CHAKRAVARTY, SATYAKI, Ph.D. The America Invents Act and innovation by small entities. (2024)
Directed by Dr. Martijn van Hasselt. 109 pp.

This dissertation studies patenting activity by small U.S. entities before and after the passage of the Leahy-Smith America Invents Act (AIA) of 2011 which changed the patenting rule in the United States from a first-to-invent to a first-inventor-to-file system. Prior to the AIA, entities had the benefit of flexibility on when to file for patents but this benefit came at a cost; it created uncertainty in an atmosphere of litigative behavior about the date of the invention. To curb litigation, the AIA intended to reward disclosure of inventions through patent filing, but this added to the already constrained budget of small entities. Contrary to the AIA's intentions, the results indicate a decline in patenting activity post-AIA for all entities, and a widened gap between small and large entities compared to the pre-AIA period. Further, a higher exposure to litigation results in a decline in patents filed, but the rate of decline stalls post-AIA. This dissertation makes two contributions. First, it provides empirical evidence about the impact of AIA on small entities' patenting behavior and empirically tests if the concerns laid out by legislators prior to the AIA's enactment hold. Second, it lays out important differences between small and large entities' portfolio of patents and their incentive to file for patent rights.

THE AMERICA INVENTS ACT AND INNOVATION BY SMALL ENTITIES

by

Satyaki Chakravarty

A Dissertation Submitted to
the Faculty of The Graduate School at
The University of North Carolina at Greensboro
in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Greensboro

2024

Approved by
Dr. Martijn van Hasselt

Committee Chair

To my birthplace and the people of Jaduguda.

APPROVAL PAGE

This dissertation written by Satyaki Chakravarty has been approved by the following committee of the Faculty of The Graduate School at The University of North Carolina at Greensboro.

Committee Chair ______________________

Martijn van Hasselt

Committee Members ______________________

Dora Gicheva

____

Albert N. Link

____

Date of Acceptance by Committee

____

Date of Final Oral Examination

ACKNOWLEDGMENTS

Whatever I write here will not be enough. I would not have the ability to write anything had I not been bestowed upon with the blessings of all the living and non-living lives I have met to date. But, I must extend a few special gratitude that comes to my mind. Also, skipping the “Dr.” as this page has too many of them.

Firstly I am forever indebted to the “director” of this dissertation, my advisor, Martijn van Hasselt, and my dissertation committee members, Dora Gicheva and Albert Link. I consider myself the luckiest Ph.D. student alive who received incessant support from their professors and had the best advisor ever. I would not have garnered enough confidence to finish this work if not for my advisors, and I wish every Ph.D. student received this kind of support. Thank you also to the tiny Ph.D. cohort, comprising Thea Liu and Hitanshu Pandit for being there. A round of applause to our fantastic Department of Economics: Ashe Peters, Jess Saunders, Stephen Holland, Nir Eilam, Eric Howard, Sebastian Laumer, Christopher Swann, Jeremy Bray, Marie Hull, and Martin Andersen. The best Department of Economics of all time.

I worked as a Research Assistant at the India Development Foundation (IDF), and this organization deserves a special place. I thank everyone affiliated to IDF for harboring me and I hope more students from India receive similar support and exposure. This dissertation also has benefited from comments and suggestions from participants from the numerous conferences and workshops I have attended, and I thank everyone who commented on my work.

Finally, my family, comprising my mother, my partner, and my father must be thanked, even though thanking them is a disservice to them. I apologize for skipping so many deserving names, but my hands are tied, and so is this page’s length.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

In this dissertation, I study the America Invents Act (AIA) and its role in changing the way small entities file for patent rights in the United States. AIA is a complex legislation with multiple objectives and it significantly modified the U.S. patent rules. To this date, its impact on patenting behavior among small entities remains understudied. Assessing this impact is important because small entities employ one-third of the U.S. labor force. Moreover, innovative small entities are critical for the U.S. economy as they take steep risks by investing their limited resources into the development of radical inventions. Theoretical predictions do not indicate if the newly adopted rule in 2011 benefited all entities, and on the contrary, the scant evidence on AIA and similar legislations passed in other countries indicate its adverse effects on small entities' and inventors' patenting activity. Here, I study the most substantial and talked about reform the AIA brought, the move of the U.S. patenting system from a first-to-invent to a first-inventor-to-file rule, and determine empirically whether it changed the quantity and quality of patents filed by small entities compared to large entities at the United States Patent and Trademark Office (USPTO).

The Leahy-Smith America Invents Act (AIA) of 2011 substantially changed the patenting rules in the United States by adopting a first-inventor-to-file (FITF) rule. A FITF rule recognizes the *first filer* of an invention as the sole inventor, as opposed to a first-to-invent (FTI) rule, which recognizes the *first inventor*, regardless of the filing date. The overarching goal of the AIA was to encourage disclosure of inventions by rewarding the filing of patents at an early stage of the invention's development, which would reduce uncertainty in its date of filing and curb any litigation that arises due to an uncertain date of invention.[1] The previous rule, FTI, provided an inventor the flexibility of inventing first without the worry of rushing to file for patents. While the FTI rule is beneficial to inventors, because they can focus solely on the invention, it also generates uncertainty on when a patent is actually filed. This uncertainty is partially responsible for attracting patent lawsuits, and it is not conducive to innovation. Along with this, there were several secondary reasons that motivated the development of a patent reform act. But the sweeping and broad reform brought forth by the AIA — recognizing the first filer — also implicitly asks an inventor to rush to the patent office immediately after invention, and entities with constrained budgets, such as small entities, may not have the luxury of doing so.[2] Small entities behave differently, have markedly different patenting strategies compared to large entities, introduce disruptive inventions, drive innovation in new directions, and are litigated at a higher rate (Abrams et al., 2019; Aghion et al., 2021; Arora et al., 2021, 2022). It is unknown how they modified their portfolio of patenting activity (i.e. their quantity and quality) post-AIA. Patents' quality can be heterogeneous and changes

[1] An invention can be kept secret, or can be disclosed to the public, either freely or through patents. Disclosure proliferates follow-on invention. But, the definition of disclosure of an invention can vary with context, and so does its intensity. From only declaring the name of a new invention, to a few mentions on certain websites without its specifics to laying down the exact steps to recreate it, everything can be counted as its disclosure. See De Rassenfosse et al. (2020) and Rantanen (2012) for a discussion on disclosure, patents, and their relation to follow-on invention.

[2] A small entity can be a person, a small business concern, or a nonprofit organization. If the entity is a small business, it is defined as having less than 500 employees. Universities and institutes of higher education are categorized under nonprofit organizations. See: https://mpep.uspto.gov/RDMS/MPEP/e8r9#/e8r9/d0e30961.html

as a strategic response to changes in incentives to patenting (Mezzanotti, 2021). In this dissertation, small entities' portfolio of patents is measured using the number of patents and their quality — using a measure of citation and compares them to large entities before and after the passage of this legislation using a difference-in-differences specification.

This dissertation makes the following contributions. First, the empirical results provide evidence in support of the longstanding concern among the U.S. Congress members wary of the detrimental effects of AIA on small inventors. Studies indicated that the divide between large and small entities' (i.e. small businesses and individual inventors) patenting activity could widen after the AIA's passage and this dissertation confirms such predictions (Abrams and Wagner, 2013; Lo and Sutthiphisal, 2009). Second, the results indicate how the benefits and costs may have changed with regard to disclosing inventions as patents for entities of varying sizes after the enactment of the legislation. This dissertation also reiterates the fact from Lerner et al. (2015) that there are multiple objectives of the legislation that require serious attention and that this dissertation only studies one of the AIA's aspects — the move from FTI to FITF, purported to be the most important aspect of the legislation. As Matal (2011a) mentions:

> [n]o part of the AIA is more significant, nor has generated more legislative discussion and debate, than the Act's changes to §102. The bill's new 102 adopts the first-to-file system of patent priority, enacts a new definition of "prior art," and creates a new grace period.

Studying the AIA is important and this dissertation can be considered among the first steps in informing the designing of future U.S. patent policy.

Since the advent of the patenting system in the United States in 1790, the country followed a different set of patenting rules, the FTI, in comparison with the rest of the major patent-granting countries, that followed the FITF. It is not known why most countries followed an FITF rule. It can be presumed that having a clear record of filing date for a patent is much more straightforward and transparent, and therefore it is easier to report it in case of a dispute. However, no country has ever mentioned this explicitly as to why they did not follow a FTI rule like the United States. Both the set of rules have their benefits and costs. The understanding and assumptions that define such benefits and costs have varied from country to country. As a substantial part of these benefits and costs are immeasurable, there is no other way but to rely on country-specific case laws and legislation. As a result, the theoretical comparisons between FITF and FTI rules do not yield a clear winner as to whether following one is strictly better than the other; rather the comparisons suggest both can be conducive to innovation under a range of parameter values. Some of the parameter conditions are: how strong the patent protection system is; or how the preferences of consumers are for a particular innovation — whether consumers prefer new and innovative products in quick succession or not, or how valuable is it for an inventor to disclose an interim invention over the final invention; and how the interim inventions are protected by the system, and when the protections are strong or weak for those. These conditions impact if FTI is a better choice for innovation than FITF (Miyagiwa, 2015; Scotchmer and Green, 1990). In addition to this, litigation is uncertain and too much of it can seriously undermine the validity of patents and as a result, discourage patenting. This also has a cascading effect on follow-on inventions (Mezzanotti, 2021; Kiebzak et al., 2016; Tucker, 2013; Lanjouw and Schankerman, 2001). On the flip side, a FITF rule limits a part of these litigations but comes at the cost of quickly filing for patents as soon as an invention is completed. Anecdotes from U.S. small entities suggest that they often invent first and then search to secure resources to file the invention as a patent. Resources are not only limited to funds but also include assistance from attorneys and agents.[3,4] Needless to say, these entities have to also worry about their idea being stolen, now more than before. However, we do not know if small entities necessarily engage in a greater secrecy or not. U.S. Congress was therefore divided when the AIA was discussed. [5] The AIA aimed to curb the increase in litigation and also included provisions to limit the increase in costs of filing patents for smaller entities and in particular individual inventors.[6] Further, to study the effects of the AIA on small entities, a study was commissioned (Lerner et al., 2015). Due to a large number of pending patent applications which resulted in a severe lack of data at the time when the study started, the authors could not provide a definite answer as to whether the AIA benefited or harmed small entities' patenting activity.[7] Therefore, we still do not know if these provisions were enough to offset the increase in costs of patenting for small entities and provide a net benefit. Thus, this dissertation studies

[3]https://www.nytimes.com/2012/02/09/business/smallbusiness/business-owners-adjusting-to-patent-system-overhaul.html

[4]In 2015, the cost of filing a patent in the United States ranged from $6000 for an "extremely simple" patent application, to more than $19,000, for a "highly complex" patent application. https://www.ipwatchdog.com/2015/04/04/the-cost-of-obtaining-a-patent-in-the-us/id=56485/

[5]Senator Feinstein's argument on the AIA being detrimental to small entities: "[t]his presents a particular hardship for independent inventors, for startups, and for small businesses, which do not have the resources and volume to employ in-house counsel but must instead rely on more-costly outside counsel to file their patents. This added cost and time directed to filing for ideas that are not productive will drain resources away from the viable ideas that can build a patent portfolio–and a business."

[6]The AIA created a new category of inventors, called "micro entity", who can avail 75 percent discount in all fees. This is discussed in detail in Section 2.1

[7]Two major limitations, as the authors note: "Only a small amount of data exists for patent activity under the FITF rules since (a) the FITF provision became effective on March 16, 2013 and (b) there was still, at that point, a major backlog of patents under the old system because the patenting process — from application to issue/abandonment — took close to 2.5 years for fiscal year 2013. The complexity of the law has led to varying interpretations from small businesses and small business investors, which has likely been reflected in varying responses."

and provides evidence of changes in patenting activity before and after the AIA in the United States with a focus on small entities. Change in patenting activity can be considered the best signal to study the participation of entities of various sizes in the patenting process. Thus, in particular, the first question this dissertation asks is: was there a change in the number of patents filed by small entities relative to large entities post-AIA? Patent counts do not indicate if an invention is truly new, or if it is written clearly enough to be replicated by experts — all of which if true, can indicate a patent is of high quality (Hall et al., 2004). Therefore, I also ask a second question for a comprehensive understanding of the total effect of the AIA on the patenting activity of small entities: did the quality of patents filed by small entities change post-AIA when compared to large entities? Studying the quantity and quality jointly can provide us with a better understanding of the total effect of AIA on small entities' patenting activity.

Early evidence indicates a reduction in disclosure of inventions by larger entities post-AIA; however, for smaller entities, the evidence remains inconclusive partly due to limitations in the data and partly due to unexplored and thus undefined scope of innovative activity in small entities (Lerner et al., 2015; Huang et al., 2020). Small entities have different incentive structures and innovative behavior. This is not represented adequately by publicly listed companies. Along with this, evidence from an AIA-like reform from Canada reported adverse effects on small firms and individual inventors upon its implementation (Abrams and Wagner, 2013; Lo and Sutthiphisal, 2009). This further warrants the identification of small entities and the study of AIA's effect on their patents.

To address these questions, I connect various datasets from the USPTO and add entity-level disambiguated identifiers from the Patentsview database to construct an entity-quarter-level dataset. I use the number of patent applications applied by different entities as an indicator of invention disclosure, which is the broadest indicator of innovative activity at the entity level. The number of patent applications captures the resulting effect of all the policy changes the AIA introduced, and this approach to the data is in line with the literature. To measure the change in the quality of patents, I use the number of citations a patent received in its first two years of issue per patent and rescale it to a percentile value for its field of invention — also called average scaled citation. These two measures used for the main set of analyses are not the only measures of quantity and quality of innovation. In recent years, the value of innovation captured through changes in stock value for public entities as well as measures of quality developed from patent texts have been widely used as a proxy for patent value (Kogan et al., 2017; Kelly et al., 2018). This measure is unfortunately only available for publicly listed companies. All the measures highlight one particular dimension of innovation and none are adequate on their own. Therefore, I will show in this dissertation a range of alternative specifications and alternative measures in an effort to highlight the nuances of patent analysis and corroborate my findings.

I use a difference-in-differences approach to estimate the change in the number of patents and scaled citations for small entities relative to large entities post-AIA. I find that on average, small entities are filing for a lower number of patents post-AIA relative to large entities, but the average quality of the patents filed has not significantly changed post-AIA. This constitutes AIA's overall effect on the quantity and quality of patent applications and is this dissertation's primary set of results.

Next, I define a secondary specification to estimate the same outcomes — the change in the number of patents and scaled citations when entities are exposed to different levels of litigation. We know that certain fields of invention are more prone to litigative behavior than others (Marco and Miller, 2019). This tests whether the AIA's objective to reduce litigation led to an increase in patenting activity for different entities especially small entities or not. According to the AIA's mechanism, I should observe an increase in confidence among inventors to publicly disclose their inventions now that litigation was supposed to drop. What I observe is over the years approaching the AIA's passage, the number of patents and scaled citations continued to fall in areas that have a larger exposure to litigation but with the AIA's passage, this fall was arrested. While this indicates a relative increase in confidence among small entities to file for patents post-AIA, this has to be investigated further as post-AIA, a landmark case — *Alice v. CLS Bank* was decided that changed the definition of eligibility of patents. This will be discussed in detail in Section 6.2 of Chapter 6. As robustness, I study patenting activity for a sub-group — publicly listed companies and I find no change in the outcomes for small publicly listed companies as compared to large ones post-AIA. This reiterates the fact that the publicly listed small entities are in reality too large and behave just as any other large publicly listed company and thus do not adequately represent a large chunk of small entities who file for patent applications.

The remainder of the dissertation is organized as follows: in Chapter 2 I detail the provisions of the AIA and compare it with other countries. In Chapter 3, I describe a simple model to understand the trade-off between patenting and keeping trade secrets and the conditions that govern it. In Chapter 4, I list the data sources, the sample selection, and the definition of the variables. Next, in Chapter 5, I explain the empirical strategy, and in Chapter 6, I discuss the inferences derived from the estimates. Finally, I conclude by summarizing my findings in Chapter 7.

Chapter 2: The America Invents Act (AIA)

2.1 Provisions of the AIA

The U.S. patent regime underwent significant changes since World War II and the most recent amendment to the 35 U.S. Code—the America Invents Act of 2011 is regarded as the most substantial one since the Patent Act of 1952, changing the regime from First-to-Invent (FTI) to First-Inventor-to-file (FITF) (Lerner, 2000).[1] Under an FTI rule, an inventor had the option to claim that they are the original inventor even if another inventor had filed a patent application for a similar invention before theirs. Under a FITF rule, the inventor who files for patent rights first is the rightful owner of the patent. Pre-AIA, the invention date could be used to claim rights that do not matter post-AIA, and instead the date of filing matters (Masur and Ouellette, 2020). While the former rule provided flexibility to an inventor as to when to file for patent rights, entities could abuse this flexibility for strategic advantages. A group of entities, commonly known as "patent trolls" or Non-Performing Entities (NPEs) is those who

[1] https://www.nytimes.com/2011/09/09/business/senate-approves-overhaul-of-patent-system.html

do not invent; but rather acquire patents and assert their rights to invalidate other patents. Such entities wait for other inventors to file for a similar patent application as theirs or start producing a product that uses an invention similar to theirs. When the producing entities apply for patent rights or start production, the NPEs file for injunctive relief and ask for hefty royalties for infringement. This action is known as a hold-up. Due to the fear of hold-ups or infringing upon other patents, inventors restrict the disclosure of their inventions. As a result, the innovation of an economy does not reach its optimum level (Tucker, 2013).

The AIA's FITF rule tries to address the problem of unnecessary lawsuits partially. While not all kinds of patent lawsuits will be taken care of by the implementation of the AIA, at least for the cases related to the invention date, the AIA tries to establish a more certain regime by not providing the opportunity to contest conflicts on the date of invention. Pre-AIA, conflicts on the date of the invention were litigated under interference proceedings. An interference proceeding would be conducted when one patent application interferes with another in the process of its filing.[2] Post-AIA, this has been changed to derivation proceedings. Now, an entity cannot claim rights if they were late to the patent office. But, if they can show that an invention already in the patent office was derived from their invention, they can file for a derivation

[2] To understand interference proceedings and its relation to the AIA, I present an example. In a pre-AIA period, an invention i_1 came into existence at t_1 but was not filed at the patent office, and an independent invention i_2, also in the pre-AIA period, was invented at t_2 ($t_2 > t_1$) and was filed as a patent application immediately. Here, invention i_2 interferes with invention i_1. Invention i_1 can still claim its patent rights or at the least negate i_2's patentability by conducting an interference proceeding, which will be conducted as a lawsuit. In the post-AIA period, in a similar situation, i_2's patent application will be upheld as the only patent application and i_1 can contest if it files for a derivation proceeding. Derivation proceeding replaced interference proceeding post-AIA and in the example, i_1 will claim that i_2 has been derived from i_1. But in the post-AIA period, invention i_2 will have stronger protection and i_1 will face higher costs because i_1 will have to show that i_2 has been derived from i_1 and not independently invented. A low count of interference proceedings does not mean that they are seldom used. Rather, it indicates that the patent applications prone to such opposition are deterred from filing.

proceeding.

To cushion small entities from bearing the extra cost of quick filing, the AIA implemented certain provisions. The AIA sets up two programs to assist small entities in filing their patent applications, called the Pro Bono Program and the Patent Ombudsman for Small Businesses. The Pro Bono Program provides the qualifiers with free legal assistance in preparing and filing patent applications. [3] The Patent Ombudsman program assists applicants when a normal application stalls. [4] The AIA defines a new category of entities, called "micro-entity" for whom the USPTO levies a discount of 75 percent on all patent filing fees. [5]

The AIA not only changed the patent rules but introduced post-grant review (PGR) and inter partes review (IPR), replacing the ex parte reexamination. Any third party can challenge the validity of a patent, using PGR in the first nine months, and then using IPR. As opposed to an ex parte reexamination, where a third party could apply for reexamination of a granted patent but not have active participation in the proceedings, PGR and IPR allow the third party to actively take part in the challenge process. The introduction of new and cheaper ways to oppose a patent's validity can caution entities who operate in areas predominantly exposed to litigation. To put numbers, court litigation can run to millions of dollars while PGR/IPR can be estimated to be at around $500,00 (Masur and Ouellette, 2020). New and cheaper opportunities to review a patent reduce the cost to both plaintiff and defendant, as filing a review and defense against it equally become cheaper. But, it is not obvious if

[3] https://www.uspto.gov/patents/basics/using-legal-services/pro-bono/patent-pro-bono-program

[4] https://www.uspto.gov/patents/ombudsman-program

[5] A micro-entity is a small entity with additional thresholds. A micro-entity has to qualify for an income threshold, which is three times the median income household. Note that institutes of higher education are considered micro-entities. The detailed definition can be found here: https://www.uspto.gov/patents/laws/micro-entity-status

it encourages or discourages patent filing.

Lastly, the AIA also sets up Post Grant Review of patents. Any third party can file a petition to challenge the validity of one or all of the claims of a granted patent.[6]

2.2 Legislative history of the AIA

While Matal (2011a), Hasford (2017), Doody (2012) explain and compare the number of changes the AIA brought to the patenting rules after its passage and their legalities, I specifically highlight and discuss the differences among Congresspersons' debates which in turn informs the benefits and costs of adopting a first-inventor-to-file rule that the AIA introduced. I separate the arguments furnished during the debates in favor and in opposition to the AIA's first-inventor-to-file rule from the Senate and House versions of the bill into two broad topics — filing first and interference proceedings; prior art and grace period.

Elements of the AIA can be traced back to as far as the 1992 Joint Hearing on the Patent System Harmonization Act as introduced by bills S. 2605 and H.R. 4978. But, the most identical version of the AIA, then called the Patent Reform Act was first introduced in the House on June 8, 2005. Apart from certain issues, which were removed in the later iterations of the bills' introduction in 2007, this 2005 version of the bill — the Patent Reform Act contained all the elements that were finally enacted in 2011 as the America Invents Act.

