the mean of the estimates of β_t, and its associated standard errors are computed using the delta method. ***, **, and * stand for statistical significance based on two-sided tests at the 1%, 5%, and 10% level, respectively.

Panel A: Citations lags				
	(1)	(2)	(3)	(4)
VARIABLES	Ln_Lag1Fcite	Ln_Lag3Fcite	Ln_Lag5Fcite	Ln_Lag7Fcite
Pre-Post-Dif ($\frac{1}{37}\sum_{\tau \in Post} \beta_t$)	-0.201 ***	-0.143***	-0.126***	-0.116***
	(0.007)	(0.006)	(0.006)	(0.006)
Month	0.001***	0.001***	0.000**	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
EarlyGrant	-0.317***	-0.218***	-0.184***	-0.169***
	(0.003)	(0.003)	(0.003)	(0.003)
OptOut	0.276***	0.212***	0.189***	0.175***
	(0.007)	(0.006)	(0.005)	(0.005)
USPC FE	Yes	Yes	Yes	Yes
Observations	990,975	772,256	604,596	482,676
Adj R-squared	0.103	0.126	0.130	0.130

Panel B: Number of citations				
	(1)	(2)	(3)	(4)
VARIABLES	Ln_Fcite3Y	Ln_Fcite5Y	Ln_Fcite7Y	Ln_Fcite10Y
Pre-Post-Dif ($\frac{1}{37}\sum_{\tau \in Post} \beta_t$)	0.038***	0.113***	0.164***	0.190***
	(0.006)	(0.007)	(0.007)	(0.007)
Month	-0.003***	-0.005***	-0.006***	-0.006***
	(0.000)	(0.000)	(0.000)	(0.000)
EarlyGrant	-0.016***	-0.034***	-0.045***	-0.058***
	(0.002)	(0.003)	(0.003)	(0.003)
OptOut	-0.071***	-0.135***	-0.158***	-0.164***
	(0.006)	(0.006)	(0.006)	(0.006)
USPC FE	Yes	Yes	Yes	Yes

Observations	1,107,656	1,107,204	1,104,670	1,089,865
Adj R-squared	0.135	0.148	0.154	0.161

Panel C: Technology similarity

VARIABLES	(1) 50th-IPC7Sim	(2) 75th-IPC7Sim	(3) 90th-IPC7Sim	(4) 95th-IPC7Sim
Pre-Post-Dif ($\frac{1}{37}\sum_{\tau \in Post} \beta_t$)	0.006***	0.006***	-0.006***	-0.014***
	(0.002)	(0.002)	(0.002)	(0.002)
Month	0.000***	0.000*	-0.000***	-0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
EarlyGrant	-0.001	-0.001	0.001	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)
OptOut	-0.002	0.010***	0.038***	0.052***
	(0.001)	(0.002)	(0.002)	(0.002)
USPC FE	Yes	Yes	Yes	Yes
Observations	1,106,975	1,106,975	1,106,975	1,106,975
Adj R-squared	0.290	0.238	0.179	0.151

Table 4. AIPA's effect on patent characteristics: before-and-after analysis

This table reports Ordinary Least Squares (OLS) regression estimates of AIPA's effect on renewal rates, patent originality, abandonment rates, and patent claims. The regressions are estimated using the following specification:

$$Outcome_{it} = \alpha_0 + \sum_{\tau \in Post} \beta_t * I\{Month_t = \tau\} + \alpha_1 Month_t + \alpha_2 EarlyGrant_i + \alpha_3 OptOut_i + TechFE + \epsilon_{it},$$

where i indicates the patent application filed in the month t. We include *Month* (a continuous variable indicating the calendar month in which the patent is filed) and a set of dummy variables indicating each month in the post-AIPA regime ($I\{Month_t = \tau\}$). Hence, the continuous variable *Month* captures a linear pre-trend. Control variables include a dummy variable (*EarlyGrant*) indicating patents that are granted before 18 months after application and another (*OptOut*) indicating patents that opt out of the 18-month disclosure requirement in the post-AIPA period. To control for technology heterogeneity, we also include technology class fixed effects (3-digit USPC classification). Standard errors are clustered by the application month. We obtain the pre-post difference by taking the mean of the estimates of β_t, and its associated standard errors are computed using the delta method. ***, **, and * stand for statistical significance based on two-sided tests at the 1%, 5%, and 10% level, respectively.

VARIABLES	(1) 3.5-Year Renewal	(2) Originality	(3) Originality2	(4) Abandon	(5) Abandon2	(6) Claims	(7) IndClaims	(8) IndClaims_ Wrd
18-Month AIPA-Effect	-0.005***	-0.005***	-0.002*	-0.010***	-0.019***	-0.165***	-0.017***	3.615***
Pre-Post-Dif ($\frac{1}{18}\sum_{\tau \in Post18} \beta_t$)	(0.002)	(0.001)	(0.001)	(0.003)	(0.002)	(0.106)	(0.013)	(0.578)
37-Month AIPA-Effect	-0.014***	-0.010***	-0.004**	-0.013***	-0.019***	-0.665***	-0.120***	8.22***
Pre-Post-Dif ($\frac{1}{37}\sum_{\tau \in Post} \beta_t$)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.106)	(0.016)	(0.717)
Control Variables								
Month	0.001***	0.001***	0.001***	-0.000***	0.002***	0.045***	0.003***	-0.223***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.003)	(0.000)	(0.019)
EarlyGrant	-0.006***	-0.049***	-0.049***			-3.033***	-0.592***	5.018***
	(0.001)	(0.001)	(0.001)			(0.047)	(0.006)	(0.333)
OptOut	-0.005**	0.015***	0.016***	-0.000	0.032***	1.787***	0.222***	5.012***
	(0.002)	(0.002)	(0.001)	(0.000)	(0.003)	(0.101)	(0.014)	(0.612)
USPC FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,107,656	1,083,275	1,086,129	1,536,346	1,536,346	1,107,656	1,100,495	1,102,757
R-squared	0.027	0.124	0.132	0.002	0.073	0.034	0.058	0.062

Table 5. AIPA's effect on knowledge diffusion (intensive margin): "twin" analysis

This table reports DID estimates of AIPA's effect on the speed of knowledge diffusion, measured as the average time between the patent application date of the focal patent and its $1^{st}/3^{rd}/5^{th}/7^{th}$ forward citation. The sample consists of U.S. patents filed during 1998-2003 and their equivalent applications filed at the EPO. A U.S. application and its equivalent EPO applications together constitute a distinct family. Only patents that have at least 1/3/5/7 forward citations (excluding self-citations) within ten years after application are included. The regressions are estimated using the following specification:

$$Outcome_{ijt} = \alpha_1 + \alpha_2 US_j + \alpha_3 US_j * PostAIPA_t + \delta W_j + Family_i + Month_t + \epsilon_{ijt},$$

where j indicates the patent application belonging to family i and filed in the year t.. Standard errors are clustered by the application month for U.S. and EP patents, separately. ***, **, and * stand for statistical significance based on two-sided tests at the 1%, 5%, and 10% level, respectively.

	(1)	(2)	(3)	(4)
VARIABLES	Ln_Lag1Fcite	Ln_Lag3Fcite	Ln_Lag5Fcite	Ln_Lag7Fcite
US	-0.394***	-0.214***	-0.150***	-0.121***
	(0.011)	(0.012)	(0.014)	(0.019)
PostAIPA#US	-0.254***	-0.294***	-0.292***	-0.274***
	(0.013)	(0.016)	(0.023)	(0.028)
Granted	-0.201***	-0.178***	-0.147***	-0.139***
	(0.015)	(0.015)	(0.021)	(0.032)
EarlyGrant	-0.393***	-0.342***	-0.298***	-0.255***
	(0.015)	(0.021)	(0.031)	(0.043)
EarlyGrant#US	0.149***	0.079**	0.059	0.031
	(0.026)	(0.035)	(0.052)	(0.064)
Family FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	454,497	318,794	239,577	188,668
Adj R-squared	0.284	0.327	0.347	0.375

Table 6. AIPA's effect on knowledge diffusion (extensive margin): "twin" analysis

This table reports DID estimates of AIPA's effect on knowledge diffusion, measured by the number of forward citations. The sample consists of 316,563 successful U.S. applications filed during 1998-2003 and their equivalent applications filed at the EPO. A U.S. application and its equivalent EPO applications, together, constitute a distinct family. The regressions are estimated using the following specification:

$$Outcome_{ijt} = \alpha_1 + \alpha_2 US_j + \alpha_3 US_j * PostAIPA_t + \delta W_j + Family_i + Month_t + \epsilon_{ijt},$$

where j indicates the patent application belonging to family i and filed in month t, and W_j represents patent characteristics such as whether the patent is granted before 18 months. The dependent variable is the natural logarithm of one plus 3/5/7/10-year forward citations (excluding self-citations). We include patent family fixed effects and application month fixed effects; hence, the impact of AIPA is identified by the interaction term $US_j * PostAIPA_t$. In Panel B, we repeat the same regressions with the dependent variables as the forward citations made by subsequent U.S. and EP patents, respectively. For brevity, only the coefficient on the interactions are reported. Standard errors are clustered by the application month for U.S. and EP patents, separately. ***, **, and * stand for statistical significance based on two-sided tests at the 1%, 5%, and 10% level, respectively.

