NBER WORKING PAPER SERIES

PATENT PUBLICATION AND INNOVATION

Deepak Hegde
Kyle F. Herkenhoff
Chenqi Zhu

Working Paper 29770
http://www.nber.org/papers/w29770

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2022

We are grateful to Luis Cabral, Wesley Cohen, Alberto Galasso, Nancy Gallini, Bernhard Ganglmair, Daniel Gross, Bronwyn Hall, Dietmar Harhoff, Joachim Henkel, Sabrina Howell, Zorina Khan, Hong Luo, Petra Moser, Alan Marco, Gaétan de Rassenfosse, Bhaven Sampat, Rosemarie Ziedonis, and seminar participants at Stanford, University of California, Berkeley, Northwestern University, University of Colorado, Boulder, and New York University. This research was supported by National Bureau of Economics Research, Innovation Policy Grant. Hegde acknowledges the support of the United States Patent and Trademark Office's Thomas Alva Edison Visiting Scholars program. All errors are our own. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the National Bureau of Economic Research.

Patent Publication and Innovation
Deepak Hegde, Kyle F. Herkenhoff, and Chenqi Zhu
NBER Working Paper No. 29770
February 2022
JEL No. D23,E02,G24,L26,O34

ABSTRACT

How does the publication of patents affect innovation? We answer this question by exploiting a large-scale natural experiment—the passage of the American Inventor's Protection Act of 1999 (AIPA)—that accelerated the public disclosure of most U.S. patents by two years. We obtain causal estimates by comparing U.S. patents subject to the law change with "twin" European patents which were not. After AIPA's enactment, U.S. patents receive more and faster follow-on citations, indicating an increase in technology diffusion. Technological overlap increases between distant but related patents and decreases between highly similar patents, and patent applications are less likely to be abandoned post-AIPA, suggesting a reduction in duplicative R&D. Firms exposed to one standard deviation longer patent grant delays increased their R\&D investment by 4% after AIPA. These findings are consistent with our theoretical framework in which AIPA provisions news shocks about related technologies to follow-on inventors and thus alters their innovation decisions.

Deepak Hegde
Stern School of Business
New York University
40 West Street
Suite 722
New York, NY 10012
dhegde@stern.nyu.edu

Kyle F. Herkenhoff
University of Minnesota
Department of Economics
4-101 Hanson Hall
1925 Fourth Street South
Minneapolis, MN 55455
and IZA
and also NBER
kfh@umn.edu

Chenqi Zhu
University of California at Irvine
4293 Pereira Dr
Irvine, CA 92697
chenqiz1@uci.edu

An online appendix is available at: http://www.nber.org/data-appendix/w29770

1 Introduction

> *"When a patent is granted and the information contained in it is circulated to the general public and those especially skilled in the trade, such additions to the general store of knowledge are of such importance to the public weal that the Federal Government is willing to pay the high price of 17 years of exclusive use for its disclosure, which disclosure, it is assumed, will stimulate ideas and the eventual development of further significant advances in the art."* – U.S. Supreme Court, Kewanee Oil Co. v. Bicron Corp. (1974; No. 73-187)

Patent publications 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] As emphasized by the U.S. Supreme Court, 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 on the completeness and rate at which patents are disclosed to the public (Scotchmer and Green (1990) and Fromer (2009)). Yet, and despite the fact that technological progress is central to endogenous growth models (e.g., Romer (1990), Lucas Jr and Moll (2014), Luttmer (2014), and Perla and Tonetti (2014)), little is known about the effects of patent publication on technology diffusion.[2]

In this study, we provide causal evidence that patent publication facilitates knowledge diffusion, reduces technology duplication, informs follow-on inventors' patenting decisions, and possibly increases R&D. Thus, our study advances what little we know about the effects of patent publication (Graham and Hegde (2015), Hegde and Luo (2018), Lück, Balsmeier, Seliger, and Fleming (2020), Furman, Nagler, and Watzinger (2021)).[3] The few studies on this topic offer contradictory conclusions: survey-based research suggests that inventors consider patents an important source of relevant technical knowledge (Ouellette (2017)) and information on rivals' R&D activities (Cohen, Goto, Nagata, Nelson, and Walsh (2002)); however, some legal scholars argue that "patent disclosures play an insignificant role in promoting R&D spillovers" (Roin (2005), p. 2027).

[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]Efforts in other fields have sought to identify the effects of peer-to-peer entrepreneur knowledge flows (e.g. Brooks, Donovan, and Johnson (2018) and Brooks, Donovan, and Johnson (2020)) and coworker knowledge flows (e.g. Herkenhoff, Lise, Menzio, and Phillips (2018), Jarosch, Oberfield, and Rossi-Hansberg (2021), Nix (2020), Gregory (2020)), often finding sizeable effects of knowledge diffusion on real output, productivity, and other aggregates.

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

We measure the effects of patent publication 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 publication on knowledge diffusion and follow-on innovation.

To guide our empirical analysis, we develop a theoretic framework in which AIPA provisions *news shocks* to inventors about recent technologies, and we derive a set of testable assumptions and implications. Our main assumption is that under AIPA, the technological know-how embedded in patents enters the stock of public knowledge faster than during the pre-AIPA regime. In this new information environment inventors see related inventions in the patent system earlier prior to filing their own patents. First, we show that inventors' search effort for prior knowledge, proxied by patent citations, increases. Better information about competing inventions reduces duplicate patent applications, which, in turn, reduces abandonments and similarity between closely related inventions. Under AIPA, inventors also draw more on recently disclosed patents and take smaller inventive steps, raising technology similarity among related, but not substitute, applications. Overall patent activity and innovation under AIPA may increase, decrease, or stay the same, depending on the net effect of the two countervailing forces of (i) lower invention costs due to superior information on the one hand, and (ii) free-riding by follow-on inventors on the other.

We provide tests of our theory using a Difference-in-Difference (DID) regression design, which we refer to 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).[4] 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 that the European twins of U.S. patents were not plausibly affected by AIPA's enactment.

This "twin" study design allows us to control for unobserved characteristics of each patent family (comprising the U.S. patent and its EP twin) and to account for quality-based selection into patenting or early disclosure using family-fixed effects. Thus, we isolate AIPA's causal 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 are published nearly simultaneously after AIPA—18 months from the filing date of the earliest twin, called the "priority date"). By

[4]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.

identifying the effect of 18-month disclosure in the U.S. for patents disclosed around the same time by the EPO, this design provides conservative estimates of patent disclosure's effects.

The "twin" study estimates show that after AIPA's enactment (i) U.S. patents' follow-on citations, our proxy for knowledge diffusion, within a ten-year window after disclosure date increase by 14.7 percent; (ii) mean delay for U.S. patents to receive 1/3/5/7 follow-on patent citations from application date drops by 25 to 29 percent; (iii) technological overlap increases between distant but related U.S. patents and decreases between highly similar U.S. patents; (iv) overall U.S. patenting increases by 6.2 percent. We also find that the increase in citations to U.S. patents is due to an uptick in citations to U.S. patents after AIPA, rather than to a decrease in citations to their European twins. This suggests a rise in knowledge diffusion stemming from earlier publication by the U.S., rather than a migration of citations from European patents to their U.S. twins, caused the increase in citations to U.S. patents after AIPA. We provide additional robustness checks including showing that self-citations, which reflect knowledge flows within patenting organizations, do not significantly rise after AIPA.

Our "twin" study's key identifying assumption is that applications disclosed by one national patent office (EPO) are not seamlessly disseminated to inventors filing for patents in another jurisdiction (the USPTO), until published by the latter. This assumption is supported by an extensive body of prior literature which has documented that patent citations and knowledge spillovers are geographically localized (Jaffe, Trajtenberg, and Henderson (1993), Jaffe and Trajtenberg (1999), Maurseth and Verspagen (2002), Bottazzi and Peri (2003), Peri (2005), Breschi and Lissoni (2001)). Moreover, Bacchiocchi and Montobbio (2006), Azagra Caro and Tur (2014), Wineburg (1988), Webster, Jensen, and Palangkaraya (2014), and de Rassenfosse, Jensen, Julius, Palangkaraya, and Webster (2019), study various patent offices (including the USPTO and EPO) and confirm search costs and "biases" that favor local inventors in patent grants and citations. We contribute three additional pieces of evidence to this literature. We show that AIPA's effects are larger when we use an alternative estimation sample of U.S. patents with twins filed at the Japanese Patent Office (JPO), rather than at the EPO. We argue that this reflects the larger search costs for JPO patents, as they are published in Japanese. Additionally, we show that citations to Patent Cooperation Treaty (PCT) patents, which are published in a single repository (by the World Intellectual Property Organization) after 18 months from filing and then seamlessly transmitted to national patent offices that participate in the PCT (including both the EPO and USPTO), show the least responsiveness to AIPA. Lastly, we find that foreign inventors filing at the USPTO experience the greatest increase in citations post-AIPA, consistent with the idea that a U.S. patent publication diffuses knowledge more effectively for inventors with fewer channels for broad international exposure. These results, taken together with prior evidence on the localization of knowledge flows, support the assumption that publication by the USPTO enhances the visibility of inventions even if they are disclosed by other national patent offices.

We next investigate heterogeneity in AIPA's treatment intensity to address the mechanisms underlying our results. First, we stratify US-EP twins by their technology class's exposure to pre-AIPA grant delays to test how variation in disclosure timeliness affects patenting outcomes. We find that patents in technology classes with the slowest pre-AIPA grant times have the greatest post-AIPA response of forward citations and citation lags, consistent with our hypothesis. Second, since AIPA allowed applicants to opt out of pre-grant publication, we use opt-out prevalence across technology fields to proxy for the value of secrecy. We find that patents in industries with greater opt-out rates (i.e., greater values of secrecy) are the most responsive to AIPA. These results suggest that AIPA's effects on knowledge diffusion are larger for inventors in contexts where barriers to knowledge flows are steeper.

Lastly, we link speedier patent disclosure to real measures of innovation by exploiting firm-specific exposure to AIPA. We do so by constructing the mean reduction in delays that each firm's patent portfolio received around AIPA. Firms that patented in technology areas with the longest patent grant delays experienced the greatest acceleration in disclosure after AIPA. We find that firms with one standard deviation greater exposure to AIPA (i.e., expediting patent publication by about 16 months) increased their R&D investment after AIPA by 4%, suggesting AIPA had significant effects on incentives to innovate. To provide a policy benchmark for the stimulative effects of AIPA, we compare the magnitude of AIPA's estimated effects to those of R&D tax credits (Rao (2016)). We show that a one standard deviation greater exposure to AIPA is comparable to reducing the user-cost of R&D by 2.6 percentage points or equivalently increasing the R&D tax credit in the U.S. by roughly 6 to 16 percentage points (i.e., increasing the R&D tax credit from 20% to 26%).[5] We provide additional suggestive evidence that AIPA increased the rate of U.S. patenting by about 6.2 percent, while decreasing patent abandonments and patent scope. Thus our findings collectively reveal that accelerated patent publication had wide-ranging effects on the speed of knowledge diffusion and innovation.

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 20^{th} 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."*[6] Our analysis confirms that AIPA indeed increased knowledge diffusion, reduced duplicative technologies and reduced patenting costs, without having a negative impact on U.S. patenting or innovation. Thus, our policy evaluation provides evidence against 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. Our results also imply that welfare analyses of patent systems should incor-

[5] The range of the equivalent R&D tax credit is determined by the calibrated interest rate and corporate tax rate.

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

porate the effects of invention disclosure, in addition to those of the monopoly rights created by patents.

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 our event study analysis of AIPA. Section 5 discusses our "twin" study design and the corresponding findings. Section 6 explores plausible mechanisms behind our results as well as various robustness tests. Section 7 analyzes the effects of AIPA on patenting and R&D. Finally, section 8 offers concluding thoughts.

2 Background

2.1 Institutional Context

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 18 months after the application date by the USPTO.[7] Since most foreign countries' patent systems already had similar 18-month publications, 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 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, 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.

Impetus and contemporaneous legislation. According to experts, efforts to accelerate patent publication in the U.S. came to fruition with the 1994 Uruguay Round negotiations on the Global Agreement on Tariffs and Trade (GATT) between the U.S. on the one hand, and other countries that already disclosed patents 18 months from application date, on the other. After these negotiations, the U.S. agreed to introduce legislation to harmonize its patent publication time-line, to

[7] Applications can be published before 18 months if applicants submit an early publication request to the USPTO. Less than 1% of applicants request publication prior to 18 months.

align with other developed countries of the world where 18-month publication after application was already in vogue. The Uruguay round negotiations also addressed other patent related matters, some of which were enacted immediately in the U.S. in 1995 (e.g., the change in patent term from 17 years after grant to 20 years from application date, and patent term adjustments for patents subject to prolonged prosecution at the USPTO were enacted for patents filed after June 8, 1995, via U.S.C. 301-307). The introduction of 18-month patent publication in the U.S. suffered delays due to intense lobbying and opposition from several groups. Famously, 26 Nobel Laureates sent a letter to the U.S. Senate opposing the 18-month disclosure requirement as harming inventors and stifling the flow of new inventions (Modigliani et al. (1999)).

On May 24, 1999, Representatives Coble and Rohrabacher introduced H.R. 1907 with the 18-month publication requirement, after brokering unanimous approval by the House Subcommittee on Intellectual Property and the Courts. The bill, dubbed the American Inventors Protection Act (AIPA) for the first time, incorporated a provision that only required publication of an application after 18 months if the applicant is also applying for protection in other countries. The bill also provided provisional rights to patentees to obtain reasonable royalties, retrospectively, after grant, if others make, use, sell, or import the invention during the period between publication and grant.[8] Eventually, on November 29, 1999, ten days after the vote in the Senate, President Bill Clinton signed the AIPA into law.

AIPA packaged other patent law changes along with the 18 month publication rule; specifically, the establishment of a patent reexamination alternative that expanded the participation of third party requester, establishment of the PTO as an agency with more control over its operating budget under the Dept of Commerce, and provision of patent term compensations for inventors who were subject to processing delays at the PTO. However, according to historians of patent law, the most significant change was the enactment of 18-month publication (see, for example, Ergenzinger Jr (2006) for a detailed account of AIPA's legislative history). Considering the hypothetical effects of these other changes, we believe it implausible for them to have caused the effects we attribute to AIPA's 18 month publication rule (e.g., AIPA's reexamination provision effects less than 0.5% of applications that are subject to reexamination each year while enhanced budgetary control at the USPTO allowed the agency to independently set its fees and plough back any surpluses to improve the agency's operations in later years). To rule out that our results are contaminated by

[8]According to AIPA the owner of a published application receives provisional rights to pursue royalties or infringement damages for the period between the date of publication and date of patent grant. To be entitled to the royalties, the published claims must be "substantially identical" to the granted claims. These rules together limit the downside of early patent disclosure to inventors. Without these protections, the inventor would face a higher risk of substitution by rivals through pre-grant disclosure. Hence, we anticipate that in the absence of these provisional rights, some inventors would prefer to delay disclosure till grant date. It is unclear how the provisional rights affect our knowledge diffusion results however: on the one hand, disclosure without provisional rights may increase citations and knowledge diffusion by lowering the cost of entry/duplication by follow-on competitors; on the other hand, disclosure without provisional rights may limit overall knowledge diffusion if inventors respond by delaying disclosure to grant date.

any anticipation effects, we remove patents filed one year before AIPA came into force (that is, filed between Nov 29, 1999 and Nov 28, 2000) and find results comparable to our main results.[9] A placebo test conducted using the pre-AIPA subsample also rules out that other spurious factors set into motion during our study period caused effects that we attribute to expedited patent publication.[10]

2.2 Literature Review

A large body of prior research has focused on studying the effects of the temporary monopoly rights conferred by patents on innovation (see Hall et al. (2014) and Williams (2017) for excellent surveys of this research). This research suggests that patents confer various benefits to their holders (e.g., Farre-Mensa, Hegde, and Ljungqvist (2020) and Kogan, Papanikolaou, Seru, and Stoffman (2017)), while plausibly blocking follow-on innovation efforts under some circumstances (see Galasso and Schankerman (2015) and Sampat and Williams (2019)). Surveys of innovative firms attempt to compare the effectiveness of patents against other appropriability mechanisms and suggest that firms prefer trade secrecy over patents to secure their R&D efforts (e.g., Levin, Klevorick, Nelson, Winter, Gilbert, and Griliches (1987) and Cohen et al. (2002)). However, Graham and Hegde (2015) find that of the companies that do patent, only about 8% of patent applicants opt out of pre-grant publication, thus revealing a preference for disclosure over secrecy for their inventions. 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 incorporating pre-AIPA data, Johnson and Popp (2003) find small positive short-run effects of pre-grant publication through this mechanism, but no long-run impact. Our research adds to this literature by examining the private and public consequences to inventors of accelerating the disclosure of their technological knowledge.

Our paper complements several contemporaneous efforts to understand the effects of invention disclosure. In particular, Furman, Nagler, and Watzinger (2018) show the introduction of patent libraries in the 1980s increased local patenting, job creation, and citations. These findings provide mechanisms through which invention publication may benefit follow-on innovation, complementing the results of Hegde and Luo (2018) and Mohammadi and Khashabi (2017), who show that AIPA accelerated technology licensing and corporate venture investments, respectively. Additionally, Kim and Valentine (2021) build on our approach to identify the real effects of AIPA by exploiting variation in realized patent grant delays of each firm's pre-AIPA patents relative to their competitors. They find positive effects of AIPA on citations and innovation inputs, including R&D, broadly consistent with our results which rely on each firms' USPC-weighted exposure to delays.[11] An important and creative contribution by Lück et al. (2020) uses data on 35 USC 102

[9] See Appendix D.4.

[10] See Appendix D.5.

[11] Kim and Valentine (2021) focus on the "Relative Spillover, defined as the log of the ratio between the patent-

rejections to provide evidence on the reduction of R&D duplication post AIPA, and de Rassenfosse, Pellegrino, and Raiteri (2020) exploits secrecy orders to measure the effects of patent secrecy on follow-on innovation. Other discussions of AIPA, such as that of Okada and Nagaoka (2020) and Baruffaldi and Simeth (2020), are descriptive in nature and compare patent citations before and after AIPA. Relative to this literature, our "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 and R&D. Furthermore, we put forth the theory of AIPA as a news shock, which allows us to interpret and rationalize our findings. In doing so, we extend recent research that investigates the effects of other policies such as compulsory secrecy orders Gross (2019) and trade secrecy laws (Ganglmair and Reimers, 2019) on innovation.

In terms of theory, some of the earliest work on information disclosure and innovation is by Horstmann, MacDonald, and Slivinski (1985), Bhattacharya, Glazer, and Sappington (1992), Anton and Yao (1994), Anton and Yao (2004), and Bhattacharya and Guriev (2006). More-recent work by Aoki and Spiegel (2009), Hopenhayn and Squintani (2011), Akcigit and Liu (2016), and Bobtcheff, Bolte, and Mariotti (2017), among others, endogenizes the timing of patent races.[12] 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 depart from these papers by simplifying the ex-ante decision of when to patent to instead model rich institutional details of the patent process. Our contribution is to integrate a multi-stage patent process into a heterogeneous firm framework with endogenous investment during the patenting process, drawing on elements from Atkeson and Burstein (2010) and Atkeson and Burstein (2019). In our theory, AIPA provides advance information about the presence of duplicates (i.e., a *news shock*) that alters the patenting, investment and abandonment decisions of inventors. Our theory thus provides a unifying framework to understand how the information environment created by patent publications influences a wide range of innovation activities.