Filing first and interference proceedings. The House and Senate separately debated on the enactment of a first-inventor-to-file (FITF) rule which would eventually replace the first-to-invent (FTI) rule. No objections were observed during the period

[6]https://www.uspto.gov/patents/ptab/trials/post-grant-review

when the Patent Reform Act was first introduced, i.e. between 2005-2007. The bulk of the opposition happened in March 2011 — six months before the AIA's passage. The arguments opposing FITF broadly spanned around its possible unintended consequences, and those are discussed in detail in the following paragraphs.

In the Senate, the opposition to FITF was led by Senator Feinstein's amendment to strike the FITF provision of the AIA down. Her initial arguments were twofold — first, FITF rewards the inventor who files first, and not the first inventor, or in her words "incentivizes a race to the patent office" which puts small entities at a disadvantage due to lack of resources, and second, this "race" would encourage "significant over filing of 'dead end' inventions", as inventors would rush to protect all of their inventions from their competitors, irrespective of the quality of their invention. Her arguments were accompanied by a supporting letter from Paul Michel, former chief judge of the Court of Appeals for the Federal Circuit, and Gregory Junemann, president of the International Federation of Professional and Technical Engineers, who reiterated Senator Feinstein's concerns and added that FITF would lead to over filing of low-quality patents which would in turn overburden the office. Though, it must be noted that this dissertation does not point towards any evidence that suggest there was overfiling of patents after the AIA. Yet, further evidence would be welcome in this regard. Coming back to Senator Feinstein's arguments, her point is understood from the following excerpts from her arguments and the letter furnished in the Senate on March 2, 2011 (S1089-S1114, 2011):

> Unfortunately, first-to-file incentivizes inventors to "race to the Patent Office", to protect as many of their ideas as soon as possible so they are not beaten to the punch by a rival. Thus, first-to-file will likely

> result in significant overfiling of these "dead end" inventions, unnecessarily burdening both the Patent and Trademark Office and inventors.

The letter adds:

> At a time when the Patent and Trademark Office has a dramatic backlog of over 700,000 patents waiting to be examined and a pendency time of some 3 years, Congress should be careful to ensure that any legislative changes will not increase patent filings that are unfruitful.

These arguments were addressed by Senator Kyl, and in anticipation of the amendment sponsored by Senator Feinstein; the sponsors and supporters of the bill — Senator Leahy, Hatch, Klobuchar, and Kyl arranged a number of letters from a spectrum of supporters of the AIA.

Senator Hatch and Senator Kyl responded to the two initial arguments made by Senator Feinstein by using descriptive statistics from interference proceedings. The interesting part regarding this is that the same statistic on interference proceedings was used for and against arguing for the FITF. Pre-AIA, interference proceedings were initiated when one patentee would allegedly accuse another of interfering in their patent filing process. Typically, multiple inventors would file the same patent application at different times — one before the other, and only one can be rewarded with patent rights, which starts a dispute between them, which when taken up formally is called an interference proceeding.[7] Historically, interferences were quite rare. For example, there were only 650 interference cases between the years 2005 to 2013 (Laplume et al., 2015). This method of dispute resolution would be discontinued as the AIA is enacted since the AIA rewards the first filer with patent rights. This would make disputes

[7] See Footnote 2 for a hypothetical example on interferences.

redundant in cases where similar or the same inventions are filed one after the other. Since interferences were quite rare, arguments in favor of discontinuing interferences repeated that interferences were this low in number because it required a hefty amount of money in terms of legal fees to show that an inventor was the first to invent, and therefore is uncertain. But the solution — filing provisional applications to establish a filing date is a lot cheaper — $100 for an application. For example, Senator Hatch mentions the following:

> Under this [first-to-invent] system, if there is a dispute, it costs applicants an average of $500,000 in legal fees to prove they were the first-to-invent. ... Unfortunately, many small businesses and independent inventors do not have the resources to engage in the process we have now.
>
> Conversely, moving to a first-inventor-to-file system would provide inventors a cost-effective and certain path to protect one's invention through the filing of a provisional application, at a much more reasonable cost of about $100.

However, arguments in favor of a first-to-invent system and interference proceedings use the same statistic, a low incidence of interference proceedings to argue that it is an insignificant issue, and therefore should not be changed. For instance, the letter accompanying Senator Feinstein's arguments says:

> [T]he reality is that this is not a significant problem under our current system. There are only about 50 "interference proceedings" a year...in other words, one-one hundredth of 1 percent of patent applications.

Additionally, Senate Majority Leader Harry Reid termed provisional applications — the cheap solution to disclose inventions as "snake oil", because drafting provisional applications would still require the expertise of patent attorneys, which would cost a lot more than just the $100 provisional application fee. He, upon consulting with small inventors in his constituency mentions the following:

> They tell me that the balm of "cheap provisionals" is snake oil, because a provisional still has to meet certain legal standards, meaning that you still have to spend a lot for patent counsel, which is the biggest single expense of filing an application. Because they can't afford to file that many applications, regular or provisional, they will have to give up on some inventions altogether.

Senator Kyl in response to Senator Feinstein's amendment to strike the FITF quotes prominent patent attorney Gene Quinn from his Feb 27, 2011 article posted on IPWatchdog.com (Quinn, 2011) and reiterates the aforementioned two points on interference proceedings. First, a move from FTI to FITF codifies the already accepted fact that filing an invention as a patent application was necessary to claim patent rights even without a FITF rule, and this codification makes the rule on filing dates more certain — thereby increasing clarity. Second, for a small inventor to gather proof and resources to argue in a legal setting that they were the first to invent was in any case a humongous task because of the efforts and costs associated with it. Through this, Senator Kyl ties the arguments on FITF and its significance. Specifically, Mr. Gene Quinn reiterates the low number of interferences to argue that the United States already followed a FITF system and it can be understood from the following quote:

> With respect to first to file, in practical effect we already have a first inventor to file system. For example, since the start of fiscal year 2005 on October 1, 2004, there have been over 2.9 million patent applications filed

> and only 502 Interferences decided. An Interference Proceeding occurs when multiple inventors file an application claiming the same invention, and is the hallmark of a first to invent system because it is possible in the United States to file a patent application second and then be awarded the patent if the second to file can demonstrate they were the first to invent. On top of the paltry 502 Interferences over nearly 7 years a grand total of 1 independent inventor managed to demonstrate they were the first to invent, and a grand total of 35 small entities were even involved in an Interference.

He also mentions the average amount required to defend a typical interference proceeding to bolster the fact that it was already an uphill battle for a small inventor.

> According to 2005 data from the AIPLA [American Intellectual Property Law Association] the average cost through an interference is over $600,000. So lets not kid ourselves, the first to invent system cannot be used by independent inventors in any real, logical or intellectually honest way, as supported by the reality of the numbers above. So first to invent is largely a "feel good" approach to patents where the underdog at least has a chance, if they happen to have $600,000 in disposable income to invest on the crap-shoot that is an Interference proceeding.

Finally, Mr. Gene Quinn comments on the concern that the use of provisional applications as a tool to disclose an invention would require extensive legal expertise and maneuvers through attorneys that would also be expensive.

> It has been brought up that an appropriate and good provisional patent application needs to be identical to a nonprovisional patent application,

> perhaps without having been spell-checked. Obviously this is a gross overstatement of the law, and not correct. It is true that a provisional patent application needs to be as complete as a nonprovisional patent application in terms of discosure, but nothing more. There are no formalities that need to be met, and it is the substance that matters. Nonprovisional patent applications exalt form over substance in large part, but a good provisional patent application needs to focus on substance.

The House debated on FITF in June 2011, and an amendment to strike the FITF down — similar to the one sponsored by Senator Feinstein was sponsored in the House by Representative Sensenbrenner. All the arguments were exactly the same as put forward in the Senate, i.e. a FITF rule would incentivize quick filing of inventions and that the provisional applications would be expensive for small inventors. Representative Kaptur, in support of the Sensenbrenner amendment, says the following:

> Since the first Congress, which included 55 delegates to the Constitutional Convention, our nation has recognized that you are the owner of your own ideas and innovations. This bill throws that out [of] the window and replaces it with a system that legalizes a rather clever form of intellectual property theft.
>
> I assure you of one thing: if this bill mistakenly passes, this debate will not be over. We will see it head straight to the courts with extended litigation for years to come, along with complete uncertainty to our markets, killing jobs, and killing innovation. I urge my colleagues to vote "no" on H.R. 1249.

This amendment was later voted down in the House.

Prior art and grace period. Prior art is defined as any evidence that an invention was already known. The boundaries of the prior art, however, are given by each nation's own definition of what could be considered as prior art. In the United States, pre-AIA, a grace period was defined as a period of one year where the following actions would not preclude the inventor from obtaining patent rights on the same invention, or in other words would not amount to being a part of the prior art, as mentioned by Senator Feinstein:

> No. 1, describing their invention in a printed publication; No. 2, making public use of the invention; or, No. 3, offering the invention for sale [within one year of the action]. This is called the grace period, and it is critical to small inventors.[8]

Post-AIA, this grace period's length was similar to pre-AIA, for one year, but from the date of disclosure. Senator Feinstein was concerned that the word "disclosure" was ambiguous and that it would need additional court time to elucidate its definition. She says:

> First, "disclosure" is not defined in the bill. This will generate litigation while the courts flesh out that term's meaning. While this plays out in the courts, there will be uncertainty about whether many inventions are patentable.
>
> This uncertainty will, in turn, chill investment, as venture capitalists will be reluctant to invest until they are confident that the inventor will be able to patent and own their invention.

[8] Also see https://www.uspto.gov/web/offices/pac/mpep/s2133.html

> Secondly, because of this lack of definition, some patent lawyers interpret "disclosure" to mean a disclosure that is sufficiently detailed to enable a person of ordinary skill in the particular art to make the invented item. In practical terms, this means a patent application or a printed publication.

Mr. Gene Quinn does agree that the word "disclosure" is not defined, but also mentions that a defined grace period does exist, in his Feb 27, 2011 article by saying the following:

> The grace period would be quite different than what we have now and would not extend to all third party activities, but many of the horror stories say that if someone learns of your invention from you and beats you to the Patent Office, they will get the patent. That is simply flat wrong.

The grace period stayed as it was defined in the bill and was not amended after Senator Feinstein's concerns.

2.3 Comparison with other countries and unintended effects

Most countries have always followed a FITF rule. To date, three countries have made a move from FTI to FITF. Since 2007, different patent reform bills have been highlighting the need to harmonize the patent system in the US with the rest of the world, and by the advent of AIA, the US is the third and last country to make such a change (Matal, 2011a; Matal, 2011b). The need to harmonize could be understood from the Senate Judiciary Committee's Report:

> (1) a patent's filing date is objective and simple to determine, whereas an invention date "is often uncertain, and, when disputed, typically requires corroborating evidence as part of an adjudication"; (2) the first-to-file system would avoid the expense and burden of interference proceedings and eliminate the need for inventors to maintain recording and document-retention systems; and (3) because many U.S. inventors and companies file for patent protection in foreign countries (which all use the first-to-file system), they "are forced [by the United States' maintenance of the first-to-invent system] to follow and comply with two different filing systems."

In 1989 and 1998 Canada and the Philippines made a similar move respectively.[9,10] Only a few studies exist on the evaluation of patenting activity by entity types around the AIA and the Canadian reform and no study evaluates the Philippines reform. A common theme from all the studies points toward no significant benefit to smaller entities after such a change.

In Canada, patenting activity skewed towards large firms compared to small and independent inventors after their reform (Lo and Sutthiphisal, 2009). Additionally, a drop in the number of patents is reported among independent inventors with no change in their patent quality (Abrams and Wagner, 2013). In the US, studies have found a decline in innovative activity, especially among publicly traded firms, after the AIA. Huang et al. (2020) use narrative R&D disclosure of publicly traded firms to measure innovative activity and find a decline, especially for the innovative firms among all publicly traded post-AIA. A narrative R&D is different from the R&D expenditure reported by firms in their annual reports. The authors use the number

[9] Canada's reform: https://www.ic.gc.ca/eic/site/cipointernet-internetopic.nsf/eng/wr04732.html
[10] Phillipines' reform: https://wipolex.wipo.int/en/text/488675

of R&D-related sentences counted from the firms' 10-K filings. Measures developed from the 10-K filings are considered to be more revealing in terms of firms' innovation strategies and activities as compared to their annual reports (Merkley, 2014). However, a decline in R&D-related sentences does not necessarily mean a decrease in the number of patents. Also, Huang et al. (2020) study a specific group of innovative entities, publicly traded firms. 10-K is not applicable for individuals and small entities but these entities participate significantly in innovative activity.

Though compared to pre-AIA, Huang et al. (2020) find a decline in R&D-related sentences for publicly traded firms post-AIA, their market returns did not significantly change (Lerner et al., 2015). Among the venture capital-backed firms (VC), Lerner et al. (2015) do not find any significant difference in their formative stage funding post-AIA. It is unknown how the other small entities, such as individual inventors, small firms not backed by VCs, and universities filed their patents before and after the AIA. In fact, the quantity and quality of patents for all entities before and after the AIA remain understudied. Neither 10-K filings nor information on VC fund disbursement completely captures the patenting activity of entities, but they do provide a direction to this dissertation. The mixed evidence of a decline or insignificant change in the innovative activity while the AIA encourages disclosure seems counter-intuitive. However, economics suggests that given different conditions we can observe either an increase or a decrease in the disclosure of inventions. AIA rewards early filers, which may result in a direct increase in patent applications. But, if an entity discloses one of its inventions, its competitors can preempt any pipeline of inventions that they may have had, which may be connected to a focal invention. Because FITF rewards the first filer as opposed to the previous rule of FTI where an inventor could furnish proof of invention date if another inventor filed for a similar invention, an entity may become secretive and only file patent applications when they are sure that their patent application will not be invalidated or their future work cannot be preempted.

In the case of the AIA, Huang et al. (2020) discusses a similar mechanism of holding back on patenting because of competitors' preemptive activity and therefore we observe a decline in narrative R&D disclosure post-AIA. Abrams and Wagner (2013) also discuss similar explanations for the decline in patents by individuals post-Canadian reform. Some of the possible explanations they discuss are fewer resources for individual inventors which makes it difficult to file quickly for multiple patents, a shift in the use of different intellectual property protection mechanisms such as trade secrets, or a shift to patenting in the US. The authors find significant evidence of an increase in Canadian inventors patenting in the US after the Canadian reform. This is one of the reasons why inferences from the Canadian case cannot readily predict the effects of the AIA. In Canada, inventors still had an option of enjoying First-to-Invent if they filed their application in the US. Therefore, it is possible that the Canadian entities shift their activity toward the US, especially because of their geographical proximity, as well as the US, being a hub for innovative activities. However, in the case of AIA, entities do not have the option to move their activity to another country where they could still enjoy the FITF rule. Therefore, concluding that reforms similar to the AIA always reduce innovative activity, especially by small entities may be misleading.

The other strategies as described in Abrams and Wagner (2013) may play out in the US, especially for small entities. Rather than shifting the patenting activity to another country, entities in the US could choose higher secrecy and file patent applications only for those inventions that they expect to prevail in the market, and it is likely that AIA had a heterogeneous impact on groups of entities. While some have become more secret than before, others have started disclosing their invention early, as was the objective of the AIA. Theoretically, as briefly mentioned earlier, its effect depends on various parameters, and empirically, it is unknown if its intended effects, i.e. incentivizing disclosure of inventions and thereby achieving certainty around the invention date were indeed achieved or not (Lerner, 2000; Vandenburg, 2013; Cerro, 2014).

2.4 Summary

The changes brought forth by AIA are multiple and all of them affect patenting activity, or the two outcomes this dissertation studies, quantity and quality of inventions in different ways. Therefore, through the number of patents — a proxy for the quantity of inventions, and citations — a proxy for the quality of inventions, I observe the aggregate or total effect of all the provisions of AIA. In this chapter, I have discussed two chief sources of insights that inform the approach to studying patenting activity. The sources are — the literature on AIA or AIA-like patent reforms in other countries, and the U.S. House and Senate hearings on AIA.

The most sweeping change that AIA brought — First-Inventor-to-File (FITF) rule targets a specific group of lawsuits. But, in general, the other policy changes aim at providing greater certainty in securing and maintaining patent rights in the post-AIA period compared to the previous regime. Therefore, the AIA affects all patent applications and not only those that were prone to litigation on the invention date. However, it is a serious challenge to decompose the overall effect of the AIA into the effects of the individual components of the policy. Through AIA's FITF rule, more certainty is brought to the U.S. patenting system, and the other components complement the FITF rule in bringing this certainty. Greater certainty should result in timely disclosure of inventions, which also translates to an increase in the number of disclosures. The rationale behind this change in policy is that the increase in benefit from increased disclosure surpasses the benefit from a flexible date of invention. The broadest and most direct measure of innovation that stems from this reasoning is the count of patent applications. Since the AIA, in all, encourages disclosure of inventions, I expect an increase in the number of patent applications post-AIA, and that should indicate an increase in disclosure of inventions post-AIA. Yet, from the two sources as described in this chapter, from time and again it has been noted there exists a concern that small entities could be adversely affected by this reform. This was observed in Canada when the country enacted a similar reform, and the same concern was raised at the U.S. House and Senate hearings quite a few times. The reason is that small entities have lower access to resources, and consequently may not be willing to disclose their pipeline of inventions to their competitors. Therefore, the AIA may or may not increase patenting activity, and it is unclear in which direction it goes. Thus, the first question this dissertation explores for small entities is — did the number of patents filed by small entities change after the AIA's enactment compared to large entities? A higher number of patent applications may be accompanied by a change in their quality. The AIA specifically mentioned that enhancing patent quality was one of its goals. It is also unknown if the quality of patent applications changed as a result of the AIA. Therefore, this is the second question this dissertation asks, did the quality of patents filed by small entities change after the AIA compared to large entities?

Chapter 3: Theoretical predictions

In this chapter, I present a simple model to show the relationship between patenting and secrecy. The AIA's main objective was to incentivize inventors to file their invention as a patent application in the patent office rather than holding on to the invention. This incentive was implemented by adopting a First-Inventor-to-File (FITF) rule over a First-to-Invent (FTI) rule that was in place prior to the AIA's passage, as has been discussed in Chapter 2. While the AIA was designed to bring certainty to the patent system and had nothing to do with the secrecy of inventions, yet looking at the incentives in a fundamental form — between disclosing an invention in the form of a patent versus keeping it secret can aid us in understanding the various outcomes that the AIA could generate. This is because an uncertain patent system — the system in which patent litigation is on the rise generates inefficiencies similar to an economy that does not have patent protection, and this was the primary reason for which the AIA was enacted.

3.1 Description of an inventor's trade-off

Producers always face a trade-off when they introduce heterogeneity in their goods in a competitive market for homogeneous goods. If we think about the heterogeneity in terms of innovation, the trade-off is either to patent their innovation or to keep it a secret, assuming that the innovation is patentable. A secret may seem to be a lucrative solution to protect an innovation. However, the flip side is that if there is even the slightest chance that the innovation will be discovered or reverse-engineered by competitors in the early stages after innovation, all effort will go in vain. Thus, one solution might be that if there is any chance that the innovation will stay a secret for at least as long as a government allows for intellectual property rights, it is better to keep a secret. Therefore, the problem can be thought of in the following way: does an inventor after innovating choose patent rights or keep the invention a secret?

The dissertation explores various choices a producer faces and models the behavior given alternatives. The following section describes the model. Then follows a section on some numerical examples where the analytical findings are put to work in a simple world.

3.2 Formal model

Each producer faces an inverse demand function given by $P = a - bq$ in a market with sufficient producers producing a homogeneous good characterized by Bertrand competition for infinite periods. One chooses between patenting their innovation or keeping it secret. If a producer chooses to patent, they get monopoly profit for n periods with certainty and after that, the patent rights expire. When the rights expire, other competitors immediately imitate the innovation and everybody starts earning zero profit. Alternatively, a producer can choose to keep their innovation a secret. There is a probability p associated with keeping it a secret. The closer p is to one, the higher the chance that nobody is going to discover the secret. Thus, p can be called a probability of secrecy. Each period's profit, whether it be profit from patenting or trade secret is discounted by r.

The game starts when producers choose to innovate. Profit from innovation at each period is given by $\pi^M = \frac{(a-c)^2}{4b}$; where c is the marginal cost of producing each unit of the good. Profit from no innovation is given by $\pi^{PC} = 0$. If producers do not innovate, the game ends. If producers choose to innovate, they choose between patenting and keeping it a secret. A secret is also referred to as a trade secret.

Lemma 3.1. *Profit from patenting is given by:*

$$\pi^P = \pi^M \left(\frac{1+r}{r} - \frac{1}{r(1+r)^{n-1}} \right) \tag{3.1}$$

Lemma 3.1 gives the profit from patenting for n periods. n can be thought of as patent length. It shows that π^M when earned every period is $\pi^M \left(\frac{1+r}{r}\right)$. This is limited by $\left(-\frac{1}{r(1+r)^{n-1}}\right)$. As patent length increases, the term reduces. At the limit

$$\lim_{n\to\infty} \left(-\frac{1}{r(1+r)^{n-1}} \right) = 0$$

Therefore,

$$\lim_{n\to\infty} \pi^P = \pi^M \left(\frac{1+r}{r} \right)$$

This is the maximum a producer can get by innovating. If patent length is finite, the profit accrued is less than this expression. The derivation of Lemma 3.1 is worked out formally in the Appendix.

Lemma 3.2. *Profit from keeping the innovation a secret is given by:*

$$\pi^{TS} = p\,\pi^{M}\left(\frac{1+r}{1+r-p}\right) \tag{3.2}$$

Lemma 3.2 gives the profit from trade secret. Profit π^M faces a probability of secrecy p and discounting r each period. Thus, if the secret is discovered in period 1, the profit accrued is $(1-p)*0$. If the secret is discovered in period 2, the profit accrued is $p(1-p)\pi^M$. If the secret is discovered in period 3, the profit accrued is $p^2(1-p)\left[\pi^M + \frac{1}{1+r}\pi^M\right]$ and so on. This goes on for infinite periods. All the possibilities are summed up to obtain the equation in Lemma 3.2. The derivation is worked out formally in the Appendix. Based on the two equations defined in Lemma 3.1 and 3.2, some propositions are defined.