Panel A: Main analyses of forward citations

VARIABLES	(1) Ln_Fcite3Y	(2) Ln_Fcite5Y	(3) Ln_Fcite7Y	(4) Ln_Fcite10Y
US	0.806***	0.986***	1.093***	1.207***
	(0.006)	(0.008)	(0.008)	(0.009)
PostAIPA#US	-0.017**	0.057***	0.108***	0.147***
	(0.008)	(0.009)	(0.010)	(0.011)
Granted	0.191***	0.242***	0.276***	0.299***
	(0.006)	(0.006)	(0.007)	(0.007)
EarlyGrant	0.419***	0.516***	0.572***	0.606***
	(0.009)	(0.010)	(0.010)	(0.011)
EarlyGrant#US	-0.468***	-0.572***	-0.634***	-0.678***
	(0.010)	(0.014)	(0.016)	(0.017)
Family FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	670,142	669,708	668,373	659,620
Adj R-squared	0.450	0.511	0.541	0.568

Panel B: Forward Citations by U.S. or EP Patents, respectively

VARIABLES	(1) Ln_Fcite3Y	(2) Ln_Fcite5Y	(3) Ln_Fcite7Y	(4) Ln_Fcite10Y
US citations				
PostAIPA#US	-0.018**	0.048***	0.094***	0.133***
	(0.008)	(0.010)	(0.011)	(0.012)
EP citations				
PostAIPA#US	-0.004	0.018***	0.036***	0.048***
	(0.002)	(0.003)	(0.003)	(0.004)

Table 7. AIPA's effect on patent similarity: "twin" analysis

This table reports DID estimates of AIPA's effect on technological overlap. The sample consists of 316,563 successful U.S. applications filed during 1998-2003 and their equivalent applications filed at the EPO. A U.S. application and its equivalent EPO applications, together, constitute a distinct family. The regressions are estimated using the following:

$$Outcome_{ijt} = \alpha_1 + \alpha_2 US_j + \alpha_3 US_j * PostAIPA_t + \delta W_j + Family_i + Month_t + \epsilon_{ijt},$$

where j indicates the patent application, belonging to family i and filed in year t. Standard errors are clustered by the application month for U.S. and EP patents, separately. ***, **, and * stand for statistical significance based on two-sided tests at the 1%, 5%, and 10% level, respectively.

	(1)	(2)	(3)	(4)
VARIABLES	50th-IPC7Sim	75th-IPC7Sim	90th-IPC7Sim	95th-IPC7Sim
US	0.023***	0.052***	0.063***	0.058***
	(0.001)	(0.001)	(0.001)	(0.001)
PostAIPA#US	0.013***	0.010***	-0.008***	-0.013***
	(0.001)	(0.001)	(0.001)	(0.001)
Granted	0.009***	0.018***	0.010***	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)
EarlyGrant	-0.001	0.006***	0.008***	0.008***
	(0.002)	(0.002)	(0.002)	(0.002)
EarlyGrant#US	-0.002	-0.017***	-0.016***	-0.007***
	(0.002)	(0.002)	(0.002)	(0.002)
Family FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	669,029	669,029	669,029	669,029
Adj R-squared	0.615	0.683	0.698	0.700

Table 8. Patenting intensity: US v. EP comparison

This table reports Ordinary Least Squares (OLS) regression analysis of patenting intensity around the enactment of AIPA. In Column 1 (2), the dependent variable is the number of applications (applications that were eventually granted) filed in each month from 1998 to 2003 at the USPTO or EPO, respectively. In Columns 3-6, the dependent variable is the monthly count of applications or granted applications by technology class (IPC 4-digit code) filed at the USPTO or EPO, respectively. Robust standard errors are reported in parentheses. ***, **, and * stand for statistical significance based on two-sided tests at the 1%, 5%, and 10% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	#Application	#Grant	#Application	#Grant	#Application	#Grant
US	11,057.943***	9,881.829***	18.014***	16.091***	18.034***	16.103***
	(449.808)	(285.928)	(0.322)	(0.278)	(0.322)	(0.277)
PostAIPA	1,497.711***	326.296***	2.439***	0.530*		
	(225.671)	(110.951)	(0.335)	(0.293)		
US#PostAIPA	2,304.165***	1,158.847***	3.686***	1.850***	3.657***	1.838***
	(549.785)	(351.614)	(0.519)	(0.432)	(0.518)	(0.431)
Fixed Effects	No	No	IPC4	IPC4	IPC4, Month	IPC4, Month
Observations	144	144	88,746	88,746	88,746	88,746
R-squared	0.938	0.963	0.703	0.631	0.704	0.631

9. Online Appendix

Figure A1. Citations to pre-grant U.S. patent applications after AIPA

The figure plots the average percentage of citations to pre-grant applications, as opposed to granted patents, by the citing patents' application year. Panel A plots the fraction of citations to pre-grant applications for citations made by examiners and applicants in all citing patents filed during 2001-2012 (and granted by June 17, 2014). Panels B-G plot the fraction of citations to pre-grant applications for applicants and examiners for each of the one-digit NBER technology class.

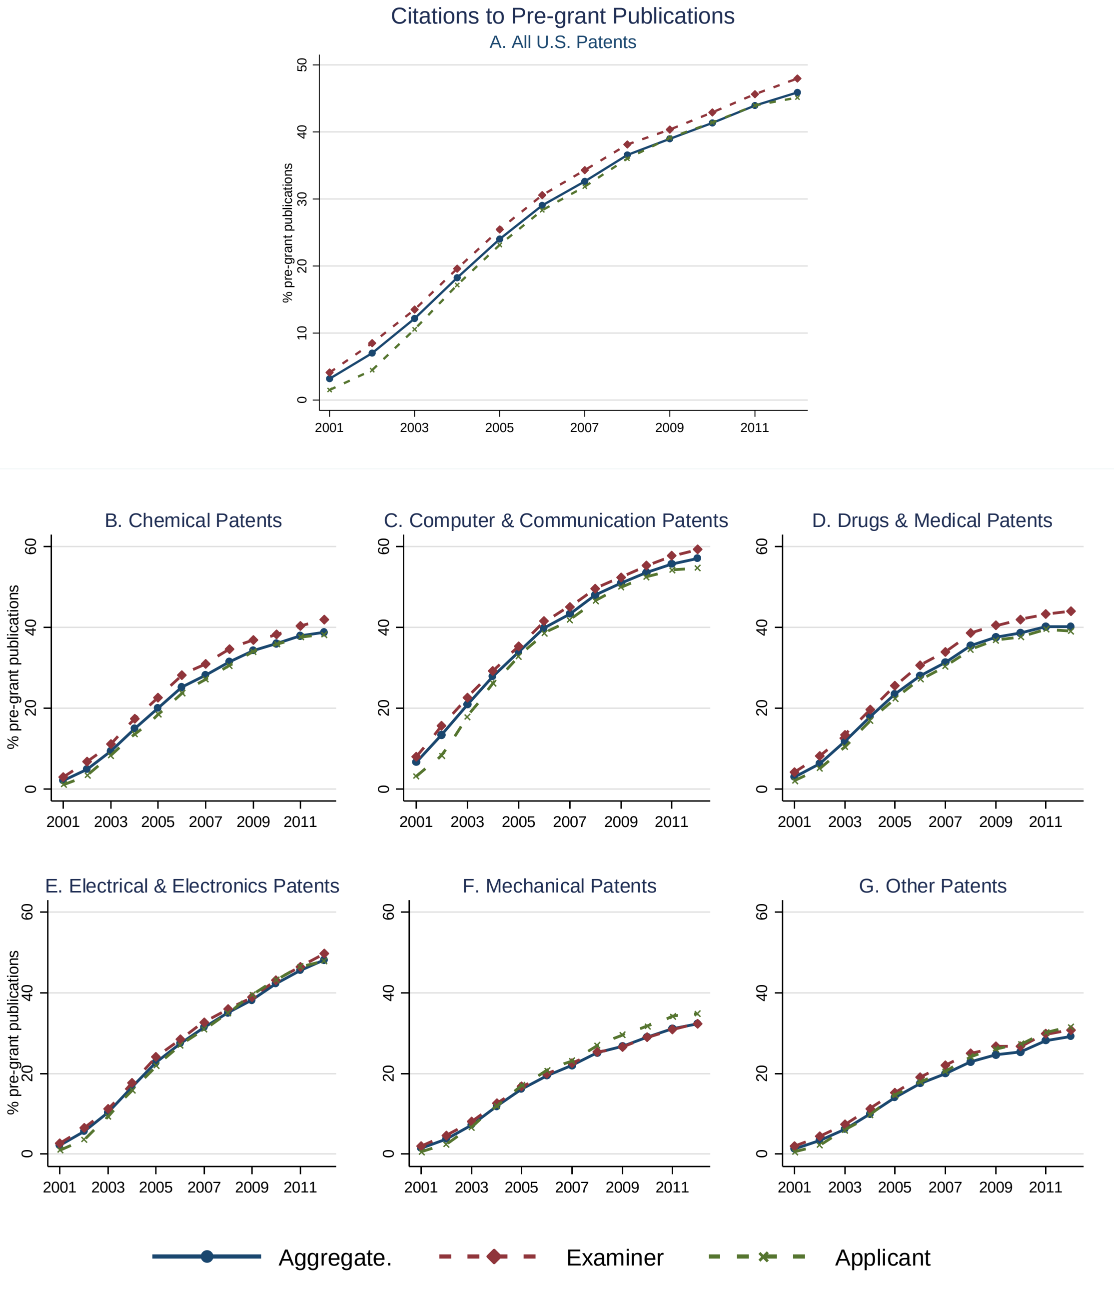

Figure A2. Cumulative forward citations before and after AIPA

The figures below plot the cumulative number of forward citations for U.S. patents filed in 2000 and 2001, respectively. The citation clock starts from the application date in Panel A, the disclosure date (publication date for patents with pre-grant publications and grant date for those without) in Panel B, and the grant date in Panel C.

