The Effect of the Leahy-Smith America Invents Act on Knowledge Transfers Through Inventor Mobility

by

İlber Atay Aslan

A Dissertation Submitted to the
Graduate School of Sciences and Humanities
in Partial Fulfillment of the Requirements for
the Degree of

Master of Arts

in

Economics

KOÇ ÜNİVERSİTESİ

June 20, 2025

The Effect of the Leahy-Smith America Invents Act on Knowledge Transfers Through Inventor Mobility

Koç University

Graduate School of Social Sciences and Humanities

This is to certify that I have examined this copy of a master's thesis by

İlber Atay Aslan

and have found that it is complete and satisfactory in all respects,

and that any and all revisions required by the final

examining committee have been made.

Committee Members:

____

Assoc. Prof. Emre Ekinci (Advisor)

____

Asst. Prof. Mustafa Kaba

____

Prof. Dr. Şerif Aziz Şimşir

Date: ______________________________

This thesis is dedicated to my first and most beloved teachers, my father and my mother.

ABSTRACT

The Effect of the Leahy-Smith America Invents Act on Knowledge Transfers Through Inventor Mobility
İlber Atay Aslan
Master of Arts in Economics
June 20, 2025

On September 16, 2011, President Barack Obama signed the Leahy-Smith America Invents Act (AIA) into law, replacing the century-old first-to-invent patent awarding system with a first-to-file rule. While prior research has documented how this reform altered firms' R&D intensity choices and patenting behavior, its impact on the labor market for inventors remains largely unexplored. This thesis addresses that gap by focusing on mobile inventors (inventors who switch firms) and examines how their mobility facilitates knowledge transfer to the hiring firm.

After identifying mobile inventors, I employ the Coarsened Exact Matching (CEM) method to construct a control sample of non-mobile inventors who have not switched firms but are similar to mobile inventors in pre-mobility characteristics. Then, using a difference-in-differences (DiD) approach, I show that hiring firms increase their use of new recruits' pre-move inventions after employing them. Moreover, the impact of inventor mobility extends beyond direct use: the tacit knowledge of the new hire diffuses throughout the hiring firm and becomes part of its broader knowledge base. Finally, I examine how the AIA affects knowledge transfers through inventor mobility. Applying a novel difference-in-differences-in-differences (DiDiD) design, I find that the AIA significantly reduces the diffusion of tacit knowledge from new hires, an unintended consequence that works against the reform's objectives.

ÖZETÇE

Leahy-Smith Amerika Buluşlar Yasası'nın Mucit Hareketliliği Yoluyla Gerçekleşen Bilgi Transferlerine Etkisi

İlber Atay Aslan
Ekonomi, Yüksek Lisans

16 Eylül 2011 tarihinde Başkan Barack Obama, Leahy–Smith Amerika Buluşlar Yasası ile ABD'de yüz yılı aşkın süredir uygulanan "ilk buluş yapan" (first-to-invent) kuralını kaldırmış ve yerine "ilk başvuran" (first-to-file) kuralını getirmiştir. Literatürdeki çalışmalar, bu reformun firmaların Ar-Ge yoğunluğu tercihleri ve patentleme davranışları üzerindeki etkilerini belgelemiş olsa da yasa değişikliğinin mucitler için iş gücü piyasası üzerindeki etkisi henüz incelenmemiştir. Bu tez, firma değiştiren mucitlere odaklanarak ve mucitlerin firmalar arası hareketliliği yoluyla işveren firmaya doğru gerçekleşen bilgi transferlerini inceleyerek literatürdeki ilgili boşluğu doldurmaktadır.

Firma değiştiren mucitler belirlendikten sonra, hareketlilik öncesi gözlemlenebilir özellikler bakımından benzer olan ancak firma değiştirmemiş mucitlerden oluşan bir kontrol grubu, Coarsened Exact Matching (CEM) adlı eşleştirme yöntemi kullanılarak oluşturulmuştur. Ardından, Farkların Farkı (Difference-in-Differences) modeliyle, işveren firmaların kendi patentlerinde yeni çalışanlarının önceki buluşlarını, onları işe aldıktan sonra daha fazla kullandıkları gösterilmiştir. Dahası, mucit hareketliliğinin etkisi bilginin doğrudan kullanımının ötesine geçmektedir: Yeni çalışanın örtük bilgisi işveren firmaya yayılmakta ve firmanın bilgi setinin bir parçası haline gelmektedir. Son olarak, Leahy–Smith Amerika Buluşar Yasası'nın mucit hareketliliği yoluyla gerçekleşen bilgi transferlerini nasıl etkilediği incelenmiştir. Farkların Farklarının Farkı (Difference-in-Difference-in-Differences) yöntemi ile Leahy–Smith Amerika Buluşlar Yasası'nın, yeni çalışanların örtük bilgisinin işveren firma içindeki yayılımını önemli ölçüde azalttığı gözlemlenmiş ve böylece reformun, amaçlananın aksine işleyen bir sonucu belgelenmiştir.

ACKNOWLEDGMENTS

I would like to express my heartfelt gratitude to my advisor, Assoc. Prof. Emre Ekinci, whose expertise, insight, support, and patience have been indispensable throughout my research. I feel privileged to be his student and will carry what I have learned from him with me for a lifetime.

I am also deeply grateful to my beloved cohort — Duygu İneci, Muhammed İkbal Oruç and Nurettin Baver Öz — for their invaluable friendship and unwavering support. This journey would have been far more difficult and far less rewarding without their presence. They truly enriched my graduate school experience, and I feel privileged to have met them.

Finally, this thesis would not have been possible without the love, encouragement and unconditional support of my mother and father.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABBREVIATIONS

AIA	America Invents Act
FTI	First-to-Invent
FTF	First-to-File
DiD	Difference-in-Differences
DiDiD	Difference-in-Difference-in-Differences
CEM	Coarsened Exact Matching
USPTO	United States Patent and Trademark Office
CIPO	Canadian Intellectual Property Office
R&D	Research and Development
IPC	International Patent Classification
ATE	Average Treatment Effect
ATET	Average Treatment Effect on the Treated

Chapter 1

INTRODUCTION

February 14, 1876, was no ordinary Valentine's Day, but rather the beginning of one of the most famous patent disputes in history. That morning, a representative for Alexander Graham Bell, a 29-year-old amateur inventor, submitted a patent application to the U.S. Patent Office in Washington for a device designed to transmit vocal sounds through electrical wires (Hounshell, 1981). Just hours later, Elisha Gray, a 41-year-old professional inventor and co-founder of the Western Electric Manufacturing Company, submitted a patent caveat to the same office for an apparatus serving precisely the same purpose. This caveat was intended to give the Patent Office formal notice and was designed to protect the idea for one year, with a full patent application required within that time (Brown, 2020).

It soon became apparent that this "talking telegraph caveat" might interfere with Bell's patent application. Gray's patent attorney even wrote to Dr. Samuel S. White, Gray's financial backer, suggesting that an interference proceeding could still be initiated if Gray promptly filed a full application, but added that his judgment was against pursuing this course of action. (Brown, 2020). Ultimately, Gray followed his attorneys' advice and decided not to pursue the matter. As a result, on March 7, Alexander Graham Bell was granted Patent No. 174,456, titled Improvement in Telegraphy - the invention we now recognize as the telephone (Boucher, 2012).

Would it have mattered if Gray's attorney had arrived at the Patent Office first on that Valentine's Day? It's hard to say, but in principle, it shouldn't have. Since its early days, the U.S. patent system operated under a first-to-invent (FTI) rule. Under this rule, when a dispute over patent priority arose, the U.S. Patent and Trademark Office (USPTO) would administer interference proceedings between applicants, evaluate various forms of evidence to establish the date of invention, and grant the patent to the party who first conceived the inventive idea (Lemley and Chien, 2003; Abrams and Wagner, 2013). Therefore, winning the race to the Patent Office was not supposed to determine the outcome under the first-to-invent rule. If Bell was indeed the first true inventor of the telephone, his attorneys' earlier arrival at the Patent Office was ultimately insignificant.

The United States remained a maverick in the intellectual property world for more than a century, standing as one of the few nations to uphold the first-to-invent rule [1]. On September 16, 2011, however, President Barack Obama signed the Leahy-Smith America Invents Act (AIA) into law, replacing the first-to-invent (FTI) system with a first-to-file (FTF) regime. Under this new system, when multiple patent applications conflict, the U.S. Patent and Trademark Office (USPTO) awards the patent to the applicant who filed an application before its competitors, regardless of which party conceived the inventive idea first. Therefore, with the transition to the first-to-file (FTF) rule, now it matters who arrives first at the Patent Office.

This thesis examines how the transition from the first-to-invent to the first-to-file rule has affected knowledge transfer through inventor mobility. In particular, I first hypothesize that inventor mobility—the movement of an inventor from one firm to another—facilitates knowledge transfer from the new hire to the employing firm. To test this claim, I employ an inventor-firm-level difference-in-differences (DiD) framework and find that hiring firms increase their use of the new recruits' pre-move inventions after employment. Moreover, I document that the impact of inventor mobility extends beyond direct use: the tacit knowledge of the new hire diffuses throughout the hiring firm and becomes part of its broader knowledge base. Second, I argue that the extent to which inventor mobility facilitates knowledge transfer declines with the enactment of AIA. Applying a novel difference-in-differences-in-differences (DiDiD) design, I show that the AIA significantly reduces the diffusion of tacit knowledge from new hires, a dark side not anticipated by the policy's architects.

My work contributes to the existing literature in several ways. First, it provides the first examination of the AIA's impact on the inventor labor market and its effect on the

[1] When Canada and the Philippines adopted the first-to-file (FTF) system in 1989 and in 1998 respectively, the U.S. was left as the last country to maintain the first-to-invent rule.

transferability of human capital through inventor mobility. This contribution is particularly important: given that a substantial share of inventive ideas originate outside the firm (Mueller, 1962), and that inventor mobility is a key mechanism through which firms access distant knowledge and novel ideas (Rosenkopf and Almeida, 2003), the detrimental effect of the AIA on knowledge transfers through inventor mobility may potentially slow aggregate inventive activity. Second, this thesis is the first empirical study to exploit variation at the inventor-firm level to examine the role of employee turnover in knowledge transfers, a methodological and a conceptual novelty I discuss in detail.