Proposition 3.3. *As $r \to 1$, profit from innovation decreases. This is given by:*

$$\frac{\partial \pi^P}{\partial r} < 0, \frac{\partial \pi^{TS}}{\partial r} < 0 \tag{3.3}$$

We observe that as the rate of interest r increases, profit decreases, *ceteris paribus.* This can be looked at in two ways — one, the direct role of interest rate as a tool that governs investment and saving, and two, the indirect role of interest rate as a discount factor. A higher rate of interest acts as an incentive to keep money in the bank. If it is on the lower side, keeping money in the bank will not accrue as much income as an investment will. Thus, borrowing and investing is a better option in the presence of a lower rate of interest. As the rate increases, it becomes costlier to invest, and hence we observe that profit falls from an increasing rate of interest. A second explanation of interest rate as a discount rate is the following: if the discount rate is high, the future is discounted heavily, as opposed to a lower discount rate that does not discount the future as much. A low discount rate means the present is close to the future in terms of profit accrued from inventing, while a high discount rate means the opposite — profit accrued from invention rapidly diminishes as time goes by.

A higher probability of secrecy, on the other hand, entails higher profit from trade secrets, *ceteris paribus.* This might be rather intuitive and needs less explanation. If the probability that a secret might be discovered by others is negligible, it might be better to opt for a trade secret, as the profit from keeping the invention secret surpasses the profit from patenting. However, as the discount rate increases, the future becomes increasingly not appealing, regardless of opting for a patent or a trade secret. If $p = 1$, a trade secret is monopoly profit summed up to infinity. This is obviously better than having monopoly profit only for finite periods — n periods. This is formally defined in Proposition 3.4. The formal proof for Proposition 3.3 and 3.4 are given in the Appendix.

Proposition 3.4. *As $p \to 1$, profit from trade secrets increases. This is given by:*

$$\frac{\partial \pi^{TS}}{\partial p} > 0 \tag{3.4}$$

An inequality has also been derived given in Proposition 3.5 which depicts the relationship between p,r, and n. It says that a trade secret is better than patenting if and only if p is greater than an expression $1 - \frac{r}{(1+r)^{n+1}-1}$. This inequality can be also read as equality, which will then depict the point of indifference between a patent and a trade secret. The inequality says that an increase in patent length increases the bar of probability of secrecy. Thus, obtaining higher profits from trade secrets becomes difficult. The proof is given in the Appendix.

Proposition 3.5. *A trade secret is better than patenting if and only if*

$$p > 1 - \frac{r}{(1+r)^{n+1} - 1} \tag{3.5}$$

We next see some numerical examples of the propositions to mimic the real world.

3.3 Numerical examples

Some numerical simulations were computed to illustrate the propositions described. Figure 3.1 and Figure 3.2 are examples of profit levels on the rate of interest keeping p fixed at various levels and on the probability of secrecy keeping r fixed at various levels respectively. For Figure 3.1 we set the probability of secrecy to be 0.2, 0.5, 0.6, 0.7, 0.8, and 0.9 and the rate of interest to be between 0 and 1 to evaluate all the π^{TS} and π^{P}. The figure thus plots π^{TS} and π^{P} with respect to r. This serves as an example for Proposition 3.1 and shows that as the rate of interest increases, profit from patenting and trade secrets falls. Note that π^{M} and n are set at 1 and 5 respectively. π^{M} only serves as a scale factor for both the profit expressions and hence does not make any difference to the depicted curves. An increase in n however will shift π^{P} upwards. In general, the patent length is 20 years. Setting n to be 20 thus shifts π^{P} upwards, making it harder for the curves to intersect. However, we can also think of each period to last for five years. In summary, n can be thought of as any finite value. Thus we chose n to be five for the curves to intersect. Relation between

π^{TS} ,π^{P} with n are discussed later.

Next, are the curves when r takes values 0.1 and 0.9 and p is set between 0 and 1. π^{TS} and π^{P} are evaluated for those values and Figure 3.2 depicts those curves. Two extremes for r are chosen only to depict two extreme values. As described in Figure 3.1, π^{M} and n are set to be 1 and 5 respectively for the same reason as before. We can see that a trade secret only becomes a better alternative for very high values of p.

If we start thinking of n as the time up to which a patent remains effective rather than the official length, it gives us insight into the relationship between n and p. Effective length can be thought of as the actual use of a patent in a market. A patent can be protected for 20 years, however, if a better substitute comes to the market, a patent can be rendered effectively useless because nobody is going to use an obsolete invention. As an example, we can think of two types of patents, patents specific to mobile phones and patents for a drug formula. A patent for a mobile phone while is protected for the same length as a drug formula, due to rapid invention in the market, substitutes are generated in less than 20 years. Thus, a mobile phone patent may in effect be used for only three to four years before becoming obsolete. On the other hand, a drug formula is difficult to replace. Once developed, it stays in the market for a long time, maybe even more than the official length of a patent (for example: Paracetamol). Given the example, the relationship becomes clearer, the lower the n, the lower the p required to achieve higher π^{TS} than π^{P}. This is shown in Figure 3.3. For two values of $n = 1, 5$, π^{P} is calculated. This is shown by the two straight lines. We can see that for a lower value of n, the minimum p is also low. Thus, a mobile phone invention, because of its short-lived nature can be released into the market without patent rights more easily than a drug formula. A drug formula can only be released if the inventor is sure about the secrecy.

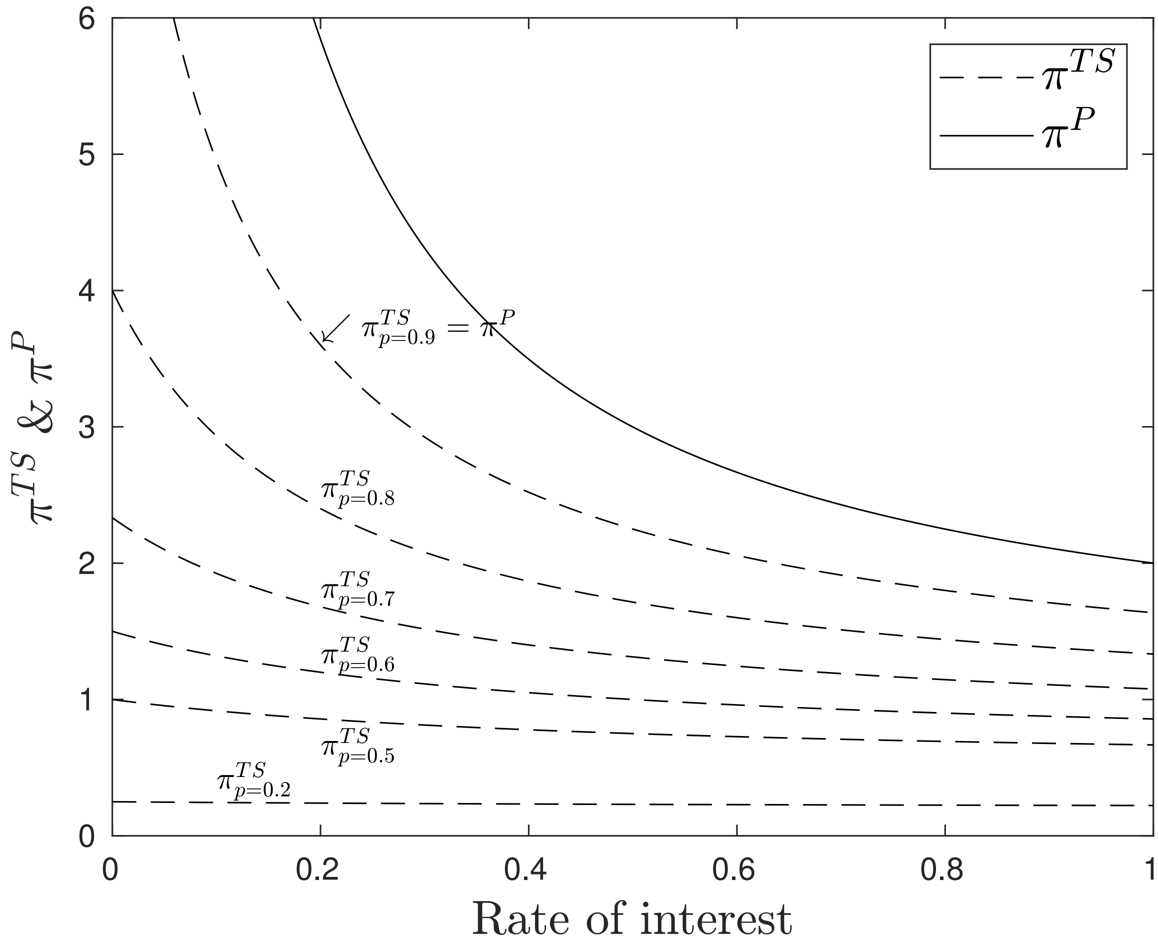

Figure 3.1. Plot of π^{TS}, π^{P} and r when $p = 0.2, 0.5, 0.6, 0.7, 0.8, 0.9$, $r \in [0, 1]$ and $n \neq \infty$

Note: π^{TS} and π^{P} are evaluated using $\pi^{M} = 1$, $n = 5$.

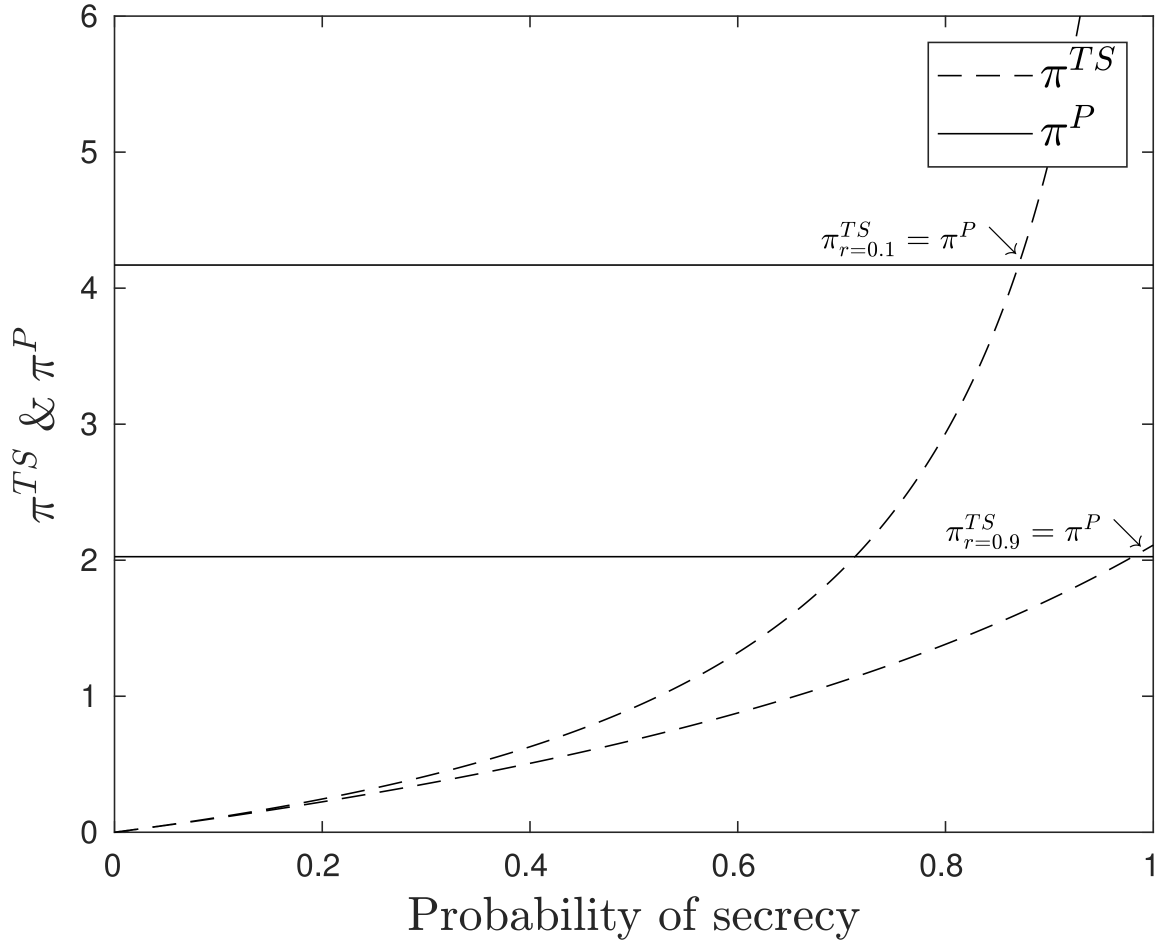

Figure 3.2. Plot of π^{TS}, π^{P} and p when $p \in [0, 1]$, r= $0.1, 0.9$ and $n \neq \infty$
Note: π^{TS} and π^{P} are evaluated using $\pi^{M} = 1$, $n = 5$.

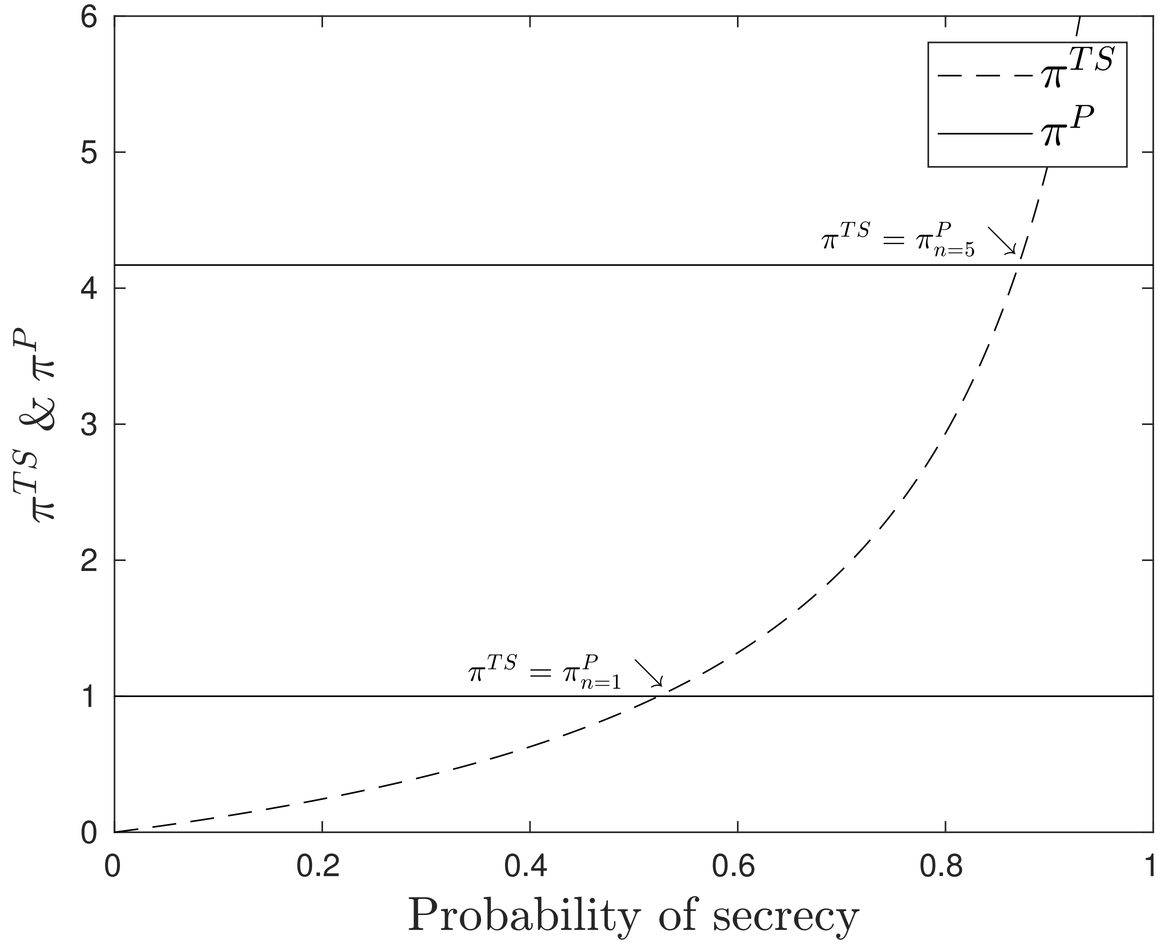

Figure 3.3. Plot of π^{TS}, π^{P} and p when $p \in [0, 1]$, $r = 0.1$ and $n = 1, 5$

Note: π^{TS} and π^{P} are evaluated using $\pi^{M} = 1$.

3.4 Model implications and relation to the AIA

The trade-off between patenting an invention or keeping it a secret involves three key variables as has been modeled — length or duration of a patent, the time discount rate, and the probability associated with keeping an invention secret. While the length of a patent is exclusively for inventors opting to patent their invention, and the probability of a patent being secret is exclusively for inventors opting to not file their invention as a patent, or in other words keeping it as a trade secret, the time discount rate influences both the patentee and the non-patentee. The implications of the model are simple and can be understood from the following points:

1. **Patent length or duration.** If a patent can be protected for long, or if its length increases, the certainty of a pipeline of income from the system of patent protection increases. Therefore, disclosing the invention as a patent is better than keeping it a secret when certainty from the system is high. This is because keeping the invention secret could result in its loss by ways of reverse engineering or it could be invented independently by others. This relationship is similar to the findings from Scotchmer and Green (1990), albeit theirs is more formal and rigorous. Figure 3.3 illustrates this, as n increases from 1 to 5, the horizontal line depicting profit from patenting shifts up. A horizontal line signifies certainty, while a curved dashed line that depicts profit from trade secrets has large variation, with possibilities of high and low profit as compared to profit from patenting.

   Pre-AIA, the rise in litigation was a concern and that brought uncertainty to the system. Without modifying the literal duration of the patent, the certainty of the system could be improved, and AIA precisely aimed for this, i.e. by incentivizing only the first filer of an invention, uncertainty around the date of invention was thought to be eliminated. Thus, the patent office could attract a larger number of patent filings instead of them being trade secrets. However, we do not know if the AIA indeed brought certainty to the system, as it could have unintended consequences, as mentioned in Chapter 2.

This is one of the crucial parameters of the model and the duration or length of a patent goes beyond its literal meaning to encompass certainty. In general, this parameter could also be interpreted as a benchmark to evaluate the patent system — how certain or consistent is the system?

2. **Probability of secrecy.** If an invention can be kept secret for long, the certainty that it will generate a stream of income for multiple periods is high. As we see in Figure 3.3, as the probability of secrecy increases (shown in the x-axis) i.e. profit from keeping an invention secret increases. In other words, with a higher probability that an invention can be kept secret, the profit from it mimics the profit from a longer length of patent. A longer patent length also meant a higher certainty from the patent system, as discussed in the previous point.

   Thus, it really depends on the AIA, as to how much certainty in patent protection it brings after its enactment. Does it enhance, or diminish it? Suppose methods of keeping a patent secret do not change drastically with AIA. In such a case, it all depends on the change AIA brings relative to the existing ways of keeping an invention secret.

3. **Discount rate.** Discount rate can be understood as the following: if there is a stream of income for multiple periods, in general, the present is more valuable

than the future — thus, the future is discounted using a *discount rate*. However, the future periods could either be valued highly, or not so highly, and this is what we understand by a low or a high discount rate respectively. From the model's Proposition 3.3, we understand that a rise in discount rate reduces the total stream of income from the invention from all periods regardless of patenting or keeping it secret — as only a few initial periods after an invention is valued in case of a high discount rate as opposed to valuing multiple periods fairly equally in case of a low discount rate. If the future is valuable, there exists a chance of extracting high value for multiple periods, and this is more lucrative for trade secrets as compared to patents. So what is better, a trade secret or a patent, if the discount rate is low? A successful trade secret could provide a stream of profit for a longer period if the future is as valuable as the present. Whereas, in case of a high discount rate, i.e. when the future is not as profitable as the present, the potential benefits from trade secrets start diminishing to the point where there is no difference between patenting and keeping the invention secret. At the extreme where the discount rate is the highest, i.e. equal to 1, or in other words the future does not hold any value for a particular invention, in such a world patenting is almost the best choice.

From the model and from Figure 3.2 we see that there is a large upside in profits from trade secrets when the future is as close as the present, i.e. has a low discount rate. But, as the discount goes closer to 1, the upside from the trade secret falls. Thus, as $r \to 1$, the future becomes increasingly non-valuable, and then patenting is almost always better than a trade secret.

Therefore, it depends on how long an invention would last once it is invented and that would govern if it would be patented or kept secret, *ceteris paribus*.

This result could be extended in multiple ways and one way this work goes is to study small entities on the AIA light. Extant literature suggests small entities (or startups) are good at solving problems while large entities are good at commercializing inventions. We also know that startups bring disruptive inventions (Arora et al., 2021). Along with this, follow-on patents are essential in developing the portfolio of patents for an inventor, and any inventor aims to file a series of patents when they file one in an area, so as to cover as much ground as possible. While we do not know if small entities are as good as large entities in filing follow-on patents as soon as they file their first patent, we know that they have limited resources. This has been shown recently to be true on the legal front as well (De Rassenfosse et al., 2021). The AIA connects on this front — because of the sweeping changes it brings in, it is hard to predict whether the AIA brings certainty to the system or not. Small entities could certainly wait for some time to test the water before filing patents, while large entities could still be filing at a rate similar to that of pre-AIA. Or, the AIA could in totality bring certainty that encourages under-resourced inventors to be confident and file their inventions as patents at a greater proportion than before. This is why law scholars were divided when the AIA was announced, as discussed in Chapter 2. This is where the empirical work comes in.