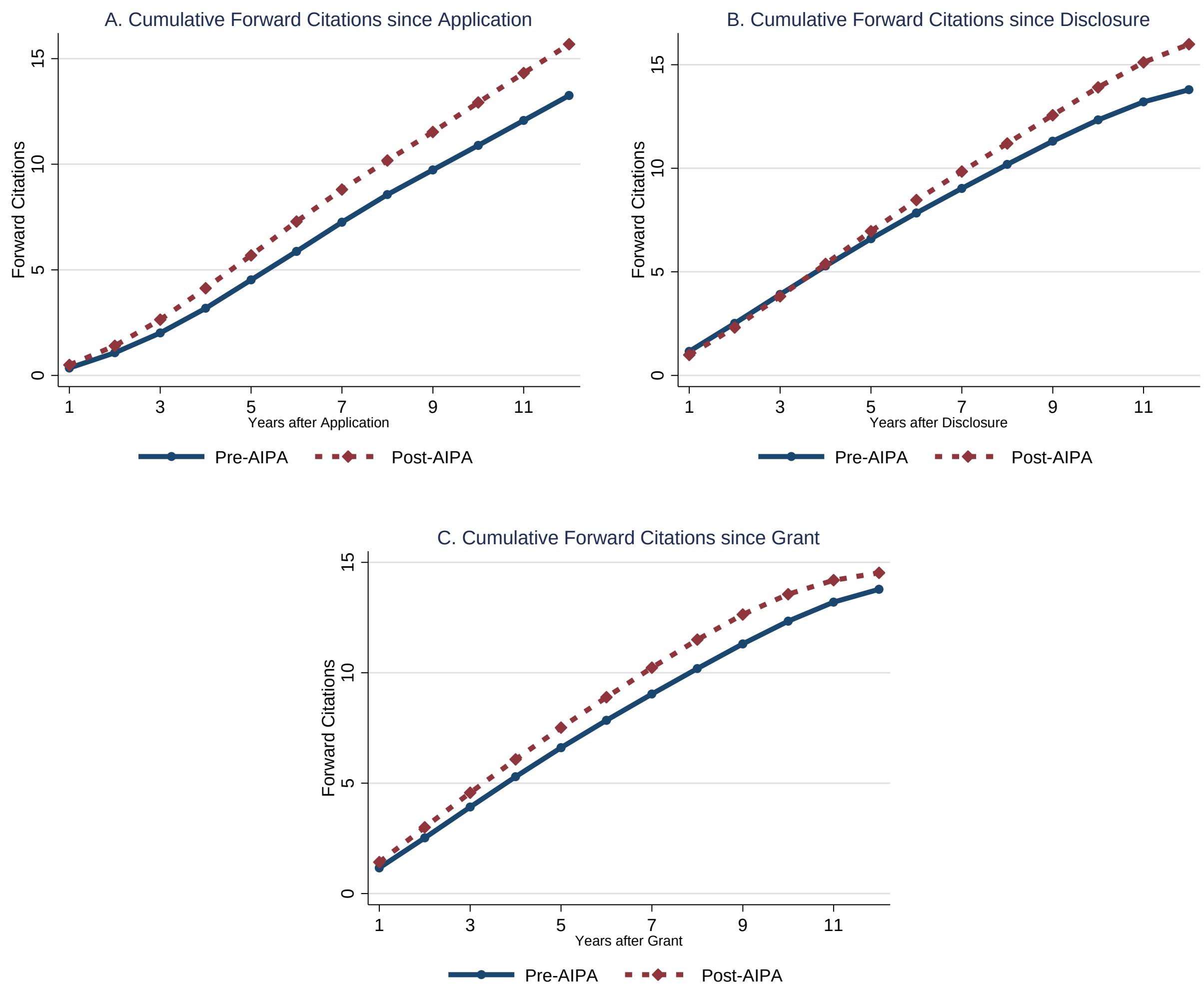

Figure A3. Technology Similarity Example

The figure illustrates an example of the similarity distribution between patent i and the next generation of patents (j=1,…,J), the 5th and 95th percentiles of patent i's similarity to the next distribution. The 5th percentile includes technologies in the next generation that are the least related to patent i. The 95th percentile includes technologies in the next generation that are most related to patent i.

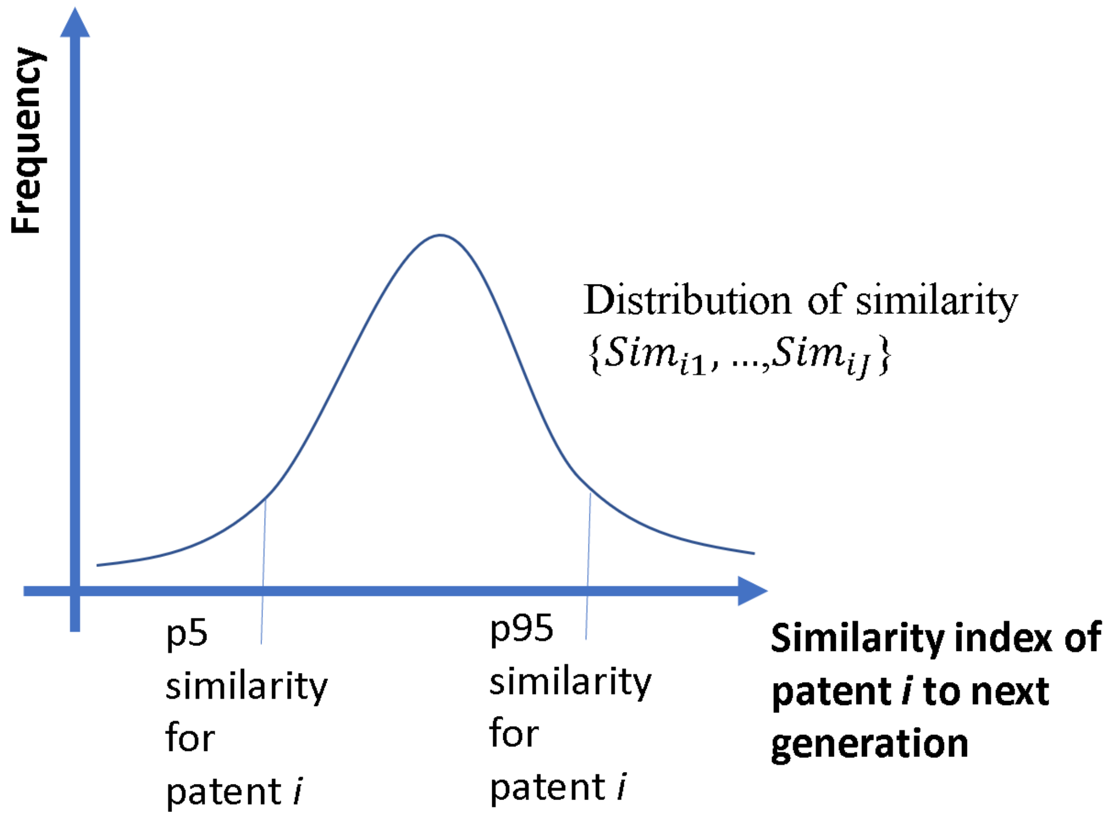

Figure A4. Comparison of patents with and without foreign parallel applications

The figures plot the citation lags, citation counts, and technology similarity for U.S. patents filed during each month during 1998-2003. We distinguish between patents with (*WF*) or without *(NF)* foreign parallel applications, which are identified from the patent family table from PATSTAT (Spring 2017 version). In Part I, we focus on citation lags, which are measured as the average number of months between the application date of a focal patent and the application dates of its 1st, 3rd, 5th, or 7th forward citations. Only patents that have accumulated the required number of forward citations within ten years after application are included. In Part II, we focus on forward citations, which are counted cumulatively 3/5/7/10 years after patent disclosure (i.e., publication date for patents with pre-grant publications or grant date for those without). In Part III, we focus on technology similarity. The vertical dashed line represents AIPA's effective date (November 29, 2000).