The remainder of this thesis is structured as follows. Chapter 2 reviews the related literature. Section 2.1 provides an overview of theoretical and empirical studies on the effects of the Leahy-Smith America Invents Act (AIA) on the innovation landscape, while Section 2.2 discusses prior evidence on the learning-by-hiring hypothesis. Chapter 3 introduces the theoretical framework and presents the two main hypotheses of the thesis. Section 3.1 details the Learning-by-Hiring Hypothesis, discusses my contributions to testing it, and highlights how this study improves upon previous approaches. Section 3.2 presents the AIA Effect Hypothesis along with the underlying mechanisms. Chapter 4 describes the data and matching methodology used in the empirical analysis. It explains how inventor mobility is identified, how the control group is constructed using Coarsened Exact Matching (CEM), and provides descriptive statistics. Chapter 5 outlines the empirical strategy. Section 5.1 introduces the Difference-in-Differences (DiD) approach used to test the first hypothesis. Section 5.2 presents the Difference-in-Difference-in-Differences (DiDiD) design used to test the second hypothesis, and Section 5.3 discusses identification and potential threats to validity. Chapter 6 presents the estimation results. Section 6.1 reports the results for the DiD specifications and the associated event study. Section 6.2 presents the results for the DiDiD specifications and the corresponding event study. Finally, Chapter 7 summarizes the main findings and concludes.

Chapter 2

LITERATURE REVIEW

In this chapter, I provide a detailed overview of the related literature. Section 2.1 reviews previous work on the first-to-invent and first-to-file rules, with theoretical and empirical findings discussed separately in subsections 2.1.1 and 2.1.2, respectively. In section 2.2, I turn to the literature on the impact of inventor mobility on knowledge transfer, commonly referred to as the learning-by-hiring phenomenon.

2.1 First-to-File vs First-to-Invent

2.1.1 Theory

Under the 1952 Patent Act, a patent applicant could obtain exclusive rights over their invention for 20 years from the filing date only if they make a full disclosure regarding the invention, such that any person skilled in the art would be able to reproduce it successfully with the aid of the patent text. In this regard, patents can be viewed as information-sharing devices that disclose information on newly developed technical knowledge to the public (Rantanen, 2012). This function of patents has strong implications on the aggregate innovative activity: As inventors disclose their technical advancements through patenting, the shared knowledge would help other innovators in their own research, reduce redundancy, and hasten the time to follow-on innovation (Scotchmer and Green, 1990). In line with this argument, many advocated for the transition to a first-to-file system on the grounds that it encourages earlier patenting and therefore, earlier disclosure (Lo and Sutthiphisal, 2009).

An early theoretical work that formulates this idea is Scotchmer and Green (1990). In this pioneering study, the authors employ a dynamic model of a patent race for an innovation consisting of three sequential and patentable stages. They show that the incentive to file a patent application and disclose interim technologies is greater under the first-to-file rule, whereas inventors of early-stage innovations are more likely to suppress patenting under the first-to-invent rule. The main mechanism is that under the first-to-invent rule, the leader faces no risk of losing ownership of the interim technology by delaying filing, as they can still establish priority even if a competitor files a patent application. As a result, the leader is more likely to postpone filing a patent application to preserve a competitive advantage in the race. This, in turn, implies that the first-to-file rule is more conducive to innovative activity, as it encourages greater disclosure of technical information, which helps other innovators and accelerates follow-on research on the final technology.

One caveat, however, is that Scotchmer and Green (1990) assume the R&D intensities—and therefore the discovery rates—of the competing firms to be constant. This assumption is inconsistent with much of the patent race literature, which models discovery rates as a function of R&D intensity levels chosen by profit-maximizing firms. To address this, Miyagiwa (2015) and Miyagiwa and Ohno (2015) employ a patent race model similar to that of Scotchmer and Green (1990), but treat R&D intensity levels as endogenous variables. Miyagiwa (2015) shows that once a firm discovers an interim technology, its incentive to conduct further research and improve the invention is weaker under the first-to-file rule, as the firm risks losing the patent to an earlier applicant if it does not file immediately. The author then analytically demonstrates that such a "file immediately" strategy is associated with a lower initial R&D intensity choice by the firm. Similarly, Miyagiwa and Ohno (2015) demonstrate that, under the first-to-invent rule, discovering the interim technology and leading the patent race is less profitable in expectation compared to the first-to-file rule. This is because, if the leader delays filing, the first-to-file rule exposes them to the additional risk of losing the patent for the interim technology to a competitor. Then, similar to Miyagiwa (2015), authors analytically show that when the expected profit from leading a patent race is lower, the firm chooses a lower R&D intensity in the first place.

In addition to the theoretical debate over which rule better incentivizes firms to conduct R&D, there has been no clear consensus on whether first-to-file or first-to-invent leads to more disclosure. While Scotchmer and Green (1990) argue that first-to-file expedites patent applications and therefore results in more disclosure, this mechanism only pertains to the disclosure of technical knowledge after a firm discovers a new technology and files a patent application. Acknowledging this, Halac et al. (2017) argue that first-to-invent does not restrict disclosure but instead encourages firms to share information prior to public disclosure, as they know they will still be rewarded for their inventions even if someone else files for a patent first. In contrast, under the first-to-invent rule, any disclosure of IP-related information before a firm obtains a patent could accelerate a competitor's filing of a conflicting application, potentially resulting in the original firm losing its patent rights (Huang et al., 2021). Therefore, the first-to-invent rule may be detrimental to disclosure prior to patent filing—a type of disclosure that is particularly important, as it helps minimize duplication of inventive efforts among firms (Yilmaz et al., 2024).

2.1.2 Empirical Studies

The United States was not entirely alone in using the first-to-invent rule in patent protection. Until the Canadian Patent Act reforms took effect in 1989, the United States' northern neighbor also operated under a first-to-invent regime. With these reforms Canada aligned its patent system with those of its major trading partners in Europe and Japan by adopting a first-to-file regime, leaving the United States as the last major holdout of the first-to-invent practice.[1] This paradigm shift in Canadian Patent Law, while the U.S. preserved the first-to-invent rule, created a quasi-natural experimental setting for researchers aiming to evaluate the consequences of the AIA on the U.S. innovation landscape.

One such study is Lo and Sutthiphisal (2009) which investigated how the shift from a first-to-invent to a first-to-file system in Canada affected firms' R&D investment and innovative output. Exploiting the 1989 Canadian reform as a natural experiment, the authors employed a difference-in-differences approach that compares Canadian firms—subject to the policy change—with U.S. firms that remained under the first-to-invent regime until 2011. They documented that, while the transition to the first-to-file rule had no significant effect on overall R&D spending or total patenting activity, it led to a decline in patent grants to independent inventors and small businesses. Abrams and Wagner (2013), another

[1]Until the patent law reforms in 1998, the Philippines also operated under a first-to-invent system. Thus, Canada's transition in 1989 is considered the last by a major industrialized nation to shift to the first-to-file system prior to the U.S. doing so.

pioneering empirical study on the AIA, further explored the potential distributional effects of the transition to a first-to-file system. Using a difference-in-differences approach similar to that of Lo and Sutthiphisal (2009), they documented that the 1989 Canadian reform reduced the share of patents granted by the Canadian Intellectual Property Office (CIPO) to Canadian individual inventors. Similarly Lerner et al. (2015) employed a difference-in-difference-in-differences (DiDiD) regression to compare patent grants for small and large firms across Canada and the U.K., before and after the 1989 reform. Authors documented that Canadian firms that patented more frequently prior to the reform—typically large firms—experienced a relative increase in patenting activity after the reform compared to smaller firms.

In addition to the detrimental effects on the inventive activity of small firms and independent inventors, several empirical studies have documented further potential harms of the policy. Using a difference-in-differences design, Huang et al. (2021) show that the enactment of the AIA is associated with a greater decline in narrative R&D disclosure among innovation-intensive firms relative to their less innovation-intensive counterparts[2]. This suggests that the policy have adversely affected the disclosure of R&D-related information prior to public disclosure, heightening the information asymmetry between innovative firms. The result presented by Huang et al. (2021) is intuitive: under the first-to-file system, firms are less likely to disclose information about their R&D activities, as doing so could benefit rivals and increase the risk of losing patent rights. In contrast, under the first-to-invent system, this risk is mitigated, since a firm that can establish priority retains the right to the patent even if a competitor files first.

A recent study Yilmaz et al. (2024) investigates the impact of the AIA on M&A markets and innovation acquisition. The authors document that the AIA led to heightened activity in the innovation acquisition market, evidenced by a significant increase in both the number of patents acquired by acquiring firms and the bid premia paid for innovative targets. Yilmaz et al. (2024) further show that this surge in demand for innovative targets coincides with a decline in R&D investments by the acquirers. This suggests that the AIA

[2]Compared with control firms, innovation-intensive firms in U.S., on average, disclosed 8% less R&D narrative information, 6% less R&D numerical information, and 9.5% less forward-looking R&D information in the post-AIA world.

may have encouraged firms to substitute away from in-house research and instead exploit external channels to acquire ownership of intellectual property. These findings are consistent with concerns that the transition to a first-to-file system can be detrimental to R&D activity of the individual firm.

2.2 Learning-by-Hiring

Firms do not generate all the knowledge required for continuous innovation internally, but often turn to external sources (Song et al., 2003). This is useful for innovative firms on several levels: First, it can be less costly and faster to source externally produced knowledge than to develop those competencies internally (Mansfield, 1988). Sourcing knowledge externally can enable firms to bypass growth constraints imposed by relying solely on internally developed resources (Singh and Agrawal, 2011). Second, a firm's past experience and existing stock of knowledge can constrain innovation, making its innovation trajectory highly path-dependent (Singh and Agrawal, 2011). Exposure to externally developed knowledge can introduce firms to new ideas, practices, and areas of expertise, potentially mitigating the path-dependent nature of their inventive activities (Song et al., 2003).