Thus, to summarize, we do not know if the AIA's FITF brings in more patenting than before from small entities. Conditions derived from the three parameters suggest multiple solutions. An inventor could increase patenting if the overall certainty of generating a stream of profit from the patent system increases. Then, the filing of patents over keeping trade secrets becomes worthwhile. However, due to firm heterogeneity, different firms (as separated by their sizes) could experience different trade-offs, and that would govern their choice of patenting.

Chapter 4: Data

4.1 Data sources

The main results of this dissertation are estimated using all the patents applied for at the USPTO between and including the years 2008 and 2016. Since not all entities release their balance sheet information, such as their R&D expenditure, assets, number of employees, etc., but such characteristics are determinants of patenting behavior, I use a subset of all patents that are applied for by the publicly traded firms in a separate analysis. For this subset of patent applications, I can control for certain entity-level characteristics (also called firm-level, since they are publicly traded firms). The data sources of patent-level and firm-level characteristics are given below:

Patents: Patent Examination Research Dataset (PatEx) is the dataset compiled by the Office of the Chief Economist (OCE) at USPTO that contains patent-level characteristics for the patents applied in the US from the 1900s to 2020. [1] In this dataset each row is a patent application which contains all the relevant information the patent application's prosecution generated till date; for example, the application number, filing date, issue date (if granted), number of claims, etc. Entity names in

[1] Patent Examination Research Dataset: https://www.uspto.gov/ip-policy/economic-research/research-datasets/patent-examination-research-dataset-public-pair

the patent database are not standardized and Patentsview bridges this gap using disambiguation algorithms to standardize patent assignee names. I sum the relevant statistics for each quarter for each assignee using the Patentsview standardized assignee names. [2]

Firms: Compustat's North America data provides quarterly financial information from the quarterly balance sheets of the publicly listed firms.[3] The Center for Research in Security Prices (CRSP) provides daily stock prices of these firms.[4] I average the daily stock prices to a quarterly-level.

While Patentsview standardizes patent assignee names, they still have to be connected to the Compustat database to get their financial information. Also, firms merge or hold subsidiaries that may individually file for patents. They may be listed as a different firm in the patent database but the patent belongs to the parent firm. Kogan et al. (2017) (henceforth referred to as KPSS) and Arora et al. (2021) (henceforth referred to as ABS) bridge this gap by standardizing the firm names, connecting them to the patent database, and connecting the Compustat firms and their subsidiaries to one standardized name. I use their databases to obtain a sub-sample of publicly listed firms' patents.

Litigation: The Office of the Chief Economist (OCE) at USPTO recently compiled a dataset on cases involving patents filed at district courts in the United States from 1963 to 2016. This dataset is called the Patent Litigation Docket Reports (PLDR).

[2]Patentsview is a collaborative project developed by the USPTO, American Institutes for Research (AIR), University of Massachusetts Amherst, New York University, University of California, Berkeley, Twin Arch Technologies, and Periscopic. See: https://patentsview.org/what-is-patentsview

[3]Compustat is accessed through Wharton Research Data Services (WRDS). https://wrds-www.wharton.upenn.edu/pages/about/data-vendors/sp-global-market-intelligence/

[4]Center for Research in Security Prices (CRSP) is accessed through Wharton Research Data Services (WRDS): https://wrds-www.wharton.upenn.edu/pages/about/data-vendors/center-for-research-in-security-prices-crsp/

From the PLDR, I create a measure of exposure to litigation at the entity-quarter level.

4.2 Sample selection

There are three types of patents in the US: utility, design, and plant. Inventions relating to new products or processes, or their improvement are utility patents. The AIA's objectives are best represented by these inventions and therefore, in this dissertation, I only study the utility patents. Plant or design patents are granted to inventions that relate to the development of new plants and new designs, which may be unaffected by the aspects of the AIA this dissertation focuses on. [5] Each row of the PatEx database is a patent application. The 2020 release of PatEx contains 16,514,638 patent application numbers. After removing design, plant, and blank patent applications, I have 15,811,897 patent applications. A patent application can appear multiple times in the dataset through continuations but will culminate in *one* granted patent.[6] A patent application appearing for the first time in the patent database is called a

[5]See USPTO patent process for details on the definitions of utility, plant, and design patents here: https://www.uspto.gov/patents/basics/patent-process-overview#step3

[6]At the USPTO, a patent application can be initially filed as a provisional or a non-provisional application. A provisional application may not contain claims or the specifics of the invention. Its primary use is to establish an effective filing date and should be followed by a non-provisional application, applied within 12 months of the provisional application's filing date. A non-provisional application is prosecuted by an examiner to determine its patentability. This type of application can further be continued as a continuation application, in-part continuation application, or divisional application. If an entity wishes to apply to more than one country, it may opt for a Patent Cooperation Treaty (PCT) application. If applied as a PCT application, the applicant has to choose the countries they wish to apply for patent rights. If an applicant wishes to file for patent protection in multiple countries, rather than applying to every country separately, the entity may choose to file the application as an international patent. This type of patent is also called a Patent Cooperation Treaty (PCT) application. The decision on grants is still given by the countries separately. For details refer here: https://www.wipo.int/pct/en/faqs/faqs.html Later, this enters the conventional application procedure, during which a new application number is assigned to the provisional or PCT application. Counting a PCT or a provisional and its conventional application counterpart as two different patent applications will result in double counting.

“parent” application, and all the connected applications, appearing later, are called its “children” applications. Multiple parent applications can be connected to multiple children applications. Using the parent and children continuation data from PatEx, I connect all the parents and their children to find the earliest application date. 13,537,926 patent applications contain a filing date.[7] I also restrict the dataset to patent applications that were applied between and including the years 2008 and 2016. I take patent applications post-2008 to avoid distortions from the financial crisis and till 2016 for comprehensive coverage of patent applications. The average grant lag of a patent is about four years and two standard deviations above the grant lag are about eight and half years at the USPTO.[8] Therefore, as years go by, the proportion of pending patent applications post-2016 will increase. For pending patent applications, information on citations is unavailable or scarce. The verdict for these applications, whether they will be granted or abandoned, is also unknown. To avoid these, I consider the patent applications that are applied on or before the end of 2016.

Restricting patent applications by year leaves me with 5,087,133 patent applications. Among these, 3,255,080 patent applications have a standardized assignee ID or an inventor ID in the Patentsview Database. The rest, 1,832,053 patent applications, do not have any ID. While these patent application numbers are unique in terms of their numbers or labels, they still can be a derivative of another application, as discussed earlier. In Appendix A.2, I provide a detailed discussion of the patent application numbers that do not have any ID and provide reasons as to why these applications

[7]The patent application that does not contain a filing date usually are filed as a PCT application. In place of the filing date, such applications receive a World Intellectual Property Organization (WIPO) publication date. The USPTO records these patent applications as National Stage Entry (NST) when they are examined at the USPTO, they receive an application number, and they do contain a filing date at this point.

[8]Grant lag winsorized at the 5 and 95 percentile cutoffs post-2000.

are either insignificant or are repeats of the patent applications already in the main sample of patent applications that can be identified with a standardized ID.

There are 581,980 unique entity IDs that have applied for patent applications between and including the years 2008 and 2016. I use the standardized Patentsview assignee IDs, wherever they are available, and the standardized inventor IDs where the assignee IDs are unavailable. For the rest of the dissertation, I refer to the entity identifiers as entity IDs. An entity can be identified as either a company, or an individual, or a Government entity. Assignee IDs are unavailable when an applicant is an individual. Each ID on average is involved in 5.6 patent applications, which includes joint work. Using the quarter of the earliest application date of each patent and the entity IDs, I sum or average all the variables at the entity-quarter level. I detail all the variables in the next section. 3,259,984 patents by 581,980 IDs are reduced to an unbalanced panel of 1,852,143 entity-quarter rows. Among these, 703,692 entity-quarters appeared only once in the sample, and therefore cannot be compared within an entity. These observations are singletons. Singletons in models with entity fixed effects where the standard errors are clustered may overstate the number of clusters and hence the statistical significance (Correia, 2015). I, therefore, report the main results without the singletons. The results of this dissertation however remain unchanged even with the inclusion of the singletons, as we will observe while discussing the results in Section 6. The final sample for the main set of analyses, therefore, is 1,148,451 entity-quarter observations. The concern is evident about entities self-selecting themselves into pre and post-AIA depending on their objectives. The unbalanced panel used in estimating the main results does not have zeroes when an entity does not file for any patent application for a given quarter. I address this issue by adding zeroes for the quarters where an entity does not patent, thereby balancing the entity-quarter panel and replicating the main analyses in Appendix A.3.

Next, I use patent number-*PERMNO* match from Kogan et al. (2017) to identify patents by publicly traded firms.[9] *PERMNO* is the permanent issue identifier as provided by CRSP data. There are 2,075 unique *PERMNO*s in Kogan et al. (2017) who had applied for patent applications between 2008 and 2016. I also use Arora et al. (2021) to match the *PERMNO*s with the Compustat identifier, *GVKEY*. There are 1,374 unique *GVKEY*s who had applied for patent applications between 2008 and 2016. For the firm-level analyses, I have 1,374 publicly traded firms' characteristics and their patent applications. Different studies employ their own assignee standardization method but Arora et al. (2021) corrects for mergers, acquisitions, and patent reassignment. From the 3,255,080 patent applications, 1,309,860 patents can be identified as having a *PERMNO* and a *GVKEY*. These observations constitute the sample for a separate set of analyses in Table 6.6.

4.3 Variables

4.3.1 Outcomes

I use two patent outcome measures as dependent variables to study the disclosure and the value of invention; log(number of patents) and scaled citations within two years of the issue of a patent.[10] Older patents may have a higher number of citations and are incomparable to the newer patents; and the number of citations may vary by technology classes and years (Mezzanotti, 2021; Lerner and Seru, 2022). Because of these two reasons, they are capped within two years of issue and adjusted by United

[9] The authors have released data updated till 2020. See: https://github.com/KPSS2017

[10] If a patent application is applied by N assignees, then each assignee receives ${\frac{1}{N}}^{th}$ of the patent.

States Patent Classification (USPC) subclass and year.[11] These values are aggregated at the assignee-quarter level.

As simple patent counts are rather discrete numbers that incorporate no weighing mechanism that could have assigned higher weight to an important patent (in terms of citations received), authors have at times transformed patent counts into citation-weighted patent counts. For example, Trajtenberg (1990) defines yearly number of patents for a particular product/technology class using citations the patents received, in the following way: $\sum_{i=1}^{n_t}(1 + C_i)$ where the total number of patents in a product/technology class is n at year t. So, for each year t, rather than reporting the number of patents for that year, the author scales each patent by the total citations received in a product class. However, in this dissertation, I explore if there is a change in filing before and after the AIA, and that requires looking at differences in raw number of patent counts filed in that period. I separately also look at an adjusted measure of citation, as explained.

Literature on innovation has routinely transformed the number of patents into the log(1+number of patents). But recently, Chen and Roth (2022), and Mullahy and Norton (2022) have questioned the addition of 1, which changes the shape of the underlying distribution. To tackle this, I only log-transform the number of patents, without the 1. Since there are no zeroes, the log of the number of patents provides a monotonic transformation and allows the interpretation of coefficients in percentage changes. The measurement of citations is slightly more complicated than the number of patents. Patents may as well receive zero citations, which cannot be log-transformed. Additionally, innovating entities have over the course of time changed their strategy of

[11] Average scaled citations are defined as the average of percentiles of citations by NBER sub cat and year.

citing other patents, and the overall citations have increased in recent years. This also varies by the various fields of inventions. These complexities have been documented by Lerner and Seru (2022), Higham et al. (2021), and Jaffe and De Rassenfosse (2019). Therefore, I calculate the percentile of citation within an NBER subcategory and year.

Using the parent application number's date of filing, I create an indicator called *Post*. This variable takes a value of 1 if the patent application was applied for after March 16, 2013, and 0 otherwise. This is the date when the AIA's FITF was enacted. But, as described in Section 2, The AIA was signed on Sept 16, 2011. Some of the other provisions were enacted at this time, while the main legislation, the adoption of the FITF rule, was implemented in March 2013. I use this date as the primary cut-off to define the pre and post-periods, but in the later part of the dissertation, and as a part of robustness, I report results using both the AIA's signing date, as well as the AIA's FITF implementation date.

Note that to determine if a patent was applied for in the pre-AIA or post-AIA period, I use the "parent" application number as opposed to the "patent" application number. This is because a patent application can claim priority to an earlier provisional or non-provisional application. Through this, an application filed later claims continuation from a previously filed application. If this happens, then the earliest application date is considered the application date for the patent application. In such cases, the patent applications filed post-AIA can still be under the purview of pre-AIA rules.

PatEx provides an indicator *Small Entity*, which takes a value of 1 if the patent application was filed by a small entity and 0 otherwise. A small entity is defined as a person, a small business concern, or a nonprofit organization, which includes universities and educational institutes. An entity is considered to be a small business concern if it meets the size standards, i.e. the total number of employees and affiliates should not exceed 500 persons at the time of application. The size is self-declared at the time of patent application and is not formally verified by the USPTO. A small entity receives a 50 percent discount on all patenting fees. An entity can file its patent as a small or a large entity, and the reasons for choosing either are unknown. [12] To keep the main results simple, I assume that an entity is small if 50 percent or more number of patents out of all patents by the entity is filed as a small entity. In Appendix A.5, I relax this assumption to redefine small entity using the proportion of patents filed as small among all patents an entity had filed for that quarter and replicate the main results.

4.3.2 Controls

The two most important controls in this dissertation are the entity and quarter-fixed effects. While quarter-fixed effects are routine, entity-fixed effects were difficult to control for due to variations in the spelling of assignees. I explain the use of Patentsview IDs in Section 4 to mitigate this problem. Entity-fixed effects allow us to estimate the average change at the entity level.

A patent application is categorized in the United States Patent Classification (USPC) System. Each patent application is assigned a USPC class and subclass. Since there are numerous USPC classes, I aggregate them into six NBER categories and 36 NBER sub-categories. The broad six categories are chemical, computer and communications, drugs and medical, electrical and electronic, mechanical, and others.[13] Since I am collapsing the patent-level statistics into the assignee-quarter level and

[12] If an entity that qualifies as a small entity transfers the rights of its patent to an undiscounted entity partly or fully, then the patent will cease to be from a small entity. However, licensing to a federal agency or using the patent as a security interest does not preclude the entity's patent from being considered as one from a small entity.

[13] A concordance between NBER categories and USPC classes are provided in Hall et al. (2005)

an assignee can have patents in various categories in each quarter, I take the mode category. If there are multiple modes of categories, I randomly choose one category for that assignee quarter.

Among the patent-level controls, literature using patent-level statistics uses the number of claims as a measure of the complexity and scope of the patent. Claims are specific statements that define the uses of a patent. In general, the higher the number of claims, the greater the patent's number of defined uses. Therefore, it becomes progressively difficult to add more and more claims to one patent application, as it adds multiple uses to one patent. Until recently, claims were thought to be a mix of complexity and quality of patents, but Marco and Miller (2019) dispel a few confusions around the significance of claims. They find that the narrow claims have a greater probability of being granted. Along with this, the claims' length and breadth are strategically written and fought with the examiner. Since the AIA changes incentives to file for patent applications substantially, all the variables that capture the patenting strategy are subject to change. This dissertation's focus is limited to studying the two broadest measures of quantity and quality of inventions, and therefore, I do not comment much on the claims of patents, and their constituents, as they deem a separate study. But in Appendix A.10, I report the gap between small and large entities' log(avg. number of claims) post-AIA. We observe a significant drop in the mean claims post-AIA by small entities. A drop in mean claims suggests that each patent post-AIA has on average lower number of defined uses. This also suggests that patent applications are now shorter in length. This is why claims are not a good covariate that should be in the list of control variables, as it would absorb significant variation in the number of patents and citations.

I use the log experience of an applicant at USPTO. I also control for the number of patents that are maintained at the 4^{th} year post-grant, or abandoned. Maintenance or renewal of patents and abandonment are predictors of the quality of patents. The higher the quality, the greater the chance that the patent would be renewed by the assignee, and on the contrary, the lower the quality the higher the chance that the patent would be abandoned by the assignee (Bessen, 2008). Next, following De Rassenfosse and Raiteri (2022) and Webster et al. (2014), I develop a measure for expertise in a field for an entity, called Revealed Technological Advantage (RTA). RTA is defined as the share of patents by an applicant over the share of patents by all applicants in a patent technology class. If an applicant applied the bulk of their patent applications in a given technology class that is higher than the share for all entities in the same technology class, the ratio will be greater than 1. Therefore, an RTA value greater than 1 signifies the high expertise of an entity in the given class, while the opposite, an RTA value lower than 1 signifies low expertise. These controls act to provide a stricter restriction to the estimates.

From Kogan et al. (2017) and Arora et al. (2021), I obtain a patent-*PERMNO-GVKEY* match. I use this information to create a subset of patents summed at the firm-quarter level. For these firms, I observe different firm characteristics that come from their quarterly balance sheets. Following the literature, I control for their quarterly log number of employees, log number of assets, log R&D expenditure, and log of firm age since establishment (Hegde and Sampat, 2009).

4.3.3 Litigation exposure

Following Mezzanotti (2021), I define exposure to litigation as a weighted average of litigation in all NBER subcategories, where the weights are the shares of patents by an entity in a subcategory at a given time.

I calculate the proportion of litigation in each subcategory till time T, where $T = 1 + 2 + ... + t$ as the number of litigated patents in subcategory c till time T (L_{ct}) over the number of patents in subcategory c till time T (P_{ct}). Formally:

$$l_{cT} = \frac{L_{cT}}{P_{cT}}$$

Next, I calculate the proportion of patents in each subcategory c by entity i at time t as the number of patents by entity i in subcategory c at time t over the number of patents by entity i at time t. Formally:

$$p_{ict} = \frac{P_{ict}}{P_{it}}$$

To obtain entity level measure of exposure to litigation (e_{ict}), I calculate the average litigation till time T for all subcategories i.e. average of l_{cT} for all the 36 NBER subcategories and weigh each of the subcategories with the proportion of patents by an entity i in those subcategories at time t (p_{ict}).[14] Formally:

$$e_{ict} = \sum_{c=1}^{36} p_{ict} \times l_{cT}$$

[14]The NBER categories and subcategories are explained in Subsection 4.3

Table 4.1. Definition of the variables

Panel A: Quarterly patent data		
Variable	**Definition**	**Source**
Log (no. patents)	Log(number of patents)	PatEx
Scaled 2 yr citations	Mean of citations' percentile by USPC subclass and year where citations are within 2 years of issue	Patentsview
Post	1 if the patent was filed on or after March 16, 2013	
Small entity (SE)	1 if the patent was filed by a small entity; 0 otherwise	
Log experience	Log experience of assignee at the USPTO	
First patent	1 if the patent was the first patent for the assignee; 0 otherwise	PatEx
NBER category	Patent classified into one of the six NBER categories: chemical, computer and communications, drugs and medical, electrical and electronic, mechanical, and others	
Renewal proportion	Proportion of patents renewed (maintained) at 4^{th} year	
Pending proportion	Proportion of patents pending as of Dec 2020	
Abandoned proportion	Proportion of patents abandoned	
Joint patent	Proportion of patents jointly applied	
Entity type	0 if company, 1 if individual, and 2 if Government	
Revealed tech. adv.	Revealed technological advantage is a ratio of ratios the proportion of patents in category c (P_{ict}) out of all patents by entity i at time t (P_{it}) over the proportion of patents in category c (P_{ct}) out of all patents by all entities at time t (P_t) Formally: $\frac{\frac{\sum P_{ict}}{\sum P_{it}}}{\frac{\sum P_{ct}}{\sum P_t}}$	PatEx
Litigation exposure	Weighted average of shares of litigation litigation till time T in category c (L_{cT}) out of all patents till time T in category c (P_{cT}) weighted by patents by entity i in category c at time t (P_{ict}) out of all patents by entity i at time t (P_{it}) Formally: $\sum \frac{P_{ict}}{P_{it}} \times \frac{L_{cT}}{P_{cT}}$ where the time till date (T) is defined as $T = 1 + 2 + ... + t$	PLDR

Panel B: Quarterly firm data		
Variable	**Definition**	**Source**
Log emp	Log number of employees	
(Log assets)/emp	Log assets per employee	Compustat
(Log R&D)/emp	Log R&D expenditure per employee	
Log age	Log age since establishment	
Log Market cap	Log(stock price × shares outstanding)	CRSP
KPSS value	Log real-KPSS value of patents	KPSS
PERMNO	CRSP firm identifier; identifies firms in KPSS	
GVKEY	Compustat firm identifier, identifies firms in Compustat; ABS provides GVKEY-PERMNO match	ABS

Panel A lists the variables used for the assignee level analyses. Panel B lists the variables used for the firm level analyses. Kogan et al. (2017) is abbreviated as KPSS, and Arora et al. (2021) is abbreviated as ABS. The variables *PERMNO* and *GVKEY* are not variables. From the assignee-quarter level patent data, the firm identifiers PERMNO and GVKEY are used to create a sub-sample of only-firm assignees.

4.4 Descriptive statistics by small and large entities

In this section, I describe a few crucial differences between small and large entities from the sample constructed for this dissertation. This helps us understand small entities in the context of this dissertation contrasting it with the literature that has previously studied small entities in a certain capacity. I additionally provide a summary of these variables in Table 4.2.

Table 4.2. Differences between small and large entities

Variable	Small	Large	Diff.
Patents	1.67	4.70	3.05***
Citations	0.48	2.06	1.58***
Abandoned	0.53	0.25	-0.28***
Cases	0.05	0.03	-0.02***
Exposure	-0.014	0.004	0.19***
RTA	3.69	1.81	-1.87***
N	488,546	517,184	

Patents and citations are average patents filed and citations received within two years of grant by small and large entities. Abandoned and cases are proportion of abandoned patents and proportion of patents involved in district court cases for small and large entities. Exposure is a standardized measure of exposure to litigation and RTA is a measure of expertise. The mean exposure and RTA are reported in this table. Column Diff. reports difference as (Large - Small). See Subsection 4.3.3 and Table 4.1 for detailed definitions. All the values are calculated using the paper's sample. P-values at 1, 5, and 10 percent are denoted by *, **, and ***.