Part I: citation lags

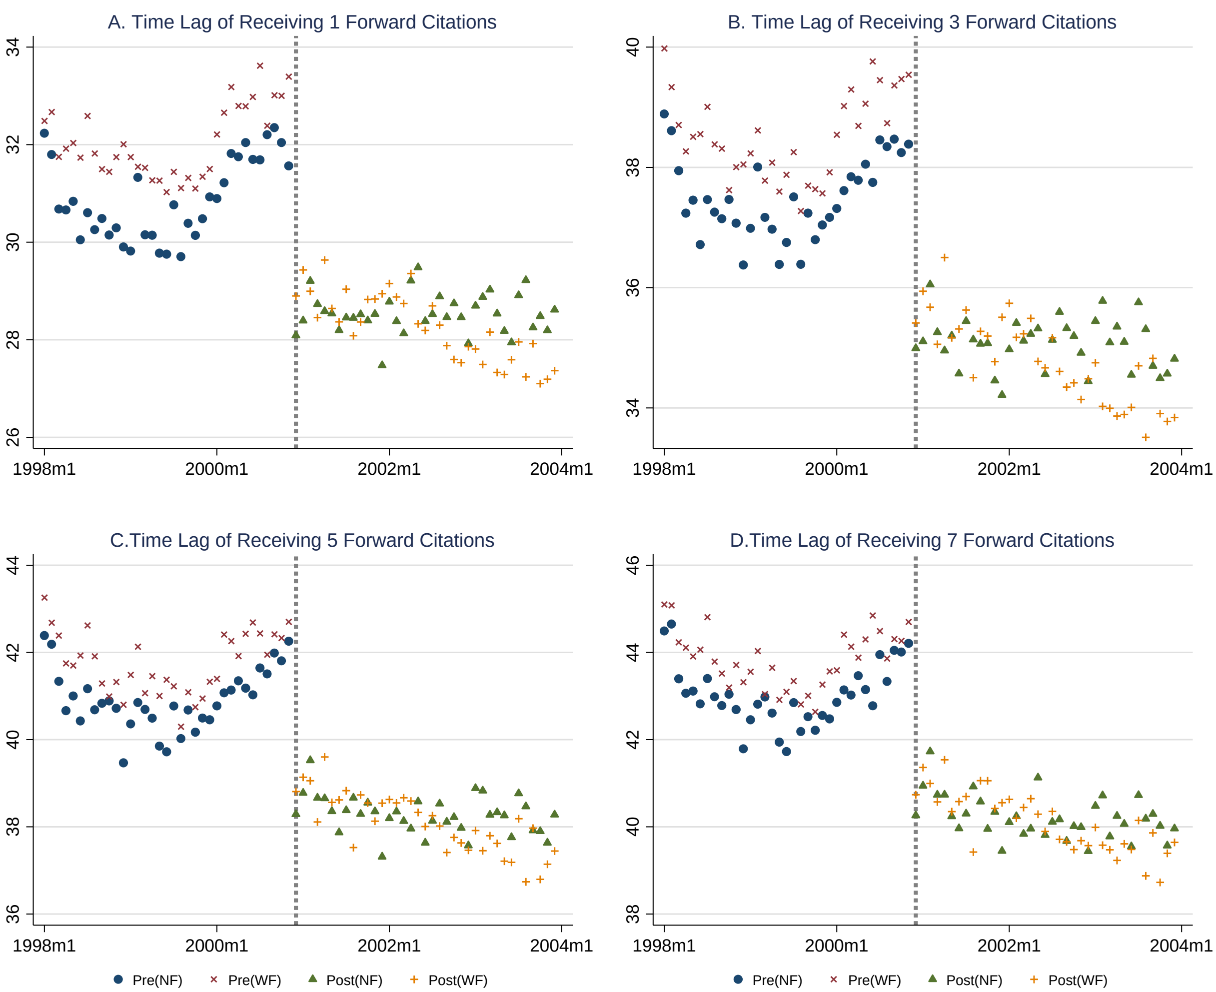

Part II: forward citations

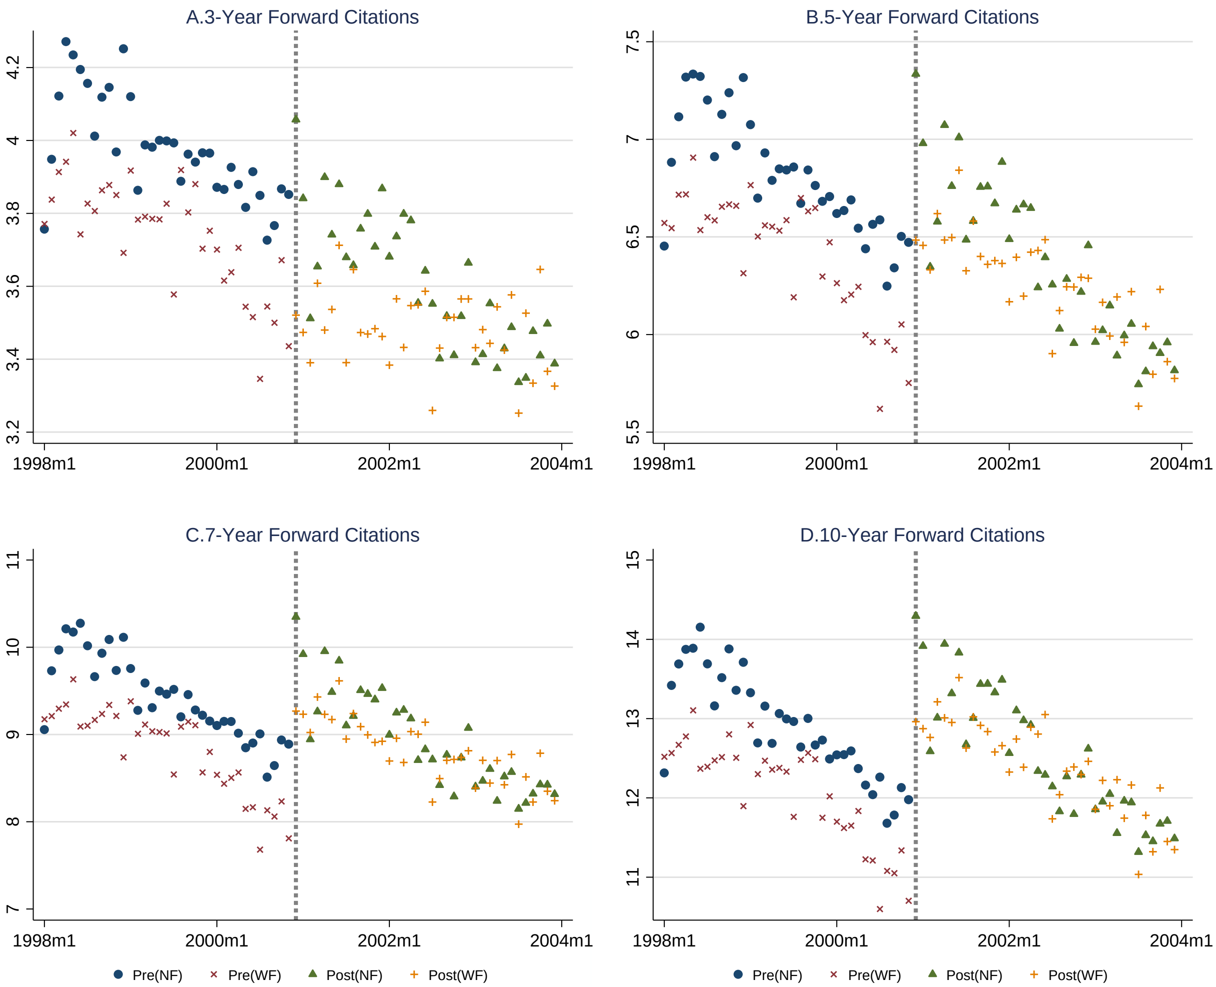

Part III: technology similarity

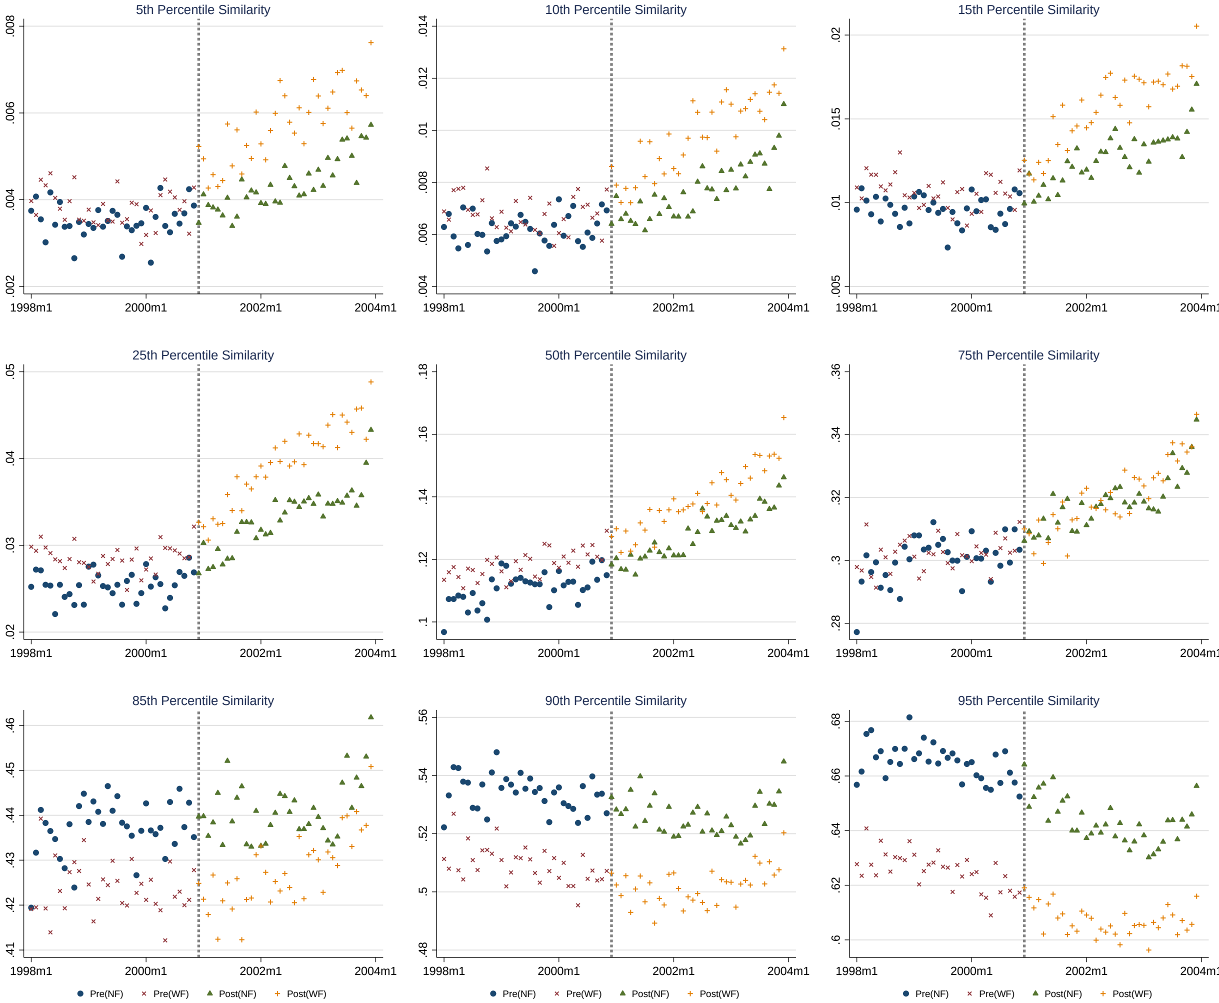

Figure A5. U.S. patentees' propensity to file for foreign or EP parallel applications

The figure below plots the percent of patents filed at the USPTO that file parallel applications in the EPO or any foreign patent office. All U.S. applications filed from 1998 to 2003 that are eventually granted by mid-2014 are included. Foreign or EP parallel applications, identified from the patent family table from PATSTAT (2017 Spring version), are required to be filed within 18 months of the application of their associated U.S. applications.