One potential channel through which firms acquire externally developed knowledge is by hiring employees from outside the organization—a mechanism commonly referred to as learning-by-hiring. Several empirical studies provide support for this channel. An early contribution by Almeida and Kogut (1999) shows that regions with spatially defined labor markets—characterized by high intraregional and low interregional labor mobility—exhibit a greater degree of localized knowledge diffusion. In a related study, Song et al. (2003) find that mobile inventors are substantially more likely than their peers at the hiring firm to base their patents on inventions from their former employers, suggesting that mobility facilitates interfirm knowledge transfer [3]. Similarly, Rosenkopf and Almeida (2003) examine firm pairs in semiconductor industry in U.S. and find that dyads with greater labor mobility between them also exhibit stronger subsequent knowledge flows, as evidenced by a higher number of inter-firm citations.

[3] In particular, mobile inventors are about seven times more likely to cite patents from their previous employers than other inventors at the hiring firm

In a more recent study, Singh and Agrawal (2011) point out two main shortcomings of the prior literature in establishing the causal effect of inventor mobility on knowledge flows. First, they argue that destination firms are more likely to use—and therefore cite—ideas of inherently higher quality, and inventors of such high-quality ideas may also have a different likelihood of being recruited. Second, they acknowledge that a hiring decisions may be endogenous: a firm might be more inclined to recruit an inventor who works in a technological domain it already plans to pursue. Authors argue that, in either case, it would be incorrect to interpret the observed cross-sectional correlation between mobility and knowledge flows as a causal relation. As a remedy to these concerns, Singh and Agrawal (2011) exploit disaggregated longitudinal citation data, utilize a stringed matching procedure to construct their sample and employ a difference-in-differences design that compares changes in citations by the destination firm to the focal patents before and after the move, relative to control patents. Using this framework, they find that inventor mobility leads to a 219% increase in hiring firms' use of recruits' prior inventions, with half of the effect driven by the recruits' own reuse of their earlier work. This suggests that, when an inventor exhibits inter-firm mobility, the hiring firm draws on the inventor's knowledge base and begins building its patents upon the new hire's pre-move inventions.

Chapter 3

THEORETICAL FRAMEWORK AND HYPOTHESIS

In this chapter, I present the two hypotheses that I empirically test. In Section 3.1, I introduce the first hypothesis, which I refer to as the learning-by-hiring hypothesis. In Section 3.2, I turn to the main hypothesis of the paper, the AIA Effect. The order of the two hypotheses is not arbitrary. The AIA Effect hypothesis is conditional on the first.

3.1 First Hypothesis: Learning-by-Hiring

As discussed under section 2.2, learning-by-hiring refers to the idea that inventor mobility facilitates knowledge transfers to the hiring firm[1]. The validity of the learning-by-hiring hypothesis is central to this paper: if inventor mobility does not serve as a mechanism for facilitating knowledge diffusion, then I cannot argue that the AIA reduces knowledge flows via inventor turnover in the first place. In below, I state my first hypothesis formally.

Hypothesis 1 (Learning-by-Hiring): Inventor mobility facilitates knowledge transfers from the new hire to the employing firm.

One central question is: How can one capture the knowledge transfers induced by inventor mobility? To address this, I follow the prior literature discussed in Section 2.2 and exploit citation flows. In particular, I use citations by the destination firm to the pre-mobility patents of the new hire as a measure of knowledge transfer from the inventor to the employing firm. Consider the following example: Suppose inventor *i* initially works at firm A and later moves to firm B. I argue that an increase in citations from firm B's patents to inventor i's pre-move patents suggests that firm B is making greater use of the inventor's prior knowledge following the hire.

[1]The source of the knowledge flow can be either the mobile inventor's prior firm, firms that she were employed before or her knowledge base itself.

However, this increase may be driven by self-citations—cases where inventor i cites her own pre-move patents after joining firm B. Since such instances do not reflect knowledge diffusion to the broader organization, I construct a second measure that excludes self-citations, counting only those citations from firm B's patents where inventor i is not listed as an inventor on the citing patent. An increase in this adjusted measure would indicate that other inventors at firm B begin building on the prior inventions of the new hire, following her mobility. This, in turn, provides evidence for the learning-by-hiring mechanism, as it suggests that these inventors learn from the newcomer and incorporate her prior work into their own research. Finally, if this pattern persists over time, it implies that the mobile inventor's tacit knowledge not only spills over to her colleagues but also becomes embedded in the broader knowledge base of the destination firm.

The learning-by-hiring hypothesis has already received empirical support (Almeida and Kogut, 1999; Rosenkopf and Almeida, 2003; Song et al., 2003; Singh and Agrawal, 2011). Among these studies, the most recent—and arguably the most robust—is the one conducted by Singh and Agrawal (2011). Despite its significant contributions, this study has two main shortcomings. First, as they note, to facilitate an unambiguous conceptualization of interfirm mobility, Singh and Agrawal (2011) restrict their sample to patents with a single inventor. In my restricted sample,[2] such patents account for only 32 percent of all patents. This limitation however, is not only about reduced sample size but also raises concerns about external validity, as there is no reason to assume that single- and multiple-inventor patents are otherwise comparable.

Second, Singh and Agrawal (2011) use a patent-level DiD model, comparing citations received by patents with a mobile inventor before and after mobility to those received by matched patents without a mobile inventor. This approach focuses on only a portion of the mobile scientist's knowledge base, represented by the her patents that find a match. This is problematic because the destination firm may draw selectively from specific domains of the inventor's expertise, increasing citations for only a subset of her patents.

This thesis addresses both of these limitations. First, I impose substantially weaker restrictions and include any patent that meets these minimal conditions, regardless of

[2] See Section 4.1.

whether it lists a single or multiple inventors. This not only provides a substantially larger sample but also helps mitigate concerns about external validity. Second, by conducting the analysis at the inventor-firm level, I consider the entire patent portfolio of each inventor in my sample. This approach allows me to capture knowledge flows from the inventor's full body of work, rather than just a selected subset.

3.2 Second Hypthesis: AIA Effect

Once I provide empirical evidence for the learning-by-hiring hypothesis, I turn to my second hypothesis, which postulates that the AIA reduced the extent to which inventor mobility facilitates knowledge transfers from the new hire to the employing firm. I refer to this hypothesis as the AIA Effect and formally define it below.

Hypothesis 2 (AIA Effect): America Invents Act decreased knowledge transfers through inventor mobility.

The main mechanism that accounts for the AIA Effect builds on the prior literature discussed in Sections 2.1 and 2.2. Miyagiwa and Ohno (2015) and Miyagiwa (2015) both suggest that the shift to the first-to-file (FTF) system weakenes firms' incentives to invest in R&D. Similarly, Huang et al. (2021) shows that the transition to FTF reduces R&D disclosures by publicly traded firms, increasing information asymmetry between firms in patent production process. This, in turn, renders in-house research riskier, as firms have less information about their competitors' research agendas or progress, exposing them to a heightened probability of duplicative R&D investment. Finally, Yilmaz et al. (2024) shows that the patent application rate for publicly traded firms between 1990 and 2018 sharply declines following the enactment of the AIA [3]. This reduction in inventive output by publicly traded firms suggests a causing reduction in inventive input as well.

Prior research provides evidence that the AIA has been detrimental to individual firms' levels of R&D investment. How would such a reduction in research activity affect learning-by-hiring? I argue that if the transition to the first-to-file (FTF) system led to a decline in

[3] They calculate the patent application rate as the fraction of public firms within a given year filing at least one patent application divided by the total number of public firms available in the market.

in-house research, it must have become more difficult for the new hire's tacit knowledge to spill over to other inventors within the destination firm and to diffuse across the organization. This is because, a reduction in the firm's R&D intensity would result in the mobile inventor collaborating less with other inventors, applying her tacit knowledge less actively, sharing fewer ideas, and demonstrating her technical skills less frequently. As the opportunities for organizational learning from the mobile inventor diminish, the overall level of knowledge transferred through her mobility would correspondingly decline.

Nevertheless, I should note that this thesis does not empirically test the potential mechanisms underlying the results. Therefore, although the R&D channel appears to be a plausible candidate to explain the empirical results I document, it may not be the actual driver of my findings. Even if the R&D channel does account for the observed effects, it may not be the sole contributor, and other mechanisms, some of which are not discussed here, may also be at play.

Chapter 4

DATA AND DESCRIPTIVE STATISTICS

For my empirical analysis, I use USPTO patent data. This dataset includes unique identifiers for each patent and for each inventor listed on a granted patent, identifiers for patent assignees, application and grant dates, International Patent Classification information for all patents and citation data indicating which patents cite which others.

Patent data is useful for testing my hypothesis on several grounds. First, using an algorithm that tracks changes in the assignee firms of an inventor's patents, the data allows me to detect instances of inventor mobility and determine—from which firm to which firm—the inventor has moved. Second, the citation information enables me to capture knowledge transfers: by law, the patent applicant and their attorney are required to cite any prior art of which they are aware (Rosenkopf and Almeida, 2003). Therefore, in principle, if Patent X cites Patent Y, this indicates that Patent X builds upon the existing knowledge embodied in Patent Y (Song et al., 2003). Moreover, since the dataset includes assignee and inventor information, I can infer which firms (inventors) own (worked in) Patent X and Patent Y, and thus make claims about the direction of knowledge flows. Finally, patent data allows me to construct a variety of inventor-level observable characteristics, which I use to match the instances of inventor mobility with pseudo-mobility instances, ensuring a balanced comparison between the control and treatment groups.

In this chapter, I describe the data analysis component of my thesis and present the accompanying descriptive statistics. Section 3.1 outlines the methodology used to detect instances of inventor mobility. Section 3.2 details the Coarsened Exact Matching (CEM) procedure employed to match treatment and control inventors and concludes the chapter.

4.1 Detecting Instances of Inventor Mobility

To test my hypothesis, I first need to identify instances of inventor mobility. Since I do not have access to any form of longitudinal employment data, I rely on an algorithm that predicts mobility events using only the available patent data. In particular, I use the method proposed by Hoisl (2007), which detects inter-firm mobility by tracking changes in the assignee IDs of an inventor's patents.

To begin my analysis, I impose several restrictions on the raw patent data. First, I eliminate any patent with multiple assignees, as otherwise, I would not be able to clearly identify the destination firm to which an inventor moves. Second, I eliminate any patent without an assignee [1]. Finally, I eliminate any patent that is assigned to a U.S. government institution. The second and third restrictions are imposed because my focus is on inter-firm inventor mobility; thus, movements involving independent inventors or government institutions fall outside the scope of this analysis. Once I obtain the sample of patents that meet these restrictions, I match each patent with its associated assignee organization, listed inventors, and filing date. Since many patents are produced by multiple inventors, this process results in multiple rows for the same patent and assignee firm.