We also build on the work by Bloom, Schankerman, and Van Reenen (2013) who use variation in federal and state R&D tax credits to identify technology diffusion. We advance this research by measuring the effect of patent publication on diffusion. 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 differ along three key dimensions. We incorporate the negative effects of technology disclosure through the outside option of the firm instead of through product market competition

weighted industry average publication lag and the patent-weighted own firm average publication lag, measured over the 20 years prior to the enactment of the AIPA." Thus Kim and Valentine (2021) use realized delays on each patent of the focal firm relative to competitors. Firm-specific realized delays may reflect endogenous characteristics of the firms (e.g. total assets, size, resources devoted to lawyers etc.). Our measure mitigates these issues by using *technology class averages* of delays, which are then weighted using the technology class shares of firms, similar to a Bartik-style instrument.

[12] See also Bryan and Lemus (2017) who endogenize the direction of innovation.

(see Section 3.4), 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.

3 Theoretical Framework

Our theoretical framework models AIPA as provisioning *news shocks* about existing patent applications to current cohorts of inventors. We use this framework to discipline our exploration of the data. In particular, we derive testable implications regarding AIPA's impact on various measures of patenting and innovation.

Consider a finite horizon economy populated by a unit measure of inventors. Let $t = 0, 1$ denote time. At date 0, inventors draw their idea quality from a distribution $F(z) : [0, \infty) \to [0, 1]$. After drawing an idea, inventors 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.[13] Pre-AIPA, inventors first invest in patent scope Δ and then subsequently learn about the presence of duplicate patents. Thus, there are two sources of risk: uncertainty over idea quality and uncertainty over the presence of a duplicate patent. The payoff at $t = 1$ to successfully patenting is given by $V^p(\cdot)$. Figure 2 illustrates the timing of the model.

Let $Z_0 > 0$ denote the stock of disclosed knowledge at $t = 0$. The profitability of a patent is determined by two factors: (i) idea quality, and (ii) inventor investment in the patent, $\Delta \in [0, \infty)$, which we will interpret as the number of claims and patent *scope*. Let $G(z, \Delta, Z_0)$ denote the profitability of a patent based on idea z, with scope Δ, and disclosed knowledge Z_0.[14]

To describe the value functions, we work backwards in time across the stages in Figure 2. In the terminal period, $t = 1$, the value of a patent is simply its profitability,

$$V^p(z, \Delta, Z_0) = G(z, \Delta, Z_0).$$

Stepping back one stage, we assume there is a duplicate patent in the system with probability $1 - q$ (we will refer to duplicates as *close* technologies).[15] If a duplicate patent is present, then the inventor abandons the patent and receives the competitive value V^c. Stepping back another stage, the inventor chooses the scope of the patent Δ at cost $c(\Delta)$, which is itself preceded by the choice to patent and the draw of idea quality. Thus, the value to an inventor at $t = 0$ is given by,

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

[13] In Section 3.4, we allow V^c to depend on own- and economy-wide ideas, thus introducing a notion of free-riding.

[14] An example includes the nested-CES function $G(z, \Delta, Z_0) = \left(\left(z^{1/\rho} + \Delta^{1/\rho}\right)^{\rho/\gamma} + Z_0^{1/\gamma}\right)^{\gamma}$, where $\rho > 1 (< 1)$ implies complements (substitutes) between z and Δ, and γ similarly controls complementarity between private and public knowledge Z_0.

[15] We discuss the case of endogenous duplication rates in Appendix A.4.

3.1 Pre-AIPA Characterization

First consider the choice of patent scope, $\Delta(z, Z_0)$, which is conditional on the idea quality. It is intuitive that if scope and public knowledge are complements ($V^p_{12}(z, \Delta, Z_0) > 0$), then scope expands when public knowledge expands. If scope and public knowledge are substitutes ($V^p_{12}(z, \Delta, Z_0) < 0$), then scope declines when public knowledge expands. Under standard regularity conditions outlined in assumption 1, lemma 1 formalizes this intuition. We relegate all proofs to Appendix A.

Assumption 1: Suppose $V^p(\cdot)$ is bound above, strictly increasing and concave in each of its arguments, and is such that $V^p(z, 0, Z_0) = 0$, and $\lim_{\Delta \to 0} V^p_2(z, \Delta, Z_0) = \infty$. Suppose $c(\cdot)$ is strictly convex, $c(0) = 0$, and it has a finite first derivative at the origin, $c_1(0) < \infty$. We further assume $c(\cdot)$ and $V^p(\cdot)$ are twice continuously differentiable in all arguments.

Lemma 1. Under assumption 1, (i) the patent scope decision is finite and unique, (ii) if public knowledge and scope are complements, then patent scope is increasing in public knowledge, and (iii) if public knowledge and scope are substitutes, then patent scope is decreasing in public knowledge.

Next consider the decision to patent. Lemma 2 establishes that with greater public knowledge, the patenting threshold declines, i.e. z^p decreases where z^p is defined below. We interpret this as a decline in the average inventive step size as public knowledge expands. To prove this lemma, we require some additional notation and assumptions. Let the *expected payoff* from patenting be given by

$$\Pi(z, \Delta, Z_0) \equiv qV^p(z, \Delta, Z_0) + (1 - q)V^c - c(\Delta).$$

Define z^p to be the minimum idea quality that enters the patent system, i.e. the value that satisfies the equality,

$$V^c = \Pi(z^p, \Delta(z^p, Z_0), Z_0).$$

To characterize how the patent threshold varies with public knowledge, we make a second assumption which prohibits corner solutions (e.g. all inventors patenting). We view this assumption as reasonable since most inventors (and firms) do not hold patents.

Assumption 2: there exist $\underline{z}$ and $\overline{z}$ such that $V^c > \Pi(\underline{z}, \Delta(\underline{z}, Z_0), Z_0)$ and $\Pi(\overline{z}, \Delta(\overline{z}, Z_0), Z_0) > V^c$.

Under assumptions 1 and 2, lemma 2 formally proves that the patenting threshold declines as public knowledge expands. Thus average inventive step size declines.

Lemma 2: Under assumptions 1 and 2, a unique interior patenting threshold z^p exists and is monotone decreasing in public knowledge Z_0.

3.2 Post-AIPA Characterization

We model the introduction of AIPA as provisioning news shocks about related inventions to follow-on inventors and rivals. 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, adjusted for duplicates.[16] Thus the available stock of public knowledge is given by $Z_1 = Z_0 + q\int_{z^{p,A}}^{\infty} zdF(z) > Z_0$. Figure 2 depicts the new timing. Post-AIPA policy functions are denoted with a scripted A, e.g. $z^{p,A}$ is the post-AIPA patent threshold rule. Likewise, value functions that differ post-AIPA are scripted by A.[17] The post-AIPA value of an inventor at $t=0$ is given by,

$$V^A(Z_0) = \int_{z^{p,A}}^{\infty} \{q\left[\max_{\Delta} V^p(z,\Delta,Z_1) - c(\Delta)\right] + (1-q)V^c\}dF(z) + V^c F(z^{p,A})$$

Lemma 4 establishes that patenting payoffs increase post-AIPA. This is due to the fact that inventors no longer have uncertainty over the presence of a duplicate when they make their decision to invest in scope. The AIPA news shock allows inventors to lower their effective costs of investing in patent scope from $c(\Delta)$ to $qc(\Delta)$. Throughout this section, we maintain assumptions 1 and 2 as hypotheses to proofs.

Lemma 4: The expected payoffs of patenting are unambiguously greater post-AIPA.

Since public knowledge expands, and payoffs to patenting increase, lemma 5 establishes that the patenting threshold declines.

Lemma 5: The patenting threshold declines post-AIPA, $z^{p,A} < z^p$.

Since the patenting threshold z^p declines under 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. However, duplication rates drop by assumption of the new information environment, implying that the closest technologies become less similar. This is one of our main testable implications.

Next, we revisit patent scope. In general, the impact of AIPA on the scope of a patent is ambiguous and depends on both the non-duplication rate q and the complementarity between scope and public knowledge. On the one hand, AIPA lowers the effective costs of investing in patent scope from $c(\Delta)$ to $qc(\Delta)$ which puts upward pressure on patent scope. On the other hand, AIPA increases public knowledge. If scope and public knowledge are complements, scope unambigiously increases post-AIPA. However, if scope and public knowledge are substitutes, then the effect of AIPA on scope is ambiguous: lower costs put upward pressure on scope, but greater public knowledge puts downward pressure on scope. Lemma 6 formalizes these mechanisms and provides a characterization of post-AIPA scope in the empirically relevant limiting case with low

[16]This can be rationalized as the new steady-state level of knowledge.

[17]Note that the function $V^p(z,\Delta,Z_1)$ is not affected by AIPA except through its arguments. Likewise, V^c is not affected by AIPA.

duplication rates $q \approx 1$.

Lemma 6: If $q \approx 1$ and $V^p_{23}(z, \Delta, Z_1) < 0$, then scope declines post-AIPA. For any $q < 1$, if $V^p_{23}(z, \Delta, Z_1) > 0$, scope increases post-AIPA.

With sufficiently low duplication rates, the cost reduction of investing in scope is minimal. If scope and public knowledge are substitutes, then greater public knowledge post-AIPA will unambiguously decrease scope. This case is particularly relevant for the U.S. data as we will discuss in Section 4.

3.3 Patent Citations

We next extend the model to derive predictions for citations. We show that if the idea distribution is Pareto (a functional form common in the knowledge diffusion literature, e.g. Lucas Jr and Moll (2014) and Perla and Tonetti (2014)), then forward citations rise post-AIPA.

We describe the key features of the environment since citations are central to our main analysis, relegating the details to Appendix A.3. Let s denote search effort of the existing stock of knowledge and let $\mu(s)$ denote the cost of searching. Our proxy for citations is search effort s. We assume greater values of search effort s imply better draws of ideas. In particular, we assume ideas are drawn from a Pareto distribution, with lower support s, i.e., $z \sim F(z, s)$. Thus greater search effort s shifts the distribution toward higher quality ideas. The pre-AIPA value at $t = 0$ is given by,

$$V(Z_0) = \left[\int_{z^p}^{\infty} \{ \max_{\Delta} qV^p(z, \Delta, Z_0) + (1-q)V^c - c(\Delta) \} f(z, s) dz + V^c F(z^p, s) \right] - \mu(s)$$

Under several additional assumptions, we show that search effort s is an increasing function of expected payoffs and therefore citations increase post-AIPA.

We note that the model's notion of a citation (search effort prior to invention) resembles backward citations; however, in Appendix A.3, we argue that in a repeated version of our model economy, backward citations from the current cohort of patents are forward citations to the previous cohort of patents. Since AIPA is a permanent news shock to all future patenting cohorts, AIPA must increase forward citations to patents disclosed after the law's enactment.

3.4 Model extensions

We consider several model extensions and their empirical implications. We discuss free-ridership and outside options that depend on idea quality in the main text, and we discuss conditions under which our main predictions continue to hold with endogeneous duplication rates in Appendix A.4.

Outside option scales with public knowledge (free-ridership). Our first discussion revolves around free-ridership. If the value of the competitive option, $V^c(Z_1)$, is an increasing function of the stock of disclosed knowledge (Z_1), there will be free ridership and AIPA may cause patenting to decrease. As we show in Section 7, our empirical results suggest this is not the case: patenting rises post AIPA. We further argue that competition should drive rents from unpatented operations to zero, regardless of the stock of disclosed knowledge.

Outside option scales with idea quality. Suppose the competitive payoff is a function of idea quality: $V^c(z)$. Then (i) predictions regarding scope are unchanged, and (ii) the patent threshold is now defined implicitly by $\Pi(z^p, \Delta, Z_0) = V^c(z^p)$. If we impose the assumption that $\Pi_1 > V_1^c$, i.e. expected payoffs from patenting better ideas increase at a greater rate than the competitive payoff, then our main predictions are unaltered. This is because $\Pi(\cdot)$ increases post-AIPA, and so the patenting threshold, z^p, must also decrease. As a result, our qualitative implications are unchanged.

3.5 Empirical Implications

In summary, modeling AIPA as provisioning news shocks yields the following testable implications about the law's effects: (i) citations increase (as inventors' search effort for ideas increases); (ii) the closest technologies increase distance (duplication declines) and patent abandonments decrease (by assumption that AIPA provisioned news shocks about related technologies); (iii) the furthest technologies decrease distance, i.e., average inventive step-size (which is proportional to z^p) declines; (iv) patent scope declines under the assumptions that (a) that public knowledge and private patent investment are substitutes and (b) pre-AIPA duplication rates are low; (v) technology enters public domain and production sooner (by assumption that AIPA is a news shock); and (vi) patent filings increase (although free-ridership can rationalize the opposite).

We use the above implications as a guide for our examination of the data. We begin with an event-study analysis surrounding AIPA, before establishing causality using our "twin" research design. We then measure the effects of AIPA on R&D expenditures using Compustat data.

4 Event Study Analysis

We begin by summarizing key innovation measures around AIPA's enactment in November, 2000. Our analysis in this section is primarily graphical, and in Appendix C, we formalize the graphical results with an event study research design following Gross, Notowidigdo, and Wang (2020).

4.1 Sample and Data

Our goal is to compare the extent and speed of knowledge diffusion, technology similarity, patent originality, application abandonment rates, and patent scope around AIPA to shed light on the effect of accelerated disclosures.

We begin with the universe of utility patent applications filed at the USPTO from 1998-2003. The application data are drawn from the agency's Patent Application Information Retrieval files and include both abandoned (rejected/withdrawn) and successful applications. We track the citations received by these applications until the beginning of 2017. 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 or EPO parallel applications; and (iii) standardized patent applicant and inventor names. After excluding reissued patents or applications that do not have information on important variables or have errors, 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, and 860,429 applications (69.5% of which were issued patents) were filed after AIPA. Table 1 provides a summary of sample observations after each step of filtering to construct the before-after comparison sample and the US-EP twin sample (more will be discussed in Section 5.2). In most analyses, we focus on granted patent applications rather than on all applications, except when we examine patent abandonment rates.[18] We discuss our main variables in detail below and list the definitions of all variables used in our analyses in Table 2.

Forward citations. The number of forward citations, after excluding self-citations, measures the extent of knowledge diffusion from the focal patent (*extensive margin*). We count citations received by each focal patent in 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 disclosure date is the time when patents become visible and, thus, *citable* by follow-on inventors.

Citation lags. 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, excluding self-citations. We use application dates as we want to capture how rapidly knowledge diffuses starting from the inception of the focal invention to the creation of follow-on inventions. Since patents become visible sooner after AIPA–at 18-months from application–one should naturally expect shorter citation lags for post-AIPA patents. Nevertheless, measuring citation lags from application dates helps assess the speed with which follow-on inventions build on prior patents.

Timing considerations in forward citations and citation lags. Our goal is to assess how the timing of patent disclosure affects the diffusion of knowledge about an invention. For this, ide-

[18] Application-level data for abandoned U.S. patents are not available pre-AIPA. Accordingly, we exclude abandoned and pending patents when analyzing citation counts, citation lags, and technology similarity to keep the pre- and post-AIPA samples comparable.

ally, we would measure knowledge diffusion starting from the invention date, which is theoretically the earliest date at which related knowledge can start diffusing.[19] However, we do not observe the invention date, so we use the patent application date to approximate it. It is important to note that our citation lag measure includes the mandated reduction in publication times caused by AIPA. Incorporating this mandated institutional acceleration of patent publication is essential to measure AIPA's effects on the speed of knowledge diffusion. Unlike the citation lag variable, our forward citation counts begin after the disclosure date. If we were to use time-since-application to measure forward citations, citation counts for post-AIPA patents will be greater than for pre-AIPA patents since post-AIPA patents have enjoyed exposure for longer due to early publication. Thus, to remove differences in exposure times from biasing our extensive margin results, we count forward citations from the date the patents are visible to the public both before and after AIPA.

Technology similarity. Technology similarity, which we also refer to as technology overlap, is measured as the cosine similarity between the focal patent and next-generation patents. Next-generation patents include all patents in the same primary technology subclass (IPC 4-digit code) as the focal patent and filed in the 19-36 month window after the focal patent.[20,21] We start the window at the 19^{th} month to ensure that the next-generation patents have had the opportunity to use the knowledge embedded in the 18-month disclosures. We stop the window at the 36^{th} month since this is roughly the average application-grant (and, thus, disclosure) lag. Presumably, patents filed within this 19-36 month window are the ones most likely to benefit from the knowledge revealed by the 18-month disclosures, although we ensure the robustness of our findings for different windows. To isolate the informational impact of patent disclosures, we exclude next-generation patents that have the same assignee as the focal patent.

The cosine similarity between the focal patent and its next-generation patents is computed 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 $j = 1, \ldots, J$ represents any patent in the next generation. $N_k = (N_{k1},\ N_{k2}, \ldots N_{k7154})$ is a vector with each element representing patent k's fraction of IPC

[19]To provide an even closer estimate of the invention date, we provide additional analysis of citation lags based on the *priority date* of the patent. The priority date reflects the date of the earliest application related to an invention as it adjusts for repeated filings of related applications through procedures such as continuations and divisionals. Appendix E shows that our main citation lag results are robust to using priority dates.

[20]Each patent typically receives multiple IPC codes. 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 primary IPC code first, but other authorities, such as the EPO, list the IPC codes alphabetically. We choose the subclasses (4-digit codes) with the highest frequency as the primary subclass for EP patents.

[21]We require the next-generation patents to have the same IPC 4-digit codes as the focal patent to reduce computational burden and to focus on patents in a related area as the focal patent. We require next-generation applications 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.

assignments in each of the 7,154 IPC main groups (IPC 7-digit code). The cosine similarity is widely used to measure the proximity of two vectors, each representing the location in a pre-defined space (e.g., Jaffe et al. (1986), Bloom et al. (2013)). Thus, for each focal patent i, we have a vector of similarities $\{Sim_{i1}, \ldots, Sim_{iJ}\}$ between patent i and its next-generation patents. We then record the similarity value at every 5^{th} percentile of this distribution, with higher percentiles in the similarity distribution (between a focal patent and its next-generation patents) corresponding to technologically close patents, and lower percentiles in the distribution corresponding to distant 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 examine the differentiation of patents as a function of the crowdedness of technological areas.[22]

4.2 Summary Statistics

Column (1) of Table 3 provides summary statistics for the 1.1 million granted U.S. patents in our sample. The average U.S. patent received 3.9 citations within three years after disclosure date (which is the 18-month publication date when available or the grant date for patents without 18-month publications), and the time it took for a patent to receive one citation was 36 months from application date. 17% of applications filed between 1998-2003 were granted before 18 months, and thus AIPA was non-binding on this subset of patents. Lastly, between 2000 and 2003, 8.6% of the successful post-AIPA applications opted out of pre-grant publication (note this statistic does not exist pre-AIPA).

4.3 Graphical Evidence

We begin by graphically examining changes in the scope and speed of knowledge diffusion around AIPA, as well as other patent characteristics pertinent to our theoretic predictions.

Forward citations. Figure 3 plots the average number of forward citations (excluding self-citations) by application month from 1998 to 2003. To illustrate AIPA's estimated impact, 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*.[23] Overall, Figure 3 shows that post-AIPA patents received more forward citations than pre-AIPA patents. The effects of AIPA are larger when we count forward citations over longer horizons.

Two patterns apparent from Figure 3 require discussion. First, citation counts show downward trends for both pre- and post- AIPA patents, most likely due to truncation of citing patents (an increasing fraction of potential citing patents are not yet granted as one approaches the end of

[22] Appendix B reports a validation exercise of our similarity measure.