A typical small entity on average files for 1.67 patents in the sample time frame, while a typical large entity on average files for 4.7 patents. While the average patents may seem low, the top patenting entities for each of these categories can file up to 4500 and 50,000 patents respectively, which illustrates the long tail of the distributions for patenting entities. The long tail is similar for both small and large entities. On similar lines, a typical small entity receives 0.48 citations while a typical large entity receives 2.06 citations within the first two years of their patents' grant. These numbers are consistent with prior literature, showing that small entities receive lower citations on average compared to large entities. However, an interesting pattern lies in the tails of the distribution of citations for small and large entities. The 95^{th} percentile cutoff of average scaled citations received by small and large entities are — 1 and 0.5 respectively. This indicates at the 95^{th} percentile, if small entities are compared with large entities, small entities receive higher citations. However, even though small entities do start receiving a greater number of citations relatively early on within the tail compared to large entities, i.e. at the 95^{th} percentile, at the 99^{th} percentile, large entities surpass the small entities and receive disproportionate citations within the entire patent dataset. Therefore, heavy-citation-receiving large entities pull the average for large entities so much that they on average surpass small entities.

Along with filing for a lower number of patents and receiving lower citations than large entities on average, small entities also do not renew, or – abandon 53 percent of their patents, which for large entities is 25 percent. This is also consistent with the literature, and the reasons cited are often times financial in nature (Bessen, 2008). Next, patents by small entities are also involved in a greater number of district court cases — about 5 percent of patents by small entities in this dissertation's sample were involved in cases between 2000 and 2016, while the same is 3 percent for large entities, as also discussed in Marco and Miller (2019) and Lanjouw and Schankerman (2001). Interestingly, if we look at average standardized exposure, we see that large entities file for patents more in the fields that have a higher exposure to being litigated. Finally, small entities are also highly specialized in their field as compared to larger entities, as shown by the average of Revealed Technological Advantage (RTA). A greater RTA means the proportion of patents by a particular entity in a particular field out of their portfolio of patents is greater than the proportion of patents by all entities in a particular field — also indicating that small entities put most of their eggs in one basket.

4.5 Descriptive evidence around the AIA

Figure 4.1 shows the number of applications filed each month at the USPTO since 1975. It also marks the months when the said amendments to the Patent Act were enforced (the figure also includes an indicator for Great Recession, which does not change patent rules in any way, but is an important event to report). We observe a rapid increase in filing of patent applications during or after each amendment was passed, for example we can see that an unusual number of applications were filed right after the implementation of Trade Related Intellectual Property Rights (TRIPS) in around 1995. After each of these spikes, we also observe either a change in growth rate, a parallel shift in monthly patents, or both. Along with the amendments to the Patent Act, certain patent lawsuits also delineate boundaries of patent rules by establishing case laws, which may affect patenting activity but such cases are not marked in the figure. The AIA came into effect as a law on September 16, 2011, which is the second red dotted line from the right. The FITF rule came into effect on March 16, 2013, as shown by the rightmost red dotted line. A large spike in patenting can be observed in the month when the AIA's FITF rule was enforced. From the Great Recession to AIA, the growth in the number of patents was positive which dampened after the AIA spike. But, we can observe an increase in noise post-AIA. I investigate this further in Figure 4.2.

Figure 4.2 furthers the investigation of the patenting activity around the AIA by separating the patents into small and large entities. After a preliminary cleaning

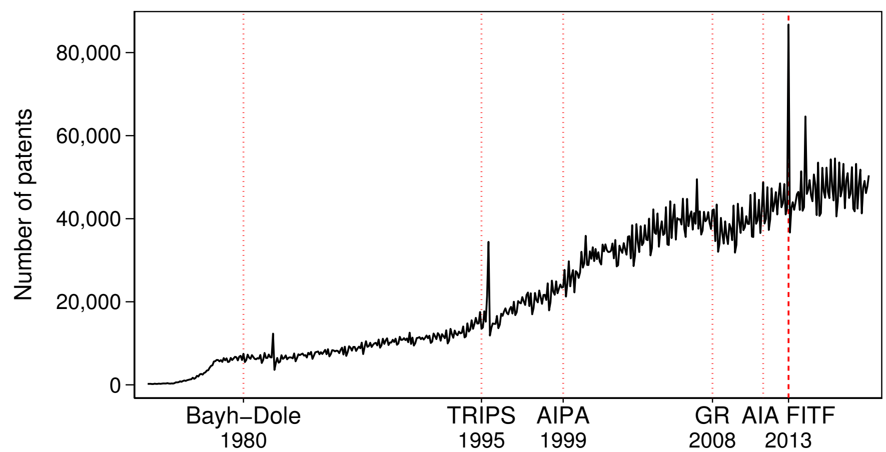

Figure 4.1. Patent applications filed each month

This figure plots total patents for each month from 1975 through 2018. The months of enforcement for the policies are the Bayh-Dole Act (Dec 1980), Trade Related Intellectual Property Rights (Jan 1995), American Inventors Protection Act (AIPA; Nov 1999), and America Invents Act (Sept 2011 and Mar 2013). The figure also indicates, for reference, the Great Recession (Sept 2008). Coverage of patents before 1981 is poor and only granted patents are observed till 1999. Post-1999, due to the introduction of pre-grant publication of patent applications, we see all barring the classified patents.

of raw patent application statistics, such as removing plant and design patents and restricting the dataset to observe patent applications only applied between 2008 and 2016, I plot Figure 4.2. It zooms in on the portion of Figure 4.1 after the Great Recession and plots the change in quarterly patents for small and large entities. I move from the monthly number of patent applications to quarterly because of two reasons. First, it irons out large monthly variations, making it easier to observe the pattern, and, second, all the empirical analyses are done at the entity-quarter level. Also, for publicly traded firms only quarterly financial statistics are available. In the estimations, I also switch from the absolute number of patents to log(number of patents) because regressing large absolute values on small $X's$ or vice-versa makes the estimates unreadable, and we are interested in studying the percentage change. But in the descriptive evidence figures, I plot changes quarter-on-quarter for the absolute number of patents.

We can observe similar trends for both small and large entities in Figure 4.2. In the USPTO database, entities are recorded as small entities or undiscounted entities. Undiscounted entities do not receive discounts in patenting fees and therefore are assumed to proxy large entities. The number of patents filed by large entities in magnitude is way higher than the number of small entities, but their change quarter-on-quarter remains similar. Similar to Figure 4.1, patent applications do not grow differently over quarters, show a rapid growth during the AIA's FITF and then the growth rate nears zero.

This fails to highlight the following: the average change pre and post-AIA within entities, and if entities who predominantly file for a high number of patents on average are different from the low patentees. The number of patents is a positively skewed distribution with a substantially long tail. A large number of small inventors and firms

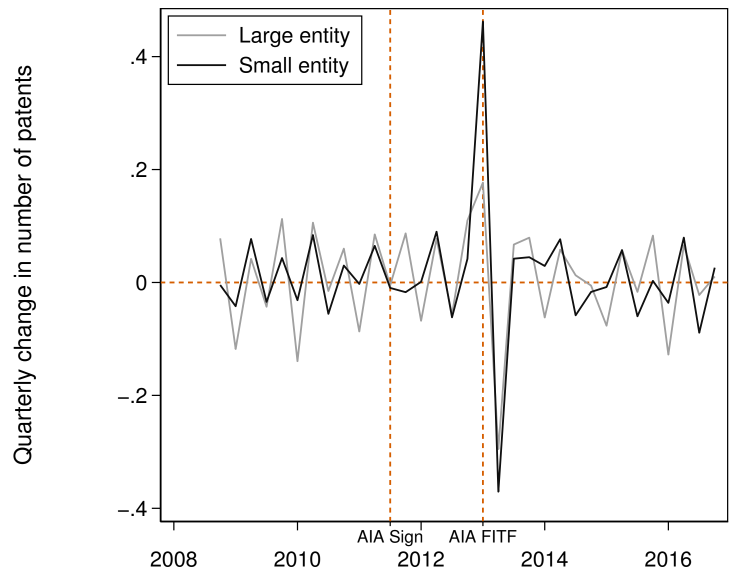

Figure 4.2. Patent applications by small and large entities
This figure plots the quarter-on-quarter change in the absolute number of patents for small and large entities from 2008 to 2016.

reside at the start of the distribution, having about one or two patents each quarter, and as we go along the x-axis of that distribution, we find large firms with around 2000 patents each quarter. Therefore, I show the main results of this dissertation considering only subsamples of entities by removing the top and bottom percentile of the distribution where they operate.

Not all patent applications are drafted equally. Some are of greater quality than others. Figure 4.3 plots the change in the number of citations per patent in the first two years after the issue of a patent quarter-on-quarter. Older patents will have an advantage here, and to limit that, I restrict the citations for the first two years after a patent's issue. Note that in the main set of analyses, I use a different measure of citations, called average scaled citations, as explained in Subsection 4.3. The quarter in the x-axis is the patent application quarter, rather than the issue quarter since a patent issued post-AIA can still be applied for before the AIA. Small entities on average have a lower number of citations, similar to a lower number of patents. One of the key differences we observe here is that the change in citations starts responding immediately after the AIA's signing, while the number of patents changes only at the FITF's implementation. Variation in citation among small and large entities is also more than the variation in the number of patents.

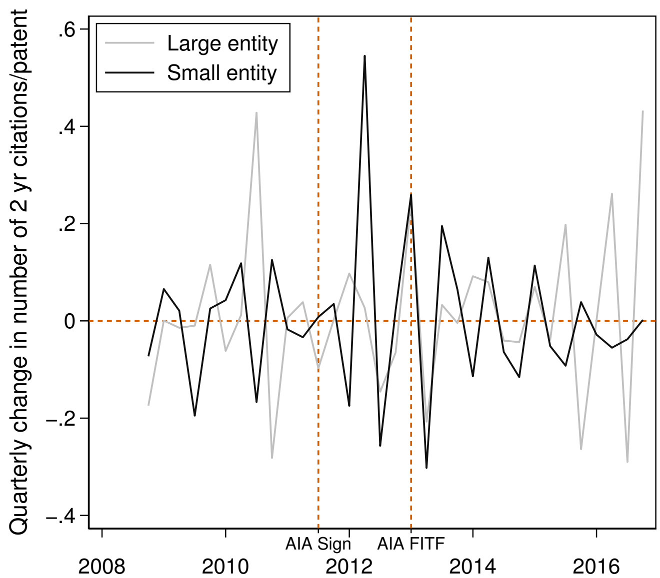

Figure 4.3. Citations/patent within two years

This figure plots the quarter-on-quarter change in the absolute number of citations per patent within two years of issue for small and large entities from 2008 to 2016.

Chapter 5: Empirical strategy

I estimate the change in the log(number of patent applications) and average scaled citations at the assignee-quarter level using an interrupted time series model. The number of patent applications and the adjusted percentile of citations proxy quantity and quality of inventions respectively. I compare the change in quantity and quality of inventions by small entities post-AIA with the inventions by small and large entities pre-AIA.

The main set of results, as presented in Tables 6.1 and 6.6 use two different samples: all patenting entities and a subset of them — the publicly traded firms involved in patenting, respectively. Next, I develop a measure of exposure to litigation, following Mezzanotti (2021) and estimate the changes in the two outcomes for only small entities. I use the measure of exposure to litigation as a continuous treatment, and the results are reported in Figures 6.4a and 6.4b.

A crucial assumption of this dissertation hinges on the comparison of small and large entities pre and post-AIA. I assume that if the AIA was not enacted, there would be no change in the rate of change of the number of patents and citations between small and large entities quarter-on-quarter. Also, by comparing discrete entity sizes i.e. small and large entities, I implicitly assume that size is an adequate measure to capture the resources the entities have at their disposal. This measure because of being discrete dampens the variation that I could have exploited in the estimation if I had a continuous measure of resources available to entities.

The main set of results i.e. the entity level and the firm level analyses, are estimated using Equation 5.1, and the full specification is as follows:

$$Y_{it} = \beta_0 + \delta(Post \times SE_i) + X_{it}^{'}\beta + \lambda_i + \lambda_t + \varepsilon_{it} \tag{5.1}$$

Here Y_{it} denotes two outcome variables, the log(number of patents) for an assignee at each quarter and the average scaled citations received by an assignee at each quarter within two years of the patents' issue. The coefficient δ captures the change in the difference in Ys between small and large entities after the implementation of the AIA. I control for the assignee i and the quarter t's baseline using assignee and quarter fixed effects, given by a range of indicator variables and their coefficients in the matrices λ_i and λ_t respectively. I also control for a range of patent quality and complexity correlates denoted by the vector X_{it}. A discussion on these measures follows after the models used in this dissertation are explained.

I estimate this equation for the full sample, i.e. all the patenting entities, and for only the publicly listed firms. For the firm-level analyses, I employ additional firm-level controls, which control the firm's resources with greater precision than only using the information derived from their patenting behavior and entity size.

While δ reports the average of Ys for small entities over all the quarters post-AIA, I also separate the effects by each quarter. This enables us to understand and verify if any pre-trend influences δ. The estimating equation is given by Equation 5.2. The full specification is:

$$Y_{it} = \beta_0 + \sum_{s \neq 0} (\beta_s \times 1[s = t] \times SE_i) + X_{it}^{'}\beta + \lambda_i + \lambda_t + \varepsilon_{it} \tag{5.2}$$

In this Equation, β_s ranges from the first quarter of 2008 to the last quarter of 2016, barring the first quarter of 2013 i.e. when the AIA was implemented, which acts as the base quarter. The other coefficients have the same interpretation as in Equation 5.1.

Next, following Mezzanotti (2021), I estimate the change in the log(patents) and avg. scaled citation for varying degrees of exposure post-AIA. The full specification is given by:

$$Y_{it} = \delta(Exposure \times Post) + X_{it}^{'}\beta + \lambda_i + \lambda_t + \varepsilon_{it} \tag{5.3}$$

Exposure is defined in Section 6.3, and is standardized. Therefore, δ reports the average change in the log(number of patents) and avg. scaled citations post-AIA when exposure to litigation increases by one standard deviation. Note that in this equation, I do not distinguish between small and large entities. I estimate this for only small entities and discuss the results. Later, in a triple-difference setup, I look at the difference in effects for small entities that are exposed to litigation post-AIA to examine if small and large entities are different when they are prone to litigation. I report this result in Appendix A.12.

Chapter 6: Results

6.1 Evidence from all small entities

In this section, I examine and report the change in two outcome variables: quantity of inventions — log(number of patent applications) and the quality of inventions — average scaled citations after the enactment of the AIA for all small entities in the sample. The results are reported in Table 6.1, and this table constitutes the main and broadest result of this dissertation. These results are also disaggregated by quarters and reported in event study forms in Table 6.2, Figures 6.1 and 6.2. In this section, I discuss those broad results, i.e. how the quantity and quality of inventions changed for small entities before and after the AIA — on average and quarterly; and compare its similarities and dissimilarities with the previous studies that focus on the AIA and AIA-like reforms in other countries. In further sections, I delve deeper into the nuances of these results. For this and all the subsequent tables, I present two columns for each Y variable, where the first column is with and the second is without controls.

The estimate δ from Equation 5.1 reports how the gap between small and large entities changes after the AIA's enactment. In Table 6.1, this is reported in the first row $SE \times Post$. Here the post is defined as the first-inventor-to-file's adoption date, i.e. March 2013, and the timeline for which the estimates are calculated is post-Great

Recession i.e. 2008, to 2016. The difference between the number of patents filed by large and small entities decreases by 4.1 percent after the AIA's implementation. This indicates a relative drop in the number of patents filed by an average small entity as compared to an average large entity post-AIA.

Table 6.1. Main results: log(no. of patents) and avg. scaled citations

	Patents		Citations	
	(1)	(2)	(3)	(4)
SE x Post	-0.0557***	-0.0412***	-0.1689***	-0.2480***
	(0.0042)	(0.0030)	(0.0568)	(0.0562)
Controls	No	Yes	No	Yes
Qtr F.E.	Yes	Yes	Yes	Yes
ID F.E.	Yes	Yes	Yes	Yes
Subcat F.E.	Yes	Yes	Yes	Yes
N	1148250	1146957	1148250	1146957

This table reports the gap in the patents and citations between small and large entities post-AIA. The number of patents are measured by log(number of patents) and citations are measured by the average scaled citations within 2 year of issue of patent. The estimates are derived from the model $Y_{it} = \delta(Post \times SE_i) + X_{it}^{'}\beta + \lambda_i + \lambda_t + \varepsilon_{it}$, where i and t denote entity and quarter respectively, X_{it} denotes a range of patent and entity level controls as defined in Table 4.1, and λ_i, λ_t control for entity and quarter fixed-effects respectively. The entity-quarter level dataset consists of all patents applied between 2008 and 2016. The indicator variable SE takes value 1 if more than 50 percent of the patents for an entity was applied as a small entity, and 0 otherwise. $Post$ takes value 1 if the patents were applied on or after the first quarter of 2013 i.e. the implementation of first-inventor-to-file rule in the United States.

Columns (1) and (3) report results without and (2) and (4) report with controls. P-values at 1, 5, and 10 percent are denoted by *, **, and ***. Standard errors are clustered at entity-level.

When we look at the differences in log(number of patents) for each quarter compared to the base quarter i.e. the first quarter of 2013 in Figure 6.1, a visual evaluation reconfirms the results from Table 6.1 — an increase in the difference between the number of patents by small entities post-AIA. While on average the gap between small and large entities increases, the quarterly estimates also fluctuate substantially. In particular, compared to Q1 2013 i.e. post-AIA's first-inventor-to-file rule's adoption, the trend for number of patents were higher pre-AIA, and after the AIA, the difference between small and large increased in every quarter. It should be reiterated that in 2011, the AIA was formally enacted. A narrower window of analysis keeping 2011's 3^{rd} quarter as the base quarter will be discussed in detail in Section 6.2.

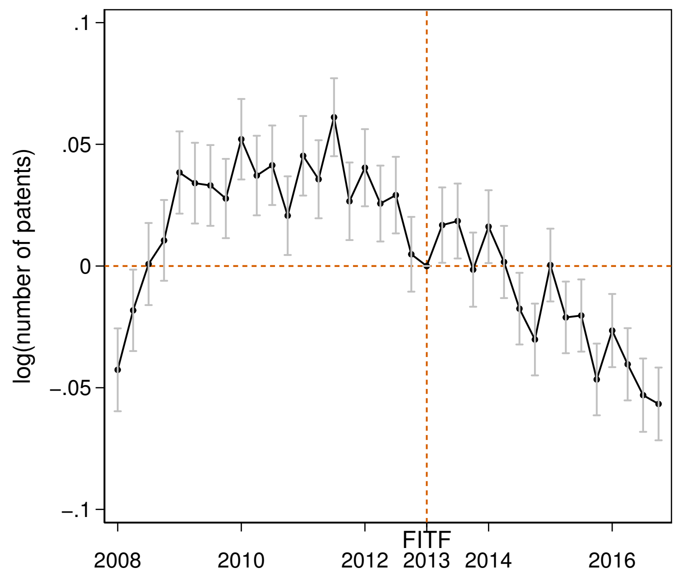

Figure 6.1. Change in small entities' patents over time
This figure plots the change in log(number of patents) by small entities for every quarter in our study timeline. The event study specification is reported in Equation 5.2, and the figure plots the β_s's estimated from $\sum_{s\neq 0} \beta_s \times 1[s = t] \times SE_i$ part of the equation, with the base quarter being the first quarter of 2013, when the AIA's FITF came in force.

This drop in patenting by small entities is similar to the one observed in Canada when a similar change in patenting rule was enacted, i.e. a first-inventor-to-file rule. Abrams and Wagner (2013) and Lerner et al. (2015) note a drop in patenting activity for all types of inventors following the enactment of such legislation, and also report a greater drop, especially among the individual inventors of Canada. The reasons for this, they note, are: one, firms have an advantage in the "race" to reach the patent office first in terms of resources compared to the individual inventors, and on the contrary, these inventors do not have adequate resources in terms of fees to attorneys and agents that are required to file for patents quickly. Two, fewer inventions by inventors and a shift to secrecy in order to protect their invention, which can also be considered as a corollary to the first reason, i.e. because of fewer resources to file for patents quickly after invention, small inventors have moved to a greater degree of secrecy than before. Three, Abrams and Wagner (2013) note that small inventors may be demoralized in filing for patents. I do find certain evidence from the results of exposure to litigation, which pushes us to think in the direction of not only "demoralization" but also "testing the waters" after such a sweeping change brought forth by the AIA. This pathway will be discussed in detail in the Subsection 6.3. The next two reasons; four and five are: they speculate that individual inventors join firms and shift their patenting activity to the U.S. This dissertation's scope and resources do not allow to test the former hypothesis of individual inventors joining firms, and the second hypothesis does not apply to the study of AIA, because the United States was the last country to adopt FITF rule. That means inventors neither have any other country where they can enjoy a rule of first-to-invent, nor do they have geographical proximity such as moving their activity from Canada to the United States.

Before we move to discuss the nuances in the drop in the number of patents by U.S. small entities, and their quality in the subsequent sections, let me comment on the other coefficients of the first two columns. The time trends are controlled using quarter-fixed effects. To minimize clutter, I plot the fixed effects in a separate graph, and is shown in Appendix A.6's Figure A.1a and Figure A.1b. The trends exhibit a pattern similar to Figures 4.1 and 4.2, i.e. for the quarters after the AIA, we observe a decline in the growth rate quarter-on-quarter and a plateauing of the number of patents for all the entities. Note that in Equation 5.1, SE and ID fixed-effects capture the overall effects for small entities and each entity's mean number of patents. The quarter fixed effects report a trend over and above the individual baselines of each entity as compared to the quarter when the AIA was enacted.