Next, I order all patents of each inventor chronologically using their filing dates[2]. Then, I identify changes in assignee firms, interpreting such changes as indicators of turnover. Consider the example provided in Table 3.1. Inventor i first had two patents assigned to firm A, with filing dates in November 2011 and December 2012. Subsequently, the same inventor was listed on a patent filed in March 2013, this time assigned to firm B. Since I have excluded any patents with multiple assignees—and given that the assignee firm typically reflects the organization under which the patent was produced (Babina et al., 2020)—I interpret the change in assignee ID from firm A to firm B as an indicator that Inventor i moved from firm A to firm B.

As the date of mobility, I take the midpoint of the last filing activity of Inventor i in

[1] Patents assigned to firms are typically produced by inventors—listed on the patents—who work in the firms' in-house R&D laboratories (Babina et al., 2020).

[2] Throughout this thesis, I rely on filing rather than grant dates, since the lag between invention and filing is, by construction, shorter than the lag between invention and grant.

Inventor ID	**Patent ID**	**Firm ID**	**Filing Date**	**Mobility**	**Move Date**
Inventor *i*	X	A	2011-11-30	No Move	NaN
Inventor *i*	Y	A	2012-12-10	No Move	NaN
Inventor *i*	Z	B	2013-03-06	Move	2013-01-22
Inventor *i*	W	A	2015-02-03	No Move	NaN
Inventor *i*	V	B	2016-05-16	No Move	NaN
Inventor *i*	Q	C	2019-07-22	Move	2017-12-18
Inventor *i*	M	C	2020-03-11	No Move	NaN

Table 4.1: Illustrative Example of Inventor Mobility Detection

her source firm (Firm B) and the destination firm (Firm A). Then, I extract the calendar year from the resulting date. This is admittedly an ad hoc estimate of the actual date of move, but it is unlikely to perform worse than more sophisticated approaches (Singh and Agrawal, 2011). Once the date of mobility is chosen, I obtain an instance of mobility, represented as a three-tuple of the form (Inventor i, Firm A ,2012), where the first element is the identification number of the mobile inventor, the second is the unique identifier for destination firm, and the third is the date of mobility [3].

There are some potential complications with this approach. For example, how should one interpret the observed change in the assignee ID from firm B back to firm A in row 4 ? I argue that such a reversal to firm A as the assignee neither indicates that the inventor never moved to firm B nor that she returned to firm A. The most convincing explanation is that inventor i worked on the patent W while still employed at firm A, subsequently moved to firm B, and only afterward did firm A file the patent application (Hoisl, 2007). In this case, although the invention was made during her employment at firm A and the patent application was filed after her move, firm A is legally obliged to list inventor i in the patent as she contributed to the invention. Based on this, I do not consider the fourth row as another instance of mobility. Another complication arises in row six, where, following the prior move to firm B, I observe a new firm (Firm C) that had not previously appeared as

[3] Notice that each such three-tuple in my data is unique. This is because, the same inventor cannot move to the same firm at the same year more than once.

an assignee in any of inventor i's patents. In that case, I argue that inventor i moved from Firm B to Firm C in December 2017, and I represent this second move as a three-tuple of the form (inventor i, Firm C, 2018).

Through this exercise, I detect approximately 1 million[4] instances of inventor mobility between 2007 and 2018. Conditional on being mobile, an inventor moves 1.633 times on average. I provide the distribution of inventor mobility frequencies in figure 3.1. Further I provide the evolution of the annual number of instances of inventor mobility between 2007 and 2018 in figure 3.2. One key takeaway from Figure 3.2 is that the total number of inventor mobility instances is greater in the 2013–2018 period compared to the 2007–2011 period. The former corresponds to the post-AIA period, and the latter to the pre-AIA period that I will consider in my empirical design. This imbalance in sample sizes can be important, particularly for the matching procedure discussed in the next section.

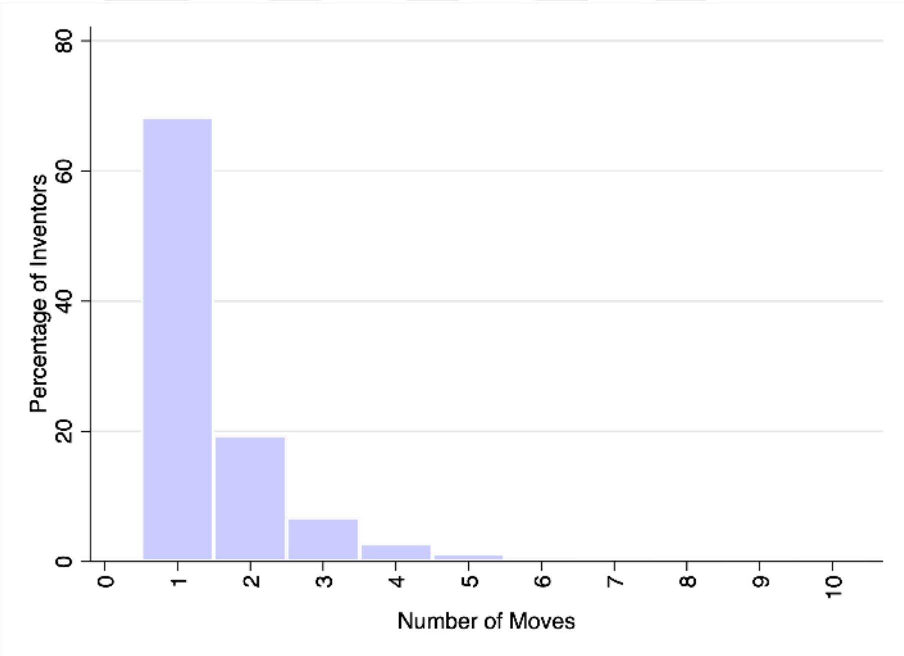

Figure 4.1: Distribution of Number of Moves for Mobile Inventors

Finally, I define an inventor as mobile if they moved at least once between 2000 and 2018. The 1 million detected instances of inventor mobility correspond to 560,610 unique mobile inventors, originating from 114,274 unique firms and moving to 131,248 unique

[4]The exact number of mobility instances is 915,485 .

destination firms. I define the remainder of inventors in the population as immobile, conditional on having at least one patent filing after January 1, 2000.

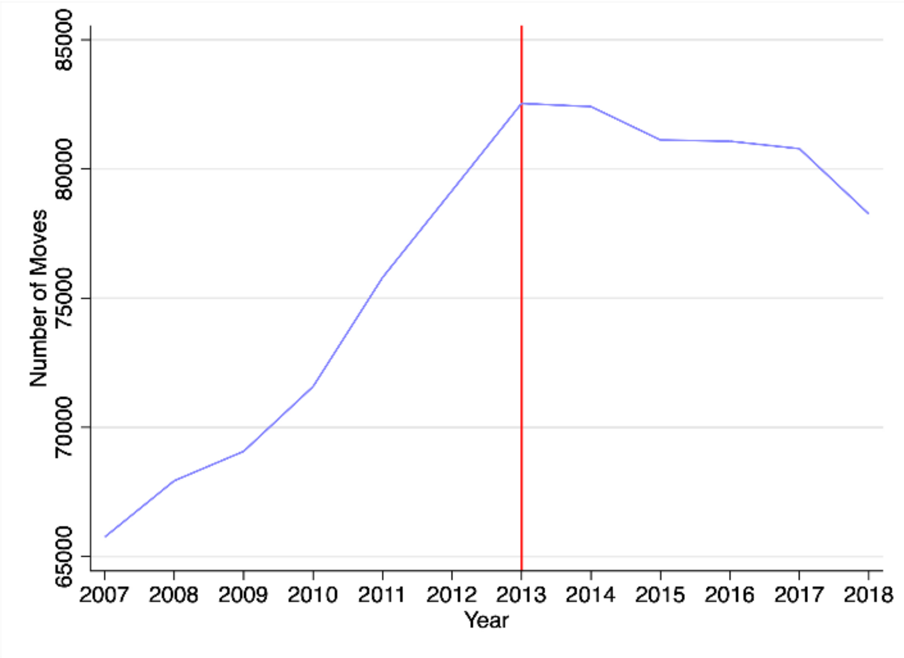

Figure 4.2: Annual Number of Instances of Inventor Mobility

Note: The red vertical line marks the year 2013, which is the year of shift to FTF

Before proceeding to the next section, I should highlight some limitations of the algorithm discussed above. First, the method is subject to a certain degree of noise. I estimate the date of mobility based on the last patenting activity at the source firm and the first patenting activity at the destination firm. While it is likely that the actual mobility date falls within this interval, taking the midpoint is only an approximation and may be misleading. For instance, if an inventor does not file a patent at the destination firm for an extended period, the midpoint approach may predict a mobility date that is too late. Second, my method identifies mobility instances conditional on the inventor having at least one patent filing. If an inventor moves between firms but never files a patent, I am unable to observe her in the patent data and, therefore, cannot capture the move. As a result, the external validity of the findings should be interpreted with caution. The estimated treatment effects apply specifically to inventors who are active in the patenting market.

4.2 Constructing the CEM Sample

Singh and Agrawal (2011) show that single-inventor patents involving a mobile inventor tend to receive more pre-move citations from the destination firm and are, on average, produced by more experienced inventors compared to control patents without a mobile inventor. This suggests that mobile and immobile inventors differ systematically on observable characteristics and raises concerns about their comparability. To address this, I employ the Coarsened Exact Matching (CEM) approach (Iacus et al., 2008), which allows me to construct treatment and control groups that are more balanced in terms of observable characteristics than the unmatched mobile and immobile inventors in the raw data. In this chapter, I discuss the mechanics of the CEM method I employ in Subsection 3.2.1 and present descriptive statistics for the CEM sample I construct in Subsection 3.2.2.

4.2.1 Mechanics

As discussed in Section 3.1, a mobility instance is defined as a three-tuple of the form (i_m, j, t), indicating that the mobile inventor i_m moved to firm j in year t [5]. The main task in the CEM procedure I follow is to match each such mobility instance to a pseudo-mobility instance of the form (i_c, j, t), where i_c is an inventor who has never exhibited inventor mobility after January 1, 2000; j is the exact destination firm to which inventor i_m moved; and t is the corresponding year of mobility. In this context, j and t respectively serve as the pseudo destination firm and pseudo mobility date for i_c, who in reality never moved.