[23] We formalize this analysis and explore different counterfactual trend assumptions in Appendix C.

our observation period–December 31, 2016–increasing undercounting of citations with time). This can lead to greater undercounting of citations for post-AIPA patents. Second, the way that we count forward citations favors pre-AIPA patents, particularly in the shorter citations windows, since their citations clock starts at grant and patents generally experience an uptick in citations upon grant both due to visibility and the finalization of property rights implied by the grant. As we increase the length of the citation counting windows, this effect of patent grant is attenuated and the cumulative effect of pre-grant patent disclosures intensifies. Therefore, the number of forward citations received by post-AIPA patents exceeds that of pre-AIPA patents five, seven, or ten years after disclosure. The figure clearly shows that the jump in citations occurs immediately in the months after AIPA's enactment, making other factors, such as contemporaneous economic conditions, less likely to be behind this increase. In Appendix C, we formalize these results using an event-study regression design with a linear pre-trend and find that post-AIPA patents receive 3.8%-19% more forward citations, on average.

Citation lag. We next graphically examine the speed of knowledge diffusion around AIPA. Figure 4 plots the average citation lag for all U.S. patents by application month from 1998 to 2003. The average time to receive the first citation ranged from 30 to 33 months before AIPA, and fell to 27 to 29 months after AIPA. The reduction in citation lag is even more pronounced at higher levels of citations. The average time to receive the first seven citations ranged from 42 to 45 months before AIPA, and fell to 39 to 41 months after AIPA. Overall, Figure 4 suggests that citation lags dropped sharply after AIPA, suggesting that timely patent disclosure accelerates knowledge diffusion. Event study analysis in Appendix C shows that after AIPA, the delay to receive one to seven forward citations decreased by 11.6%-20.1%, relative to the predicted delay based on a linear pre-trend.

Technology similarity. Figure 5 plots the monthly average similarity for all U.S. patents filed from 1998 to 2003. As described in Section 4.1, similarity at lower (higher) percentiles proxies for technological overlap with technologically remote (close) patents. We observe a large increase in similarity between technologically remote patents (5^{th}-15^{th} percentiles of similarity) and technologically moderate patents (25^{th}- 85^{th} percentiles of similarity) and a sharp drop in similarity among the top 5% or 10% closest patents (90^{th} -95^{th} percentiles of similarity). These findings are formalized in the event study analysis in Appendix C and are consistent with our theoretical predictions.[24]

Patenting threshold: renewals and originality. Our model predicts that the post-AIPA patenting threshold decreases due to reduced uncertainty in the patenting process and the positive externality from the larger stock of public knowledge from recently disclosed applications. Thus,

[24] We discuss alternative measures of similarity in Appendix G.

patentees 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.

Panel A of Figure 6 plots the monthly average 3.5-year renewal rate.[25] Patents that renew after grant are considered more valuable than those that do not. We find that the renewal rate went up before AIPA but gradually decreased after AIPA. The regression analysis reported in Appendix C confirms this graphic evidence suggesting that, indeed, inventors pursue patenting for less-valuable ideas in the post-AIPA period.

Panel B of Figure 6 plots the monthly average originality. Originality is measured as one minus the Herfindahl dispersion of backward citations to previously patents made by the focal patent across different technology classes (e.g., Trajtenberg, Henderson, and Jaffe (1997)). Intuitively, patents that refer to a broader class of prior art are more original. We find that patent originality rose steadily before AIPA and gradually fell afterward. The regression estimates, reported in Appendix C, confirm that inventors patent less-original ideas after AIPA.

Patent scope: claims. Our model predicts that patent scope declines after AIPA under the assumption that public knowledge and private patent investment are substitutes and duplication rates are low. Following the prior literature, such as Kuhn and Thompson (2019), we measure patent scope by 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 indicate a broader scope, and a greater number of words indicate that claims are defined with greater precision and clarity, thus narrower scope. Panels C through E of Figure 6 plot the monthly average patent scope. We find that patent scope decreased across the three measures post-AIPA relative to the pre-trend (as in Kuhn and Thompson (2019), we interpret more words in claims as indicative of more narrowly delineated scope). The regression estimates in Columns 3-5 of Table C.3 reported in Appendix C confirm these results.

Patent application abandonments. Our model implies 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 G and H of Figure 6 plot application abandonment rates before and after AIPA.[26] Once we account for the increasing trend of abandonments before AIPA, we find that abandonment rates declined after AIPA. The regression estimates in Column 6 of Table C.3 in Appendix C confirm the graphical evidence and

[25]As the sample patents were granted by mid-2014, four years before the record date of renewals (April 2018), there is no truncation errors in the computation of 3.5-year renewal rates.

[26]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, Mowery, and Graham (2009)). Therefore, we examine all abandonments and abandonments that are not followed by continuation filings in Panels D and E, respectively.

suggest up to a 5.2% (=0.013/0.247, the pre-AIPA average) decrease in abandonments relative to the pre-AIPA period.

While the sharpness of the jumps that coincide with AIPA's enactment suggests that these differences are due to AIPA, the magnitude of these differences may be contaminated by other confounding 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 knowledge diffusion. We address these concerns in the following section.

5 "Twin" Difference-in-Differences Analysis

5.1 Empirical Design

To identify AIPA's causal effects, we use a difference-in-differences framework that compares patents filed in the USPTO and their equivalent applications filed in the EPO before and after AIPA's enactment. As discussed in Section 2, AIPA mandated all post-AIPA U.S. applications with parallel foreign applications to be published 18 months after filing. In contrast, patents filed at the EPO were always published in 18 months, thus providing us with an ideal control group as the EPO parallel applications (equivalent, or "twin," applications) protect the same underlying invention as those filed at the USPTO. All inventors are required to disclose equivalent foreign applications at the time of filing in a particular country's patent office, and examiners further ascertain their equivalence status. We verify that the "twins" in our study are indeed equivalent in Appendix D.2.

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 sizable sample of twin applications. Further, the comparison of patents in two relatively comparable jurisdictions that cover the same technology sharpens our identification of AIPA's effects. To ensure robustness of our results, and to further elucidate the mechanisms behind our results, we later examine US-Japan twin patents.

Our main regression specification is summarized below:

$$Outcome_{ijt} = \alpha_1 + \alpha_2 \mathbf{I}(US_j) + \alpha_3 \mathbf{I}(US_j) \times \mathbf{I}(Post\ AIPA_t) + \delta W_j + Family_i + Month_t + \epsilon_{ijt} \qquad (1)$$

where j designates a patent in family i filed in month t. $\mathbf{I}(US_j)$ indicates whether the patent application is at the USPTO (denoted 'US (d)' in the tables, where (d) denotes dummy), and $\mathbf{I}(Post\ AIPA_t)$ indicates whether the patent application is submitted after AIPA's effective date (denoted 'Post AIPA (d)' in the tables), which itself is not separately identified due to the month fixed effects. The variable of interest is the interaction between $\mathbf{I}(US_j)$ and $\mathbf{I}(Post\ AIPA_t)$ (denoted 'Post AIPA × US (d)' in the tables), which captures AIPA's effect on patenting and innovation outcomes. W_j represents control variables, which include whether the patent is granted (Granted (d)) and whether it is granted before 18 months (Early Grant (d)). 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 the U.S. and early grant dummies 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 or twin fixed effects.[27]

Our identification strategy rests on the assumption that information barriers across regions (between the U.S. and Europe) come in the way of knowledge diffusion. Inventors from one region may not search as keenly for knowledge disclosed in other national jurisdictions either because the search is likely to reveal less relevant knowledge or because other factors, such as language differences, render searches in foreign jurisdictions more costly. In Section 6, we provide evidence to support this assumption of cross-country barriers. In particular, one of our tests repeats our twin analysis for U.S. and Japanese equivalent filings, and we show that AIPA's effects are more pronounced when we use Japanese twins of U.S. applications as the control group. We argue this reflects the higher informational barriers between U.S. and Japan, than between U.S. and Europe, as Japanese patent equivalents are published in the Japanese language. Hence, we believe our strategy of using EP twins provides conservative estimates of AIPA's effects since it captures only the marginal effect of disclosure by the USPTO for identical inventions that are disclosed simultaneously by the EPO.

We check, and rule out, that the propensity to file for EP parallel applications changed after AIPA because of the mandated disclosure requirement in the U.S., as we do not find any noticeable change in the proportion of U.S. patents with EP parallel applications (or other foreign applications) during 1998-2003, as shown in Appendix Figure D.1.

5.2 Sample Selection and Summary Statistics

Our US-EP twin sample is constructed based on the patent family table from PATSTAT (simple DOCDB patent family), which records the complete set of equivalent patent applications filed across different national patent offices. As described in Table 1, we start with 1,536,346 U.S. applications filed between January 1, 1998 and December 31, 2003. We require these U.S. patents

[27]The online Appendix D explores alternative fixed effects in Table D.2 and alternative specifications for pretrends in Table D.4.

to be granted and matched to EP equivalents in order to be included in the twin sample. We further: (i) require EP parallel applications are filed within 18 months of their associated U.S. applications according to the Patent Cooperation Treaty (1970); and (ii) exclude EP applications that are international PCT filings with the EPO designated as the receiving office.[28] These steps result in a sample of 316,563 U.S. patents with 354,227 EP equivalents with an average twin family size of 2.12. The family size is slightly greater than two, as one U.S. patent may be matched to multiple EP equivalents either because the EPO requested an amended application or because EPO requested a single U.S. application be split into two to adhere to EPO's application specifications under certain circumstances. [29]

Column 2 of Table 3 reports sample averages of key variables used in the analysis for U.S. patents with EP equivalents. On average, U.S. patents with EP equivalents receive marginally more citations at the 3/5/7/10 year horizons when compared to the pooled sample of U.S. patents over the same time period. U.S. patents with EP equivalents receive 4.1 citations 3-years after disclosure versus 3.9 citations among the broader sample of all U.S. patents. At the 10-year horizon, U.S. patents with EP equivalents have 14.4 forward citations compared to 13.2 for all U.S. patents. Citation lags are comparable between U.S. patents with EP equivalents and all U.S. patents. U.S. patents with EP equivalents are marginally less similar to prior cohorts, and they include more claims (19.1 vs. 18.24 for all US patents). The *signs* of these differences are not surprising as inventions seeking patent protection in multiple countries are presumably more valuable, but we argue that the limited *magnitude* of these differences between U.S. patents with EP equivalents and all U.S. patents implies relatively high external validity of our twin analysis. If anything, we expect the "twin" analysis to provide lower-bound estimates of AIPA effects as it identifies the effects of disclosure by the USPTO for inventions that are simultaneously disclosed elsewhere.

5.3 Difference-in-Differences Results

Forward citations. We first examine AIPA's effect on the extent of knowledge diffusion in the twin sample. We supplement the USPTO's data on citations made by U.S. patents with PATSTAT data on citations made by patents filed in other jurisdictions. We require the citing U.S. patent to be granted to avoid double counting the forward citations made by corresponding pre-grant publications after AIPA. For EP applications, citations data are obtained from PATSTAT, and include citations from both granted patents and pre-grant publications. Following Harhoff, Hoisl, and Webb (2006), we adjust for patent equivalents when counting forward citations for EP patents. Specifically, if a future EP patent cites a U.S. patent but not its EP equivalent, it is counted as

[28] International PCT filings are identified based on the kind code of "W" in PATSTAT.

[29] In untabulated tests, we find that our results are identical when we limit the sample to twins that had a one-to-one correspondence between U.S. and EPO equivalents.

one forward citation for the EP equivalent.[30]

Figure 7 plots the monthly average forward citations for U.S. and EP twins, normalized to 1 in January, 1998 for comparability. Figure 7 reveals comparable pre-trends for both U.S. and EP forward citations prior to AIPA.

Panel A of Table 4 reports the DID estimates for forward citations using equation (1). Consistent with the graphic evidence, we observe a significant positive coefficient on 'Post AIPA × US (d)' 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 generations of patents, which, in turn, stimulate further follow-on innovation.

One may be concerned that the more forward citations received by U.S. patents simply reflect a migration of citations. In other words, 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, thereby boosting forward citations received by post-AIPA U.S. patents through a substitution effect. If this were 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. We investigate this concern in Table D.4 included in Appendix D. Overall, we find that rather than a drop in citations, EP equivalents receive more forward citations after AIPA, although the increase is economically small and only statistically significant at the 10% level when we count the citations in a ten-year window after application. These results suggest that the increase in U.S. patents' forward citations after AIPA is unlikely to be driven by an EP-to-US substitution.

Since detailed information about the underlying inventions covered by US-EP twin patents is always publicly available through EP pre-grant publications, greater knowledge diffusion associated with timely U.S. patent disclosures suggests the existence of search frictions across patent offices. Such frictions may have arisen from search costs, language barriers, or a lack of other channels that facilitate knowledge diffusion across national patent office jurisdictions. Given these search frictions, we expect AIPA to cause a larger increase in knowledge diffusion in the U.S. than in Europe.[31] To test this expectation, we examine forward citations made by future U.S. patents and EP patents separately in Panel B of Table 4. The coefficients on 'Post AIPA × US (d)' are significantly positive for 5/7/10-year forward citations made by future U.S. patents, as well

[30] Appendix F undoes the Harhoff et al. (2006) adjustments and shows very similar results.

[31] Since twin patents may have already diffused within the U.S. due to the presence of an equivalent EP patent, AIPA's effects should be greater for applications without twins—that is those filed at, and disclosed by, the USPTO alone. While not as clean as the twin analysis, our event study analysis in Appendix C yields moderately larger results, supporting the idea that the twin analysis provides conservative estimates of AIPA. Likewise, analysis on US-JP twins, for which language barriers are arguably higher, yields larger effects of AIPA (see Section 6 below).

as by future EP patents. More importantly, the magnitude is much larger for forward citations made by future U.S. patents, and the difference between the two is statistically significant at the conventional level. This evidence suggests that pre-grant publications increase knowledge diffusion by reducing search frictions across patent offices.

Citation lag. As in Section 4, we use citation lags to proxy for the speed of knowledge diffusion. Figure 8 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 citation lags 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.

In Table 5, we estimate equation (1) to test the impact of AIPA on the pace of knowledge diffusion. Consistent with the graphic evidence, we find significant drops in the four citation lag measures. The economic magnitude is large. 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.

Technology similarity. Lastly, we study the impact of pre-grant publications 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 9, 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 (5^{th} percentile to 50^{th} percentile). The difference-in-difference estimate shrinks in magnitude at the 75^{th} percentile. At the highest percentiles of similarity (90^{th} percentile and 95^{th} percentile), we observe a drop in technology overlap among U.S. patents after AIPA.

The regression results estimating equation (1) in Table 6 confirm this graphic finding, revealing an increase in similarity at the 50^{th} and 75^{th} percentiles to be about 12.0% (=0.013/0.108) and 3.4% (=0.010/0.289), respectively. These patterns reverse at the upper percentiles. We find that at the 95^{th} percentile, there is a 2.1% (=0.013/0.611) reduction in similarity. Note that the estimated reduction of duplicative patenting is likely to be underestimated since *all* percentiles – including the 95^{th} percentile – are pushed upwards by increased knowledge spillovers after AIPA.

Figure 10 plots the coefficients on the interaction term in equation (1) using the similarity measure at every 5^{th} percentile as the dependent variable. The coefficient initially increases, reaches a plateau around 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.

Since our twin sample has only granted U.S. patents with EP equivalents, the design is not suitable for identifying AIPA's effects on patent abandonment. Likewise, since by definition twins have the same, or nearly identical, claims and backward citations, our DID analysis cannot be used to test predictions regarding AIPA's effects on patenting scope and originality. Hence, we have focused on examining AIPA's effects on citations, citations lag and similarity with our "twin" design. We refer readers back to Section 4.3 for evidence, based on our event study, regarding AIPA's effects on patenting threshold, patent scope and abandonments.

6 Mechanisms and Robustness Checks

In this section, we investigate the mechanisms driving our results and the robustness of our findings. We first provide evidence supporting a key assumption of our "twin" based identification strategy—that information barriers between national jurisdictions come in the way of knowledge flows across borders. Specifically, we analyze US-JP twins, PCT filing status, and twins filed foreign inventors and find that AIPA eased flows where the impediments were high. We then explore AIPA's heterogeneous effects and show that the law's estimated impact on the extent and the speed of knowledge diffusion is stronger for (i) patents in technological fields that suffered the longest grant delays pre-AIPA; and (ii) technological fields in which inventors value secrecy for their patents the most. These findings suggest accelerated disclosure had the greatest impact on citations in contexts where we expect barriers for free knowledge diffusion to be the higher. In a setting where we expect the barriers to be non-existent — that is, for within-firm knowledge flows, which we proxy using self-citations — we show that AIPA had no systematic impact. Taken together, these findings suggest that the mechanism underlying the post-AIPA increase in citation count and speed is enhanced knowledge diffusion. Finally, we report a battery of checks to establish the robustness of our "twin" analysis.

6.1 Cross-country barriers to knowledge diffusion

Our "twin" analysis is inspired by a well-established literature which documents that knowledge flows, proxied by patent citations, tend to be localized within city, state, and national boundaries (Jaffe et al. (1993), Jaffe and Trajtenberg (1999), Maurseth and Verspagen (2002), Bottazzi and Peri (2003), Peri (2005), Breschi and Lissoni (2001)). Several studies including Wineburg (1988), Bacchiocchi and Montobbio (2006), Azagra Caro and Tur (2014), Webster et al. (2014), and de Rassenfosse et al. (2019) document search costs and biases, favoring local inventors in patent grants and citations, in various national jurisdictions including the U.S., Europe, Japan, and China. Nevertheless, we conduct three exercises here aimed at testing the existence of barriers to cross-country knowledge diffusion.

US-JP twins. Our first exercise highlights the role of language barriers by repeating our "twin" analysis for Japanese twins of U.S. patent applications. The Japanese Patent Office (JPO) publishes its applications in Japanese, so inventions published at 18 months by the JPO are rather inaccessible to follow-on inventors in other countries. In contrast, EPO publishes most of its applications in English (85% in our sample) or languages based on the English alphabet (German and French) that are more readily translated. Thus, our prior is that AIPA, and the consequent accelerated publication of US-JP twins in English by the USPTO, had greater effects on knowledge diffusion for US-JP twins than US-EP twins.

Table 7 reports the results of this analysis with Panel A focusing on citation counts and Panel B on citation lags. We begin by augmenting equation (1) with country-specific linear trends, which purge the predictable upward linear trend in citations exhibited by Japanese patents (see Appendix Figure D.2). For comparability, we repeat our US-EP twin analysis with country-specific linear trends, yielding similar point estimates to the baseline estimates without controlling for the trends. Specifically, in Panel A of Table 7, Columns (1) and (2) illustrate the impact of AIPA on 3-year forward citations for US-EP and US-JP twins, respectively. While Column (1) implies no significant difference of 3-year forward citations for US patents relative to their EP equivalents, Column (2) shows that U.S. patents receive significantly more citations post-AIPA relative to their JP equivalents. The last row of Panel A ('Difference in AIPA Effect') tests equality of the coefficients. We find that the coefficients differ at the 5% level. Columns (3) through (8) yield similar results, with the US-JP twins exhibiting significantly greater post-AIPA citations than US-EP twins.

Panel B of Table 7 compares the AIPA impact on citation lags across US-JP and US-EP twins. We find significant reductions in citation lags at all horizons for both US-EP and US-JP twins. In particular, U.S. patents in US-EP pairs exhibit a 20% reduction in time delay to receive 1 citation whereas U.S. patents in US-JP pairs exhibit a 37% reduction. This difference is significant at the 1% level. Columns (3) through (8) repeat the analysis for different citation thresholds. While the point estimates broadly support our argument of greater barriers for US-JP twins, the differences fall outside of standard significance levels.