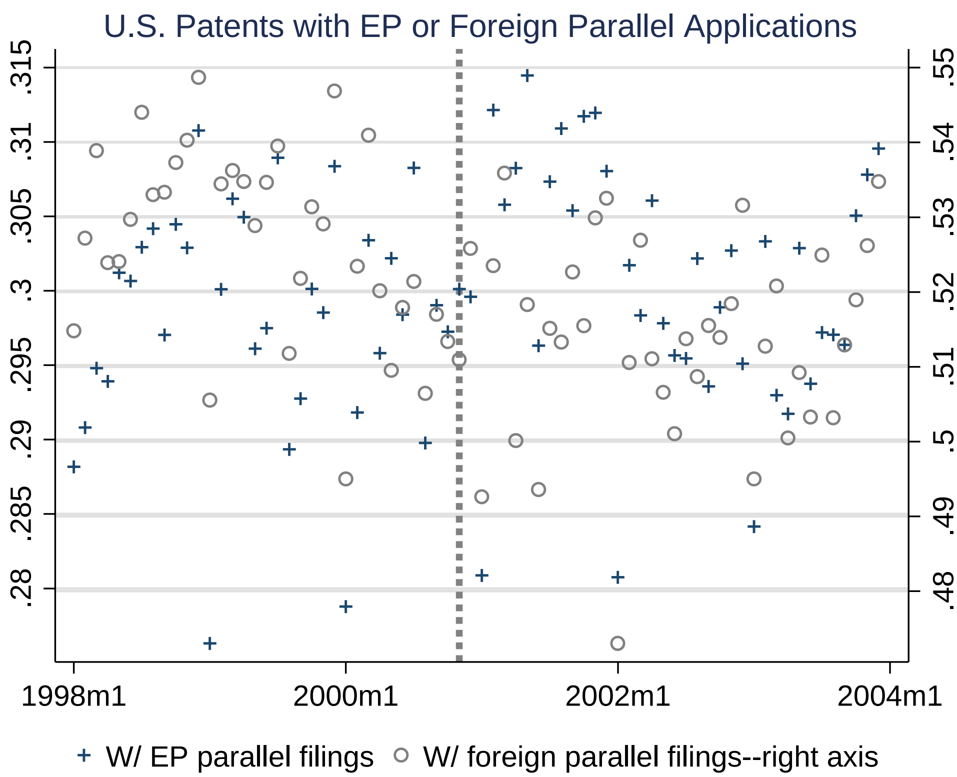

Figure A6. Adjusted citations to U.S. patents and EPO “twins” before and after AIPA

The figures plot the monthly average number of forward citations (excluding self-citations) to U.S. patents and their equivalent “twins” at the European Patent Office (EPO). Forward citations are counted cumulatively 3/5/7/10 years after patent disclosure (i.e., publication date for patents with pre-grant publications and grant date for those without). We adjust the citation count by the average number of citations received by patents of the same technology class (NBER 2-digit code) filed in 1998. We then take the natural logarithm of one plus the adjusted citation count. The solid (dashed) line represents an OLS regression line fit to all pre-AIPA U.S. (EP) patents’ monthly adjusted citations. The vertical line represents the effective date of AIPA.

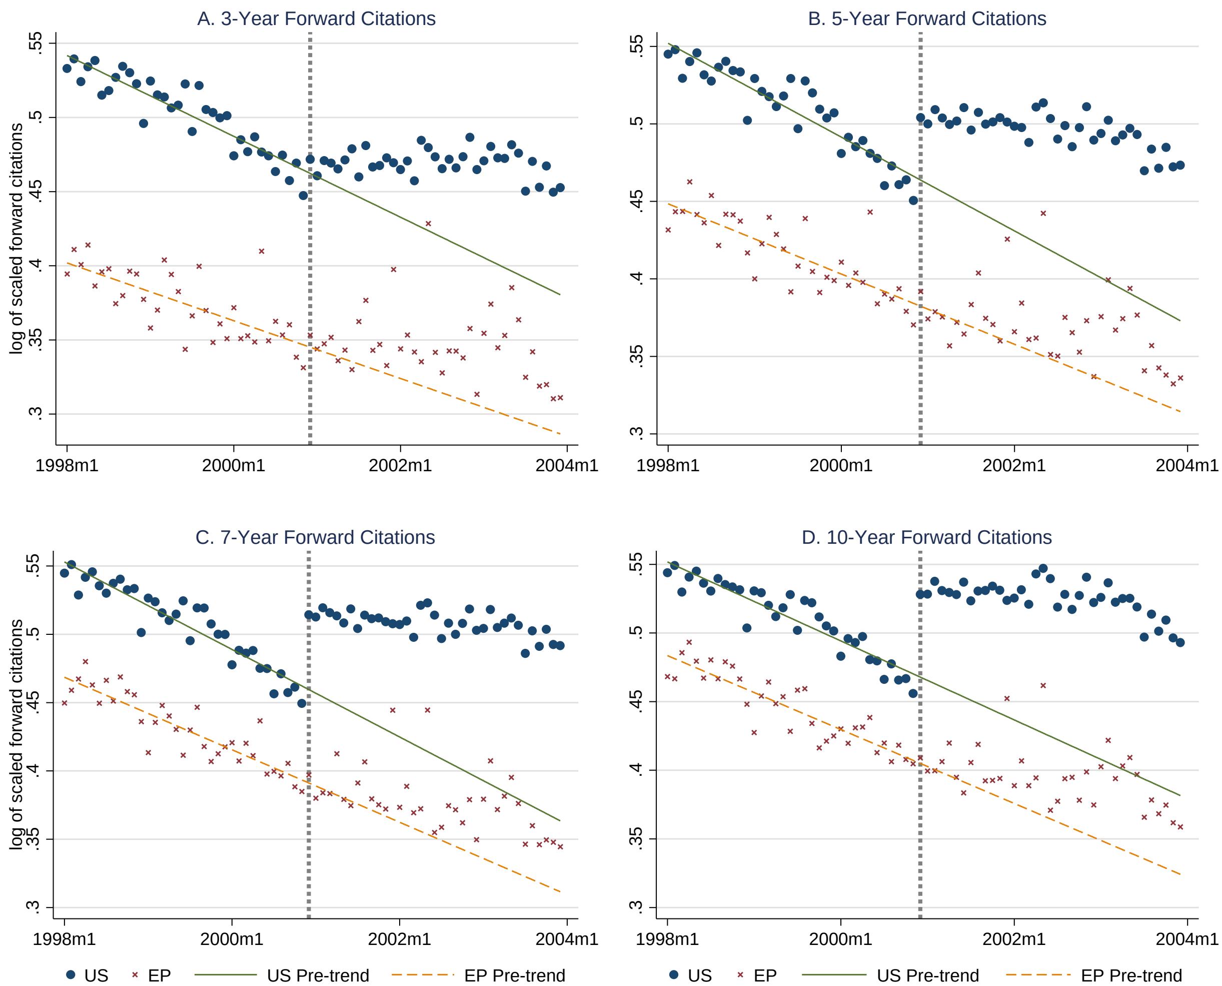

Figure A7. Patenting intensity (by technology class)

The figures in Panel A plot the number of patent applications filed during each month of 1998-2003 by one-digit IPC technology class at the USPTO and EPO, respectively. The figures in Panel B plot the number of patent applications that are eventually granted. We define technology class by the one-digit IPC code rather than by the one-digit NBER technology code because EPO patents are not assigned an NBER technology code and there is no one-to-one mapping between the IPO code and the NBER code. The vertical dashed line represents AIPA's effective date (November 29, 2000).