To match the actual and pseudo-mobility instances, I use discrete buckets based on the following observable characteristics: (1) total pre-mobility citations received by the inventor (25 bins), (2) total pre-mobility patent filings by the inventor (18 bins), (3) years of tenure at the mobility date[6] (6 bins), and (4) the inventor's three-digit IPC code (exact match)[7]. I provide the exact bins I used in table 3.2.

---

[5] I use subscript m to denote mobile and c to denote immobile.

[6] Tenure is calculated as the year difference between the date of mobility and the Filing Date of inventor's first patent.

[7] The IPC is a hierarchical patent classification system, with sections ranging from A to H, further divided into subsections, classes, groups, and subgroups. For each inventor, I identify their primary IPC code as the

Once I construct the binned observable characteristics for each three-tuple in my sample, I collect them into a vector. In the first stage, I match an actual and a pseudo-mobility instance if they share the same vector of binned observable characteristics. In the second stage, I impose an additional condition and retain a matched pair only if both inventors received the same number of total pre-mobility citations from the destination firm specifically. I impose this final condition at the second stage because, because otherwise, for each immobile inventor, I would have to compute the citations received from any destination firm, up to any calendar year from 2007 to 2018, which is computationally infeasible. However, since the matching procedure imposes an exact match on the number of citations received from the destination firm, separating the process into two stages does not alter the resulting set of matched units.

Table 4.2: Bin Definitions for Pre-Move Variables

Variable	**Bin Definitions**
Pre-move Citations	0, 1–2, 3–4, 5–7, 8–12, 13–25, 26–35, 36–50, 51–65, 66–80, 81–100, 101–125, 126–150, 151–175, 176–200, 201–250, 251–300, 301–400, 401–500, 501–600, 601–700, 701–800, 801–900, 901–1000, ≥1001
Pre-move Patent Filings	0–1, 2–3, 4–6, 7–12, 13–20, 21–30, 31–40, 41–50, 51–60, 61–70, 71–80, 81–90, 91–100, 101–200, 201–300, 301–400, 401–500, ≥501
Inventor Tenure	0–1, 2–3, 4–6, 7–10, 11–15, ≥16
IPC Code	Exact Match
Pre-move Citations (Dest.)	Exact Match

To illustrate this approach, consider an actual example from the data. fl:zy_ln:gorgol-1

one most frequently appearing among the patents on which they are listed as an inventor.

is a mobile inventor who moved to firm A in 2011[8]. Prior to his move, fl:zy_ln:gorgol-1's patents had received a total of 80 citations, he had three successful patent applications, his tenure was 11 years, and the majority of his patents were classified under IPC code G06. Therefore, his pre-move unbinned vector of observables was (2011, 80, 3, 11, G06). Based on the discrete buckets defined in table 4.2, his binned vector becomes (2011, 66–80, 2–3, 11–15, G06).

Now consider fl:ni_ln:coleman-1, an inventor who never exhibited mobility. Right before fl:zy_ln:gorgol-1 moved to firm A, fl:ni_ln:coleman-1 had the following vector of observables: (2011, 79, 2, 13, G06), which also maps to the same binned vector: (2011, 66–80, 2–3, 11–15, G06). Since both inventors have identical binned vectors, I match them in the first stage of the CEM procedure. In the second stage, I check criterion (2), which involves comparing the number of pre-mobility citations each inventor received from the destination firm A. If the two inventors received different numbers of pre-move citations from firm A, the match is broken. However, in this case, both fl:zy_ln:gorgol-1 and fl:ni_ln:coleman-1 received the same number of citations from firm A, which is zero. Therefore, I keep the match. As a result, one treatment unit becomes (fl:zy_ln:gorgol-1, Firm A) and its associated control unit becomes (fl:ni_ln:coleman-1, Firm A). Notice that once I perform the match, I do not include the move date of the mobile inventor in the vector, as the combination of the inventor and destination firm identifiers is sufficient to uniquely define each instance of mobility. Therefore my units are inventor-destination firm pairs.

4.2.2 Descriptive Statistic for the CEM Sample

Using the Coarsened Exact Matching (CEM) method described in the previous subsection, I identify at least one match for 155,081 instances of inventor mobility, with move dates from 2007 to 2018, excluding the years 2012 and 2013.[9] Many of these mobility events are matched to multiple control units. In such cases, I allow for multiple control units and incorporate CEM weights into the regression analysis. Mobility instances without any match

[8]Specifically, to the firm with the unique identifier 061b02fd-95e6-474c-87d0-062217c24ffd.

[9]I exclude mobility instances occurring in 2012 and 2013 to eliminate anticipation effects.

are excluded from the sample. Table 3.3 presents descriptive statistics for the matched mobility instances and those dropped from the analysis. I document that mobility instances with matched controls differ systematically from unmatched ones across all observable characteristics. In particular, on average, unmatched instances involve mobile inventors who received more pre-move citations from the destination firm, more cited overall, have a greater number of pre-move patent applications and grants, and exhibit higher levels of experience.

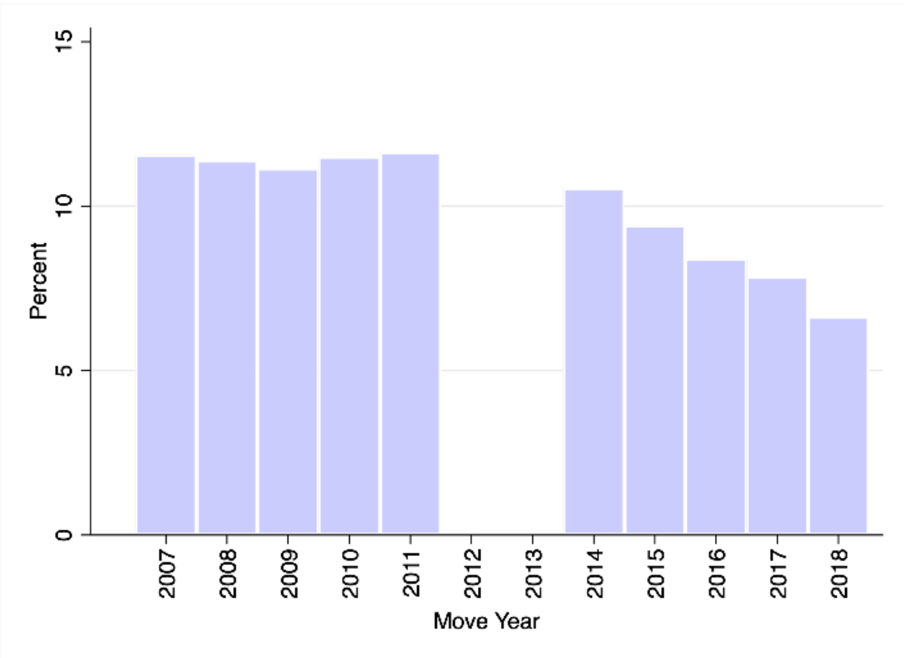

Figure 4.3: Distribution of Move Years in The CEM Sample for Actual Mobility Instances

This systematic difference between matched and unmatched mobility instances has two important implications for my analysis. First, it raises concerns about the external validity of my results: Unmatched mobility instances involve more experienced and productive inventors. These inventors are also more likely to facilitate organizational learning through mobility (Varshney, 2023). Hence, their exclusion implies that I omit a subgroup with potentially stronger treatment effects, where the treatment is mobility itself. As a result, in testing the first hypothesis, the estimated effect of mobility on knowledge transfers (ATET) is likely to underestimate the average treatment effect (ATE). Second, the disparity documented in Table 3.3 suggests that the matching method I employ is effective: Immobile inventors are cited less, have lower inventive output and experience (Singh and Agrawal, 2011). Therefore, it is expected that the units they are matched to have significantly lower inventive output and experience compared to the units they did not match with.

Table 4.3: Descriptive Statistics for Matched and Unmatched Mobility Instances

	Matched		**Unmatched**	
	Mean	Std. dev.	Mean	Std. dev.
Cumulative Pre-Mobility Citations (Dest.)	0.171	1.276	1.108	26.03
Cumulative pre-move citations	36.42	94.59	185.5	1167.4
Pre-MobilityPatents	3.324	6.348	13.53	35.61
Pre-Mobility Applications	5.375	7.710	18.69	43.99
Pre-Mobility Inventor Tenure	7.658	8.270	12.49	10.22

Table 4.4: Descriptive Statistics for Control and Matched Treatment Units

	Treatment		**Control**	
	Mean	Std. dev.	Mean	Std. dev.
Cumulative Pre-Mobility Citations (Dest.)	0.171	1.276	0.107	0.773
Cumulative Pre-Mobility Citations	36.42	94.59	18.72	65.53
Pre-Mobility Patents	3.324	6.348	1.830	4.444
Pre-Mobility Applications	5.375	7.710	2.811	5.065
Inventor age	7.658	8.270	3.166	6.670
Number of Treatment Units=155,081				
Number of Control Units= 515,606				
Total Units= 670,687				

Finally, in Table 3.4, I present the descriptive statistics for the treatment and control units. Although the control units are statistically more similar to the treatment units they are matched with, the CEM sample remains unbalanced. On average, treatment units involve mobility instances where inventors have a higher number of pre-move citations from the destination firm, receive more overall citations, have submitted more pre-move patent applications and been granted more patents, and possess greater experience than their matched controls. Although this is not ideal, it reflects a data limitation for which a remedy is likely infeasible. Moreover, as I will discuss in Chapters 5 and 6, the imbalance in the CEM sample does not raise concerns about identification. On the contrary, the exact matching imposed on pre-move citations received from destination firms supports the validity of the parallel trends assumption in both the DiD and DiDiD designs.

4.3 Panel Data

Once I obtain the treatment and control units using the CEM approach, I construct the panel data structure that I use in my regression analysis. In this panel, my units are inventor–destination firm pairs. As the first dependent variable, I count the annual number of citations received by inventor i's patents from patents assigned to destination firm j, in citation year t. As the citation years, I consider the time period between 2003 to 2022, inclusive. Since I have a total of 670,687 inventor–destination firm pairs, my panel has a total of $670,687 \times 20 = 13,413,740$ observations. As the second dependent variable, I count the annual number of citations received by inventor i from the destination firm j, conditional on inventor i not being listed as an inventor in the citing patent.