Overall, we find larger effects of AIPA on our knowledge diffusion measures for US-JP twins than on US-EP twins, suggesting information frictions, worsened by language differences, impede the free flow of knowledge across national boundaries. These results are consistent with accelerated patent publications by the USPTO lowering these barriers.

PCT Filings. Our second exercise estimates AIPA's effects in our twin analysis for patent families filed under the international Patent Cooperation Treaty. PCT patents are published in a single repository (by the International Bureau at the World Intellectual Property Organization) after 18 months from filing and then seamlessly transmitted to national patent offices that participate in the PCT (including both the EPO and USPTO). These patents are easily and readily searchable by U.S. inventors, and so arguably exhibit fewer barriers of EPO-to-USPTO knowledge flows. We, therefore, expect AIPA's accelerated patent disclosures to have minimum impacts on the diffusion of these inventions.

Panel A of Table 8 estimates equation (1) with a full set of interactions with a PCT twin indicator ('PCT (d)'). The resulting triple difference estimator compares U.S. vs EP twin citations before and after AIPA among PCT and non-PCT twins. Columns (1) through (4) report results for forward citations. In general, we find that PCT twins are less responsive to AIPA relative to non-PCT ones. Adding the relevant interaction terms ('Post AIPA $\times$ US $\times$ PCT (d)' plus 'Post AIPA $\times$ US') implies that PCT twins exhibit very little difference in U.S. patent citations relative to their EP twins, pre- and post-AIPA. Columns (5) through (8) repeat this analysis for citation lags and reveal a smaller AIPA impact on citation lags for PCT twins. We view these estimates as further supporting our assumption that cross-country barriers impede the free flow of knowledge across borders.

Foreign inventors. Our third exercise differentiates AIPA's impact on foreign and U.S.-based inventors. Our prior is that foreign inventors' ideas are less likely to diffuse in the U.S., thus making the effects of AIPA more pronounced for U.S. twins filed by these inventors. To implement this triple difference analysis, we augment equation (1) with a full set of interactions with a foreign inventor indicator ('Foreign Inventor (d)').

Panel B of Table 8 reports the results. The coefficient of interest is the triple interaction 'Post AIPA $\times$ US $\times$ Foreign Inventor (d)', which estimates the differential AIPA effect for foreign inventors. Columns (1) through (4) report the forward citation results, which show a significant larger impact of AIPA on foreign inventors' patent citations. Thus, AIPA exhibits stronger effects for ideas that are least likely to have diffused within the U.S. through alternative mechanisms such as local investor meetings and industry conferences. Columns (5) through (8) report the results for citation lags. Although two of the estimates are not significant at conventional levels, we find a more pronounced negative effect on citation lags for foreign inventors, further supporting our hypothesis that AIPA provisioned relevant news to follow-on inventors, particularly in settings where other information diffusion mechanisms are less effective.

6.2 Heterogeneous Effects

Our analysis of US-JP twins, PCT filing status, and foreign inventors suggests that national boundaries impede knowledge flows and that AIPA eased flows where the impediments were high. Here, we further probe AIPA's heterogeneous effects, stratifying our sample along two dimensions which plausibly correlate with other types of barriers for knowledge flows: (1) exposure to patent grant delays and (2) the value of secrecy, proxied by opt-out intensity in the technology class.

Exposure to grant delays. We first stratify US-EP twins by their technology field's pre-AIPA grant delays. Our prior is that patents in fields with the greatest pre-AIPA grant delays react more to AIPA since AIPA accelerated their disclosure the most. We construct pre-AIPA grant delay as the average time delay from application to grant, in years, for U.S. patents filed in the three-year pre-AIPA period for each main IPC technology class (4-digit). The average pre-AIPA grant delay is 31 months in our sample and it varies a lot across technology classes. For example, patents in HO4L, "Transmission of Digital Information," have an average delay about 48 months while the delay is 19 months for patents in "Artificial flowers; Wigs; Masks; Feathers." We implement the triple difference estimator by augmenting equation (1) with a full set of interactions with pre-AIPA grant delay exposure.

Panel A of Table 9 reports our results. Consistent with our prior, we find that if pre-AIPA grant delays are one standard deviation larger (0.61), forward citations increase a further 2.4-4.0% and citation lags reduce by a further 5.1-8.5%. This result supports our hypothesis that timely patent disclosure accelerates knowledge diffusion.

Exposure to opt-outs. As explained in Section 2, AIPA allowed patent applicants to opt out of pre-grant publication by forgoing foreign protection for their inventions. Our assumption is that applicants opt-out of early disclosure for inventions that benefit more from secrecy, thus the value of secrecy can be proxied by opt-out intensity. The average opt-out rate is 8.6%, as reported in Table 3, but some technology classes (e.g., B21L, Making Metal Chain) have virtually no opt-outs while others (e.g., H05C, Electric Circuits) have an opt-out ratio as high as 25%. We expect that patents in technology fields with the greatest opt-out prevalence (and thus having the greatest value of secrecy) will be the most responsive to accelerated patent publication. Since opt-outs became available after AIPA, we construct opt-out ratio as the percentage of opt-out U.S. patents filed during the 3-year post-AIPA period (2000-2003) in each 4-digit IPC class. We implement the triple difference estimator by augmenting equation (1) with a full set of interactions with opt-out ratios.

Panel B of Table 9 reports our results. We find that for patents in fields with opt-out rates one standard deviation higher (0.046) than the mean, the estimated post-AIPA effects on forward citations are 1.2-1.9% greater and on citation lags are 2.9-4.6% greater. These results suggest that accelerated patent publication matters the most for knowledge diffusion in fields of higher secrecy.

6.3 Placebo test using self-citations

Throughout our analyses, we have interpreted an increase in citations post-AIPA as reflecting increased knowledge diffusion after AIPA. Here, we use self-citations data to further test the plausibility of this interpretation. Since self-citations – defined to be citations by follow-on inventors to their own prior patents – are not affected by barriers to knowledge diffusion, they should not respond to AIPA.

Table 10 estimates specification (1) for self-citation counts and the share of self-citations among total citations. Panel A of Table 10 shows that 18-month publication does not have a systematic effect on self-citations: the 3-year and 5-year estimates are negative, the 7-year estimate is not significant at conventional levels, and the 10-year estimate is significant and positive but very small relative to the baseline estimates in Table 4. Furthermore, in Panel B, we find a systematic negative effect of AIPA on the *share* of total citations that are self-citations. We view this as corroborating evidence that the post-AIPA increase in citations reflects the effects of disclosure-related knowledge flows occurring across firm boundaries, and that the increase is not due to correlated shocks that may have increased citations overall, including self-citations. Even while the placebo test supports our interpretation of the citations result, we note that this placebo is not perfect – AIPA may have had an effect on self-citations if it altered the nature of firms' R&D allocations and invention strategies with respect to its own previous-generation inventions (which, in turn, may respond to the broader information environment as discussed in our theory section). For example, a decreasing fraction of self-citations is also consistent with firms being more likely to use recent external knowledge as a substitute to own-R&D post-AIPA, as predicted by our theory.

6.4 Other Robustness Checks

We conduct a number of robustness tests to ensure the validity of our measures and regression-based results. Specifically, we provide evidence that supports our twin study's representativeness and validity; we conduct a placebo exercise in the pre-AIPA period; we use U.S. twins granted within 18 months of filing, and thus unaffected by AIPA's 18-month publication rule, as a placebo; we adjust for continuation and divisional patents that may induce noise into our citation lag measures; we exclude software patents to mitigate confounding macro trends in patenting; and we explore alternative fixed effects, including applicant fixed effects. We find our main results are quantitatively and qualitatively robust across the wide variety of checks.

Representativeness of "twin" analysis. We examine the representativeness of the estimates delivered from our US-EP "twin" study design by comparing them to the estimates from event-study design on the universe of U.S. patents (Tables C.1 - C.3 in Appendix C). On the one hand, we expect that AIPA's effects are understated by the "twin" design as it identifies the marginal effect of disclosure by the USPTO for applications that are simultaneously disclosed by other national offices (EPO). AIPA's effects should arguably be higher for applications without twins–that is, for U.S. patent applications without foreign equivalents. On the other hand, since applications with foreign equivalents are likely to be more valuable than those without (consistent with comparative descriptives on patent claims, citations and renewals presented in Table 3), one may expect early disclosures to have a greater impact on knowledge diffusion for twin patents. The comparison suggests the plausibility of both hypotheses: the event-study estimates of AIPA's effects are larger for the extensive margin but smaller at the intensive margin than "twin" study estimates. Nevertheless, the estimates obtained by the two designs are broadly comparable, suggesting that AIPA's effects on the population of patentees are not too far off from those obtained by the US-EP "twin" analysis.

Placebo AIPA effective date. In Appendix D.5, we use an earlier fictitious AIPA enactment date (July, 1999, roughly the mid-point of our pre-AIPA sample period) to rule out spurious effects driving the results. The results show no impact on citations, citation lags, or technology similarity between before this fictitious date and AIPA's actual enactment in November, 2000. This evidence suggests that our results are not biased by spurious trends or possible anticipation effects of AIPA's actual enactment.

U.S. patents granted within 18 months. In Appendix D.6, we conduct another placebo test in which we restrict US-EP twins to only include twins where the U.S. patent is granted within 18 months, and thus AIPA was not binding. We find much more muted effects in this subsample: no effect of AIPA on forward citations and a slight decrease in citation lags (about 20% of the magnitude we see in the entire twin sample).

Continuation applications. In Appendix E, we address the concern about the influence of continuation filings, which allow patent applicants to submit related applications that claim priority to previously field applications, on our results. We first analyze citation lags based on priority dates, the filing date of the earliest related patent application, which is closer to the actual invention date. This analysis yields very similar effects of AIPA on citation lag reductions. We also repeat our baseline analysis after excluding 63,365 (9.45%) continuation patents from the US-EP twin sample, as identified from PATSTAT's application continuation table. The results estimated from this "clean" sample are very similar to our baseline.

Alternative similarity measures. In Appendix G, we analyze alternative similarity measures, including Jaccard similarity based on 7-digit IPC codes, cosine similarity based on 4-digit IPC codes, and cosine similarity based on text-converted vectors as constructed by Google. Figure G.1 plots the results along with our main measure (cosine similarity based on 7-digit IPC codes) and Table G.1 reports the regression results. The general pattern that similarity rises among distant but related patents is present in all versions of similarity considered. The reduction in similarity among the closest related patents is also present in all measures except for Google text similarity, presumably due to the difficulty of converting large bodies of text into sparse vectors. We further discuss the pros and cons of various similarity measures in detail in Appendix G.

Excluding software-related patents. To alleviate concerns that our results reflect contemporaneous macro-economic trends related to the dot-com stock-market bubble that peaked in March 2000, Appendix I repeats our main analysis after excluding software patents as defined by Graham and Vishnubhakat (2013). We find results very similar to the ones obtained from our main sample that does not impose any technology-field sample exclusions.

Alternative fixed effects and trends. In Table D.2 of Appendix D.3, we repeat our "twin" analysis with alternative fixed effects, including applicant fixed effects. We find very similar results across a wide range of specifications. The fact that our twin analysis results are robust to applicant fixed effects provides strong evidence that selection on inventor-specific characteristics is unlikely to explain our main findings. We further include country-specific technology trends in Table D.3 of Appendix D.3 to address concerns regarding the differences in industry composition and sectoral trends across patent offices; we find very similar results.

7 AIPA, Patenting, and Innovation

In this section, we examine AIPA's broader impact on measures of innovation. We test two predictions: (i) AIPA increased patenting overall, (ii) AIPA increased real measures of innovation proxied by R&D investments. We find evidence that both of these predictions hold in the data.

Patenting intensity. First, we examine AIPA's effect on overall U.S. patenting. Our theory predicts that patenting activities increase after AIPA due to a richer information environment that decreases the cost of patenting.[32]

To empirically examine AIPA's effect on patenting, in Figure 11, we plot the monthly count of patent applications in the USPTO and EPO (without imposing the twin requirement) and find that they both grew steadily before AIPA and that the growth rate slowed slightly after AIPA while remaining positive. We formalize this graphical analysis by running difference-in-differences

[32]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.

regressions that compare the total number of applications and eventually granted applications filed in each month from 1998 to 2003 at the USPTO and EPO, respectively, in Columns (1) and (2) of Table 11. We find that the number of U.S. applications (eventually granted) increased by 2,304 (1,158) per month after AIPA, relative to EP applications. This increase is economically considerable, equivalent to 11.9% (7.9%) of the pre-AIPA average (19,384 applications and 14,621 grants, respectively). As a robustness check, we also conduct the analysis at the patent-office × technology-class × month level in Columns (3)-(6) and find confirmative evidence that patenting intensity increased after AIPA.

R&D intensity. Second, to shed direct light on the effect of early disclosure on innovation measures other than patents, we examine AIPA's effect on firms' R&D investments.[33] Although all pre-grant publications are disclosed in 18 months after the filing date, the (expected) advancement in disclosure caused by AIPA varies across technology classes due to different levels of grant delays. Specifically, we exploit this between-class variation and calculate a firm-specific measure of exposure to AIPA's information shock as the weighted average of delays across all technology classes that the firm patents in. The weights are the share of patents filed in each class by the firm during the three years before AIPA. Firms with a larger weighted average delay are more affected by AIPA and hence are considered to have greater exposure to AIPA. Firms with zero patents or weighted average delay smaller than 18 months are considered to have zero exposure to AIPA.

We then estimate AIPA's effect on the R&D investments (as measured by log R&D or R&D intensity) as a function of the firm-specific AIPA exposure.[34] To minimize R&D intensity outliers, we analyze three samples: firms with annual sales greater than \$10m, \$20m, and \$50m.

The results are reported in Table 12. Across the six specifications, we find a consistent increase in R&D investment after AIPA for firms with greater exposure to AIPA. Columns (1) through (3) estimate the elasticity of R&D to AIPA exposure. Across each of our samples, the standard deviation of exposure to AIPA is approximately 1.4 years (Appendix H.1 contains summary statistics for the samples). Consider column (3) which restricts firms to those with greater than \$50m in sales. A one standard deviation faster patent disclosure post-AIPA is associated with a 3.95% (=1.4×.0282) increase in R&D investment. These results are robust across samples, and they imply that pre-grant publication had a large stimulative effect on R&D.

[33] We focus on U.S. publicly listed companies for this analysis due to the availability of their R&D data. We obtain the R&D expenses and other firm variables from Compustat for all companies listed in the U.S. from 1998 to 2003, six years centered around AIPA's enactment. We exclude firms in finance (SIC code: 6000-6999) and utility industry (SIC code: 4900-4999) and firms with missing values for the variables. We further require firms to have at least one observation before and one observation after AIPA in the estimation sample to ensure comparability. We report the sample size and summary statistics for key variables in Appendix Table H.1. The link between patents and public companies is obtained from Kogan et al. (2017).

[34] We explore the possibility of differential pre-trends between high and low exposure firms in Appendix H.2. We find very similar pre-trends of R&D among firms in the top and bottom quintiles of AIPA exposure, suggesting that our results are not driven by the correlation of unobservables and exposure to AIPA.

Columns (4) through (6) of Table 12 repeat our estimation using R&D intensity as the dependent variable. The coefficients are interpretable as semi-elasticities. We find that greater AIPA exposure is positively related to R&D intensity; however, the results are sensitive to sample selection. To minimize R&D intensity outliers and provide conservative estimates, we focus on column (6), which restricts firms to those with greater than \$50m in sales. A one standard deviation faster patent disclosure post-AIPA is associated with a 0.27 percentage point (=1.4×.00193) increase in R&D intensity. This represents a 4% increase in R&D intensity relative to the sample average R&D intensity rate of 6.6%. Thus, the implied elasticities across columns (3) and (6) are consistent and imply large positive effects of AIPA on R&D.

To interpret the magnitude, we compare our results to Rao (2016) who measures the impact of R&D tax credits on R&D spending using confidential firm-level IRS data. Rao estimates the semi-elasticity of R&D intensity with respect to the user cost of R&D is -0.104 (Table 4, Column (2) of Rao (2016)). Thus, a one standard deviation faster patent disclosure has an equivalent effect to a 2.6 percentage point (=.0027/.104) decrease in the user cost of R&D. Rao (2016) reports an interquartile range of R&D user costs of 5 percentage points, placing AIPA's effect on R&D user costs well within an empirically plausible range.

Using the formulas and parameter values in Rao (2016) (setting the input price growth rate as well as the firm-specific claw-back and carry-forward terms to zero), we calculate that a 2.6 percentage point reduction in the user-cost of R&D is equivalent to increasing the R&D tax credit from 6 to 16 percentage points (i.e., increasing the R&D tax credit from 20% to 26%). Figure 12 graphically illustrates the R&D tax credit equivalent of AIPA for various interest rates and corporate tax rates.

Taken together with our previous evidence on enhanced knowledge diffusion and reduced duplicative patenting post-AIPA, this analysis suggests that accelerated patent publication had a net positive effect on both innovation inputs (R&D) and outputs (patents).

R&D Mechanisms. In Appendix H.3, we explore the heterogeneous effects of AIPA on R&D along the dimensions of technology cycle, value of secrecy, and technology similarity. Specifically, we compute a firm's heterogeneous AIPA exposure as the average application-grant delay weighted by its patent portfolio in the subset of technology classes with above/below median: (1) length of technology cycle (measured as the median citation lag for the technology class during the three-year pre-AIPA period), (2) opt-out rates, and (3) technology similarity. We find that industries with shorter technology cycles – in which faster disclosure of knowledge post-AIPA arguably should have the greatest effect – are also the industries that increase R&D the most post-AIPA. Likewise, consistent with our analysis in Section 6, AIPA has larger effects on R&D investments in areas with high opt-out rates, our proxy for the value of secrecy. Lastly, we find that the AIPA effect on R&D is stronger in areas with low similarity, where knowledge diffusion is more limited to begin with. We note that the ability to stratify along more granular dimensions of heterogeneity is limited by our small sample of public firms. We therefore leave more detailed exploration of these mechanisms to future researchers with administrative Census bureau data, such as the Survey of Industrial Research and Development.

Summary. Economists have long used patents as proxies for innovation and patent citations as proxies for knowledge diffusion (e.g., Griliches, Pakes, and Hall (1987), Jaffe et al. (1993)). A recent body of careful empirical work has validated this usage by providing evidence that patents can be causally linked to increases in productivity, revenue growth, R&D investment, and better market valuation for their holders (e.g., Kogan et al. (2017), Farre-Mensa et al. (2020)). In the context of AIPA, Hegde and Luo (2018) find that patents are about 30 percentage points more likely to be licensed before grant after AIPA, which was interpreted as a reduction in the search costs of finding technology inputs after AIPA. Our study shows that accelerated patent publication had substantial effects on patenting, R&D, and citations by follow-on inventors. We also provided corroborative evidence that the mechanism behind these effects is enhanced knowledge diffusion. Combining our evidence with the prior literature, we believe it is plausible that the accelerated development of disclosed technologies by follow-on inventors (a subset of whom license the technology) is one of the mechanisms driving an increase in innovation outcomes post-AIPA.

8 Concluding thoughts

In this study, we provide causal estimates of the effects of patent publication on various measures of innovation. We make progress by exploiting the passage of AIPA — a law which expedited the disclosure of U.S. patent applications by nearly two years, on average — as a natural experiment. To guide our empirical analysis, we develop a theoretic framework in which AIPA provisions *news shocks* to inventors about recent technologies, and we derive a set of testable assumptions and predictions.