The AIA is likely to induce a strategic change, as we observe with the number of patents. This may as well be observed in other measures that are derivatives and details within a patent document, and the inclusion of those variables as a covariate in the model can absorb some of the variation in Y and therefore are "bad controls". Such a control, which is also an important measure of complexity is "claims" that is often used in the literature on the economics of patents. As discussed in Section 4.3.2, the use of claims as a strategy is documented by Marco and Miller (2019). A change in incentives to patent due to the AIA's introduction can not only drop the number of patents as a response, but each of those patents could also be shorter and of a different quality. And, this is what is observed when I estimate change in log(claims) using the same specification as for the log(number of patents), i.e. Equation 5.2. This result is reported in Appendix A.5. We can see from the figure that the number of claims per patent does fall after the AIA's enactment. Since it is a per-patent measure, and since we know that the number of patents also drops after AIA, it must be the case that the number of claims also significantly drops post-AIA. This measure is deemed a separate study, and the focus of this dissertation is to estimate the broadest change in quantity and quality of inventions by small entities post-AIA. Therefore, I do not discuss the change in claims further. However, this result stays as an addendum and assists in the mechanism that will be discussed in the next sections.

Now, I discuss the second outcome, the average scaled citations within two years of the issue of the patent. From Table 6.1, we observe results similar to that of the log(patents). Specifically, the estimate of δ reported in columns (3) and (4) of Table

6.1 indicates a drop in the number of citations within two years of the issue of a patent per patent of about 19.8 percent for small entities post-AIA relative to large entities. This is rather an alarming drop, but this drop is driven by the years which are substantially far away from the AIA's dates, unlike the results from the log(number of patents), which stays robust with a shorter window, different specifications, and additional controls. This compels us to think if this drop in citations is indeed due to the AIA, or due to other changes, unrelated to the legislation. In an event study setup, shown in Table 6.2, and Figure 6.2, this becomes clearer, that the years 2008 and 2016 in particular are driving this gap. In particular, the yearly coefficients for citations starting from 2009 till 2015 all remain statistically insignificant in Table 6.2.

While Abrams and Wagner (2013) and Lerner et al. (2015) report a reduction in the number of patents in Canada, they do not find any appreciable difference in the quality, as also measured by citation, for small entities during the Canadian reform. This dissertation's results show a similar trend among the U.S. inventors. However, De Rassenfosse (2013) show that trade-offs between quantity and quality of patents do exist. This result is not prominent at least when we study patenting activity around the AIA's enactment.

Citations are a tricky and complicated proxy to measure the quality of patents. Over time, strategic citations have become prevalent, and the overall number of citations has increased per patent (Lerner and Seru, 2022). Because of this, I report results with alternative definitions of citations in Appendix A.7 and A.8. In Appendix A.7, rather than using citations within two years of the patent's issue as an outcome and then transforming it into its scaled version, I use citations within one year of the issue date and show the results in two ways: scaled citations by NBER subcategory and year, and quarterly citations per patent. In Appendix A.8, I separate citations

Table 6.2. Event study results: log(no. of patents) and avg. scaled citations

	Patents		Citations	
	(1)	(2)	(3)	(4)
SE x 2008	0.0116	-0.0180***	0.1992*	0.3074***
	(0.0072)	(0.0052)	(0.1127)	(0.1115)
SE x 2009	0.0312***	0.0236***	0.1495	0.2285**
	(0.0069)	(0.0050)	(0.1081)	(0.1073)
SE x 2010	0.0214***	0.0302***	-0.0252	0.0466
	(0.0065)	(0.0048)	(0.1045)	(0.1036)
SE x 2011	0.0157***	0.0300***	0.0154	0.0544
	(0.0060)	(0.0045)	(0.1005)	(0.0997)
SE x 2012	0.0035	0.0124***	-0.0195	-0.0012
	(0.0054)	(0.0041)	(0.0913)	(0.0906)
SE x 2014	-0.0323***	-0.0192***	-0.0213	-0.0329
	(0.0051)	(0.0038)	(0.0842)	(0.0836)
SE x 2015	-0.0579***	-0.0353***	-0.1248	-0.1822**
	(0.0055)	(0.0041)	(0.0856)	(0.0851)
SE x 2016	-0.0803***	-0.0546***	-0.3803***	-0.4214***
	(0.0059)	(0.0043)	(0.0868)	(0.0861)
Controls	No	Yes	No	Yes
ID F.E.	Yes	Yes	Yes	Yes
Subcat F.E.	Yes	Yes	Yes	Yes
year	Yes	Yes	Yes	Yes
N	1148250	1146957	1148250	1146957

This table reports the gap between small and large entities for each year before and after the AIA. The number of patents are measured by log(number of patents) and citations are measured by the average scaled 2 citations within 2 year of issue of patent. The estimates are derived from the model $Y_{it} = \beta_0 + \sum_{s\neq 0}(\beta_s \times 1[s = t] \times SE_i) + X_{it}^{'}\beta + \lambda_i + \lambda_t + \varepsilon_{it}$, where i and t denote entity and quarter respectively, X_{it} denotes a range of patent and entity level controls as defined in Table 4.1, and λ_i, λ_t control for entity and quarter fixed-effects respectively. The entity-quarter level dataset consists of all patents applied between 2008 and 2016. The indicator variable SE takes value 1 if more than 50 percent of the patents for an entity was applied as a small entity, and 0 otherwise. The year 2013 acts as the base year, when the AIA's FITF came into effect.

Columns (1) and (3) report results without and (2) and (4) report with controls. P-values at 1, 5, and 10 percent are denoted by *, **, and ***. Standard errors are clustered at entity-level.

added by the examiner and applicants and re-calculate the average scaled 2-year citations. All of the exercises show that citations for small entities compared to large entities do not significantly change post-AIA.

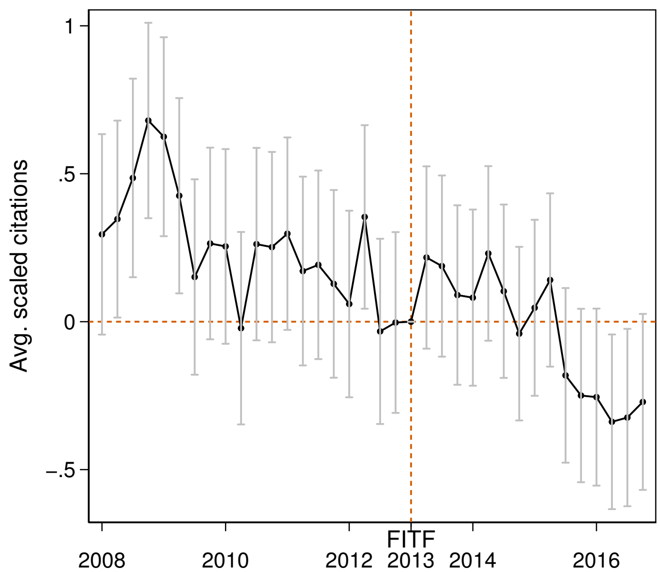

Figure 6.2. Change in small entities' citations over time
This figure plots the change in 2 yr. scaled citations by small entities for every quarter. The number of citations is transformed into the percentile within the NBER subclass and quarter to create scaled citations. The event study specification is reported in Equation 5.2, and the figure plots the β_s's estimated from $\sum_{s \neq 0} \beta_s \times 1[s = t] \times SE_i$ part of the equation, with the base quarter being the first quarter of 2013, when the AIA's FITF came in force.

The AIA was implemented in stages, as explained in Subsections 2.1 and 2.4. But, these stages are fundamentally different from each other in terms of changes in the benefits and costs of patenting. Therefore, they have to be evaluated individually. In the next section — Section 6.2, I use the enactment of AIA as the base quarter and discuss how the results stay the same and do not change qualitatively.

6.2 Narrow window of analyses: 2009-2014

In this section, I compare the results discussed in Section 6.1 for a narrower window of analyses (2009-2014), and with a different comparison quarter — the implementation quarter of the AIA — September 2011, or quarter 3 of 2011.

This section serves two purposes: it illustrates an anticipation for the AIA's first-inventor-to-file rule and rules out any effect that could be attributed to the recent Supreme Court judgments on *Alice v. CLS Bank* decided on June 2014, and *Bilski v. Kappos* decided on June 2010.[1, 2] This case was on patentable subject matter, or patentability of abstract ideas. Prior to these cases, any "useful results" were patentable, and the Supreme Court did not intervene in the patent litigation process, leaving it to the Federal Circuit. *Alice* held four patents which were challenged by *CLS Bank* on their patentability and were later invalidated by the Supreme Court. Feng and Williams (2023) and Lemley and Zyontz (2021) study this case in detail. Lemley and Zyontz (2021) study *Alice v. CLS Bank* and the related cases, and their relation to invalidated patents after the case's decision. They found that small and individual inventors were affected adversely due to the case's decision, and the proportion of invalidations increased. Feng and Williams (2023) confirm that *Alice*-like patents (patent similar to the ones in dispute in *Alice v. CLS Bank*) were invalidated at a greater rate post-*Alice v. CLS Bank*. This increase in invalidation, as the authors point out, was linked to greater scrutiny by examiners. The activity of Patent Assertion Entities (PAE) and overall litigation fell in *Alice*-like areas (patent technology classes that are similar to the patents in *Alice v. CLS Bank*) which in essence weeded out

[1] Bilski v. Kappos, 561 U.S. 593 (2010)
[2] Alice Corp. v. CLS Bank Int'l, 573 U.S. 208 (2014)

some of the costs to the patenting system that arise due to excessive litigation. Did this come at the cost of reducing the patenting activity of small entities, the central question of this dissertation? Feng and Williams (2023) find that startups in *Alice*-like areas or industries were not affected due to the Supreme Court decision.

Table 6.3 and Figures 6.3a and 6.3b report pre and post-AIA averages, and event study results respectively for a shorter window of analyses, 2009-2014, and considers the quarter of the AIA's enactment as the base quarter for comparison — as opposed to the AIA's first-inventor-to-file's quarter. Table 6.3 serves as a robustness, and shows that mean drop for small entities post-AIA occurred over the course of the AIA's implementation, and the coefficients of column 2 are close to that of the main results from column 2 of Table 6.1. Specifically, we observed a 4.4 percent drop in the longer timeframe and a 4.2 percent drop for small entities in the shorter timeframe. For citations, we observe no statistically significant changes post-AIA. This confirms that the tail years, i.e. 2008, 2015, and 2016 were driving the results.

It is possible that the Great Recession of 2008 induced a different patenting strategy, and thus an increase in citations around 2008-2009, but it is at this point unclear and requires further exploration of its mechanisms. For the years 2015 and 2016, it is likely that the *Alice*'s decision is influencing a drop in citations. Lemley and Zyontz (2021) note that the case, on one hand, may increase certainty in what can and what cannot be patented, they also have spurred multiple inconsistencies in subsequent patent invalidity judgments, increasing confusion. Because of this, it is possible that we are observing a strategic drop in average citations over all patent categories. Entities are deliberately keeping their inventions secret, and are disclosing only those inventions that are of poor quality, on average.

Figure 6.3a provides a cleaner view of the widening gap between small and large entities' number of patents post-AIA. While the gap remained insignificant compared to Sept 2011, even after *Bilski v. Kappos* (a case similar to *Alice*) which was decided in June 2010, it started dropping after the AIA's enactment, showing that at the least, part of the results can be attributed to the AIA rather than only to the Supreme Court decisions. While we do observe a significant change in the number of patents by small entities, the citations do not move away from zero, and the estimates exhibit a greater imprecision, shown by the wide standard errors, as compared to the number of patents' confidence intervals. In the two figures, I mark the two important events of the AIA, its enactment, and the implementation of first-inventor-to-file. The second vertical dotted line compares the results with Figure 6.1.

While this dissertation's objective is to show certain broader results related to small entities' quantity and quality of patents around the AIA, dividing innovating entities into small and large may not be enough. As Lerner and Seru (2022) and Jaffe and De Rassenfosse (2019) note, the use of citation to proxy quality of inventions has its own pitfalls and has been knowingly or unknowingly misused in the literature. Among the distribution of patentees, the entities who file routinely file for a considerably higher number of patents than the average are qualitatively and strategically different from the others. Abrams et al. (2019) show that a group of small entities aggregate, and sell their "weak" patents to non-practicing entities, who in turn use those to invalidate other patents. To disassociate different strategies by patent portfolios of entities, I provide event studies for different subsamples within the distribution of low to heavy patentees in Appendix A.9. I loosely define "low patentees" as the entities who fall between [0-75) percentile in the distribution of the number of patents held by entities in their portfolio, and "high patentees" as those who fall above the 90^{th} percentile in the same distribution. We do observe interesting behavior among the two types, and in particular, the drop in patenting activity is driven by small entities who are between the 90-99 percentile of the distribution of number of patents, as given by Figure A.4b. As opposed to this, the patentees having a lower number of patents in their portfolio, on average, spring back after the AIA's first-inventor-to-file implementation. The top 1 percent are not statistically different from the large.

Table 6.3. Main results with shorter timeframe

	Patents		Citations	
	(1)	(2)	(3)	(4)
SE x Post	-0.0409***	-0.0380***	-0.0958	-0.0903
	(0.0052)	(0.0040)	(0.0860)	(0.0853)
Controls	No	Yes	No	Yes
Qtr F.E.	Yes	Yes	Yes	Yes
ID F.E.	Yes	Yes	Yes	Yes
Subcat F.E.	Yes	Yes	Yes	Yes
N	457448	456944	457448	456944

This table reports the gap in the patents and citations between small and large entities post-AIA. The number of patents are measured by log(number of patents) and citations are measured by the average scaled citations within 2 year of issue of patent. The estimates are derived from the model $Y_{it} = \delta(Post \times SE_i) + X_{it}^{'}\beta + \lambda_i + \lambda_t + \varepsilon_{it}$, where i and t denote entity and quarter respectively, X_{it} denotes a range of patent and entity level controls as defined in Table 4.1, and λ_i, λ_t control for entity and quarter fixed-effects respectively. The entity-quarter level dataset consists of all patents applied between 2009 and 2014. The indicator variable SE takes value 1 if more than 50 percent of the patents for an entity was applied as a small entity, and 0 otherwise. $Post$ takes value 1 if the patents were applied on or after the third quarter of 2011 i.e. the AIA's enactment quarter — as opposed to the implementation of first-inventor-to-file rule that was considered in Table 6.3.

Columns (1) and (3) report results without and (2) and (4) report with controls. P-values at 1, 5, and 10 percent are denoted by *, **, and ***. Standard errors are clustered at entity-level.

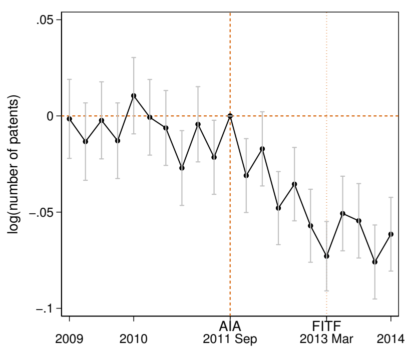

(a) Patents

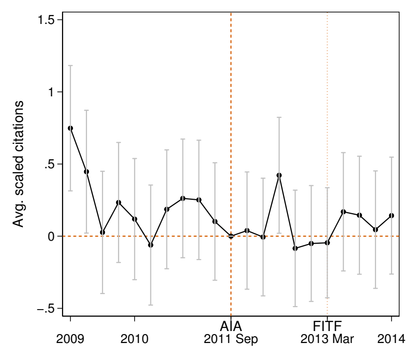

(b) Citation

Figure 6.3. Narrow window of analyses: 2009-2014

These figures plot changes in log(number of patents) and average scaled citations by small entities for a shorter timeline, the first quarter of 2009 to the first quarter of 2014. Also, the base quarter is the AIA's signing quarter (September 2011) rather than the quarter of the first-inventor-to-file's enactment (March 2013). The event study specification is reported in Equation 5.2, and the figure plots the β_s's estimated from $\sum_{s\neq 0} \beta_s \times 1[s=t] \times SE_i$ part of the equation.

6.3 Evidence from the litigation exposure

In this section, I focus on the variation from exposure to litigation and report results for small entities and all entities separately. In Appendix A.12, I provide triple difference results, comparing small entities — exposed to litigation — post-AIA.

We observe that in addition to the gap in the number of patents between small and large entities widening post-AIA; within small entities, those operating in patent categories with heavier exposure to litigation do decrease their patenting even more compared to all small entities post-AIA. Specifically, a one-standard-deviation increase in exposure to litigation results in about a 15.7 percent reduction in the number of patents for all entities, and 42 percent among small entities, given by the coefficient of $Exposure \times Post$ from Tables 6.4 and 6.5 respectively.

It should also be noted that while on average citations did not significantly change for small entities post-AIA, among those exposed to litigation filed for patents that received lower citations on average, as shown in Figure 6.4b, and columns 3 and 4 of Table 6.5. While the number of patents, and citations did drop for all entities, for small entities, we observe an increased effect, and this result is confirmed when I estimate the averages using a triple difference model reported in Appendix A.12.

Post-AIA, we can infer that the entities experiencing heavy exposure to litigation are in a way more "discouraged" to patent in those areas. The mechanism which can explain this result is complex. While the AIA tried to increase certainty, the establishment of the Patent Trial and Appeal Board (PTAB) made it easier to oppose patents' claims, and on top of this, the Supreme Court Cases, from *Bilski* to *Alice* exacerbated the "discouraging effect" of filing patents in areas with higher exposure to litigation.

Lemley and Zyontz (2021) report an increase in patent invalidations, and Marco et al. (2017) report a dramatic increase in litigation from 2010. In Figure 6.4, I separate the pre and post-difference in means for each quarter and compare it with the quarter when the AIA's FITF was enacted, only for small entities. We note that the exposed entities were anticipating the change and were lowering their activity till the AIA's passage, and this decline sustains, especially for small entities, post-AIA's first-inventor-to-file. For the number of patents for all entities, we can see a change in the rate of decline, and it nears zero in the post-AIA period. A drop can also be observed for the citations. When we compare the event studies for only small entities versus all entities, we see that the exposed small entities are the most adversely affected among all. This result is a combination of the Supreme Court cases and the AIA. At this point, due to the overlapping nature of events, it is difficult to disentangle the effects-driven due to the AIA, and due to the cases, and this dissertation can only inform the average change seen during the period. With a different and nuanced identification strategy, this route can be separately explored in future studies.

Table 6.4. Exposure to litigation for all entities

	Patents		Citations	
	(1)	(2)	(3)	(4)
Exposure x Post	-0.1165***	-0.1139***	-0.0101***	-0.0093***
	(0.0232)	(0.0226)	(0.0023)	(0.0021)
Exposure	0.2975***	0.2799***	0.0228***	0.0211***
	(0.0557)	(0.0530)	(0.0048)	(0.0045)
Controls	No	Yes	No	Yes
Qtr F.E.	Yes	Yes	Yes	Yes
ID F.E.	Yes	Yes	Yes	Yes
Subcat F.E.	Yes	Yes	Yes	Yes
N	1148250	1146957	1148250	1146957

This table reports difference-in-differences estimates for litigation exposure post-AIA, given by $Exposure \times Post$ for all entities. The number of patents are measured by log(number of patents) and citations are measured by the average scaled 2 citations within 2 year of issue of patent. The estimates are derived from the model $Y_{it} = \delta(Exposure \times Post) + X_{it}^{'}\beta + \lambda_i + \lambda_t + \varepsilon_{it}$, where i and t denote entity and quarter respectively, X_{it} denotes a range of patent and entity level controls as defined in Table 4.1, and λ_i, λ_t control for entity and quarter fixed-effects respectively. The entity-quarter level dataset consists of all patents applied between 2008 and 2016. The indicator variable SE takes value 1 if more than 50 percent of the patents for an entity was applied as a small entity, and 0 otherwise. $Post$ takes value 1 if the patents were applied on or after the first quarter of 2013 i.e. the implementation of first-inventor-to-file rule in the United States.

Exposure is a standardized measure and is defined in Subsection 4.3.3.

Columns (1) and (3) report results without and (2) and (4) report with controls. P-values at 1, 5, and 10 percent are denoted by *, **, and ***. Standard errors are clustered at entity-level.

Table 6.5. Exposure to litigation within small entities

	Patents		Citations	
	(1)	(2)	(3)	(4)
Exposure x Post	-0.3490**	-0.3709**	-0.0514**	-0.0417**
	(0.1674)	(0.1746)	(0.0254)	(0.0207)
Exposure	1.1380***	1.0951***	0.1333***	0.1116***
	(0.2112)	(0.2125)	(0.0287)	(0.0237)
Controls	No	Yes	No	Yes
Qtr F.E.	Yes	Yes	Yes	Yes
ID F.E.	Yes	Yes	Yes	Yes
Subcat F.E.	Yes	Yes	Yes	Yes
N	407900	406918	407900	406918

This table reports difference-in-differences estimates for litigation exposure post-AIA, given by *Exposure* $\times$ *Post* only for small entities. The number of patents are measured by log(number of patents) and citations are measured by the average scaled 2 citations within 2 year of issue of patent. The estimates are derived from the model $Y_{it} = \delta(Exposure \times Post) + X'_{it}\beta + \lambda_i + \lambda_t + \varepsilon_{it}$, where i and t denote entity and quarter respectively, X_{it} denotes a range of patent and entity level controls as defined in Table 4.1, and λ_i, λ_t control for entity and quarter fixed-effects respectively. The entity-quarter level dataset consists of all patents applied between 2008 and 2016. The indicator variable *SE* takes value 1 if more than 50 percent of the patents for an entity was applied as a small entity, and 0 otherwise. *Post* takes value 1 if the patents were applied on or after the first quarter of 2013 i.e. the implementation of first-inventor-to-file rule in the United States.

Exposure is a standardized measure and is defined in Subsection 4.3.3.

Columns (1) and (3) report results without and (2) and (4) report with controls. P-values at 1, 5, and 10 percent are denoted by *, **, and ***. Standard errors are clustered at entity-level.