Panel A. Number of patent applications

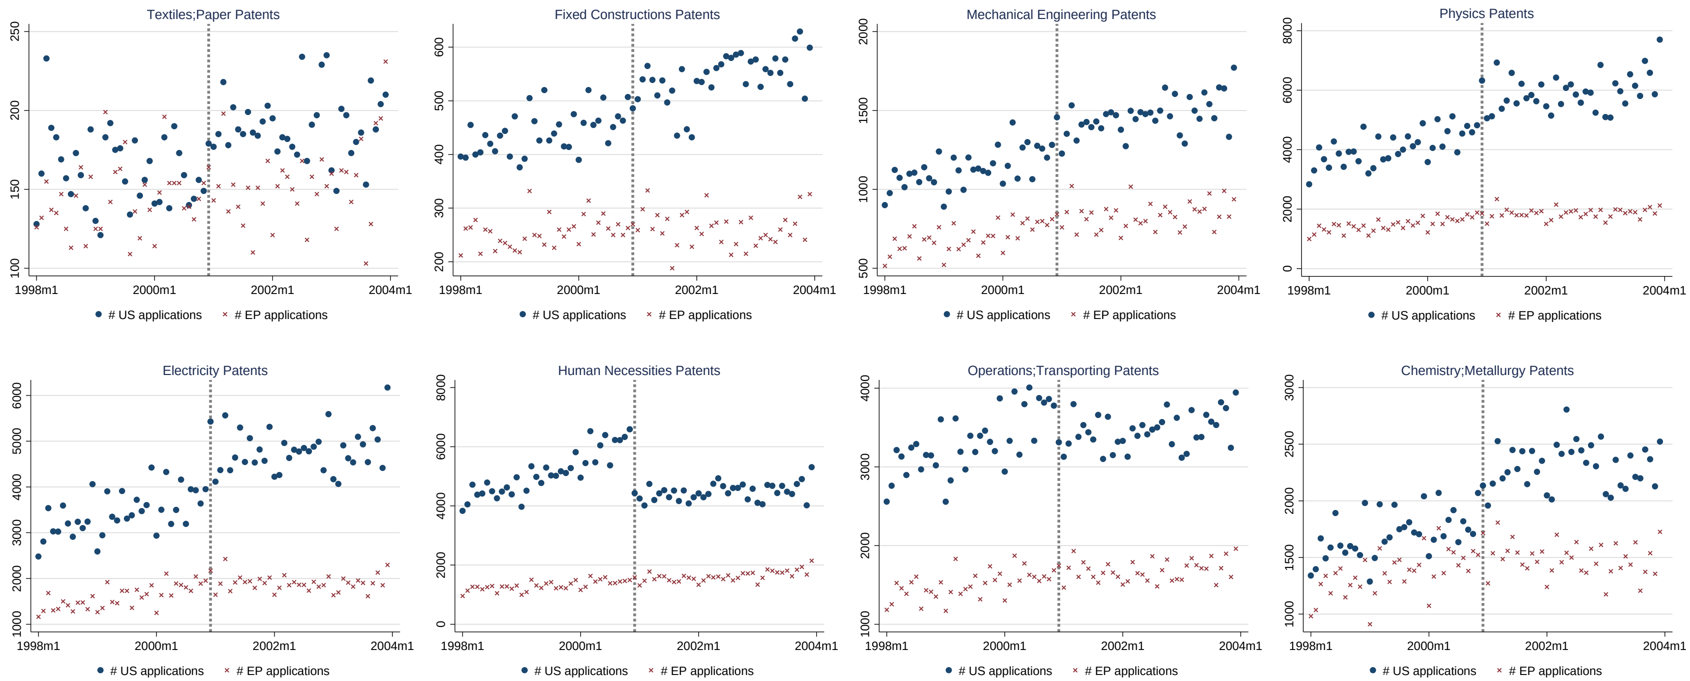

Panel A. Number of granted patent applications

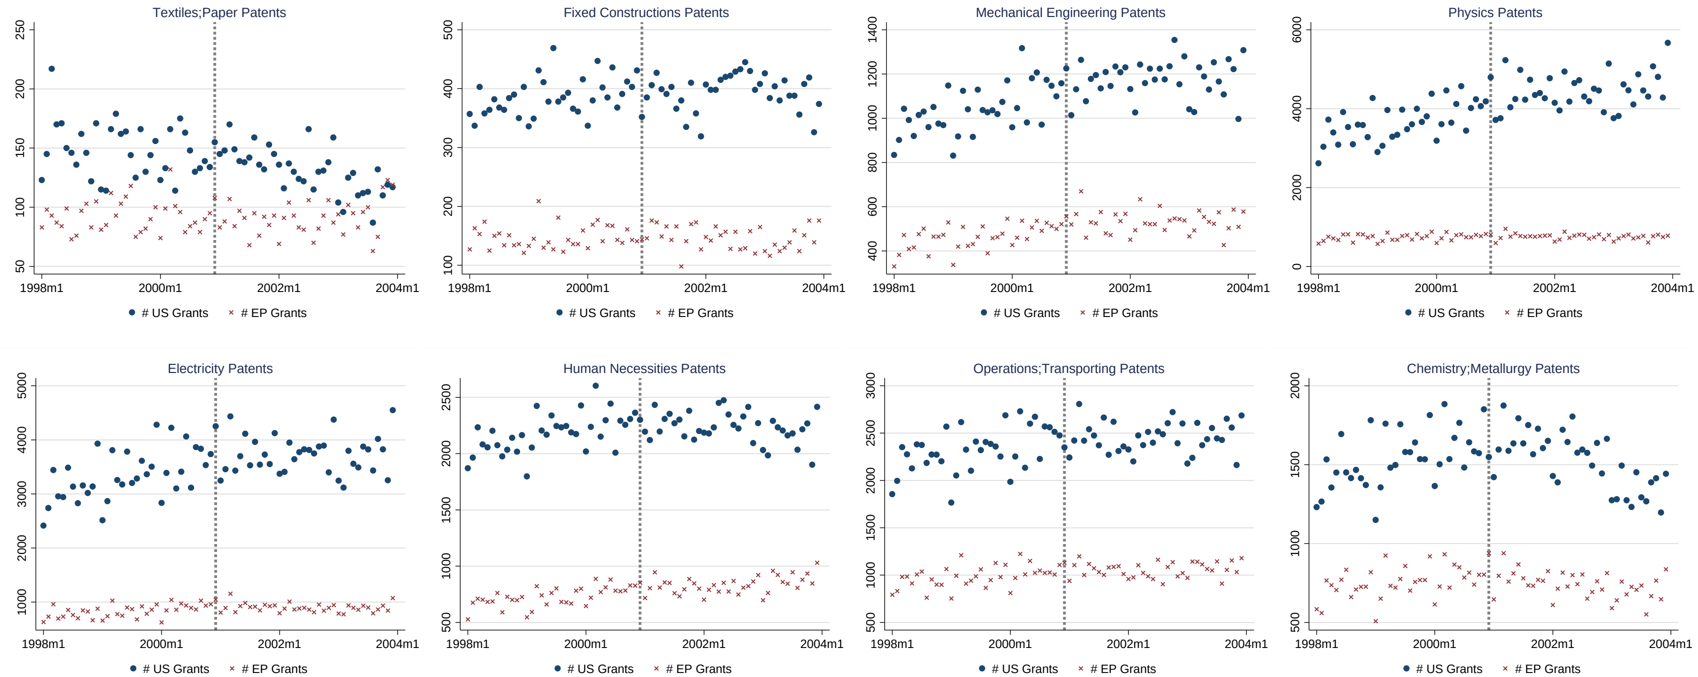

Table A1. Citations to post-AIPA abandoned patents

This table reports the forward citations to 235,530 abandoned applications that were filed in the USPTO from November 29, 2000 to December 31, 2003. We exclude 27,113 abandoned applications due to missing data on technology classes. Applications without foreign parallel applications had the option to opt out of the pre-grant publication requirement, and, on average 10.32% of applications in the sample exercised the option. For the subsample of abandoned applications that are disclosed, we report their forward citations received in 3/5/7/10 years in the right panel of the table.

				Abandoned applications with pre-grant publications			
Tech Class	**#ABN-Pub**	**#ABN-Secrecy**	**%ABN-Secrecy**	**Fcite3Y**	**Fcite5Y**	**Fcite7Y**	**Fcite10Y**
Chem.	29,774	1,906	6.02%	0.106	0.527	1.256	3.006
Cmp. & Comm	56,601	8,100	12.52%	0.090	0.669	2.249	6.969
Drug & Med.	50,892	3,904	7.12%	0.095	0.509	1.351	4.475
Elec.	28,669	2,862	9.08%	0.148	0.729	1.712	3.876
Mech.	27,844	3,421	10.94%	0.111	0.499	1.106	2.463
Other	41,750	6,922	14.22%	0.088	0.425	0.988	2.542
Total	235,530	27,115	10.32%	0.102	0.563	1.522	4.293

Table A2. Validation of patent classification-based patent similarity measure

This validation exercise of the technological similarity measure defined in Section 4.1 draws on the institutional feature that citations made by the EPO are classified based on the relationship between the cited and citing patents. The most important types of citations are X- and Y-citations, which account for 21% and 20% of all the citations categorized in PATSTAT, respectively. An X- or Y-citation indicates that at least one claim in the citing patent cannot be considered novel or does not involve an inventive step, either taking the cited patent alone or combining it with other cited documents. Thus, X- and Y-citations are particularly relevant to the citing patent. In comparison, other types of citations provide mainly general background information. For example, the most common citations (comprising 49% of all citations) are A-citations, which merely define the general state of the art. Therefore, if our similarity measure captures the fundamental proximity of the patented invention, we should observe a higher *Sim* (defined in Section 4.1) for X- and Y-citations than for other types of citations.

For comparison, we also match the cited or the citing patent involved in an X- or Y-citation to a random patent that is not cited by or citing any patents in the citation pair in question. The matching is based on application quarter, IPC 4-digit code, and grant status. We then construct the *Sim* for the fake citation, where the matched patent replaces either the cited or the citing patent. We expect *Sim* to be even lower for these fake citations.

To reduce computational burden, we restrict the cited patents to be filed from 1998 to 2003 and citing patents from 1998 to 2009. This restriction yields a sample of 141,582 X-citations and 50,445 Y-citations from PATSTAT. We group them together (labeled as *Important Citations*). We then keep the citing patent in the important citations constant and construct its similarity to its other types of citations (labeled as *Background Citations*) and to a patent matched to the cited patent (labeled as "*Fake Citations"*). Given that the EPO makes only four backward citations per application, on average, the requirement for the citing patent to have at least one important citation and one background citation that are both filed from 1998 to 2003 reduces the sample substantially. For this reason, we summarize and compare *Sim* for the whole sample of important citations and the reduced sample, separately. In a similar vein, we also compare *Sim* across important citations, background citations, and fake citations, while keeping the cited patent fixed.

The results are reported below. Holding the citing patent fixed, the average *Sim* for the entire sample of important citations is 0.594, more than three times as large as that of faked citations. When we focus on the subsample in which the citing patent has at least one important citation and one background citation, we observe a slight decrease in *Sim* from important citations to background citations, a further decrease from background citations to fake citations. The pattern is similar when we compare *sim* while keeping the cited patent constant, as reported in Panel B of this table. Collectively, the above evidence demonstrates that our similarity measure exhibits variations that are consistent with the technological overlap identified by patent examiners at the EPO.