Consistent with the framework's implications, our event study analysis implies that AIPA (i) increased the magnitude and the pace of knowledge diffusion associated with U.S. patents; (ii) increased overlap between technologically distant patents; decreased overlap between similar patents; (iii) lowered inventive steps and patent scope; (iv) decreased patent abandonments; and (v) increased U.S. patenting.

We then establish the causal effects of AIPA by comparing U.S. patents subject to AIPA's accelerated 18 month disclosure with "twin" European patents which were always disclosed at 18 months. This "twin" study design allows us to control for unobserved characteristics of each patent family (comprising the U.S. patent and its EP twin) and to account for quality-based selection into patenting or early disclosure using family-fixed effects. This "twin" analysis largely aligns with our event-study analysis: (i) U.S. patents' follow-on citations increase; (ii) U.S. patents' time lag to citations decrease; and (iii) technological overlap increases between distant but related U.S. patents and decreases between highly similar U.S. patents.

Lastly, we link speedier disclosure of inventions to real measures of innovation, such as R&D, by exploiting firm-specific exposure to AIPA. Firms exposed to one standard deviation longer patent grant delays increased their R&D investment by 4% after AIPA, suggesting that AIPA had significant effects on incentives to innovate.

Overall, we provide causal evidence that patent publication matters and that rules governing the timing of publication can have a profound impact on follow-on innovation. Early disclosure appears to promote knowledge diffusion, to lower patenting costs, and to reduce technology duplication. We also find that overall, patent publication's effect on reducing the cost of innovation through knowledge diffusion dominates possible losses from free-ridership, resulting in a net increase in patenting and R&D investments. Our findings imply that welfare analyses of patents should incorporate their information disclosure effects.

References

Ufuk Akcigit and Qingmin Liu. The role of information in innovation and competition. *Journal of the European Economic Association*, 14(4):828–870, 2016.

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

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

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

Andrew Atkeson and Ariel Burstein. Aggregate implications of innovation policy. *Journal of Political Economy*, 127(6):2625–2683, 2019.

Andrew Atkeson and Ariel Tomas Burstein. Innovation, firm dynamics, and international trade. *Journal of political economy*, 118(3):433–484, 2010.

Joaquín Azagra Caro and Elena M Tur. Differences between examiner and applicant citations in the european patent office: a first approach. 2014.

E Bacchiocchi and F Montobbio. Technological opportunities or absorptive capacity?: an estimation of the rate of diffusion and decay of technical knowledge using patent citations. 2006.

Stefano H Baruffaldi and Markus Simeth. Patents and knowledge diffusion: The effect of early disclosure. *Research Policy*, 49(4):103927, 2020.

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

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

Nicholas Bloom, Mark Schankerman, and John Van Reenen. Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393, 2013.

Catherine Bobtcheff, Jérôme Bolte, and Thomas Mariotti. Researcher's dilemma. *The Review of Economic Studies*, 84(3):969–1014, 2017.

Laura Bottazzi and Giovanni Peri. Innovation and spillovers in regions: Evidence from european patent data. *European economic review*, 47(4):687–710, 2003.

Stefano Breschi and Francesco Lissoni. Knowledge spillovers and local innovation systems: a critical survey. *Industrial and corporate change*, 10(4):975–1005, 2001.

Wyatt Brooks, Kevin Donovan, and Terence R Johnson. Mentors or teachers? microenterprise training in kenya. *American Economic Journal: Applied Economics*, 10(4):196–221, 2018.

Wyatt Brooks, Kevin Donovan, and Terence R Johnson. From micro to macro in an equilibrium diffusion model. Technical report, Working Paper, 2020.

Kevin A Bryan and Jorge Lemus. The direction of innovation. *Journal of Economic Theory*, 172:247–272, 2017.

Wesley M Cohen, Akira Goto, Akiya Nagata, Richard R Nelson, and John P Walsh. R&d spillovers, patents and the incentives to innovate in japan and the united states. *Research policy*, 31(8-9):1349–1367, 2002.

Gaétan de Rassenfosse, Paul H Jensen, T'Mir Julius, Alfons Palangkaraya, and Elizabeth Webster. Are foreigners treated equally under the trade-related aspects of intellectual property rights agreement? *The Journal of Law and Economics*, 62(4):663–685, 2019.

Gaétan de Rassenfosse, Gabriele Pellegrino, and Emilio Raiteri. Do patents enable disclosure? evidence from the invention secrecy act. *Evidence from the Invention Secrecy Act (March 26, 2020)*, 2020.

Edward R Ergenzinger Jr. The american inventor's protection act: A legislative history. *Wake Forest Intell. Prop. LJ*, 7:145, 2006.

Joan Farre-Mensa, Deepak Hegde, and Alexander Ljungqvist. What is a patent worth? evidence from the us patent "lottery". *The Journal of Finance*, 75(2):639–682, 2020.

Jeanne C Fromer. Claiming intellectual property. *The University of Chicago Law Review*, pages 719–796, 2009.

Jeffrey L Furman, Markus Nagler, and Martin Watzinger. Disclosure and subsequent innovation: Evidence from the patent depository library program. Technical report, National Bureau of Economic Research, 2018.

Jeffrey L Furman, Markus Nagler, and Martin Watzinger. Disclosure and subsequent innovation: Evidence from the patent depository library program. *American Economic Journal: Economic Policy*, 13(4): 239–70, 2021.

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

Bernhard Ganglmair and Imke Reimers. Visibility of technology and cumulative innovation: Evidence from trade secrets laws. *ZEW-Centre for European Economic Research Discussion Paper*, (19-035), 2019.

Stuart Graham and Deepak Hegde. Disclosing patents' secrets. *Science*, 347(6219):236–237, 2015.

Stuart Graham and Saurabh Vishnubhakat. Of smart phone wars and software patents. *Journal of Economic Perspectives*, 27(1):67–86, 2013.

Victoria Gregory. *Firms as learning environments: Implications for earnings dynamics and job search.* Federal Reserve Bank of St. Louis, Research Division, 2020.

Zvi Griliches, Ariel Pakes, and Bronwyn Hall. The value of patents as indicators of inventive activity. *Economic policy and technological performance*, page 97, 1987.

Daniel P Gross. The consequences of invention secrecy: Evidence from the uspto patent secrecy program in world war ii. Technical report, National Bureau of Economic Research, 2019.

Tal Gross, Matthew J Notowidigdo, and Jialan Wang. The marginal propensity to consume over the business cycle. *American Economic Journal: Macroeconomics*, 12(2):351–84, 2020.

Bronwyn Hall, Christian Helmers, Mark Rogers, and Vania Sena. The choice between formal and informal intellectual property: a review. *Journal of Economic Literature*, 52(2):375–423, 2014.

Dietmar Harhoff, Karin Hoisl, and Colin Webb. European patent citations-how to count and how to interpret them. *University of Munich and CEPR (London), University of Munich, and OECD*, 2006.

Deepak Hegde and Hong Luo. Patent publication and the market for ideas. *Management Science*, 64(2): 652–672, 2018.

Deepak Hegde, David C Mowery, and Stuart JH Graham. Pioneering inventors or thicket builders: Which us firms use continuations in patenting? *Management Science*, 55(7):1214–1226, 2009.

Kyle Herkenhoff, Jeremy Lise, Guido Menzio, and Gordon M Phillips. Production and learning in teams. Technical report, National Bureau of Economic Research, 2018.

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

Ignatius Horstmann, Glenn M MacDonald, and Alan Slivinski. Patents as information transfer mechanisms: To patent or (maybe) not to patent. *Journal of Political Economy*, 93(5):837–858, 1985.

Adam B Jaffe and Manuel Trajtenberg. International knowledge flows: Evidence from patent citations. *Economics of innovation and new technology*, 8(1-2):105–136, 1999.

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

Adam B Jaffe et al. Technological opportunity and spillovers of r&d: Evidence from firms' patents, profits, and market value. *American Economic Review*, 76(5):984–1001, 1986.

Gregor Jarosch, Ezra Oberfield, and Esteban Rossi-Hansberg. Learning from coworkers. *Econometrica*, 89(2):647–676, 2021.

Daniel KN Johnson and David Popp. Forced out of the closet: The impact of the american inventors protection act on the timing of patent disclosure. *RAND Journal of Economics*, pages 96–112, 2003.

Jinhwan Kim and Kristen Valentine. The innovation consequences of mandatory patent disclosures. *Journal of Accounting and Economics*, 71(2-3):101381, 2021.

Leonid Kogan, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics*, 132(2):665–712, 2017.

Jeffrey M Kuhn and Neil C Thompson. How to measure and draw causal inferences with patent scope. *International Journal of the Economics of Business*, 26(1):5–38, 2019.

Richard C Levin, Alvin K Klevorick, Richard R Nelson, Sidney G Winter, Richard Gilbert, and Zvi Griliches. Appropriating the returns from industrial research and development. *Brookings papers on economic activity*, 1987(3):783–831, 1987.

Robert E Lucas Jr and Benjamin Moll. Knowledge growth and the allocation of time. *Journal of Political Economy*, 122(1):1–51, 2014.

Sonja Lück, Benjamin Balsmeier, Florian Seliger, and Lee Fleming. Early disclosure of invention and reduced duplication: An empirical test. *Management Science*, 66(6):2677–2685, 2020.

Erzo GJ Luttmer. *An assignment model of knowledge diffusion and income inequality*. Federal Reserve Bank of Minneapolis, Research Department, 2014.

Per Botolf Maurseth and Bart Verspagen. Knowledge spillovers in europe: a patent citations analysis. *Scandinavian Journal of Economics*, 104(4):531–545, 2002.

Fraco Modigliani et al. An open letter to the u.s. senate. *https://eagleforum.org/patent/nobel_letter.html*, 1999.

Ali Mohammadi and Pooyan Khashabi. Embracing the sharks: The impact of information disclosure on cvc investments. In *Academy of Management Proceedings*, volume 2017, page 15980. Academy of Management Briarcliff Manor, NY 10510, 2017.

Emily Nix. Learning spillovers in the firm. Technical report, Working paper, 2020.

Yoshimi Okada and Sadao Nagaoka. Effects of early patent publication on knowledge dissemination: Evidence from us patent law reform. *Information Economics and Policy*, 51:100852, 2020.

Lisa Larrimore Ouellette. Who reads patents? *Nature biotechnology*, 35(5):421–424, 2017.

Giovanni Peri. Determinants of knowledge flows and their effect on innovation. *Review of economics and Statistics*, 87(2):308–322, 2005.

Jesse Perla and Christopher Tonetti. Equilibrium imitation and growth. *Journal of Political Economy*, 122(1):52–76, 2014.

Nirupama Rao. Do tax credits stimulate r&d spending? the effect of the r&d tax credit in its first decade. *Journal of Public Economics*, 140:1–12, 2016.

Benjamin N Roin. The disclosure function of the patent system (or lack thereof). *Harvard Law Review*, 2005.

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

Bhaven Sampat and Heidi L Williams. How do patents affect follow-on innovation? evidence from the human genome. *American Economic Review*, 109(1):203–36, 2019.

Suzanne Scotchmer and Jerry Green. Novelty and disclosure in patent law. *The RAND Journal of Economics*, pages 131–146, 1990.

Manuel Trajtenberg, Rebecca Henderson, and Adam Jaffe. University versus corporate patents: A window on the basicness of invention. *Economics of Innovation and new technology*, 5(1):19–50, 1997.

Elizabeth Webster, Paul H Jensen, and Alfons Palangkaraya. Patent examination outcomes and the national treatment principle. *The RAND Journal of Economics*, 45(2):449–469, 2014.

Heidi L Williams. How do patents affect research investments? *Annual Review of Economics*, 9:441–469, 2017.

Arthur Wineburg. Japanese patent system: A non-tariff barrier to foreign businesses. *J. World Trade*, 22:11, 1988.

Figure 1: Thomas Edison's Light Bulb Patent

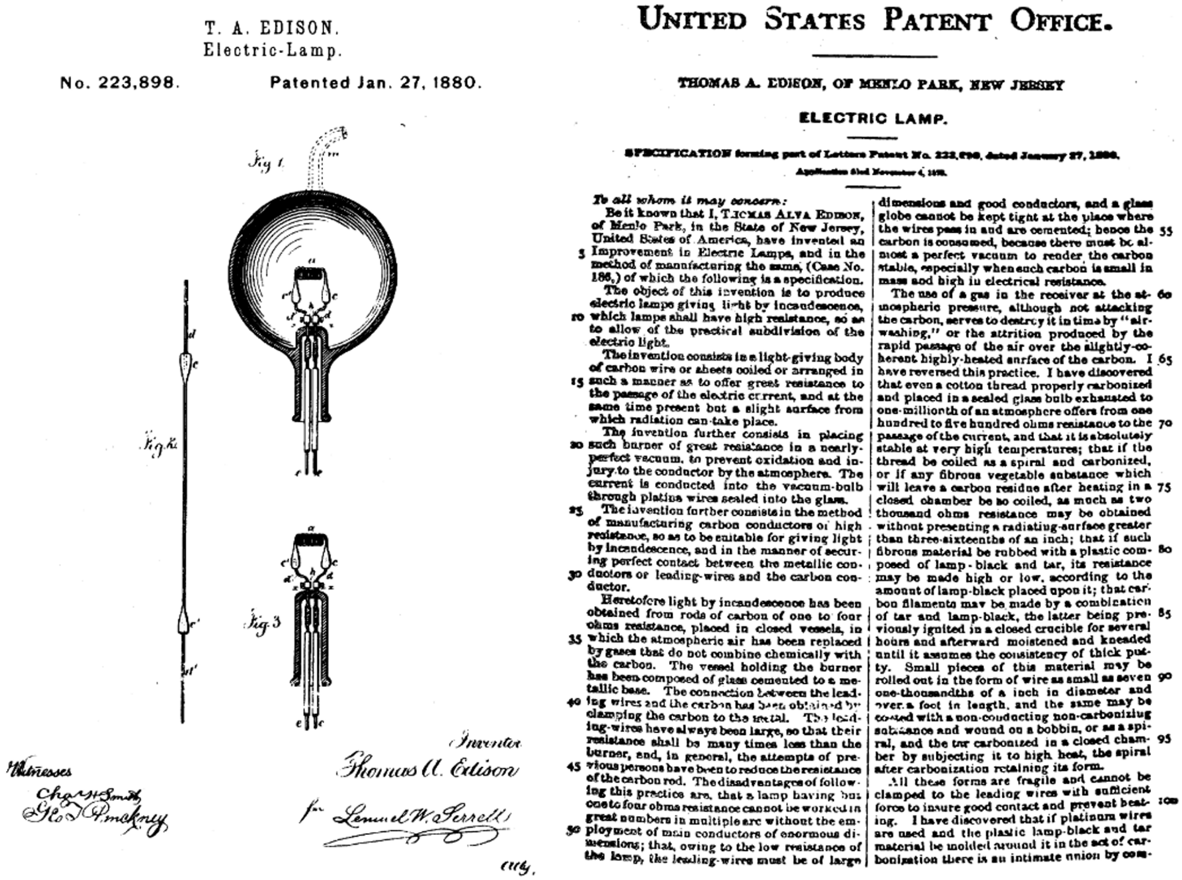
T. A. EDISON.
Electric-Lamp.

No. 223,898. Patented Jan. 27, 1880.

Fig. 1

Fig. 2

Fig. 3

Witnesses
Chas H Smith
Geo T Pinckney

Inventor
Thomas A. Edison
per Lemuel W. Serrell
atty.

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: Model Timing

This figure depicts the timing of our model's stages (1) Pre-AIPA and (2) Post-AIPA.

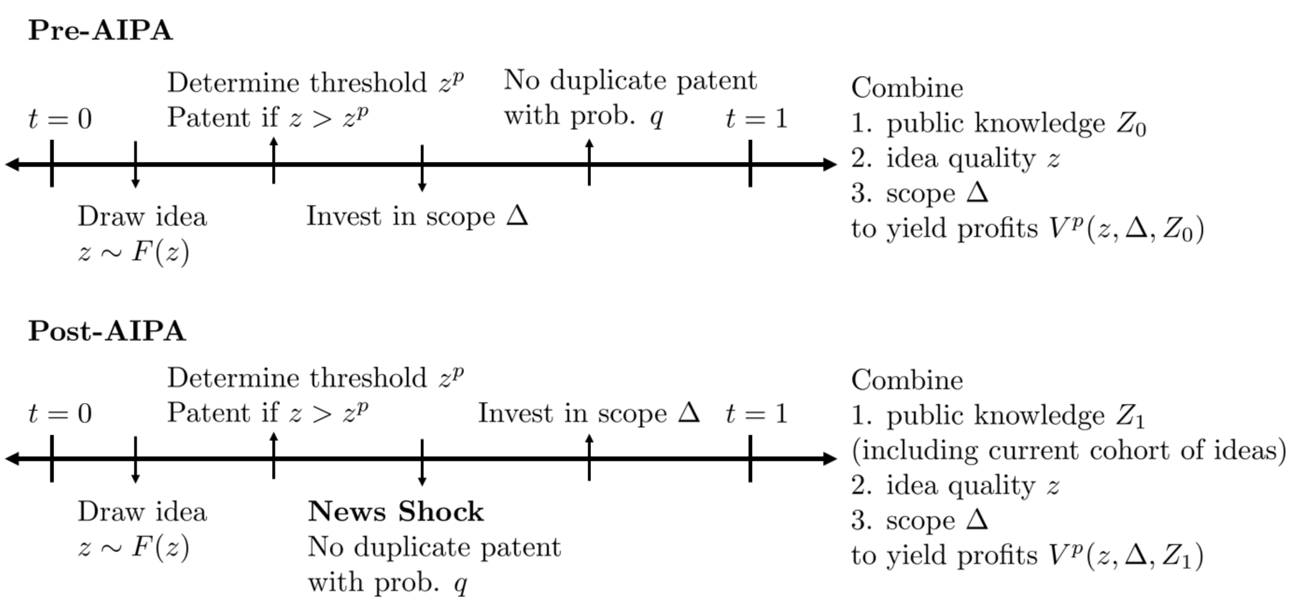

Figure 3: Citations to U.S. patents before and after AIPA

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).