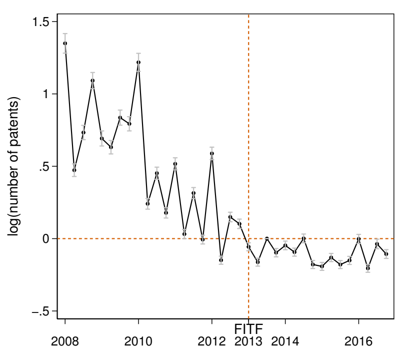

(a) Exposure to litigation: Patents

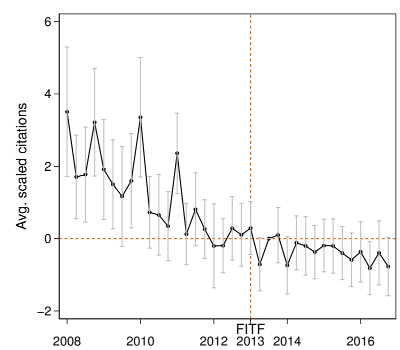

(b) Exposure to litigation: Citations

Figure 6.4. Small entities' exposure to litigation

These figures plot changes in log(number of patents) and average scaled citations for an exposed vs. an under-exposed to litigation only for a small entity by quarter. The estimates are derived from the model $Y_{it} = \beta_0 + \sum_{s \neq 0}(\beta_s \times 1[s = t] \times Exposed) + X_{it}^{'}\beta + \lambda_i + \lambda_t + \varepsilon_{it}$, where i and t denote entity and quarter respectively, X_{it} denotes a range of patent and entity level controls as defined in Table 4.1, and λ_i, λ_t control for entity and quarter fixed-effects respectively. The Appendix Table A.5, Figure A.7b report these results for all entities.

6.4 Evidence from the Compustat firms

Small entities cover about 30 percent of the patents in the main sample. But, among the publicly traded firms, only about one percent of the patents are from small entities. While this sample enables me to control for a greater number of entities' characteristics, such as assets, R&D expenditure, and the number of employees; I inevitably lose all the individual patentees and a large proportion of unlisted firms. Patents by individuals and small firms are important when studying the AIA because of the results from the Canadian reform, as documented by Lerner et al. (2015), Abrams and Wagner (2013), and Lo and Sutthiphisal (2009). The overarching conclusion all the authors report is that small entities, and particularly individual patentees decrease their patenting activity in Canada, and the authors observe an increased gap between large and small entities after an AIA-like reform in Canada. Publicly traded firms do not adequately represent small entities. Table 6.6 reports results from two different models given by equations 5.1 and 5.3 in rows $SE \times Post$ and $Exposure \times Post$ respectively. In this table, I also add a measure of the value of patents derived from the stock market as an alternative measure of patent quality (value) — KPSS value from Kogan et al. (2017).

The results from row $SE \times Post$ show and reiterate that the publicly traded small firms are similar to large firms in terms of their patent strategies, and we do not observe any significant difference between them. They face similar incentives to patent, while the other non-publicly traded small entities, which comprise unlisted small firms and individuals, face a different set of incentives and constraints to patent their invention. From the next row, $Exposure \times Post$, we observe results similar to Table 6.4 i.e. entities patenting in areas where litigation is prevalent reduce their activity post-AIA. An interesting finding is that the drop in patenting by listed firms is about 3.6 percent, while the same coefficient when computed for only small entities is about 37 percent. No change post-AIA can be observed for the KPSS value.

It should be noted that the publicly listed small entities are still sufficiently large to be listed as compared to the unlisted firms and individuals. They also are significantly less budget-constrained than the other small entities. This can be one explanation as to why this sample of small entities does not exhibit a different behavior as compared to larger ones post-AIA.

Table 6.6. Results for listed firms

	Patents		Citations		KPSS	
	(1)	(2)	(3)	(4)	(5)	(6)
SE x Post	-0.0757	-0.0216	1.0795	1.8093	-0.0118	0.0072
	(0.0579)	(0.0594)	(1.7096)	(1.5000)	(0.0889)	(0.0919)
Exposure x Post	-0.0298***	-0.0356***	-0.0356	-0.0566**	-0.0002	-0.0015
	(0.0099)	(0.0124)	(0.0229)	(0.0276)	(0.0012)	(0.0013)
Controls	No	Yes	No	Yes	No	Yes
Qtr F.E.	Yes	Yes	Yes	Yes	Yes	Yes
ID F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Subcat F.E.	Yes	Yes	Yes	Yes	Yes	Yes
N	18406	17512	18406	17512	18406	17512

This table reports the gap in the patents and citations between *listed* small and large firms, and by exposure to litigation post-AIA. The number of patents are measured by log(number of patents), citations are measured by the average scaled 2 citations within 2 year of issue of patent, and exposure is a standardized weighted average of proportion of litigation by NBER subcategory and quarter where the weights are proportion of patents filed by an entity in a subcategory out of all patents filed by the entity in a quarter. For row *SE x Post* the estimates are derived from the model $Y_{it} = \delta(Post \times SE_i) + X^{'}_{it}\beta + \lambda_i + \lambda_t + \varepsilon_{it}$, and for row *Exposure x Post* the estimates are derived from the model $Y_{it} = \delta(Exposure \times Post) + X^{'}_{it}\beta + \lambda_i + \lambda_t + \varepsilon_{it}$ where i and t denote entity and quarter respectively, X_{it} denotes a range of patent and entity level controls as defined in Table 4.1, and λ_i, λ_t control for entity and quarter fixed-effects respectively. The entity-quarter level dataset consists of all patents applied between 2008 and 2016. The indicator variable SE takes value 1 if more than 50 percent of the patents for an entity was applied as a small entity, and 0 otherwise. $Post$ takes value 1 if the patents were applied on or after the first quarter of 2013 i.e. the implementation of first-inventor-to-file rule in the United States. Exposure is standardized and is defined in Subsection 4.3.3.

Columns (1) and (3) report results without and (2) and (4) report with controls. P-values at 1, 5, and 10 percent are denoted by *, **, and ***. Standard errors are clustered at entity-level.

Chapter 7: Concluding remarks

This dissertation studies the quantity and quality of inventions by small entities relative to large entities around a recent and substantial change in the patenting rule of the United States, the Leahy-Smith America Invents Act (AIA) of 2011. The AIA brought forth a range of changes in the patenting rules in the United States to put a check on litigation and provide ease in the filing of patents, especially for smaller entities. However, the AIA's Congressional hearings highlighted that small entities could be disproportionately impacted through the AIA's first-inventor-to-file rule, as they were required to file as quickly as possible upon invention, which adds additional costs to their already constrained resources. Therefore, the AIA's passage required a mandated study to assess the impact of implementing FITF on small entities. Unfortunately, the authors note that the study was premature and warranted further exploration into the questions. Through this dissertation, I contribute the following: one, I shed light on the innovative activities of a relatively understudied group of entities, small entities before and after the AIA's enactment, acting as an update to Lerner et al. (2015); two, I show that small and large entities face different trade-offs when filing patents; and three, this dissertation acts as the first step in understanding the AIA's impacts and addresses some of the legislators' concerns that could be used to design future patent reforms.

I study the impact of the AIA in two parts: first, estimate the total change in patenting activity by small entities before and after the AIA, as measured by the number of patents and their citations; and second, estimate the change in patenting activity for entities exposed to litigation. I find that on average, the gap in patents' quantity and quality between small and large entities was widening even before the enactment of AIA. Entities with greater exposure to litigation were reducing their patenting activity significantly, and among this group, small entities report an even more pronounced gap in patenting. This hints that resource-constrained entities are exercising more caution than larger ones. And among them, especially the entities that operate in areas previously plagued by litigation. If certain entities are more cautious about disclosing their invention than before, we might expect a drop in follow-on inventions in the years to come. The reasons for a cautious move can be many, and a part of the reasons were discussed by Huang et al. (2020) and Abrams and Wagner (2013). Because an FITF rule recognizes the first filer of an invention as the sole inventor, an entity has to file quickly after invention, but also has to ensure that the patent document is as complete as possible. An incomplete document can cause more harm than good. Second, this might also signal an entity's competitors about the portfolio of inventions that they are developing. Hence, an entity is more likely to ensure greater secrecy as long as they are not fully ready to disclose their invention. This behavior is accentuated by the addition of the post-grant review. A post-grant review can question the validity of any granted patent, and therefore, an entity needs to ensure that questions on the validity of their inventions do not arise, or at least are minimized. A drop in follow-on inventions may be deemed as an unintended consequence of the legislation, but this dissertation's scope may not be adequate to provide evidence on the same. There are multiple aspects of the AIA that remain to be studied, and therefore, determining the AIA to be singularly "good" or "bad" may be premature, and such a conclusion, if drawn, must be thoroughly examined, since legislation as complex as the AIA is unworthy of a singular label, and requires examinations through multiple lenses.

This dissertation acts as a first step in analyzing the broadest aspect of the AIA with a focus on small entities. But, a range of questions remain unanswered. We do not know if the AIA indeed resulted in eliminating litigation and "bad patents". While this dissertation suggests that the additional support small and micro entities received was insufficient to counter the increase in resource requirements the AIA invoked, how was the support used by these entities? Also, if the entities patent at a lower rate than before post-AIA, are they also inventing at a lower rate or are they inventing at the same rate but keeping those secret, and later secretly engaging in licensing deals with large manufacturers? In the years to come, an examination of the follow-on patents is required to assess if the AIA's blanket changes in the patenting rules were too harsh and excessively dampened disclosure of the invention or if it hit the sweet spot in dampening the costs of patenting arising from litigation and "bad patents" and simultaneously proliferating invention disclosure through patents.

References

Abrams, D. S., U. Akcigit, G. Oz, and J. G. Pearce (2019). The patent troll: Benign middleman or stick-up artist? Technical report, National Bureau of Economic Research.

Abrams, D. S. and R. P. Wagner (2013). Poisoning the next apple-the america invents act and individual inventors. *Stan. L. Rev. 65*, 517.

Aghion, P., C. Antonin, and S. Bunel (2021). The power of creative destruction. In *The Power of Creative Destruction*. Harvard University Press.

Arora, A., S. Belenzon, M. Marx, and D. Shvadron (2021). (when) does patent protection spur cumulative research within firms? Technical report, National Bureau of Economic Research.

Arora, A., S. Belenzon, and L. Sheer (2021). Matching patents to compustat firms, 1980–2015: Dynamic reassignment, name changes, and ownership structures. *Research Policy 50*(5), 104217.

Arora, A., A. Fosfuri, and T. Roende (2022). Caught in the middle: The bias against startup innovation with technical and commercial challenges. Technical report, National Bureau of Economic Research.

Bessen, J. (2008). The value of us patents by owner and patent characteristics. *Research Policy 37*(5), 932–945.

Cerro, M. (2014). Navigating a post america invents act world: How the leahy-smith america invents act supports small businesses. *J. Nat'l Ass'n Admin. L. Judiciary 34*, 193.

Chen, J. and J. Roth (2022). Log-like? ates defined with zero outcomes are (arbitrarily) scale-dependent. *arXiv preprint arXiv:2212.06080*.

Correia, S. (2015). Singletons, cluster-robust standard errors and fixed effects: A bad mix. *Technical Note, Duke University 7*.

De Rassenfosse, G. (2013). Do firms face a trade-off between the quantity and the quality of their inventions? *Research Policy 42*(5), 1072–1079.

De Rassenfosse, G., P. H. Jensen, A. P. T'Mir Julius, and E. Webster (2021). Is the patent system a level playing field: The effect of patent attorney firms.

De Rassenfosse, G., G. Pellegrino, and E. Raiteri (2020). Do patents enable disclosure? Evidence from the Invention Secrecy Act.

De Rassenfosse, G. and E. Raiteri (2022). Technology protectionism and the patent system: Evidence from china. *The Journal of Industrial Economics 70*(1), 1–43.

Doody, P. A. (2012). Comprehensive legislative history of the leahy-smith america invents act: compiled, edited, and abridged. Technical report, Pillsbury Winthrop Shaw Pittman.

Feng, J. and P. Williams (2023). Abstract patents and innovation: Evidence from Alice v. CLS Bank. *Available at SSRN 4447719*.

Hall, B. H., S. Graham, D. Harhoff, and D. C. Mowery (2004). Prospects for improving us patent quality via postgrant opposition. *Innovation policy and the economy 4*, 115–143.

Hall, B. H., A. Jaffe, and M. Trajtenberg (2005). Market value and patent citations. *RAND Journal of economics*, 16–38.

Hasford, S. (2017). *America Invents Act Primer*. Elsevier.

Hegde, D. and B. Sampat (2009). Examiner citations, applicant citations, and the private value of patents. *Economics Letters 105*(3), 287–289.

Higham, K., G. De Rassenfosse, and A. B. Jaffe (2021). Patent quality: Towards a systematic framework for analysis and measurement. *Research Policy 50*(4), 104215.

Huang, R., L. Li, L. Y. Lu, and H. Wu (2020). The impact of the leahy-smith america invents act on firms' r&d disclosure. *European Accounting Review*, 1–38.

Jaffe, A. B. and G. De Rassenfosse (2019). Patent citation data in social science research: Overview and best practices. *Research handbook on the economics of intellectual property law*.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2018). Measuring technological innovation over the long run. Technical report, National Bureau of Economic Research.

Kiebzak, S., G. Rafert, and C. E. Tucker (2016). The effect of patent litigation and patent assertion entities on entrepreneurial activity. *Research Policy 45*(1), 218–231.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics 132*(2), 665–712.

Lanjouw, J. O. and M. Schankerman (2001). Characteristics of patent litigation: a window on competition. *RAND journal of economics*, 129–151.

Laplume, A. O., E. Xavier-Oliveira, P. Dass, and R. Thakur (2015). The organizational advantage in early inventing and patenting: Empirical evidence from interference proceedings. *Technovation 43*, 40–48.

Lemley, M. A. and S. Zyontz (2021). Does alice target patent trolls? *Journal of Empirical Legal Studies 18*(1), 47–89.

Lerner, J. (2000). The government as venture capitalist: the long-run impact of the sbir program. *The Journal of Private Equity 3*(2), 55–78.

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

Lerner, J., A. Speen, and A. Leamon (2015). The leahy-smith america invents act: A preliminary examination of its impact on small business. *Bella Research Group*.

Lo, S.-t. and D. Sutthiphisal (2009). Does it matter who has the right to patent: first-to-invent or first-to-file? lessons from canada. Technical report, National Bureau of Economic Research.

Marco, A. C. and R. D. Miller (2019). Patent examination quality and litigation: Is there a link? *International Journal of the economics of business 26*(1), 65–91.

Marco, A. C., A. Tesfayesus, and A. Toole (2017). Patent litigation data from us district court electronic records (1963-2015).

Masur, J. S. and L. L. Ouellette (2020). *Patent Law: Cases, Problems, and Materials*. Independently published (October 28, 2020).

Matal, J. (2011a). A guide to the legislative history of the america invents act: Part i of ii. *Fed. Cir. BJ 21*, 435.

Matal, J. (2011b). A guide to the legislative history of the american invents act: Part ii of ii. *Fed. Cir. BJ 21*, 539.

Merkley, K. J. (2014). Narrative disclosure and earnings performance: Evidence from r&d disclosures. *The Accounting Review 89*(2), 725–757.

Mezzanotti, F. (2021). Roadblock to innovation: The role of patent litigation in corporate r&d. *Management Science 67*(12), 7362–7390.

Miyagiwa, K. (2015). The 2011 America Invents Act: does it undermine innovation? *Journal of Economics & Management Strategy 24*(2), 211–227.

Mullahy, J. and E. C. Norton (2022). Why transform y? a critical assessment of dependent-variable transformations in regression models for skewed and sometimes-zero outcomes. Technical report, National Bureau of Economic Research.

Quinn, G. (2011, February). Senate to vote on patent reform, first to file fight looms. IPWatchdog.

Rantanen, J. (2012). Peripheral disclosure. *U. pItt. l. rev. 74*, 1.

S1089-S1114, . C. R. (2011). (statement of Sen. Dianne Feinstein).

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

Trajtenberg, M. (1990). A penny for your quotes: patent citations and the value of innovations. *The Rand journal of economics*, 172–187.

Tucker, C. E. (2013). Patent trolls and technology diffusion.

Vandenburg, E. P. (2013). America invents act: How it affects small businesses. *Idaho L. Rev. 50*, 201.

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

Chapter A: Appendix

A.1 Proof

Proof of Lemma 3.1. Profit in each period is given by:

$$\pi^M = \max_P (a - bq)q - cq = \frac{(a-c)^2}{4b}$$

If patent rights are for n periods;

$$\begin{aligned}
\pi^P &= \pi^M + \frac{\pi^M}{(1+r)} + \frac{\pi^M}{(1+r)^2} + ... + \frac{\pi^M}{(1+r)^n} \\
&= \pi^M \left(1 + \frac{1}{(1+r)} + \frac{1}{(1+r)^2} + ... + \frac{1}{(1+r)^n}\right) \\
&= \pi^M \left(\frac{1+r}{r} - \frac{1}{r(1+r)^{n-1}}\right)
\end{aligned}$$

□

Proof of Lemma 3.2.

$$\begin{aligned}
\pi^{TS} &= (1-p)*[0] + p(1-p)\pi^M + p^2(1-p)\left[\pi^M + \frac{1}{(1+r)}\pi^M\right] \\
&+ p^3(1-p)\left[\pi^M + \frac{1}{(1+r)}\pi^M + \frac{1}{(1+r)^2}\pi^M\right] \\
&+ ... + p^n(1-p)\left[\pi^M + \frac{1}{(1+r)}\pi^M + \frac{1}{(1+r)^2}\pi^M + ... + \frac{1}{(1+r)^{n-1}}\pi^M + ...\right] + ... \\
&= \pi^M(1-p)\Bigg[p + \left(p^2 + \frac{p^2}{(1+r)}\right) + ... \\
&+ \left(p^n + \frac{p^n}{(1+r)} + \frac{p^n}{(1+r)^2} + ...\right) + \frac{p^n}{(1+r)^{n-1}} + ...\Bigg] + ... \\
&= p\,\pi^M(1-p)\Bigg[\left(1 + p + ... + p^{n-1} + ...\right) + ... + \left(\frac{p}{(1+r)} + ... + \frac{p^{n-1}}{(1+r)} + ...\right) + ... \\
&+ \frac{p^{n-1}}{(1+r)^{n-1}} + ...\Bigg] + ... \\
&= p\,\pi^M(1-p)\Big[S_1 + S_2 + ... + S_n + ...\Big]
\end{aligned} \tag{A.1}$$

Now,

$$\begin{aligned}
S_1 &= 1 + p^2 + ... + p^{n-1} + ... = \lim_{n\to\infty}\frac{1-p^n}{1-p} = \frac{1}{1-p} \\
S_2 &= \frac{1}{1+r}\left(p + p^2 + ... + p^{n-1} + ...\right) = \frac{1}{(1-p)}\lim_{n\to\infty}\frac{p-p^n}{(1+r)} = \frac{1}{(1-p)}\frac{p}{(1+r)} \\
&\vdots \\
S_n &= \frac{p^{n-1}}{(1+r)^{n-1}} + \frac{p^n}{(1+r)^{n-1}} + \frac{p^{n+1}}{(1+r)^{n-1}} + ... \\
&= \frac{1}{(1-p)}\lim_{n\to\infty}\frac{p^{n-1}-p^{n+n+1}}{(1+r)^{n-1}} = \frac{1}{(1-p)}\left(\frac{p}{1+r}\right)^{n-1} \\
&\vdots
\end{aligned}$$

$Thus,$

$$S_1 + S_2 + ... + S_n + ... = \frac{1}{(1-p)}\left[1 + \frac{p}{1+r} + \left(\frac{p}{1+r}\right)^2 + ... + ...\right] = \left(\frac{1+r}{1+r-p}\right)$$

$Hence,$

$$\begin{aligned}\pi^{TS} &= p\,\pi^M (1-p)\frac{1}{(1-p)}\left(\frac{1+r}{1+r-p}\right)\\ &= p\,\pi^M\left(\frac{1+r}{1+r-p}\right)\end{aligned}$$

Note: It is assumed that $p < 1 + r$ for the sum to converge. The assumption impies that the probability of secrecy cannot be unity or the rate of interest cannot be zero, which is not unfair to assume. □

Proof of Proposition 3.3. Taking partial derivative of π^P and π^{TS} with respect to r:

$$\begin{aligned}\frac{\partial \pi^P}{\partial r} &= \pi^M\left(\frac{(r+1)^{1-n}}{r^2} - \frac{r+1}{r^2} + \frac{1}{r} + \frac{n-1}{r\,(r+1)^n}\right) < 0\\ \frac{\partial \pi^{TS}}{\partial r} &= p\,\pi^M\left(\frac{1}{1+r-p} - \frac{(1+r)}{(1+r-p)^2}\right) < 0\end{aligned}$$

$$\frac{1}{(1+r-p)} < \frac{(1+r)}{(1+r-p)} \implies \left(\frac{1}{1+r-p} - \frac{(1+r)}{(1+r-p)^2}\right) < 0 \implies \frac{\partial \pi^P}{\partial r} < 0$$

□

Proof of Proposition 3.4. Taking partial derivative of π^{TS} with respect to p,

$$\frac{\partial \pi^{TS}}{\partial p} = \pi^M\frac{(1+r)^2}{(1+r-p)^2}$$

$$\frac{(1+r)^2}{(1+r-p)^2} > 0 \implies \frac{\partial \pi^{TS}}{\partial p} > 0$$

□

Proof of Proposition 3.5. If trade secret is better than patenting, $\pi^{TS} > \pi^{P}$. Simplifying the inequality:

$$\begin{aligned}
p\,\pi^{M}\left(\frac{1+r}{1+r-p}\right) &> \pi^{M}\left(\frac{1+r}{r} - \frac{1}{r(1+r)^{n-1}}\right)\\
p\left(\frac{1+r}{1+r-p}\right) &> \frac{r(1+r)^{n}-r}{r^{2}(1+r)^{n-1}}\\
\frac{1}{\frac{1+r}{p}-1} &> \frac{(1+r)^{n}-1}{r(1+r)^{n}}\\
\frac{1+r}{p}-1 &< \frac{r(1+r)^{n}}{(1+r)^{n}-1}\\
\frac{1+r}{p} &< \frac{r(1+r)^{n}}{(1+r)^{n}-1}+1\\
\frac{1}{p} &< \frac{r(1+r)^{n-1}}{(1+r)^{n}-1}+\frac{1}{1+r}\\
p &> \frac{(1+r)^{n+1}-(1+r)}{r(1+r)^{n}+(1+r)^{n}-1}\\
p &> \frac{(1+r)^{n+1}-(1+r)}{(1+r)^{n+1}-1}\\
p &> \frac{(1+r)^{n+1}-1}{(1+r)^{n+1}-1} - \frac{r}{(1+r)^{n+1}-1}\\
p &> 1 - \frac{r}{(1+r)^{n+1}-1}
\end{aligned}$$

□

A.2 Patent applications without assignee IDs

The Subsection 4.2 describes the sample selection this dissertation. In the second paragraph, I end up with 5,087,133 patent applications, with or without entity identifiers. Among these, 1,832,053 did not have either assignee ID or inventor ID from the Patentsview database. In this section of the Appendix, I argue that the patent application without the IDs is not a systematic error of the disambiguation algorithm. Rather, most of these patent applications are a derivative of another patent application already considered in the main sample. 87 percent of the patent applications which do not have an ID are either connected to the patent applications in the main sample through a parent or a child application. Among the 87 percent, 99.2 percent are either PCT or provisional applications. These applications are not examined if they are not converted into a non-provisional application within a given time. They either end up being abandoned or are marked as "pending" throughout their life in the USPTO patent database. A summary of the types of patent applications among those who do not have an ID is provided in Table A.1. Each row reports an application type, and if they are connected to the main sample through parent or child applications. The main concern here is the utility patents, which amount to 13,203 patent applications. Utility patents may end up being examined, but the other patents will not, and therefore do not pose a threat to the main results of the dissertation. The proportion of utility patents out of the total missing is minuscule and will not disturb the estimates even if they were in the main sample.