As the right-hand-side variables, I first construct a binary indicator, Mobile, which takes a value of one if the inventor–destination firm pair involves a mobile inventor, and zero otherwise. Second, I construct another binary indicator, Post, which takes a value of one for the inventor–destination firm pair in the years following the associated mobility date of the pair. Third, for the DiDiD specification, I define a binary variable, Move 2013, which captures whether the instance of mobility occurred before or after the shift to the FTF rule. Move 2013 takes a value of one for inventor–destination firm pairs with a mobility date after 2013, and zero otherwise. Finally, following Singh and Agrawal (2011), I construct the variable Tenure, which measures the inventor's tenure and is computed as the difference between the panel year and the year of the inventor's first patent filing. I provide the variable definations in table 4.5.

Table 4.5: Variable Definitions

Variable	**Definition**
Dependent Variables:	
`Citation`	Number of citations received by the inventor's patents from the destination firm in a given year.
`No Self Citation`	Number of citations received by the inventor from the destination firm in a given year, excluding patents where the inventor is also listed on the citing patent.
Explanatory Variables:	
`Mobile`	Binary indicator equal to 1 if the inventor–destination firm pair involves a mobile inventor, and 0 otherwise.
`Post`	Binary indicator equal to 1 for years following the inventor's move to the destination firm; 0 otherwise.
`Move 2013`	Binary indicator equal to 1 if the inventor moved after 2013 (post-AIA); 0 otherwise.
`Tenure`	Years since the inventor's first patent filing, calculated as the difference between the current panel year and the inventor's first filing year.
# of Observations:	13,413,740
Unit of Observation:	Inventor–destination firm–year.
Time Span:	Panel covers citation years from 2003 to 2022.

Chapter 5

EMPIRICAL STRATEGY

5.1 Difference-in-Differences Design for Testing the First Hypothesis

To test the first hypothesis, I use a Difference-in-Differences (DiD) design based on the panel data described in Section 4.3. This framework exploits the fact that I observe citations received by both actual and pseudo mobility instances before and after the mobility event, by the associated destination firm. Therefore, although factors other than mobility may contribute to differences in post-mobility citations between the treatment and control units, the DiD design allows me to disentangle these effects by accounting for differences in citations prior to the mobility date. In this sense, the pre-mobility differences in citations received by the treatment and control units from the destination firm serve as a benchmark, allowing me to attribute any excess post-mobility difference between the two groups to the effect of mobility. To implement this idea, I estimate the following Difference-in-Differences (DiD) specification:

$$\begin{aligned} citation_{ijt} = \beta_0 + \beta_1 \cdot mobile_i + \beta_2 \cdot post_{ijt} + \beta_3 \cdot (mobile_i \times post_{ijt}) \\ + \alpha_{ij} + \lambda_t + \delta_{t-\text{earliest filing}(i)} + \epsilon_{ijt} \end{aligned} \tag{5.1}$$

where i indexes the inventor—either mobile or immobile— denotes j the destination firm, and t is the citation year. The dependent variable $citation_{ijt}$ denotes the number of citations to the patents of inventor i by destination firm j in year t. The first binary indicator in the regression, $mobile_i$, takes a value of one if the inventor associated with the inventor–destination firm pair (i, j) is mobile, and zero otherwise. Note that $mobile_i$ is defined at the inventor level and does not vary across destination firms j. This is because, if inventor i is classified as mobile, all inventor–destination firm pairs involving i take the value of one, regardless of the specific firm j to which the inventor moves. The coefficient β_1 captures systematic differences between mobile and immobile inventors in terms of the citations they receive from the destination firm, which exist irrespective of mobility. I expect the estimated β_1 to be positive because as shown in Table 4.4, mobile inventors receive more citations from the destinaiton firms, even before the mobility. Next, I include another binary indicator, $post_{ijt}$, which takes a value of one for the inventor–destination firm pair in citation years following the specified date of mobility, and zero otherwise. Note that $post_{ijt}$ equals one for both actual mobility instances and the matched pseudo pairs during the periods after the mobility date. Therefore, as argued by Singh and Agrawal (2011), β_2 reflects the counterfactual change in the citation rate at the time of the move, had the move not actually occurred.

As the final variable, I include the interaction term $mobile_i \times post_{ijt}$, which equals one only for actual mobility instances in the periods following the move. The associated parameter, β_3, is the coefficient of interest, as it captures the causal effect of mobility on the destination firm's use of the newly hired inventor's prior inventions, conditional on the validity of the identification assumptions. A positive and statistically significant estimate of β_3 would indicate that mobility itself leads to an increase in the destination firm's use of the new hire's knowledge.

In addition to the binary indicators and the DiD interaction term discussed above, I include several fixed effects in the model specification. First, I include a matched-pair fixed effect α_{ij}, which controls for any time-invariant characteristics specific to the inventor–destination firm pairs that are matched together. As a result, the coefficient of interest in equation 5.1. is identified just using the variation within matched pairs over time, which is desirable as I use the CEM sample. Further, I include time fixed effects λ_t to capture year-specific shocks that may influence citation rates. Finally, to account for life-cycle effects[1], I include the tenure fixed effect $\delta_{t-\text{earliest filing}(i)}$, which is computed as the difference between the current citation year and the inventor's earliest patent filing year. In addition to my preferred model 5.1, I also explore alternative specifications that vary in their fixed effects structure, including models with inventor fixed effects , firm fixed effects and both, presented in Equations 5.2, 5.3 and 5.4, respectively.

[1]Inventors are likely to be more productive at particular stages of their careers relative to other years of tenure.

$$\begin{aligned} citation_{ijt} = \beta_0 + \beta_2 \cdot post_{ijt} + \beta_3 \cdot (mobile_i \times post_{ijt}) \\ + \gamma_i + \lambda_t + \delta_{t-\text{earliest filing}(i)} + \epsilon_{ijt} \end{aligned} \tag{5.2}$$

$$\begin{aligned} citation_{ijt} = \beta_0 + \beta_1 \cdot mobile_i + \beta_2 \cdot post_{ijt} + \beta_3 \cdot (mobile_i \times post_{ijt}) \\ + \theta_j + \lambda_t + \delta_{t-\text{earliest filing}(i)} + \epsilon_{ijt} \end{aligned} \tag{5.3}$$

$$\begin{aligned} citation_{ijt} = \beta_0 + \beta_2 \cdot post_{ijt} + \beta_3 \cdot (mobile_i \times post_{ijt}) \\ + \theta_j + \gamma_i + \lambda_t + \delta_{t-\text{earliest filing}(i)} + \epsilon_{ijt} \end{aligned} \tag{5.4}$$

In these specifications, γ_i denotes the inventor fixed effect, and θ_j represents the destination firm fixed effect. I include inventor fixed effects to better address unobserved heterogeneity across inventors, recognizing that inventors may differ in terms of productivity or the relevance of their expertise to the destination firm [2]. The inclusion of θ_j is motivated by similar reasoning, as there may be firm-level unobservable characteristics that explain their propensity to cite prior inventions.

Finally, I estimate each of the above specifications using an alternative dependent variable, No Self Citations, which is also described in table 4.5. Specifically, in the baseline specification, I run the following DiD regression:

$$\begin{aligned} no\_self\_citation_{ijt} = \beta_0 + \beta_1 \cdot mobile_i + \beta_2 \cdot post_{ijt} + \beta_3 \cdot (mobile_i \times post_{ijt}) \\ + \alpha_{ij} + \lambda_t + \delta_{t-\text{earliest filing}(i)} + \epsilon_{ijt} \end{aligned} \tag{5.5}$$

where the dependent variable $no\_self\_citation_{ijt}$ denotes the number of citations to the patents of inventor i by destination firm j in year t, conditional on inventor i not being listed as an inventor on the citing patent. I use this alternative dependent variable because, after an inventor moves to a new firm, she is likely to cite her own prior patents, which would mechanically increase the number of citations from the hiring firm to her past inventions. Clearly, such a rise in citations does not necessarily reflect a knowledge

[2] Note that, when I include the inventor fixed effect γ_i, I omit $mobile_i$, since the fixed effect absorbs all time-invariant inventor-specific characteristics.

transfer from the inventor to the firm, but rather captures the inventor's own reuse of her previous work. A rise in the variable No Self Citation, however, would suggest that other inventors at the destination firm increasingly use the new hire's prior inventions in their own research following her mobility, thereby providing evidence for learning-by-hiring. Therefore, in Equation 5.5, conditional on the validity of the identification assumptions, the parameter β_3 captures the causal effect of inventor mobility on knowledge transfers from the new hire to the employing firm.

5.2 Difference-in-Difference-in-Differences Design for Testing the Second Hypothesis

The second hypothesis postulates that the AIA reduces the extent to which inventor mobility facilitates knowledge transfers from the new hire to the employing firm. To test this hypothesis, I employ a Difference-in-Difference-in-Differences (DiDiD) design. My empirical strategy relies on comparing the causal effect of inventor mobility on knowledge transfers before and after the enactment of the policy. Put simply, the DiDiD equation I use contrasts the average knowledge transfers facilitated by inventor mobility instances that occurred before and after the transition to first-to-file. The DiDiD specification I estimate is as follows:

$$\begin{aligned} no\_self\_citation_{ijt} = \beta_0 &+ \beta_1 \cdot Move2013_{ij} + \beta_2 \cdot Mobile_i + \beta_3 \cdot Post_{ijt} \\ &+ \beta_4 \cdot (Move2013_{ij} \times Mobile_i) + \beta_5 \cdot (Move2013_{ij} \times Post_{ijt}) \\ &+ \beta_6 \cdot (Mobile_i \times Post_{ijt}) + \beta_7 \cdot (Move2013_{ij} \times Mobile_i \times Post_{ijt}) \\ &+ \alpha_{ij} + \lambda_t + \delta_{t-\text{earliest filing}(i)} + \epsilon_{ijt} \end{aligned} \tag{5.6}$$

where i indexes the inventor—either mobile or immobile—j denotes the destination firm, and t is the citation year. As in Equation 5.5, the dependent variable $no\_self\_citation_{ijt}$ denotes the number of citations to the patents of inventor i by destination firm j in year t, conditional on inventor i not being listed as an inventor on the citing patent. The right-hand-side variables $Mobile_i$ and $Post_{ijt}$ are defined the same as in the DiD specifications described in Section 5.1. The variable that is new to the previous framework is

$Move2013_{ij}$, which takes a value of one for inventor–destination firm pairs with mobility date after 2013, and zero otherwise. Thus, $Move2013_{ij}$ essentially captures whether the mobility instance involving inventor i and destination firm j occurred before or after the transition to the first-to-file rule.