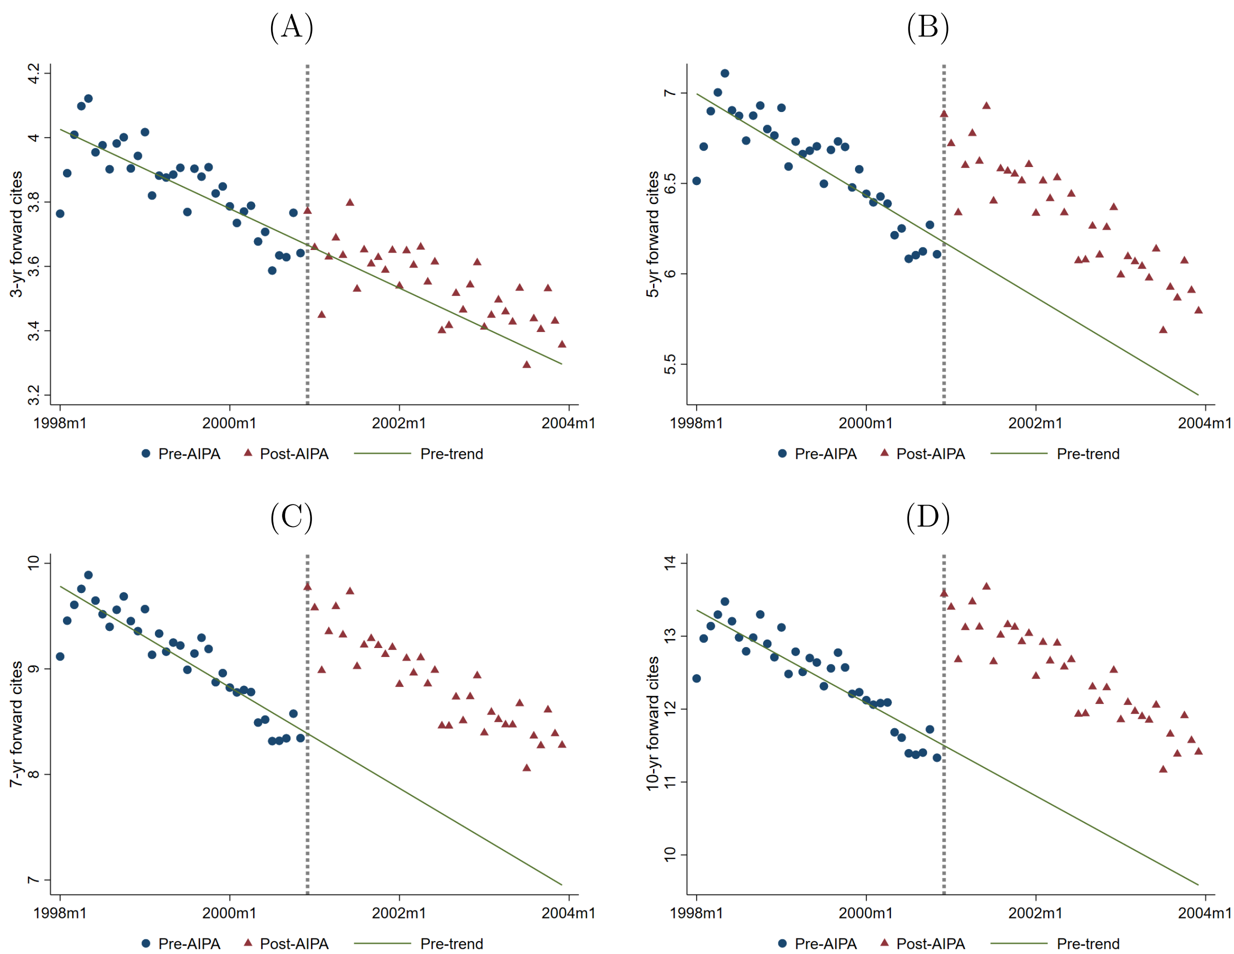

Figure 4: Citation lags of U.S. patents before and after AIPA

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 first or first 3, 5, 7 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).

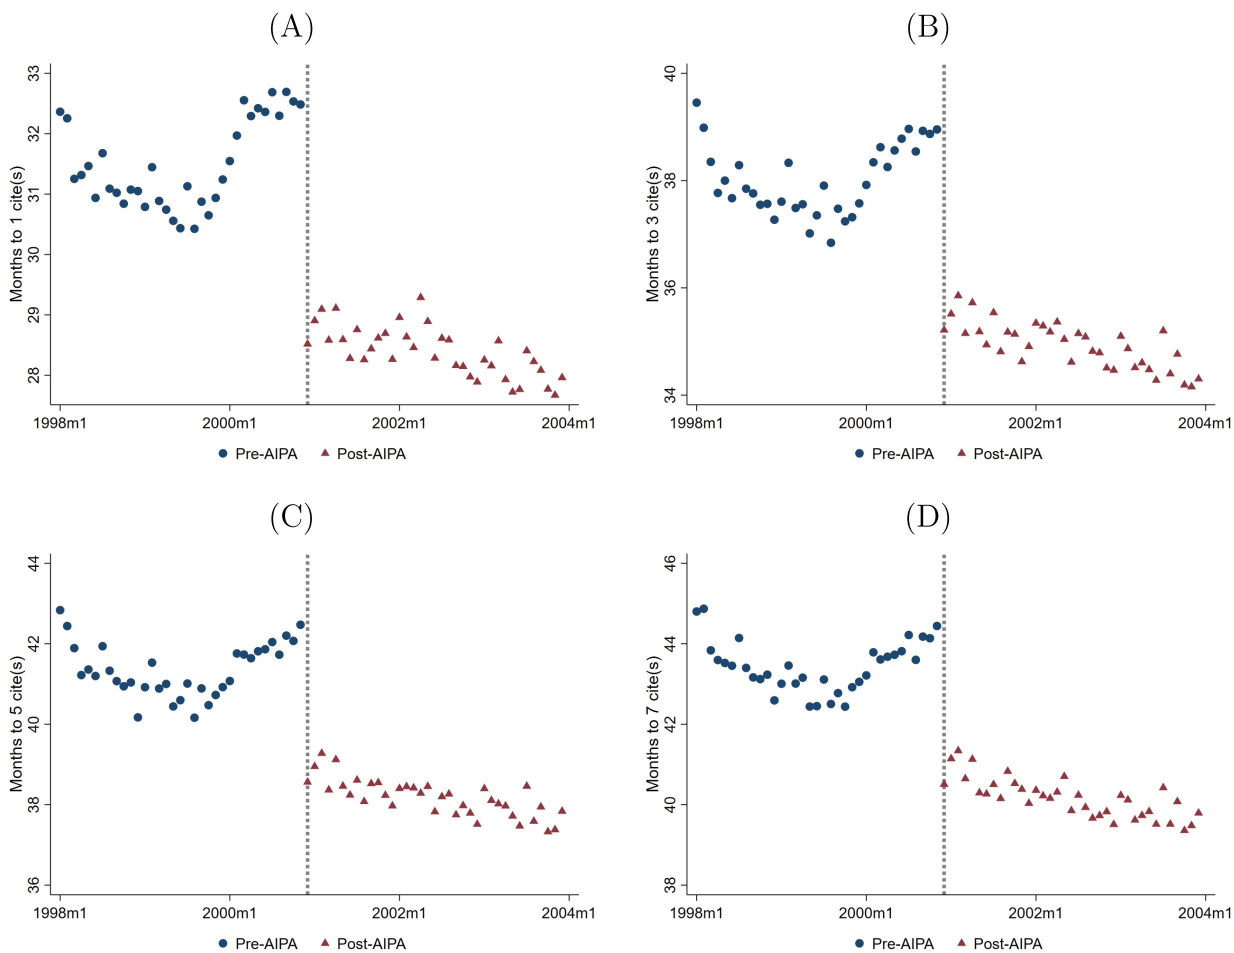

Figure 5: Technology similarity of U.S. patents before and after AIPA

The figures plot the monthly average technology similarity between U.S. patents filed during 1998-2003 and their "next-generation" patents. Similarity is measured as the pair-level cosine similarity, 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).

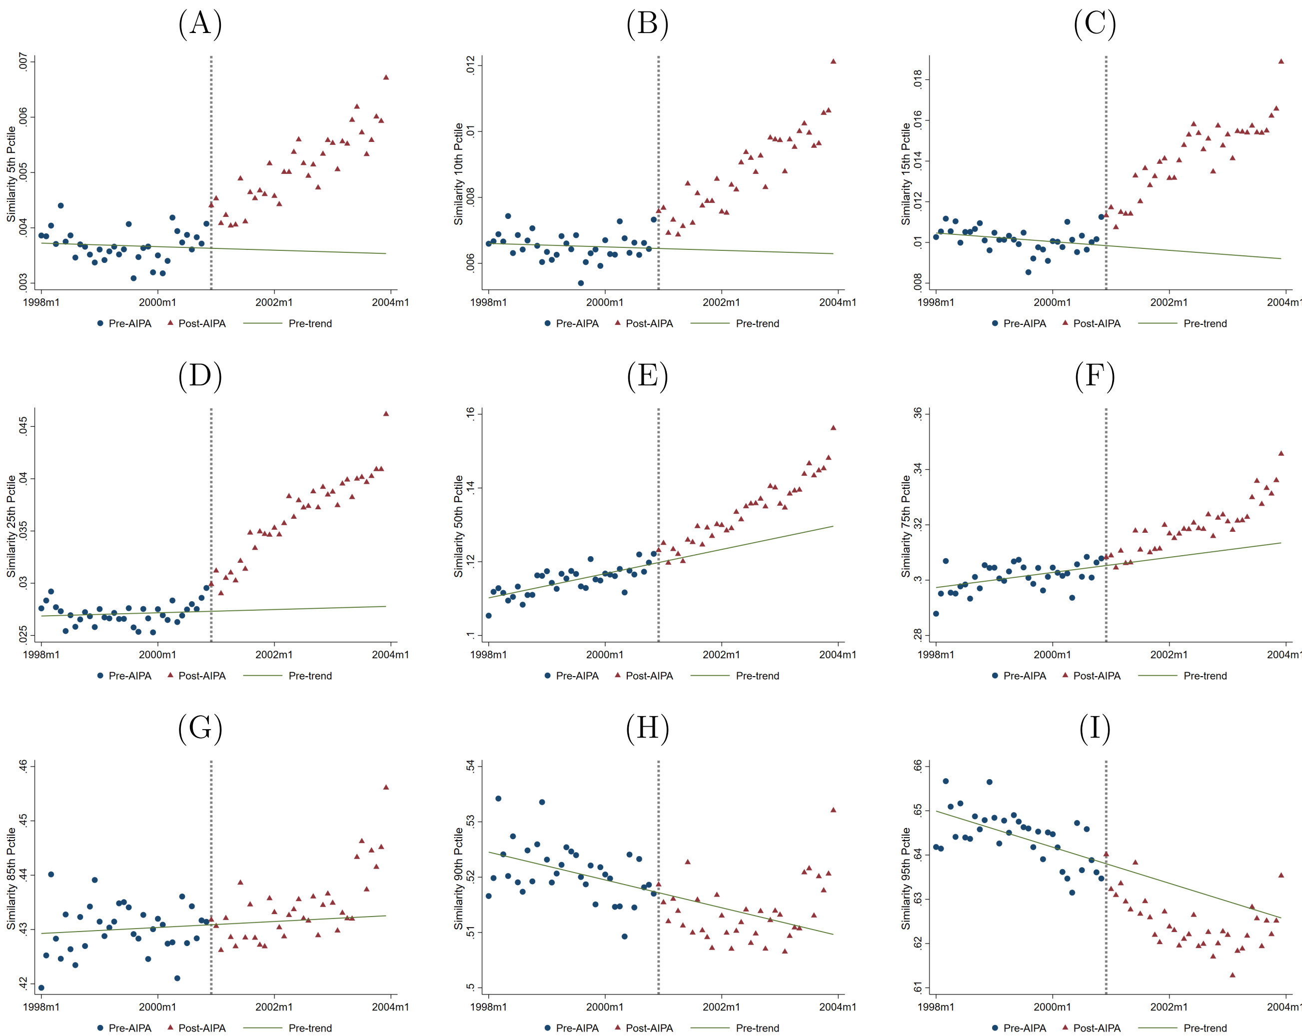

Figure 6: U.S. patent renewal, originality, claims, and abandonment

The figures plot the average renewal rates in Panel A and originality index in Panel B by application month during 1998-2003. Panels C-E plot the average number of total allowed claims, independent claims, and average words per independent claim for issued patents. Panel F plots the monthly average abandonment rates. 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).

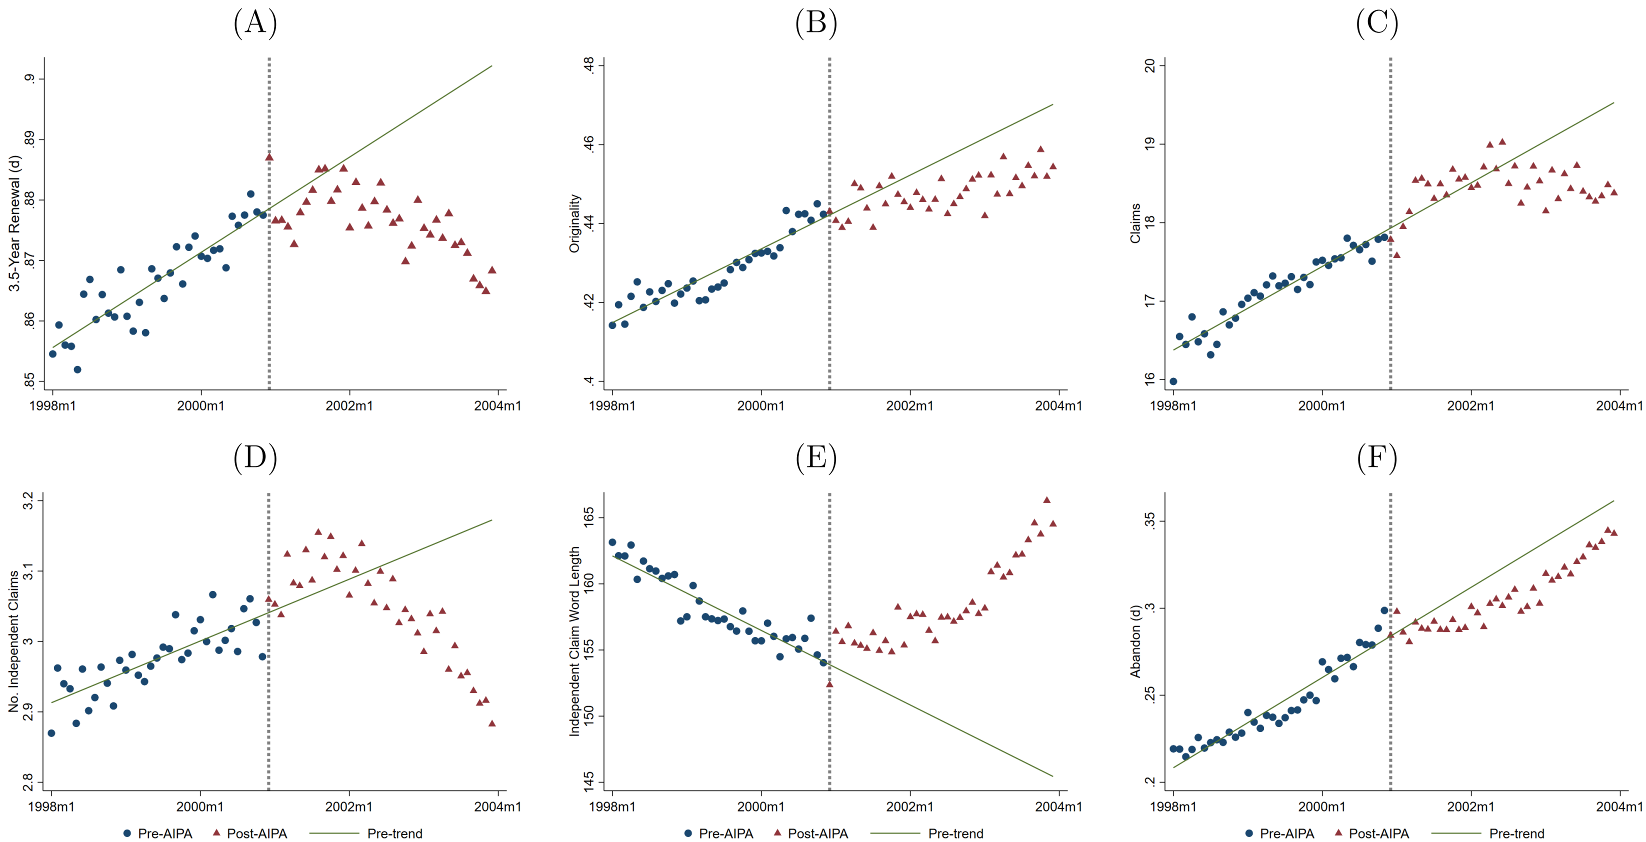

Figure 7: Citations to US-EP “twins” before and after AIPA

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). We normalize the average by its value at the beginning of the sample period. Citations data are obtained from the USPTO and PATSTAT. The vertical line represents AIPA’s effective date (November 29, 2000).

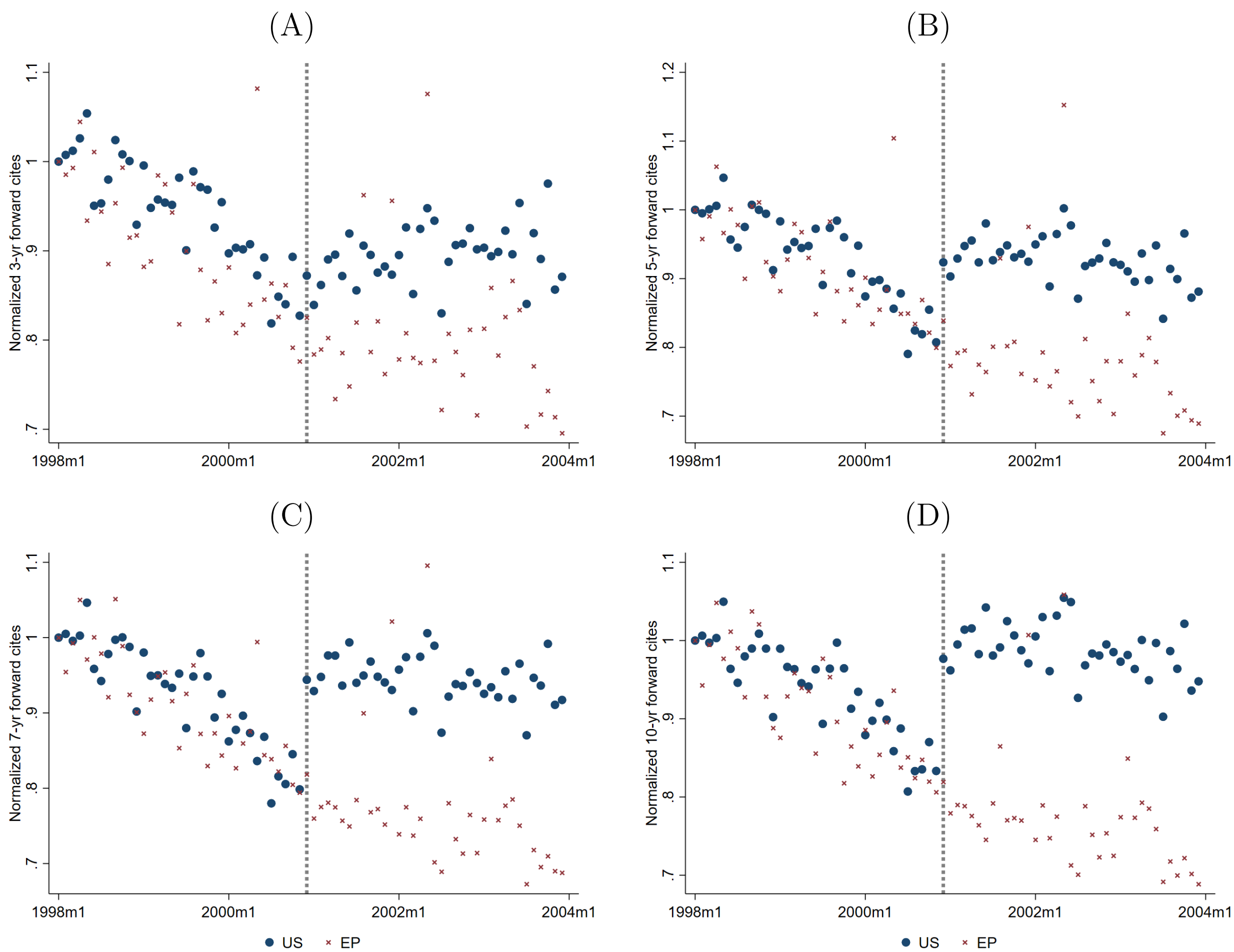

Figure 8: Citation lags of US-EP "twins" before and after AIPA

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).