Table A.1. Patent applications with missing IDs

Patent application type	Connected	Not-connected	Total
Utility	9,080	4,123	13,203
PCT	901,286	209,622	1,110,908
Provisional	679,304	25,644	704,948
Re-issue	2,146	24	2,170
Re-examination	815	7	822
Missing	1	1	2
Total	1,592,632	239,421	1,832,053

A.3 Results from a balanced panel

In Subsection 4.2, I briefly mention a concern regarding selection that may arise from the use of an unbalanced panel. While the sample is intrinsically not unbalanced, because I do observe all the patents each entity files for each quarter and there are no missing observations for any particular quarter for a given entity; in a definitional sense of an unbalanced panel, the main sample *is* unbalanced.

Different entities may find it favorable to choose between the pre and post-AIA periods to file their patent application which may relate to their objectives and characteristics. If the entities in the groups small and large for the before and after periods are vastly different, estimates showing the change in their patenting activity before and after the AIA may also contain bias. One way to tackle the issue would be to control for enough of the varying entity characteristics which explain their choice between the two periods if any. If we assume that the control variables adequately capture their strategies, the estimates would be consistent. This is one reason why I estimate the model with different samples and variables.

Another way to tackle this is to force the unbalanced panel to be balanced. In the quarters when an entity did not file for patent applications, I put zeroes in the number of patents and citations column. The problem with arranging the data in this way is that it adds zeroes to the time period when an entity was not established. This changes the effect size, and because small entities on average appear at significantly lower rates than larger ones, their representation is negatively skewed.

Table A.2 reports results when a balanced panel is forced. We observe that small entities still file for patents at a lower rate than larger ones. Their citations per patent though reported to be significantly higher, is still near zero and not economically significant.

Table A.2. Results from a balanced panel

	Patents		Citations	
	(1)	(2)	(3)	(4)
SE x Post	0.4876***	-0.1664***	0.2382**	-0.1958*
	(0.0832)	(0.0395)	(0.0951)	(0.1052)
Individual	-5.1391***		-2.7560***	
	(0.0415)		(0.0475)	
Government	5.3494***		0.0816	
	(0.7196)		(0.8222)	
Constant	6.3705***	2.4955***	3.1808***	1.1168***
	(0.0361)	(0.0101)	(0.0412)	(0.0268)
Qtr F.E.	Yes	Yes	Yes	Yes
ID F.E.	No	Yes	No	Yes
N	1147959	1148250	1147959	1148250

This table reports results from a balanced entity-quarter panel. Entities may or may not file for patent applications in all quarters. I add zeroes to the quarters where the entity did not file for patents, as opposed to the sample in Table 6.1 which does not have zeroes. The later method is prevalent in the literature. Rest of the descriptions of this table is similar to Table 6.1.

Columns (1) and (3) report results without and (2) and (4) report with controls. P-values at 1, 5, and 10 percent are denoted by *, **, and ***. Standard errors are clustered at entity-level.

A.4 Results from Poisson regression

Table A.3. Results from Poisson model

	Patents		Citations	
	(1)	(2)	(3)	(4)
SE x Post	-0.1117***	-0.0436**	-0.4202***	-0.3433***
	(0.0221)	(0.0195)	(0.1043)	(0.1173)
Controls	No	Yes	No	Yes
Qtr F.E.	Yes	Yes	Yes	Yes
ID F.E.	Yes	Yes	Yes	Yes
Subcat F.E.	Yes	Yes	Yes	Yes
N	1148250	1146957	384502	384147

This table reports results from Poisson regression model for the outcomes number of patents, and number of citations' gaps between small and large entities post-AIA. The estimates are derived from the model $Y_{it} = \delta(Post \times SE_i) + X_{it}^{'}\beta + \lambda_i + \lambda_t + \varepsilon_{it}$, where i and t denote entity and quarter respectively, X_{it} denotes a range of patent and entity level controls as defined in Table 4.1, and λ_i, λ_t control for entity and quarter fixed-effects respectively. The entity-quarter level dataset consists of all patents applied between 2008 and 2016. The indicator variable *SE* takes value 1 if more than 50 percent of the patents for an entity was applied as a small entity, and 0 otherwise. *Post* takes value 1 if the patents were applied on or after the first quarter of 2013 i.e. the implementation of first-inventor-to-file rule in the United States.

Columns (1) and (3) report results without and (2) and (4) report with controls. P-values at 1, 5, and 10 percent are denoted by *, **, and ***. Standard errors are clustered at entity-level.

A.5 Alternate definition of small entities

An entity can file as a small or undiscounted entity. The reason for this choice is unclear. The main set of results assumes an entity to be small if it is ever claimed to be small. But, it is possible that an entity grew over time to be large and be misrepresented as a small entity because of the assumption. Therefore, Table A.4 reports results when the variable small entity is not an indicator variable. Rather, it is the proportion of times an entity claimed to be small out of total patents filed for that quarter. This value ranges between 0 and 1 and is a continuous measure of an entity being small and large in each quarter.

I re-estimate Equation 5.1 considering the proportion of small entity for each quarter in place of the indicator variable small entity. The coefficients of SE $\times$ Post do not qualitatively change compared to Table 6.1.

Table A.4. Results from an alternate definition of small entities

	Patents		Citations	
	(1)	(2)	(3)	(4)
SE x Post	-0.0730***	-0.0416***	-0.1835***	-0.2760***
	(0.0047)	(0.0034)	(0.0640)	(0.0635)
Controls	No	Yes	No	Yes
Qtr F.E.	Yes	Yes	Yes	Yes
ID F.E.	Yes	Yes	Yes	Yes
Subcat F.E.	Yes	Yes	Yes	Yes
N	1148250	1146957	1148250	1146957

This table reports results from the model similar to the main results, as reported in Table 6.1 but with an alternative definition of small entities. Here *SE* ranges from 0 to 1, and the proportion is calculated as the number of patents filed as a small entity over the total number of patents the entity is involved in inventing. For simplicity, the definition used in the Table 6.1 *SE* takes value 1 if more than 50 percent of the patents for an entity was applied as a small entity, and 0 otherwise. All other parts of the model is same as in Table 6.1.

Columns (1) and (3) report results without and (2) and (4) report with controls. P-values at 1, 5, and 10 percent are denoted by *, **, and ***. Standard errors are clustered at entity-level.

A.6 Quarter fixed-effects from the main table

Figures A.1a and A.1b plot the quarter fixed-effects from Table 6.1.

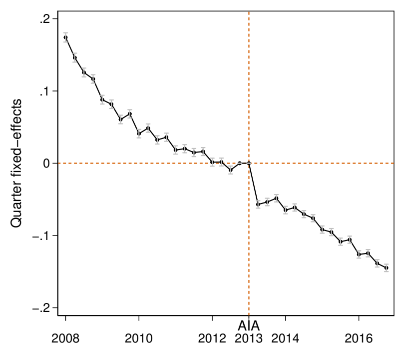

(a) Number of patents' quarter fixed-effects

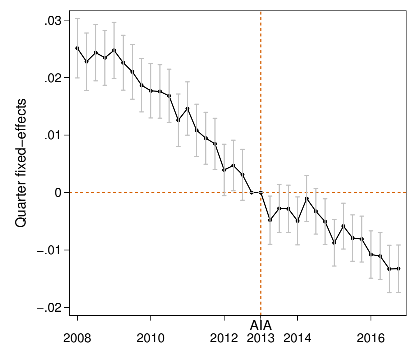

(b) Number of citations' quarter fixed-effects

Figure A.1. Quarter fixed effects from Table 6.1

A.7 Number of citations within one year of issue

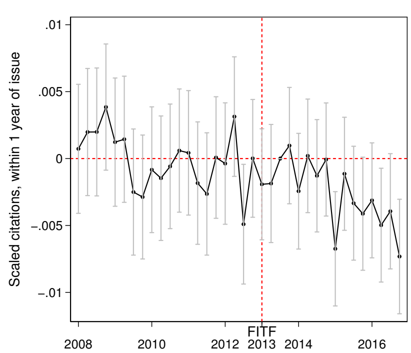

(a) Change in small entities' 1 yr. scaled citations over time

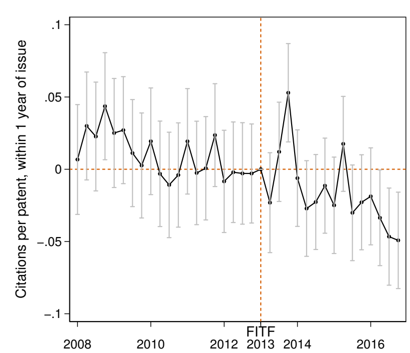

(b) Change in small entities' number of 1 yr. citations per patent over time

Figure A.2. Citations within a year of issue, scaled and absolute measure

A.8 Results from applicants' and examiners' citation additions

The Figures A.3 report results by separating examiner and applicant added citations. In the main dissertation, I use total citations and convert it into percentiles for each NBER subclass and year combination. While the examiner and applicant added citations are not different from zero, before and after the AIA, an interesting finding is an increase in applicant added citations, while a drop in examiner added citations, and the estimates are statistically significant towards the later part of the post-AIA period. The effects are prominent after 2014, which may indicate an effect that could be attributed to *Alice*. Examiner added citations are stronger predictor of patent value compared to the applicant added citations (Hegde and Sampat, 2009). Since *Alice* questions the validity of abstract patents, entities are citing granted patents at a greater rate to prove their own patents' validity.

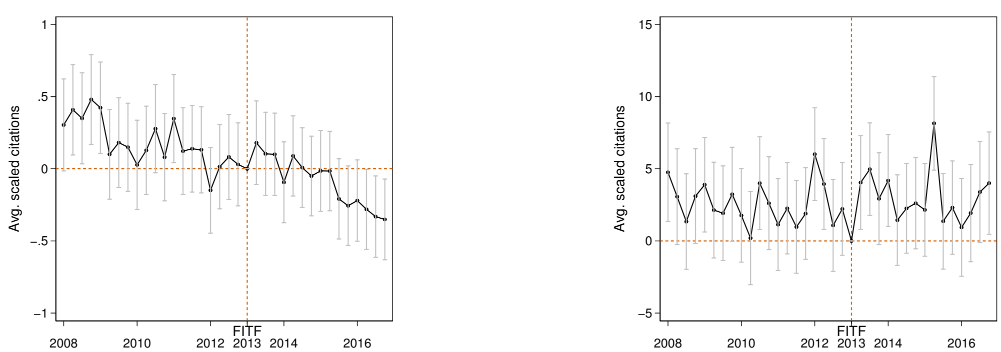

(a) Change in small entities' scaled citations (examiner added)

(b) Change in small entities' scaled citations (applicant added)

Figure A.3. Citations, examiner and applicant added

A.9 Patentees' subsample

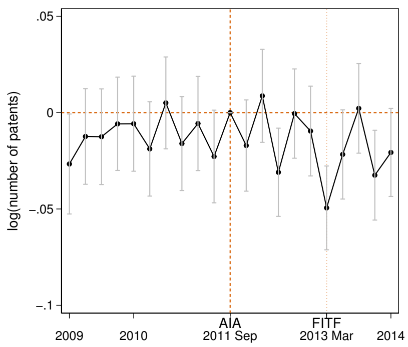

(a) [0-75) percentile

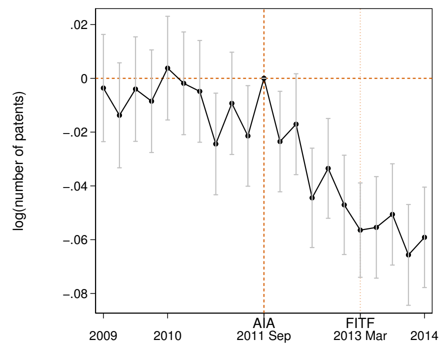

(b) [0-99) percentile

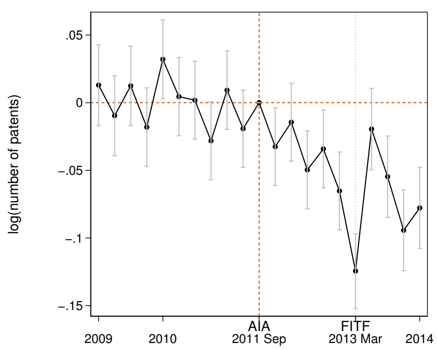

(c) [75-100] percentile

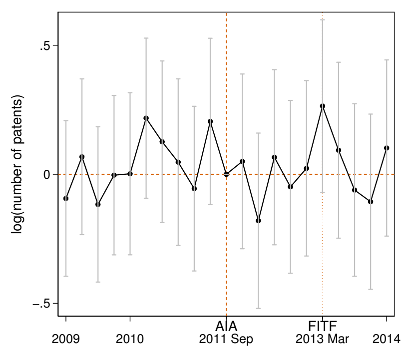

(d) [99-100] percentile

Figure A.4. Subsamples by light and heavy patentees

A.10 Results considering from number of claims as outcome

Figure A.5 reports changes in log(number of claims per patent) post-AIA for small entities. When we contrast and compare Figure 6.3a with Figure A.5, we observe similar outcomes, which shows that not only did the number of patents drop post-AIA, they also became shorter in length, scope, and its use. This requires a separate study, and therefore is not part of the main dissertation.

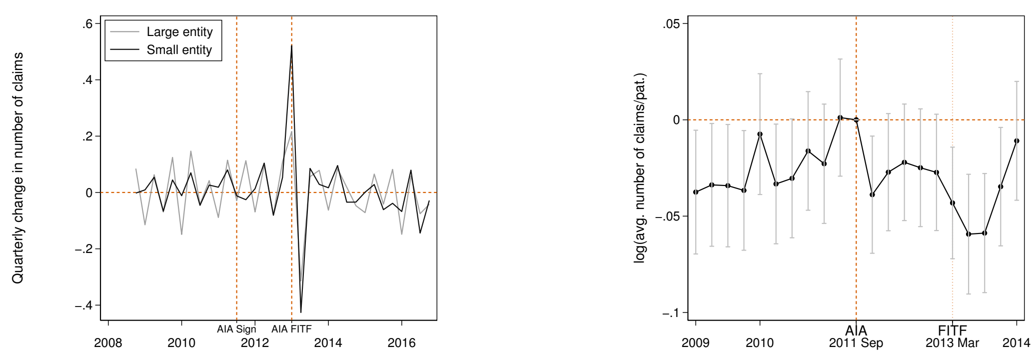

(a) Claims, Quarter-on-quarter

(b) Change in log(average claims per patent)

Figure A.5. Change in small entities' claims

Figure (a) reports the change in number of claims quarter-on-quarter for small and large entities. Figure (b) plots the change in log(number of claims) by small entities between 2009 and 2014. The base quarter is the AIA's signing quarter (September 2011). The event study specification is reported in Equation 5.2, and the figure plots the β_s's estimated from $\sum_{s\neq 0} \beta_s \times 1[s = t] \times SE_i$ part of the equation.

A.11 Cases filed per month

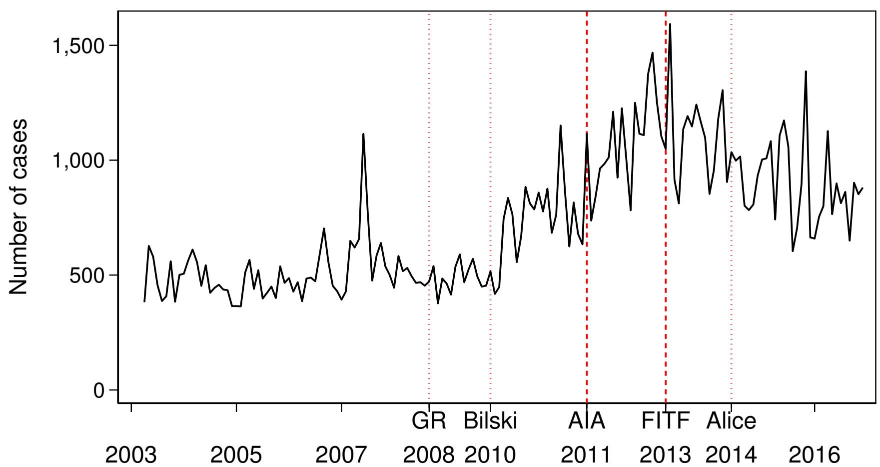

Figure A.6. Court cases filed each month
This figure plots total cases filed at the District Courts each month from 2003 through 2017. The important Supreme Court cases are marked: *Bilski v. Kappos* (Nov 2009), and *Alice v. CLS Bank* (June 2014)

A.12 Litigation exposure for all entities

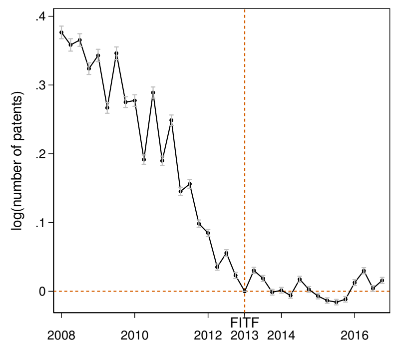

(a) Exposure to litigation: Patents

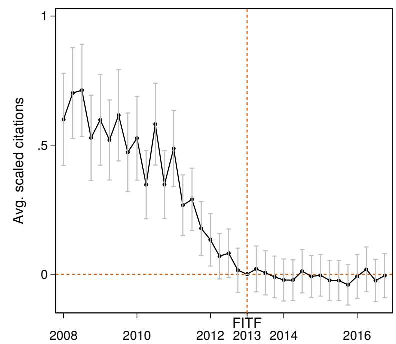

(b) Exposure to litigation: Citations

Figure A.7. Exposure to litigation: all entities

Table A.5. DDD estimates for exposure to litigation

	Patents		Citations	
	(1)	(2)	(3)	(4)
SE x Exposure x Post	-0.2340***	-0.2673***	-0.0417***	-0.0333***
	(0.0068)	(0.0049)	(0.0027)	(0.0027)
SE x Exposure	0.8165***	0.7966***	0.1092***	0.0915***
	(0.0066)	(0.0048)	(0.0027)	(0.0026)
Exposure x Post	-0.0971***	-0.0945***	-0.0074***	-0.0072***
	(0.0011)	(0.0008)	(0.0004)	(0.0004)
SE x Post	-0.1189***	-0.1048***	-0.0179***	-0.0200***
	(0.0030)	(0.0022)	(0.0012)	(0.0012)
Exposure	0.2481***	0.2327***	0.0164***	0.0158***
	(0.0015)	(0.0011)	(0.0006)	(0.0006)
Controls	No	Yes	No	Yes
Qtr F.E.	Yes	Yes	Yes	Yes
ID F.E.	Yes	Yes	Yes	Yes
Subcat F.E.	Yes	Yes	Yes	Yes
N	1148250	1146957	1148250	1146957

This table reports triple difference estimates for small entities — exposed to litigation — post-AIA, given by $SE \times Exposure \times Post$. The number of patents are measured by log(number of patents) and citations are measured by the average scaled citations within 2 year of issue of patent. The estimates are derived from the model $Y_{it} = \delta_1(SE_i \times Exposure \times Post) + \delta_2(SE_i \times Exposure) + \delta_3(Exposure \times Post) + \delta_1(SE_i \times Post) + X_{it}^{'}\beta + \lambda_i + \lambda_t + \varepsilon_{it}$, where i and t denote entity and quarter respectively, X_{it} denotes a range of patent and entity level controls as defined in Table 4.1, and λ_i, λ_t control for entity and quarter fixed-effects respectively. The entity-quarter level dataset consists of all patents applied between 2008 and 2016. The indicator variable SE takes value 1 if more than 50 percent of the patents for an entity was applied as a small entity, and 0 otherwise. *Post* takes value 1 if the patents were applied on or after the first quarter of 2013 i.e. the implementation of first-inventor-to-file rule in the United States.

Exposure is standardized and is defined in Subsection 4.3.3.

Columns (1) and (3) report results without and (2) and (4) report with controls. P-values at 1, 5, and 10 percent are denoted by *, **, and ***. Standard errors are clustered at entity-level.