Panel A: Compare *Sim* between Important, Background, and Fake References (Keep Citing Patent Fixed)

	1	2	3	4	5	6	7	8	
	All Important References			**Require the citing patent in the important reference to have at least one background references filed from 98 to 03**					
	Important	Fake	Dif(1-2)	Important	Background	Dif(4-5)	Fake	Dif(4-7)	Dif(5-7)
#	191,358	191,358		32,141	32,141		32,141		
Mean	0.594	0.139	0.455***	0.615	0.592	0.024***	0.147	0.468***	0.444***
S.E.	0.001	0.001	0.001	0.002	0.002	0.002	0.002	0.002	0.002

Panel B: Compare *Sim* between Important, Background, and Fake References (Keep Cited Patent Fixed)

	All Important References			**Require the cited patent in the important reference to be cited as background references by patents filed from 98 to 09**					
	Important	Fake	Dif(1-2)	Important	Background	Dif(4-5)	Fake	Dif(4-7)	Dif(5-7)
#	191,490	191,490		133,982	133,982		133,982		
Mean	0.594	0.160	0.434***	0.596	0.579	0.017***	0.162	0.434***	0.417***
S.E.	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001

Table A3. Before-and-after analyses of citation lags with non-linear trends

This table reports the before-and-after analyses of AIPA's effect with non-linear pre-trends. The regressions are specified as follows,

$$Outcome_{it} = \alpha_0 + \sum_{\tau \in Post} \beta_t * I\{Month_t = \tau\} + f(Month_t) + \alpha_2 EarlyGrant_i + \alpha_3 OptOut_i + TechFE + \epsilon_{it},$$

where i indicates the application filed in month t. *Month* is the calendar month in which the patent is filed. The specification is the same as the ones in Table 3, except that the linear pre-trend is replaced by non-linear trends specified by second-order polynomials of *Month* in Panel A and third-order in Panel B. Standard errors are clustered by the application month. *Pre-Post-Dif* ($\frac{1}{37}\sum_{\tau \in Post} \beta_t$), which is computed as the average β_t with its associated standard errors computed using the delta method. ***, **, and * stand for statistical significance based on two-sided tests at the 1%, 5%, and 10% level, respectively.

Panel A: With second-order polynomials of *Month*

	(1)	(2)	(3)	(4)
VARIABLES	Ln_Lag1Fcite	Ln_Lag3Fcite	Ln_Lag5Fcite	Ln_Lag7Fcite
Pre-Post-Dif ($\frac{1}{37}\sum_{\tau \in Post} \beta_t$)	-0.414***	-0.330***	-0.315***	-0.300***
	(0.025)	(0.022)	(0.022)	(0.023)
Month	-0.155***	-0.135***	-0.137***	-0.133***
	(0.018)	(0.015)	(0.015)	(0.016)
Month^2	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
EarlyGrant	-0.316***	-0.218***	-0.183***	-0.168***
	(0.003)	(0.003)	(0.003)	(0.003)
Optout	0.276***	0.212***	0.189***	0.174***
	(0.007)	(0.006)	(0.005)	(0.005)

Panel B: With third-order polynomials of *Month*

	(1)	(2)	(3)	(4)
VARIABLES	Ln_Lag1Fcite	Ln_Lag3Fcite	Ln_Lag5Fcite	Ln_Lag7Fcite
Pre-Post-Dif ($\frac{1}{37}\sum_{\tau \in Post} \beta_t$)	-0.163***	-0.124***	-0.117***	-0.110***
	(0.008)	(0.007)	(0.007)	(0.007)
Month	-3.339***	-2.709***	-2.594***	-2.471***
	(0.264)	(0.227)	(0.226)	(0.241)
Month^2	0.007***	0.006***	0.005***	0.005***
	(0.001)	(0.000)	(0.000)	(0.000)
Month^3	-0.000***	-0.000***	-0.000***	-0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
EarlyGrant	-0.316***	-0.218***	-0.183***	-0.168***
	(0.003)	(0.003)	(0.003)	(0.003)
Optout	0.276***	0.212***	0.189***	0.174***
	(0.007)	(0.006)	(0.005)	(0.005)

Table A4. U.S. patents and EPO equivalents—matching and sample selection

The table describes the sample selection process. We start with all U.S. applications filed between January 1, 1998 and December 31, 2003, regardless of grant status. We match these U.S. applications to applications filed at the European Patent Office (EPO) according to the simple family member table in PATSTAT, where family is defined as patents that share the same priority. We further split the matched U.S. applications into four sub-groups, based on whether or not the U.S. applications and EPO parallel applications are granted. U.S. applications that are not matched to any EPO application are split into two groups depending on their grant status. Abandoned U.S. patent applications filed before AIPA cannot be matched to EPO applications since the USPTO does not publish such applications.

Note that, in total, 403,292 granted U.S. applications are identified with EP equivalents, among which 264,651 (138,641) U.S. patents have granted (ungranted) EP equivalents. By "ungranted," we mean that the patent application is still pending, already withdrawn by its applicant, or rejected by its patent examiner. Nearly 34.38% (138,641/403,292) of granted U.S. applications do not get granted EP counterparts, and the discrepancy of grant statuses across the two patent systems is driven largely by "Computers and Communications" patents. 22,076 (41,506) ungranted U.S. applications are matched with granted (ungranted) EP equivalents. These ungranted U.S. applications are all published through the 18-month disclosure. The majority of ungranted U.S. applications cannot be matched to any EP applications because they are kept secret (pre-AIPA applications or post-AIPA applications that opt out of the 18-month publication requirement). We exclude ungranted U.S. applications, even if they can be matched to EP equivalents, to construct a balanced pre- and post-AIPA sample. To deal with the concern that distinct grant statuses in the USPTO and EPO might drive the results, we also check the robustness of our results using U.S. patents and EP equivalents that are both granted. Our results hold in this smaller sample.

	Grant Status	Whole Sample		% By Technology Class					
Match Outcome	US & EPO	Freq.	Percent	Chem.	Cmp. & Comm	Drug & Med.	Elec.	Mech.	Other
Matched	Grant & Grant	264,651	17.23	19.12	18.21	15.79	15.34	17.50	13.67
	Grant & No-grant	138,641	9.02	16.04	24.71	15.96	20.01	12.09	10.97
	No-grant & Grant	22,076	1.44	19.46	13.42	32.85	7.20	12.68	14.40
	No-grant & No-grant	41,506	2.7	15.44	24.23	28.32	10.42	9.65	11.95
Unmatched	Grant & N/A	704,377	45.85	9.41	28.14	8.10	23.35	15.00	15.76
	No-grant & N/A	365,095	23.77	11.85	24.17	17.65	13.17	12.74	20.42
	Total	1,536,346	100	12.57	24.86	13.30	18.67	14.45	15.95

Table A5. Summary statistics for sample of U.S. patents and EP "twins"

This table reports the univariate analyses for the main variable of interest. The *DID* is calculated as the difference in means of (US Post-US Pre) –(EP Post-EP Pre). *Pre* refers to applications filed before AIPA's effective date, November 29, 2000, while *Post* refers to applications filed after AIPA. Detailed definitions can be seen in Table 1.

Variable	**US Pre**	**US Post**	**EP Pre**	**EP Post**	**DID**	**T-Stat**
Fcite3Y	4.904	4.625	0.901	0.793	-0.171	-5.439
Fcite5Y	8.409	8.312	1.510	1.272	0.141	2.613
Fcite7Y	11.601	11.783	1.984	1.646	0.521	6.855
Fcite10Y	15.936	16.800	2.510	2.078	1.296	11.830
Fcite3Y(US)	4.329	4.026	0.240	0.165	-0.228	-8.059
Fcite5Y(US)	7.479	7.282	0.424	0.281	-0.055	-1.112
Fcite7Y(US)	10.399	10.382	0.573	0.374	0.182	2.596
Fcite10Y(US)	14.334	14.820	0.727	0.479	0.734	7.166
Fcite3Y(EP)	0.170	0.200	0.415	0.462	-0.017	-3.484
Fcite5Y(EP)	0.293	0.338	0.721	0.729	0.037	5.080
Fcite7Y(EP)	0.384	0.440	0.957	0.933	0.079	8.592
Fcite10Y(EP)	0.491	0.563	1.219	1.145	0.145	12.778
Lag1Fcite	31.368	25.220	46.688	44.626	-4.086	-24.471
Lag3Fcite	37.974	32.482	46.328	45.815	-4.979	-26.756
Lag5Fcite	41.197	35.954	46.270	46.201	-5.174	-23.321
Lag7Fcite	43.217	38.103	46.813	46.537	-4.838	-18.324
5th-IPC7Sim	0.004	0.005	0.004	0.004	0.001	9.125
10th-IPC7Sim	0.007	0.009	0.007	0.007	0.002	9.011
15th-IPC7Sim	0.011	0.014	0.010	0.011	0.004	13.398
25th-IPC7Sim	0.026	0.036	0.021	0.022	0.008	18.860
50th-IPC7Sim	0.108	0.125	0.080	0.085	0.011	12.601
75th-IPC7Sim	0.289	0.303	0.229	0.236	0.007	4.955
85th-IPC7Sim	0.410	0.410	0.342	0.345	-0.002	-1.505
90th-IPC7Sim	0.496	0.486	0.424	0.427	-0.012	-7.989
95th-IPC7Sim	0.611	0.591	0.544	0.542	-0.017	-11.818
# Patents	151177	165386	165075	189152		

Table A6. DID analysis with linear pre-trends

This table reports the DID analyses of AIPA's effect on citation lags (Panel A), citation counts (Panel B), and technology similarity (Panel C) using the US-EP twin sample. The regressions are specified as follows:

$$Outcome_{ijt} = \alpha_0 + \alpha_1 Month_t + \alpha_2 US_j + \delta W_j + \sum_{\tau \in Post} \beta_t I\{Month_t = \tau\} + \sum_{\tau \in Post} \gamma_t US_j * I\{Month_t = \tau\} + \epsilon_{ijt},$$

where j indicates the patent application belonging to family i and filed in month t, and W_j represents patent characteristics such as *Granted*—whether the EP patent is granted (all U.S. patents in this sample are granted by sample construction)—and *EarlyGrant*—whether the patent is granted before 18 months. *Month* is the calendar month in which the patent is filed, and *US* indicates whether the patent is filed in the USPTO. We control for a linear pre-trend and include a set of dummy variables indicating each month in the post-AIPA period to identify the impact of AIPA. We compute the *AIPA-Effect* by taking the mean of γ_t, and its associated standard errors are computed using the delta method. Standard errors are clustered by the application month for U.S. and EP patents, separately. ***, **, and * stand for statistical significance based on two-sided tests at the 1%, 5%, and 10% level, respectively.