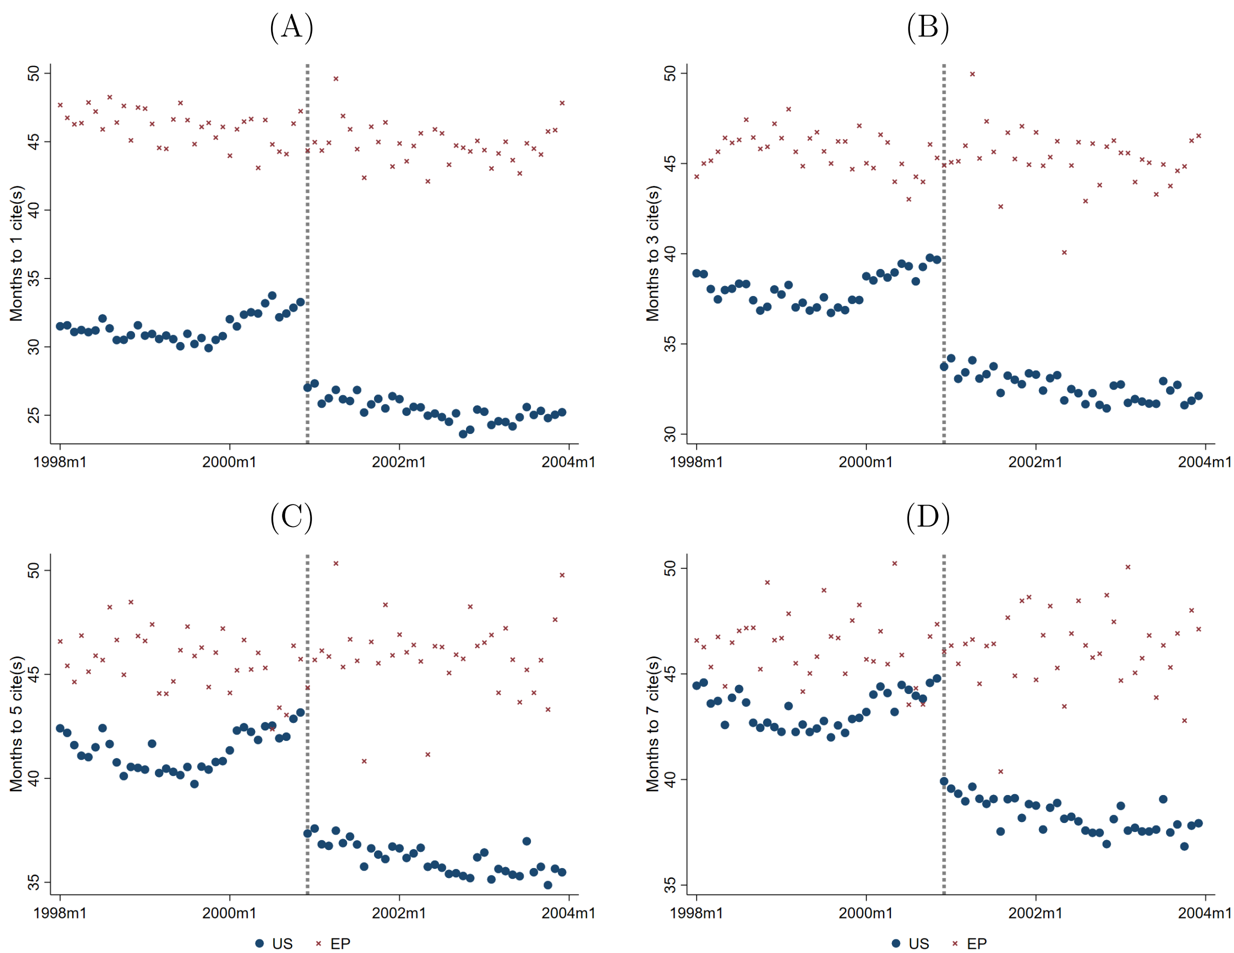

Figure 9: Similarity of US-EP "twins" before and after AIPA

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).

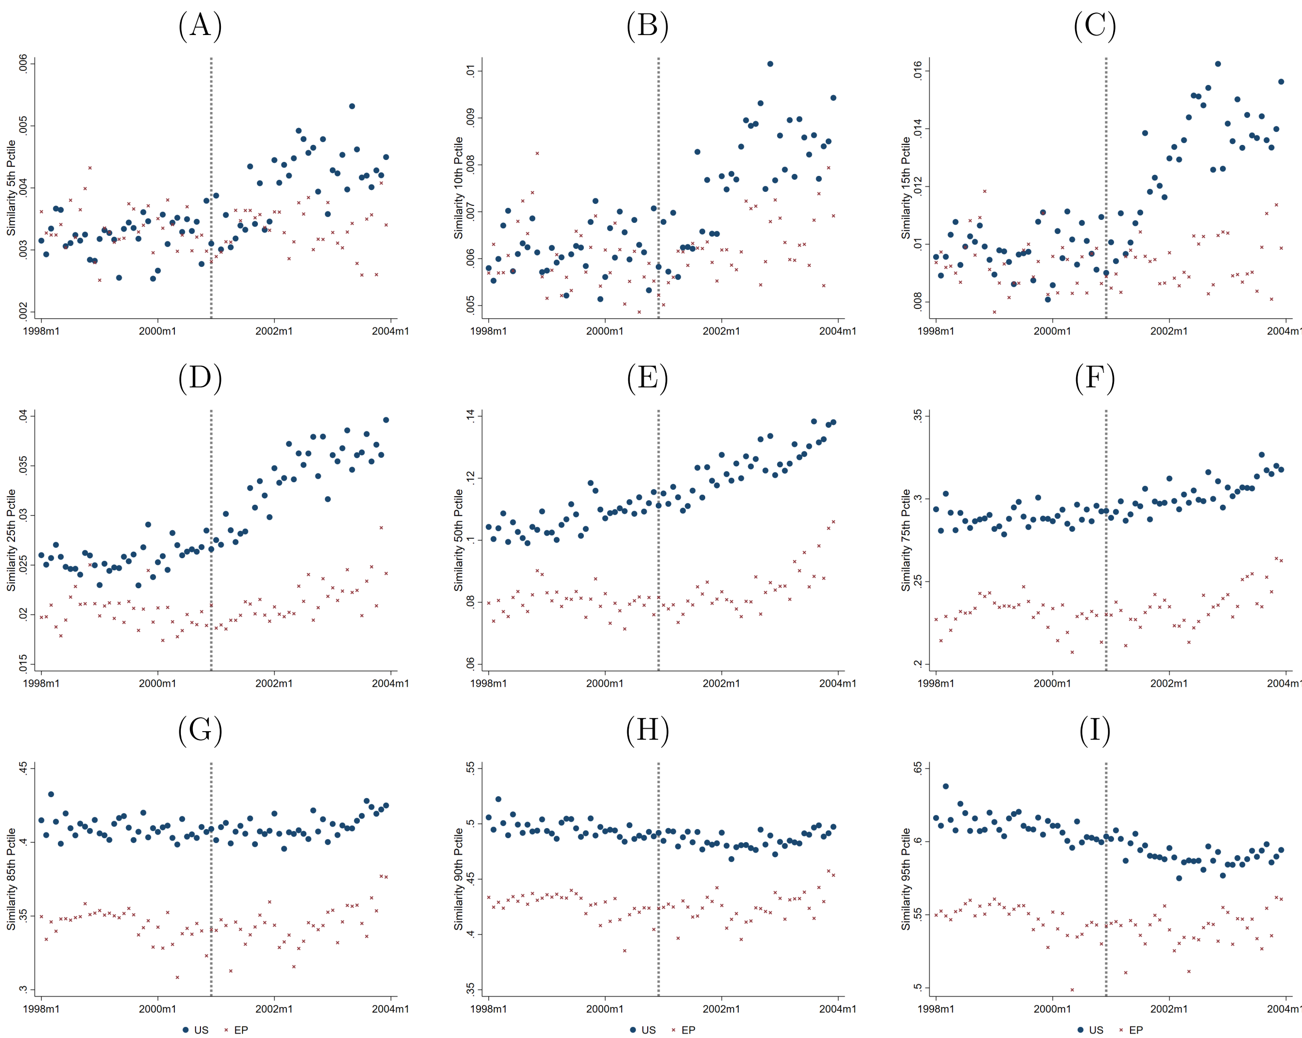

Figure 10: AIPA's effect on technology similarity at different percentiles (ventiles)

This figure plots the estimated AIPA effect on technology similarity measured at different percentiles (ventiles) between focal patents and next-generation patents. The estimated AIPA effect is the coefficient on the interaction term 'US × Post AIPA (d)'. Refer to Table 6 notes for a description of the regression specification.

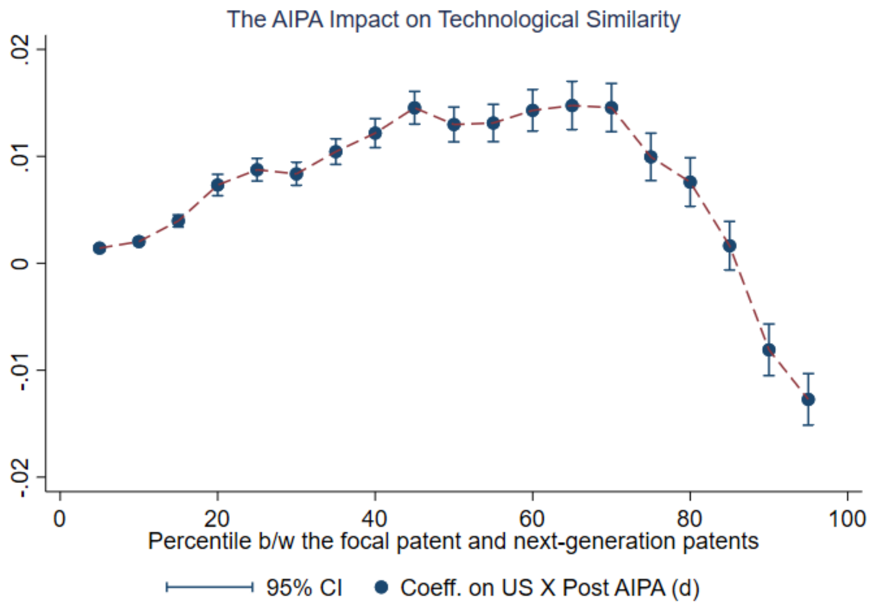

Figure 11: Patenting Intensity (U.S. v. EP)

The figures below 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).

(A)

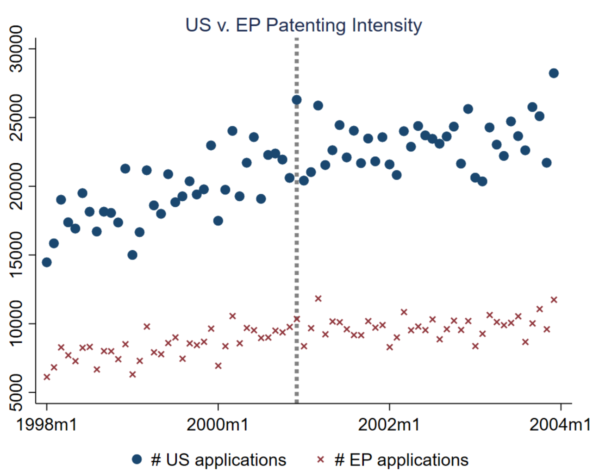

(B)

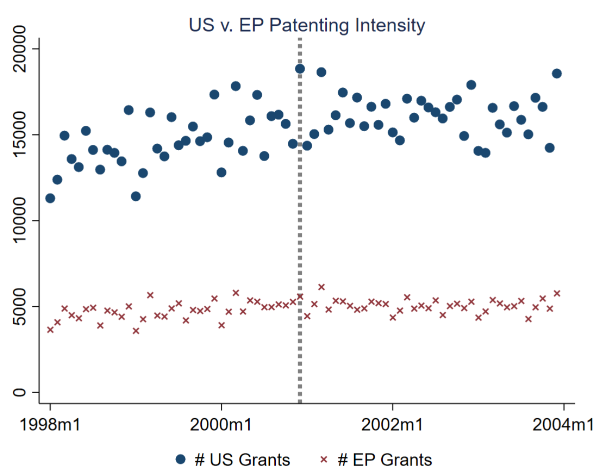

Figure 12: Compare the impact of R&D tax credit and AIPA on R&D

We compute the effect of AIPA on the R&D to sales ratio. Rao (2016) reports the effect of R&D user costs on R&D to sales ratios. We use Rao (2016)'s R&D user cost formula $\rho = \frac{(r+\delta)(1-\tau-c)}{1-\tau}$ *where* r *is the interest rate,* δ *is the depreciation rate,* τ *is the corporate tax rate, and* c *is the R&D tax credit rate. Following Rao (2016), we set* $\delta = 0.15$, $c = 0.25$ *and we vary the interest rate and corproate tax rates on the axes. We set the carry-forward provisions to zero.*

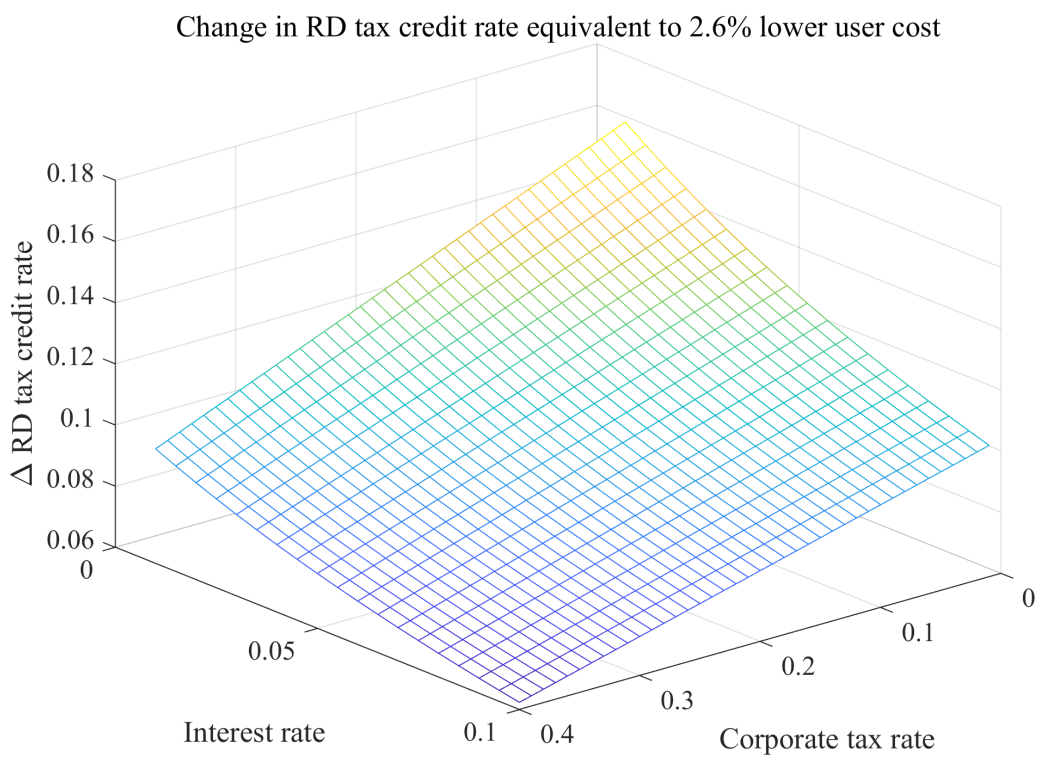

Table 1: Sample Counts

	# US applications	# EP equivalents	Avg Twin Size
U.S. applications filed between Jan. 1, 1998 and Dec. 31, 2003	1,536,346		
Exclude US ungranted patent applicants (abandoned or pending)	1,107,710		
Exclude US patent applicants that dont match to EP equivalents	403,292	467,587	2.16
Require EP filings are non-PCT and filed w/in 18mo of US equivalent	316,563	354,227	2.12

Table 2: Variable definitions

Variable	**Definition**
Opt Out (d)	Dummy variable, equal one if the patent application is filed after the enactment of AIPA and opts out of the pre-grant publication requirement.
Early Grant (d)	Dummy variable, equal one if the patent application is granted 18 months after application.
3-yr forward cites	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).
X-yr forward cites	The number of forward citations received within X 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).
Months to 1 cite	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).
Months to X cites	The average time lag to receive the first X forward citations conditional on having at least X forward citations within ten years of application (unit: month).
Similarity X^{th} Pctile	The Xth percentile of the pair-wise cosine similarity based on the distribution 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 primary 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-yr Renewal (d)	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.
Claims	Total number of claims allowed at grant.
Independent Claims	The number of independent claims allowed at grant.
Independent Claim Word Length	The average number of words per independent claim.
Abandon (d)	Dummy variable, equal one if the application is abandoned.

Table 3: Summary Statistics

This table reports the summary statistics of key variables of interest for the universe of U.S. patents (patent applications filed in the USPTO from 1998 to 2003 and granted by mid-2014) and U.S. patents with EP equivalents. The sample composition is indicated in the table header. For more details on the variable definitions, please refer to Table 2. We use (d) to denote dummy indicator variables.