In this specification, the parameter of interest is β_7, which is the coefficient on the triple interaction term $(Move2013_{ij} \times Mobile_i \times Post_{ijt})$. An intuitive interpretation I prefer for the effect that β_7 captures is based on the seminal paper by Olden and Men (2022), which provides a comprehensive presentation of the triple-difference estimator, explaining its mechanics and identification assumptions. In particular, Olden and Men (2022) shows that the triple-difference estimator can be computed as the difference between two underlying difference-in-differences estimators. In the context of Equation 5.6, this implies that the DiDiD estimator for β_7 is equivalent to the difference between two DiD estimators for β_3 in Equation 5.2: one estimating the effect of mobility on knowledge transfers for the pre-2013 period and the other for the post-2013 period. This implies, β_7 in the DiDiD Equation 5.6 captures the effect of the AIA on knowledge transfers through inventor mobility. A negative estimate of β_7 would indicate that the AIA reduced the extent to which inventor mobility facilitates knowledge transfers from the new hire to the employing firm. However, the common interpretation of the DiDiD estimator as the difference between two DiD estimates—and thus the basis for a causal interpretation—may not hold when covariates play an important role in identification. This calls for additional caution in interpreting the results.

5.3 Identification

In this section, I further discuss the identification strategy underlying the empirical models I present. In particular, Subsection 5.3.1 outlines the main identification assumption of the DiD and DiDiD models and explains how I assess their validity. Subsection 5.3.2 then addresses potential threats to identification and discusses how my empirical design mitigates them.

5.3.1 Paralel Trends

The main identifying assumption for the DiD framework is that, in the absence of treatment, the average outcomes for the treatment and control groups would have followed parallel paths over time. In the context of this study, the parallel trends assumption for the DiD model implies that, had there been no mobility, the citations received by mobile and immobile inventors from the destination firm would have followed parallel paths. Since the counterfactual citation trajectories for treated units in the absence of mobility in the post-move period are unobservable, the most credible test of the parallel trends assumption is to show that citation flows from the destination firm to mobile and immobile inventors exhibit parallel paths up to the date of mobility.

Similar to the DiD framework, the DiDiD estimator requires a parallel trend assumption for the estimated effect to have a causal interpretation (Olden and Men, 2022). Since the triple difference estimator is constructed as the difference between two underlying difference-in-differences estimators, the validity of the parallel trend assumption for each underlying DiD implies the validity of the assumption for the DiDiD estimator (Olden and Men, 2022). However, the DiDiD estimator does not require that both DiD components individually satisfy the parallel trend assumption. As Olden and Men (2022) argue, the difference between two biased DiD estimators can still be unbiased, as long as the bias is the same for both estimators. Therefore, the parallel trend assumption for the DiDiD estimator is that, in the absence of mobility, the differential in citation flows to mobile and immobile inventors from the destination firm in the post-2013 period to trend similarly to the differential in citation flows to mobile and immobile inventors in the pre-2013 period.

In Chapter 6, I assess the validity of the parallel trends assumptions for both the DiD and DiDiD models using event study designs. This approach introduces leads into the regression equation and therefore allows me to analyze the pre-treatment trends for the control and the treatment group. For the DiD specification, I simply check whether the trends in the outcome variable are the same for the treatment and control groups during the pre-mobility period. For the DiDiD framework, I evaluate the validity of the parallel trends assumptions for the two underlying DiD components individually.

5.3.2 Threats to Identification

I should outline and discuss several threats to identification, as well as the extent to which my empirical strategy mitigates them. One such concern is the presence of anticipation effects, for which there is supporting evidence in the existing literature. For example, when examining the effects of the Canadian reform, Abrams and Wagner (2013) documents a spike in patent applications just before the effective date of the FTF system on October 1, 1989. Similarly, Yilmaz et al. (2024) finds a surge in patent applications to the USPTO right before the effective date of the transition to FTF on March 16, 2013. These empirical findings suggest strong anticipation effects, as firms rush to the patent office to exploit the FTI system before its abandonment. This raises the concern that firms anticipating the transition to FTF rule may also alter their hiring behavior, or that their citation flows may change as a result of the rush to the patent office.

Table 5.1: Legislative Timeline of the America Invents Act (AIA)

Date	**Event**	**Vote**
Feb. 3, 2011	AIA (S. 23) is reported by the Senate Judiciary Committee	*15-0*
March 8, 2011	AIA (S. 23) is passed by the Senate	*95-5*
April 14, 2011	AIA (H.R. 1249) is reported by the House Judiciary Committee	*32-3*
June 23, 2011	AIA (H.R. 1249) is passed the House[a]	*304-117*
Sept. 8, 2011	AIA (H.R. 1249) is passed by Senate	*89-9*
Sept. 16, 2011	AIA (H.R. 1249) is signed into law	n.a.
	………………………	
March 16, 2013	The FTF system became effective	n.a.

Avoiding such contamination from anticipation effects requires a careful consideration of the AIA's enactment and the effective dates of its provisions. The AIA was reported by the Senate Judiciary Committee on February 3, 2011, and signed into law on September 16, 2011. However, its provisions were not implemented uniformly, but rather phased in over a two-year period (Yilmaz et al., 2024). The shift to the first-to-file (FTF) system became effective on March 16, 2013. Since I use citation years rather than exact dates, I treat the enactment as beginning in early 2012 (i.e., end of 2011) and the FTF implementation as effective throughout 2013. I therefore argue that firms anticipated the transition during 2012 and 2013. Hence, to avoid confounding effects driven by anticipatory behavior, I exclude from the sample all inventor mobility instances that occurred in those two years.

Another threat to identification is the potential endogeneity of firms' hiring decisions. Firms may be more inclined to recruit inventors who are already working in technological domains that the firm intends to pursue (Singh and Agrawal, 2011). If this is the case, any subsequent increase in citations—relative to control inventors—may not reflect the effect of mobility per se, but rather a shift in the firm's technological focus. To mitigate this concern, I restrict the sample to matched inventor pairs whose patents predominantly fall under the same three-digit IPC code. This ensures that the matched mobile and immobile inventors are comparable in terms of technological expertise. Consequently, if a firm's hiring decision is influenced by a directional change in its technological focus, this effect should be similarly present across both groups, helping to isolate the impact of mobility itself.

A further concern regarding endogeneity is the heterogeneity in unobservable characteristics of inventors that correlates with both their likelihood of being mobile as well as the citations they receive from the destination firm. For example, it is very likely that inventors of higher ability/productivity are not only more likely to be mobile but also more likely to receive citations from the destination firm, irrespective of whether they move. To mitigate this concern, I use inventor-level fixed effects, which capture the unique inventor characteristics that are time-invariant. Second, I argue that the CEM approach helps reduce bias arising from unobserved heterogeneity between mobile and immobile inventors. While CEM matches these two groups based on observable characteristics, unobservables such as ability or productivity are likely to be correlated with those observables [3]. Therefore, matching indirectly reduces differences in unobservable characteristics as well. Most

[3]For example, high-ability inventors tend to have more patents and receive more citations than low-ability inventors.

importantly, as long as the unobservable characteristics that differ across mobile and immobile inventors are constant over time, the difference-in-differences method will cancel them out by construction.

A final concern for the identification of the DiDiD desing is the presence of any contemporaneous policy change or shock that also affects citation flows. If this is the case, my results would reflect the joint effect of the AIA and the contemporaneous shock on knowledge transfers through inventor mobility. To the best of my knowledge, there was no major policy change in 2013 that could have influenced knowledge transfers through hiring. Even if such a shock did exist, it is implausible that it had an effect as direct and targeted as the AIA, making it unlikely to fully explain my findings. The same rationale applies to provisions of the AIA other than the change in the priority rule.

Chapter 6

RESULTS

In this chapter, I present the empirical results of my study. Section 6.1 includes two subsections: Subsection 6.1.1 presents the DiD estimation results used to test Hypothesis 1, while Subsection 6.1.2 provides the results of the event study design, which assesses the validity of the parallel trend assumption and explores the persistence of the effect of mobility on knowledge flows. Section 6.2 also consists of two subsections. Subsection 6.2.1 presents the DiDiD estimation results for testing Hypothesis 2, and Subsection 6.2.2 examines the event study design that verifies the parallel trend assumption for the triple-difference framework.

6.1 Testing the First Hypothesis

6.1.1 Regression Results

Table 6.1 reports the regression results for the DiD specifications with various fixed effects structures, corresponding to equations 5.1 through 5.4. As a preliminary step however, I regress the dependent variable, Citation, on the binary indicator Mobile, using only the post-move subsample. The positive and statistically significant coefficient on Mobile in column (1) implies that mobile inventors receive more citations in the post-move period compared to their matched immobile counterparts. This estimate reflects the result one would obtain by simply comparing citation flows to mobile and immobile inventors after the move.

The DiD results in column (2) illustrate why such an approach is misleading. Even after including the interaction term Mobile $\times$ Post, the coefficient on Mobile remains positive and significant. This suggests that mobile inventors systematically receive more citations from destination firms even before the move, pointing to the presence of pre-existing differences and selection into mobility. Nevertheless, since the coefficient on the interaction term Mobile × Post is positive and statistically significant, the results provide evidence that the hiring firm increases its use of the new hire's pre-move inventions and knowledge following the mobility.