Panel A: Citation Lags

	(1)	(2)	(3)	(4)
VARIABLES	Ln_Lag1Fcite	Ln_Lag3Fcite	Ln_Lag5Fcite	Ln_Lag7Fcite
Month	-0.001**	-0.000	-0.000	-0.000
	(0.001)	(0.000)	(0.000)	(0.000)
US	-0.406***	-0.203***	-0.125***	-0.092***
	(0.010)	(0.008)	(0.008)	(0.010)
Granted	-0.045***	-0.017**	0.004	0.006
	(0.009)	(0.008)	(0.010)	(0.013)
EarlyGrant	-0.227***	-0.159***	-0.122***	-0.101***
	(0.012)	(0.013)	(0.017)	(0.025)
EarlyGrant#US	0.033*	0.012	0.000	-0.003
	(0.018)	(0.017)	(0.020)	(0.027)
CommonDif ($\frac{1}{37}\sum_{\tau \in Post} \beta_t$)	0.008	0.007	0.017	0.001
	(0.021)	(0.016)	(0.016)	(0.018)
AIPA-Effect($\frac{1}{37}\sum_{\tau \in Post} \gamma_t$)	-0.148***	-0.171***	-0.173***	-0.153***
	(0.010)	(0.007)	(0.008)	(0.010)
Observations	454,497	318,794	239,577	188,668
Adj R-squared	0.080	0.048	0.033	0.027

Panel B: Citation Counts

	(1)	(2)	(3)	(4)
VARIABLES	Ln_Fcite3Y	Ln_Fcite5Y	Ln_Fcite7Y	Ln_Fcite10Y
Month	-0.001**	-0.002**	-0.002***	-0.002***
	(0.001)	(0.001)	(0.001)	(0.001)
US	0.853***	1.039***	1.149***	1.264***
	(0.008)	(0.010)	(0.010)	(0.011)
Granted	0.064***	0.107***	0.143***	0.183***
	(0.004)	(0.004)	(0.005)	(0.006)
EarlyGrant	0.107***	0.151***	0.186***	0.216***
	(0.006)	(0.008)	(0.009)	(0.011)

EarlyGrant#US	-0.254***	-0.334***	-0.385***	-0.435***
	(0.008)	(0.012)	(0.014)	(0.016)
CommonDif ($\frac{1}{37}\sum_{\tau\in Post}\beta_t$)	0.028	0.025	0.035	0.038*
	(0.018)	(0.021)	(0.021)	(0.021)
AIPA-Effect($\frac{1}{37}\sum_{\tau\in Post}\gamma_t$)	-0.036***	0.030***	0.078***	0.113***
	(0.008)	(0.010)	(0.010)	(0.010)
Observations	670,142	669,708	668,373	659,620
Adj R-squared	0.215	0.256	0.281	0.305

Panel C: Technology Similarity

	(1)	(2)	(3)	(4)
VARIABLES	50th-IPC7Sim	75th-IPC7Sim	90th-IPC7Sim	95th-IPC7Sim
Month	-0.001**	-0.000	-0.000	-0.000
	(0.001)	(0.000)	(0.000)	(0.000)
US	-0.406***	-0.203***	-0.125***	-0.092***
	(0.010)	(0.008)	(0.008)	(0.010)
Granted	-0.045***	-0.017**	0.004	0.006
	(0.009)	(0.008)	(0.010)	(0.013)
EarlyGrant	-0.227***	-0.159***	-0.122***	-0.101***
	(0.012)	(0.013)	(0.017)	(0.025)
EarlyGrant#US	0.033*	0.012	0.000	-0.003
	(0.018)	(0.017)	(0.020)	(0.027)
CommonDif ($\frac{1}{37}\sum_{\tau\in Post}\beta_t$)	-0.000	0.006	0.013***	0.010*
	(0.002)	(0.004)	(0.005)	(0.005)
AIPA-Effect($\frac{1}{37}\sum_{\tau\in Post}\gamma_t$)	0.011***	0.007***	-0.012***	-0.017***
	(0.001)	(0.002)	(0.002)	(0.002)
Observations	454,497	318,794	239,577	188,668
Adj R-squared	0.010	0.013	0.012	0.011

10. Theoretical Appendix

Our proofs assume a general nested-CES function form:

$$V^p(z, \Delta, Z_0) = \left((z^\alpha + \Delta^\alpha)^{\frac{\rho}{\alpha}} + Z_0^\rho\right)^{\frac{1}{\rho}}$$

The results in the text are a special case where $\alpha = 1$.

Single-Crossing and Regularity Assumptions (A1): there exist $\underline{z}$ and $\overline{z}$ such that $V^c > V^p(\underline{z}, \Delta(\underline{z}, Z_0), Z_0)$ and $V^p(\overline{z}, \Delta(\overline{z}, Z_0), Z_0) > V^c$, and $V^p(\cdot)$ is differentiable and monotone increasing in each of its arguments.

Lemma 1: Under A1, a unique interior patenting threshold z_p exists and is monotone decreasing in public knowledge Z_0.

Proof: Under assumption $A1$, z_p is unique, and under monotonicity, in conjunction with $\frac{\partial V^c}{\partial Z_0} = 0$, z_p declines whenever Z_0 increases.

Under the functional form assumption $c(\Delta) = \frac{\Delta^\gamma}{\gamma}$, where $\gamma > 1$, we can characterize the investment choice.

Lemma 2: Suppose that A1 holds, $\gamma > \alpha$, $\gamma > \rho$, and suppose that $\rho > 1$. Then, patent scope, Δ, is decreasing in the stock of public knowledge Z_0.

Proof: The firm optimization problem is given by,

$$\max_\Delta q\left((z^\alpha + \Delta^\alpha)^{\frac{\rho}{\alpha}} + Z_0^\rho\right)^{\frac{1}{\rho}} + (1-q)V^c - \frac{\Delta^\gamma}{\gamma}$$

This is a concave programming problem in Δ. Taking first-order conditions:

$$\frac{q}{((z^\alpha + \Delta^\alpha)^{\frac{\rho}{\alpha}} + Z_0^\rho)^{1-\frac{1}{\rho}}} = \frac{\Delta^{\gamma-\alpha}}{(z^\alpha + \Delta^\alpha)^{\frac{\rho}{\alpha}-1}} \tag{1}$$

Define the left-hand side of (1) as $f(z) \equiv \frac{q}{((z^\alpha+\Delta^\alpha)^{\frac{\rho}{\alpha}}+Z_0^\rho)^{1-\frac{1}{\rho}}}$. $f(0) > 0$, $\lim_{z\to\infty} f(z) = 0$, and $f'(z) < 0$ by hypothesis. Denote $g(z) = \frac{\Delta^{\gamma-\alpha}}{(z^\alpha+\Delta^\alpha)^{\frac{\rho}{\alpha}-1}}$. $g'(z) = \Delta^{\gamma-1}(z^\alpha + \Delta^\alpha)^{-\frac{\rho}{\alpha}}[(\gamma-\alpha)\left(\frac{z}{\Delta}\right)^\alpha + (\gamma-\rho)]$, which is guaranteed to be positive when $\gamma > \rho$ and $\gamma > \alpha$. $g(0) = 0$, $\lim_{z\to\infty} g(z) > 0$ (under the hypothesis), and, thus, there is a unique interior Δ. $\frac{\partial f}{\partial Z_0} < 0$; thus, as Z_0 increases, Δ declines. QED.

Lemma 3: Under the assumptions of Lemma 2, if $Z_1 > Z_0$ and $q \approx 1$, then post-AIPA investment declines.

Proof: Post-AIPA, the first-order conditions yield

$$\frac{1}{\left((z^{\alpha}+\Delta^{\alpha})^{\frac{\rho}{\alpha}}+Z_1^{\rho}\right)^{1-\frac{1}{\rho}}}=\frac{\Delta^{\gamma-\alpha}}{(z^{\alpha}+\Delta^{\alpha})^{\frac{\rho}{\alpha}-1}} \tag{2}$$

As long as $q \approx 1$, $\frac{1}{\left((z^{\alpha}+\Delta^{\alpha})^{\frac{\rho}{\alpha}}+Z_1^{\rho}\right)^{1-\frac{1}{\rho}}} < \frac{q}{\left((z^{\alpha}+\Delta^{\alpha})^{\frac{\rho}{\alpha}}+Z_0^{\rho}\right)^{1-\frac{1}{\rho}}}$, and post-AIPA investment declines. QED.