	(1) All US Patents #=1,107,710		(2) US Patents with EP Equivalents #=316,563	
Variable	Mean	S.D.	Mean	S.D.
3-yr forward cites	3.885	7.599	4.103	8.189
5-yr forward cites	6.779	13.079	7.261	14.29
7-yr forward cites	9.472	18.518	10.236	20.455
10-yr forward cites	13.191	26.718	14.433	29.938
Months to 1 cite	36.4	21.561	36.788	21.685
Months to 3 cites	39.667	20.45	39.888	20.579
Months to 5 cites	41.69	19.875	41.853	20.005
Months to 7 cites	43.631	19.318	43.747	19.466
Similarity 5th Pctile	0.005	0.034	0.005	0.031
Similarity 10th Pctile	0.009	0.048	0.008	0.044
Similarity 15th Pctile	0.013	0.062	0.013	0.058
Similarity 25th Pctile	0.033	0.101	0.031	0.095
Similarity 50th Pctile	0.126	0.215	0.117	0.2
Similarity 75th Pctile	0.311	0.319	0.296	0.302
Similarity 85th Pctile	0.432	0.334	0.41	0.319
Similarity 90th Pctile	0.517	0.326	0.49	0.314
Similarity 95th Pctile	0.634	0.299	0.601	0.292
Early Grant (d)	0.17	0.376	0.134	0.341
3.5-yr Renewal	0.872	0.334	0.889	0.314
Originality	0.439	0.276	0.441	0.276
Claims	18.24	17.2	19.086	17.708
Independent Claims	3.06	2.57	3.052	2.679
Independent Claim Word Length	160.252	104.049	156.637	102.877
Opt-Out Rate (2000-2003)	0.086	0.280	–	–
Abandon (d) (1,536,346 Applications)	0.279	0.449	–	–

Table 4: AIPA's effect on knowledge diffusion (extensive margin): "twin" analysis

This table reports DID estimates of AIPA's effect on the extent of knowledge diffusion. The sample consists of U.S. patents filed during 1998-2003 and their equivalent applications filed at the EPO. The regressions are specified as follows:

$$Outcome_{ijt} = \alpha_1 + \alpha_2 \boldsymbol{I}(US_j) + \alpha_3 \boldsymbol{I}(US_j) \times \boldsymbol{I}(Post\ AIPA_t) + \delta W_j + Family_i + Month_t + \epsilon_{ijt}$$

*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 $\boldsymbol{I}(US_j) \times \boldsymbol{I}(Post\ AIPA_t)$. 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 coefficients on the interactions are reported. Standard errors are clustered by application month for U.S. and EP patents, separately. ***, **, and * indicate 1%, 5%, and 10% significance levels, resp. (d) denotes dummy indicator variable.*

Panel A: Main Analyses of Forward Citations

	(1) Log 3-Yr. Forward Cites	(2) Log 5-Yr. Forward Cites	(3) Log 7-Yr. Forward Cites	(4) Log 10-Yr. Forward Cites
Post AIPA × US (d)	-0.017**	0.057***	0.108***	0.147***
	(0.008)	(0.009)	(0.010)	(0.011)
US (d)	0.806***	0.986***	1.093***	1.207***
	(0.006)	(0.008)	(0.008)	(0.009)
Granted (d)	0.191***	0.242***	0.276***	0.299***
	(0.006)	(0.006)	(0.007)	(0.007)
Early Grant (d)	0.419***	0.516***	0.572***	0.606***
	(0.009)	(0.010)	(0.010)	(0.011)
Early Grant × US (d)	-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

	(1) Log 3-Yr. Forward Cites	(2) Log 5-Yr. Forward Cites	(3) Log 7-Yr. Forward Cites	(4) Log 10-Yr. Forward Cites
Citations by U.S. Patents				
Post AIPA × US (d)	-0.018**	0.048***	0.094***	0.133***
	(0.008)	(0.010)	(0.011)	(0.012)
Citations by E.P. Patents				
Post AIPA × US (d)	-0.004	0.018***	0.036***	0.048***
	(0.002)	(0.003)	(0.003)	(0.004)

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. The sample consists of U.S. patents filed during 1998-2003 and their equivalent applications filed at the EPO. The dependent variable is the average time between the patent application dates of the focal patent and its first X (=1/3/5/7) forward citations. 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 specified as follows:

$$Outcome_{ijt} = \alpha_1 + \alpha_2 \boldsymbol{I}(US_j) + \alpha_3 \boldsymbol{I}(US_j) \times \boldsymbol{I}(Post\ AIPA_t) + \delta W_j + Family_i + Month_t + \epsilon_{ijt}$$

*where j indicates the patent application belonging to family i and filed in the year t. Standard errors are clustered by application month for U.S. and EP patents, separately. ***, **, and * indicate 1%, 5%, and 10% significance levels, resp. (d) denotes dummy indicator variable.*

	(1) Log Months to 1 Cite	(2) Log Months to 3 Cites	(3) Log Months to 5 Cites	(4) Log Months to 7 Cites
Post AIPA × US (d)	-0.254***	-0.294***	-0.292***	-0.274***
	(0.013)	(0.016)	(0.023)	(0.028)
US (d)	-0.394***	-0.214***	-0.150***	-0.121***
	(0.011)	(0.012)	(0.014)	(0.019)
Granted (d)	-0.201***	-0.178***	-0.147***	-0.139***
	(0.015)	(0.015)	(0.021)	(0.032)
Early Grant (d)	-0.393***	-0.342***	-0.298***	-0.255***
	(0.015)	(0.021)	(0.031)	(0.043)
Early Grant × US (d)	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 patent similarity: "twin" analysis

This table reports DID estimates of AIPA's effect on technological similarity. The sample consists of U.S. patents filed during 1998-2003 and their equivalent applications filed at the EPO. The regressions are specified as follows:

$$Outcome_{ijt} = \alpha_1 + \alpha_2 \boldsymbol{I}(US_j) + \alpha_3 \boldsymbol{I}(US_j) \times \boldsymbol{I}(Post\ AIPA_t) + \delta W_j + Family_i + Month_t + \epsilon_{ijt}$$

*where j indicates the patent application, belonging to family i and filed in year t. Standard errors are clustered by application month for U.S. and EP patents, separately. ***, **, and * indicate 1%, 5%, and 10% significance levels, resp. (d) denotes dummy indicator variable.*

	(1) Similarity 50th Pctile	(2) Similarity 75th Pctile	(3) Similarity 90th Pctile	(4) Similarity 95th Pctile
Post AIPA × US (d)	0.013***	0.010***	-0.008***	-0.013***
	(0.001)	(0.001)	(0.001)	(0.001)
US (d)	0.023***	0.052***	0.063***	0.058***
	(0.001)	(0.001)	(0.001)	(0.001)
Granted (d)	0.009***	0.018***	0.010***	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)
Early Grant (d)	-0.001	0.006***	0.008***	0.008***
	(0.002)	(0.002)	(0.002)	(0.002)
Early Grant × US (d)	-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 7: US-EP versus US-JP twin analyses

*This table compares DID estimates of AIPA's effect on the extent and the speed of knowledge diffusion across US-EP twins and US-JP twins. The dependent variable is the natural logarithm of one plus 3/5/7/10-year forward citations (excluding self-citations) in Panel A and the time lag of forward citations in Panel B. Columns (1), (3), (5), and (7) use US-JP twins that do not have EP equivalents while Columns (2), (4), (6), and (8) use our main sample (i.e., US-EP twins). Differences in AIPA effects across the two samples are tested using two-tailed t-tests, and reported at the bottom of the panels. All regression specifications are the same as our main specifications except that we add country-specific linear time trends in all regressions. For brevity, only the coefficients on the interactions ('Post AIPA × US(d)') are reported. Standard errors are clustered by application month for U.S., EP, and JP patents, separately. ***, **, and * indicate 1%, 5%, and 10% significance levels, resp. (d) denotes dummy indicator variable.*

Panel A. Forward citations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log 3-Yr. Forward Cites		**Log 5-Yr. Forward Cites**		**Log 7-Yr. Forward Cites**		**Log 10-Yr. Forward Cites**	
	US-JP	US-EP	US-JP	US-EP	US-JP	US-EP	US-JP	US-EP
Post AIPA × US (d)	0.042*	-0.025	0.129***	0.070**	0.175***	0.114***	0.207***	0.141***
	(0.025)	(0.030)	(0.029)	(0.035)	(0.032)	(0.038)	(0.033)	(0.041)
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Family FE	Y	Y	Y	Y	Y	Y	Y	Y
Month FE	Y	Y	Y	Y	Y	Y	Y	Y
Observations	669,877	247,903	668,218	247,805	659,154	247,071	642,640	243,208
R-squared	0.710	0.738	0.743	0.770	0.762	0.784	0.778	0.797
Difference in AIPA Effect	0.068**		0.059**		0.061*		0.066**	
	(0.028)		(0.029)		(0.032)		(0.033)	

Panel B. Citation lags

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log Months to 1 Cite		**Log Months to 3 Cites**		**Log Months to 5 Cites**		**Log Months to 7 Cites**	
	US-JP	US-EP	US-JP	US-EP	US-JP	US-EP	US-JP	US-EP
Post AIPA × US (d)	-0.366***	-0.196***	-0.278***	-0.236***	-0.275***	-0.242***	-0.259***	-0.211***
	(0.021)	(0.026)	(0.023)	(0.021)	(0.027)	(0.022)	(0.032)	(0.023)
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Family FE	Y	Y	Y	Y	Y	Y	Y	Y
Month FE	Y	Y	Y	Y	Y	Y	Y	Y
Observations	317,853	138,280	149,423	65,699	81,400	34,544	49,424	19,989
R-squared	0.616	0.674	0.635	0.712	0.637	0.725	0.641	0.737
Difference in AIPA Effect	-0.170***		-0.042		-0.032		-0.048	
	(0.030)		(0.031)		(0.034)		(0.041)	

Table 8: Heterogeneous effects of AIPA: PCT Filing Status and Inventor Country of Origin

*This table reports the heterogeneous AIPA effects by Patent Cooperation Treaty (PCT) filing status in Panel A and inventors' country of origin in Panel B on forward citations, excluding self-citations, (Columns (1)-(4)) and citation lags (Columns (5)-(8)). When the US-EP patent family has affiliated PCT filings, we set the PCT filing indicator to be one ('PCT Filing (d)'). We use the first inventor's country of residence and code it as one if the inventor resides outside the US ('Foreign Inventor (d)'). Each family has the same PCT filing status and first inventor, hence these standalone terms are absorbed by family fixed effects. The regression specification is the same as the ones used to obtain our main results except that we add interaction terms here. We report only the key coefficients for brevity. Standard errors are clustered by application month for U.S. and EP patents, separately. ***, **, and * indicate 1%, 5%, and 10% significance levels, resp.*

Panel A: PCT Filing Status

	(1) Log 3-Yr. Forward Cites	(2) Log 5-Yr. Forward Cites	(3) Log 7-Yr. Forward Cites	(4) Log 10-Yr. Forward Cites	(5) Log Months to 1 Cite	(6) Log Months to 3 Cites	(7) Log Months to 5 Cites	(8) Log Months to 7 Cites
Post AIPA × US × PCT Filing (d)	-0.126***	-0.111***	-0.108***	-0.105***	0.172***	0.068***	0.014	0.020
	(0.011)	(0.012)	(0.013)	(0.013)	(0.016)	(0.021)	(0.034)	(0.037)
Post AIPA × US (d)	0.027***	0.090***	0.137***	0.170***	-0.289***	-0.283***	-0.271***	-0.257***
	(0.006)	(0.008)	(0.009)	(0.009)	(0.010)	(0.010)	(0.010)	(0.010)
Observations	670,108	669,608	668,081	658,352	317,853	149,423	81,400	49,424
Adj R-squared	0.491	0.557	0.590	0.619	0.312	0.362	0.370	0.380

Panel B: Foreign Inventors

	(1) Log 3-Yr. Forward Cites	(2) Log 5-Yr. Forward Cites	(3) Log 7-Yr. Forward Cites	(4) Log 10-Yr. Forward Cites	(5) Log Months to 1 Cite	(6) Log Months to 3 Cites	(7) Log Months to 5 Cites	(8) Log Months to 7 Cites
Post AIPA × US × Foreign Inventor (d)	0.122***	0.114***	0.108***	0.097***	-0.164***	-0.083***	-0.026	-0.020
	(0.014)	(0.015)	(0.016)	(0.018)	(0.019)	(0.023)	(0.038)	(0.038)
Post AIPA × US (d)	-0.100***	-0.024**	0.025**	0.072***	-0.140***	-0.233***	-0.274***	-0.261***
	(0.009)	(0.011)	(0.011)	(0.012)	(0.015)	(0.017)	(0.032)	(0.033)
Observations	669,877	668,218	659,154	642,640	317,853	149,423	81,400	49,424
Adj R-squared	0.469	0.532	0.565	0.592	0.286	0.330	0.340	0.354

Table 9: Heterogeneous effects of AIPA: Pre-AIPA Grant Delays and Opt-Out Exposure

*This table reports the heterogeneous AIPA effects by pre-AIPA application-grant time delay in Panel A and opt-out prevalence in Panel B on forward citations, excluding self-citations, (Columns (1)-(4)) and citation lags (Columns (5)-(8)). We compute the average time delay (in years) by IPC main technology class (4-digit) for US patents filed in the three-year pre-AIPA period ('Application Grant Delay'). We measure the prevalence of opt-outs by computing the percentage of US patents that are filed during the 3-year post-AIPA period (2000-2003) in each 4-digit IPC class ('Opt-Out Ratio'), since opt-outs are only available post-AIPA. Both 'Application Grant Delay' and 'Opt-Out Ratio' vary by technology class and do not not vary over time, hence the standalone term is absorbed by family fixed effects. The regression specification is the same as the ones used to obtain our main results except that we add interaction terms here. We report only the key coefficients for brevity. Standard errors are clustered by application month for U.S. and EP patents, separately. ***, **, and * indicate 1%, 5%, and 10% significance levels, resp.*

Panel A: Pre-AIPA Grant Delay Exposure

	(1) Log 3-Yr. Forward Cites	(2) Log 5-Yr. Forward Cites	(3) Log 7-Yr. Forward Cites	(4) Log 10-Yr. Forward Cites	(5) Log Months to 1 Cite	(6) Log Months to 3 Cites	(7) Log Months to 5 Cites	(8) Log Months to 7 Cites
Post AIPA × US × Application Grant Delay	0.039***	0.060***	0.065***	0.061***	-0.084***	-0.124***	-0.140***	-0.087***
	(0.010)	(0.010)	(0.010)	(0.010)	(0.016)	(0.015)	(0.018)	(0.021)
Post AIPA × US (d)	-0.107***	-0.084***	-0.050**	-0.002	-0.067*	-0.010	0.034	-0.071
	(0.023)	(0.023)	(0.024)	(0.024)	(0.036)	(0.037)	(0.045)	(0.053)
Observations	669,470	667,812	658,749	642,237	317,660	149,339	81,344	49,398
Adj R-squared	0.461	0.521	0.553	0.578	0.281	0.321	0.330	0.339

Panel B: Opt-Out Exposure

	(1) Log 3-Yr. Forward Cites	(2) Log 5-Yr. Forward Cites	(3) Log 7-Yr. Forward Cites	(4) Log 10-Yr. Forward Cites	(5) Log Months to 1 Cite	(6) Log Months to 3 Cites	(7) Log Months to 5 Cites	(8) Log Months to 7 Cites
Post AIPA × US × Opt-Out Ratio	0.404***	0.392***	0.325***	0.258**	-0.668***	-0.816***	-0.996***	-0.623*
	(0.111)	(0.110)	(0.116)	(0.122)	(0.179)	(0.197)	(0.295)	(0.354)
Post AIPA × US (d)	-0.049***	0.024**	0.073***	0.116***	-0.207***	-0.237***	-0.221***	-0.230***
	(0.010)	(0.011)	(0.012)	(0.013)	(0.016)	(0.017)	(0.028)	(0.032)
Observations	669,809	668,150	659,086	642,572	317,837	149,417	81,396	49,420
Adj R-squared	0.462	0.524	0.556	0.582	0.280	0.319	0.328	0.339

Table 10: Placebo analysis with self-citations

*This table reports DID estimates of AIPA's effect on self-citations. Panel A uses the logged citation count (one plus the number of self-citations and then take natural logarithm) while Panel B uses the percentage of self-citations relative to total forward citations as the dependent variable. The regression specification is the same as the ones used in our main analyses. Standard errors are clustered by application month for U.S. and EP patents, separately. ***, **, and * indicate 1%, 5%, and 10% significance levels, resp. (d) denotes dummy indicator variable.*

Panel A. Self-citation count

	(1) Log 3-Yr. Forward Self-Cites	(2) Log 5-Yr. Forward Self-Cites	(3) Log 7-Yr. Forward Self-Cites	(4) Log 10-Yr. Forward Self-Cites
Post AIPA $\times$ US (d)	-0.091***	-0.038***	0.001	0.034***
	(0.004)	(0.006)	(0.006)	(0.006)
Observations	669,877	668,218	659,154	642,640
Adjusted R-squared	0.242	0.288	0.324	0.353

Panel B. Self-citation percentage

	(1) Self-Cite Share of Total 3-Yr. Forward Cites	(2) Self-Cite Share of Total 5-Yr. Forward Cites	(3) Self-Cite Share of Total 7-Yr. Forward Cites	(4) Self-Cite Share of Total 10-Yr. Forward Cites
Post AIPA $\times$ US (d)	-0.037***	-0.016***	-0.008***	-0.004***
	(0.002)	(0.002)	(0.002)	(0.001)
Observations	669,877	668,218	659,154	642,640
Adjusted R-squared	0.218	0.264	0.293	0.314

Table 11: Patenting intensity

*This table reports Ordinary Least Squares (OLS) regression analysis of patenting intensity around AIPA's enactment. In Columns (1) and (2), the dependent variable is the number of applications and granted applications filed in each month from 1998 to 2003 at each patent office (USPTO or EPO). In Columns (3)-(6), we conduct the analysis at patent-office × technology-class × month level. The dependent variable is the monthly count of applications or granted applications by technology class (primary IPC 4-digit code) filed at each patent office. 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. (d) denotes dummy indicator variable.*

	(1) #Application	(2) #Grant	(3) #Application	(4) #Grant	(5) #Application	(6) #Grant
Post AIPA × US (d)	2,303.815***	1,158.371***	3.694***	1.857***	3.641***	1.834***
	(549.774)	(351.626)	(0.516)	(0.436)	(0.515)	(0.436)
US (d)	11,061.429***	9,883.629***	18.498***	16.504***	18.541***	16.529***
	(449.807)	(285.939)	(0.324)	(0.282)	(0.324)	(0.282)
Post AIPA (d)	1,498.057***	326.659***	2.532***	0.564*		
	(225.647)	(110.988)	(0.337)	(0.300)		
Fixed Effects	No	No	IPC4	IPC4	IPC4, Month	IPC4, Month
Observations	144	144	86,797	86,797	86,797	86,797
R-squared	0.938	0.963	0.740	0.655	0.741	0.656

Table 12: AIPA's effect on R&D investment: U.S. public companies

*This table reports the AIPA's impact on R&D investments by U.S. public companies. The dependent variable, indicated in the table header, is the level of R&D investment (natural logarithm of R&D expenses in Columns (1)-(3)) and R&D intensity (R&D expenses scaled by total sales in Columns (4)-(6)). Both measures are winsorized at the 1% level. 'Post AIPA (d)' is a dummy indicator equal one if the fiscal year ends after the effective date of AIPA (November 29, 2000) and zero otherwise. 'Firm Exposure to AIPA' is the firm-specific exposure to the AIPA shock measured by the median application-grant delay (in years) by USPC technology class weighted by the firm's patent share in each class. The portfolio weight is based on the firm's portfolio of patents filed from January 1, 1998 to November 28, 2000. Exposure is set to be zero for firms without any patent applications. Standard errors are clustered by firm. ***, **, and * stand for statistical significance based on two-sided tests at the 1%, 5%, and 10% level, respectively. (d) denotes dummy indicator variable.*

	(1) Log R&D	(2) Log R&D	(3) Log R&D	(4) R&D to Sales Ratio	(5) R&D to Sales Ratio	(6) R&D to Sales Ratio
Sample:	Sales≥10m	Sales≥20m	Sales≥50m	Sales≥10m	Sales≥20m	Sales≥50m
Post AIPA (d)	-0.0309	-0.0343	-0.0469*	0.00244	0.00971	-0.00116
	(0.0208)	(0.0221)	(0.0252)	(0.0122)	(0.0122)	(0.00249)
Post AIPA × Firm Exposure to AIPA	0.0297***	0.0288***	0.0282***	0.00550**	0.00355**	0.00193**
	(0.00871)	(0.00922)	(0.0107)	(0.00267)	(0.00181)	(0.000934)
Size	0.364***	0.369***	0.369***	0.0311*	0.0267*	0.00339
	(0.0212)	(0.0237)	(0.0308)	(0.0180)	(0.0152)	(0.00531)
Tobin's Q	-0.0133***	-0.0141***	-0.0141***	-0.00401*	-0.00342*	-0.00387***
	(0.00365)	(0.00403)	(0.00516)	(0.00240)	(0.00188)	(0.00125)
Return on Assets	-0.134***	-0.116***	-0.0776	-0.0781***	-0.0549***	-0.0343**
	(0.0343)	(0.0395)	(0.0582)	(0.0171)	(0.0174)	(0.0160)
Loss	0.0811***	0.0775***	0.0753***	0.0300***	0.0273***	0.0218***
	(0.0140)	(0.0149)	(0.0179)	(0.00534)	(0.00546)	(0.00292)
Leverage	-0.193***	-0.198***	-0.181**	0.00430	-0.0236	-0.0179
	(0.0663)	(0.0715)	(0.0860)	(0.0674)	(0.0585)	(0.0159)
Controls	Y	Y	Y	Y	Y	Y
Observations	17,978	16,406	13,693	17,978	16,406	13,693
R-squared	0.967	0.968	0.969	0.737	0.684	0.863