Table 6.1: DiD Estimation Results for the Effect of Mobility on Knowledge Transfers

	(1)	(2)	(3)	(4)	(5)
Dependent Variable:	Citations	Citations	Citations	Citations	Citations
Fixed Effect:	Pair FE	Pair FE	Firm FE	Inventor FE	Inv. & Firm FE
Sample:	Post-Move	Full Sample	Full Sample	Full Sample	Full Sample
`Mobile`	0.549***	0.012***	0.012***	—	—
	(0.019)	(0.003)	(0.002)	—	—
`Post`	—	0.162***	0.210**	0.171***	0.175***
	—	(0.007)	(0.0048)	(0.005)	(0.007)
`Mobile×Post`	—	0.537***	0.536***	0.456***	0.463***
	—	(0.022)	(0.023)	(0.019)	(0.02)
Number of Observations	13,413,740	13,413,740	13,413,740	13,413,740	13,413,740
R^2	0.127	0.080	0.0142	0.272	0.2765
F-Statistic	818.78	1731.20	1752.57	2045.12	1978.65

Notes: Standard errors in parentheses. I use robust standard errors clustered at inventor level. Results are robust to clustering at the matched-pair level. I use CEM weights in each regression. All specifications include year and tenure fixed effects. Column (1) Uses the post-move subsample and includes matched pair fixed effects. Remaining columns use the full sample. Column (2) includes firm fixed effects. Column (3) includes inventor fixed effects. Column (4) includes both inventor and firm fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

In columns (2) through (5), I explore different fixed effects structures, while including year and tenure fixed effects in each specification. The model estimated in column (2) uses matched-pair fixed effects, which absorb the time-invariant differences across matched pairs. As a result, the coefficient of interest is identified just using the variation within matched pairs over time. This approach is particularly appropriate given that the matched pairs were constructed using the Coarsened Exact Matching (CEM) method, and is therefore my main specification. The remaining specifications are used to address various concerns regarding identification.

In column (3) I use destination firm fixed effects to control for firm-level unobservable characteristics. However, as expected, the results do not change over column (2) [1]. In columns (4) and (5) I use inventor fixed effects to address the concern that mobile inventors may differ from immobile inventors in unobservable characteristics, even after the CEM procedure. In line with this concern, the estimated coefficient on Mobile × Post in columns (4) and (5) is significantly lower than in column (2), suggesting that my main specification may be overestimating the effect of mobility [2]. Nevertheless, in either specification, the estimated coefficient on Mobile× post is statistically significant and positive. The lowest estimated effect across all specifications indicates that mobility leads to an increase of 0.46 citations per year from the destination firm. Given that the average mobile inventor in my sample receives only 0.03 citations per year from the destination firm prior to the move, my estimations are economically significant.

The estimation results in Table 6.1 show that the hiring firm increases its use of the new hire's knowledge following her employment, as evidenced by the increase in citation flows to her prior patents. However, as discussed in detail in Sections 3.1 and 5.1, this increase in citations may partly reflect the new hire's own reuse of her previous work, rather than genuine knowledge transfer to the firm. To disentangle these two effects, I re-estimate the specifications from Table 6.1 using No Self Citation as the dependent variable. The coefficient on Mobile × Post remains positive and statistically significant across all specifications. However, the estimated effect of mobility on citation flows from the destination firm is substantially smaller compared to the results in Table 6.1. In particular, the estimated coefficient on Mobile × Post in column (4) of Table 6.2 is only 60% (0.274/0.456) of the corresponding estimate in Table 6.1. This suggests that the earlier estimates in table 6.1. were partly driven by the mobile inventor's self-use of her prior inventions, rather than by knowledge spillovers to other inventors at the hiring firm.

[1] The result is expected because, I consider the citations from the same destination firm for each matched mobile and immobile inventors.

[2] The most likely explanation is that mobile inventors differ from immobile inventors in unobservable characteristics such as inherent ability, which are positively correlated with both the likelihood of being mobile and the number of citations they receive from the destination firm. As a result, once these time-invariant characteristics are captured by the inventor fixed effects, the estimated effect of mobility on citation flows declines significantly.

Table 6.2: DiD Estimation Results for the Effect of Mobility on Knowledge Transfers - Excluding Self Citations

	Dependent Variable: No Self Citation			
	(1)	(2)	(3)	(4)
Fixed Effect:	Pair FE	Firm FE	Inventor FE	Inv. & Firm FE
`Mobile`	0.006***	0.006***	—	—
	(0.003)	(0.001)	—	—
`Post`	0.174***	0.216***	0.184**	0.188***
	(0.007)	(0.005)	(0.006)	(0.006)
`Mobile×Post`	0.348***	0.347***	0.274***	0.280***
	(0.020)	(0.021)	(0.018)	(0.018)
Number of Observations	13, 413, 740	13, 413, 740	13, 413, 740	13, 413, 740
R^2	0.0812	0.0124	0.009	0.274
F-Statistic	1469.28	1716.71	1752.57	1780.49

Notes: Standard errors in parentheses. I use robust standard errors clustered at inventor level. Results are robust to clustering at the matched-pair level. I use CEM weights in each regression. All specifications include year and tenure fixed effects. Column (1) includes matched pair fixed effects. Column (2) includes firm fixed effects. Column (3) includes inventor fixed effects. Column (4) includes both inventor and firm fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The results for the main specification in column (1) show that mobility leads to an increase of 0.348 citations per year from inventors at the destination firm, excluding the new hire herself. When inventor fixed effects are included, the estimated effect declines to 0.274. Compared to the average annual citation count received by mobile inventors from the destination firm in the pre-move period (0.03), even the smaller estimate implies more than a nine-fold increase and is therefore economically significant. I therefore conclude that, following the mobility of an inventor, other inventors at the destination firm significantly increase their use of the new hire's pre-move inventions. This suggests that inventor mobility facilitates knowledge flows from the new hire to the employing firm, providing support for my first hypothesis.

6.1.2 Event Study

For the results in Subsection 6.1.1 to be interpreted causally, the parallel trends assumption must hold. However, since my model involves heterogeneous treatment timing, I cannot simply plot the outcome variable for the control and treatment units during the pre-move period and visually assess the validity of the parallel trends assumption. Therefore, I employ an event study design, which allows me to examine whether the two groups exhibit parallel trends during the pre-treatment periods when no treatment has yet occurred.

To conduct the event study, I normalize each inventor–destination firm pair's mobility year to $\tau = 0$, which marks the year of the move. I consider a symmetric event window of 10 years before and 10 years after mobility. For each year relative to the mobility event from $\tau = -10$ to $\tau = +10$, I construct event-time indicators of the form $D_{ij\tau}$, Following the recommendations in Miller (2023), I use end-cap dummies to handle observations that fall outside the event window. In particular, I define a single indicator $D_{ij,\tau<-10}$ to group all pre-event observations earlier than year -10, and another indicator $D_{ij,\tau>+10}$ for all post-event observations beyond year +10. Finally, I exclude the event time $\tau = -1$ from the regression, treating it as the omitted category. This normalizes the event-time coefficients relative to the year immediately preceding the move, which is a common practice in event study designs (Miller, 2023). I provide the event study specifications in Equations (6.1) and (6.2). In both models, I use the matched-pair fixed effects together with the year fixed effects. My results are robust to the use of inventor fixed effects.

$$citation_{ijt} = \alpha_{ij} + \lambda_t + \sum_{\tau=-k}^{-2} \delta_\tau D_{ij\tau} + \sum_{\tau=0}^{m} \delta_\tau D_{ij\tau} + \varepsilon_{ijt} \tag{6.1}$$

$$no\_self\_citation_{ijt} = \alpha_{ij} + \lambda_t + \sum_{\tau=-k}^{-2} \delta_\tau D_{ij\tau} + \sum_{\tau=0}^{m} \delta_\tau D_{ij\tau} + \varepsilon_{ijt} \tag{6.2}$$

I present the event study graphs in Figure 6.1 and Figure 6.2. In both specifications, all the lead coefficients are statistically indistinguishable from zero. This indicates that citation counts exhibit parallel trends between the treatment and control groups in the pre-move period, thereby validating the main identification assumption.

The event study graphs also provide important insights into the dynamics of the treatment effect. As shown in Figure 6.2, the effect of mobility on citation flows gradually increases over eight consecutive post-move years before beginning to decline. This pattern suggests that the diffusion of knowledge across the hiring firm unfolds gradually and takes time. Furthermore, the treatment effect remains significant even ten years after the move, suggesting that the knowledge of the new hire diffuses across the firm and becomes a part of the organization's broader knowledge base.

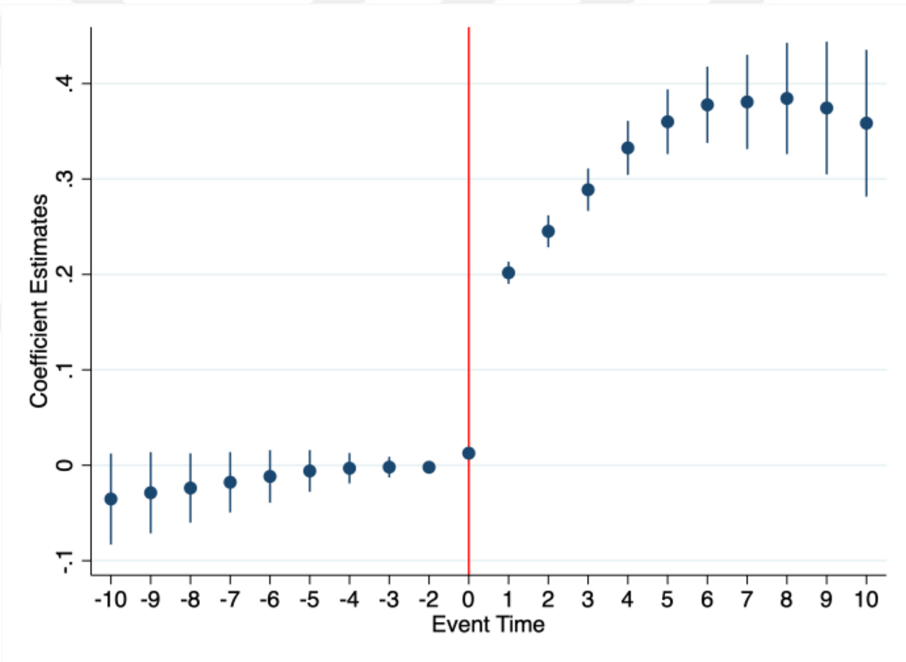

Figure 6.1: Event Study for Checking the Parallel Trends with Citation as the Dependent Variable

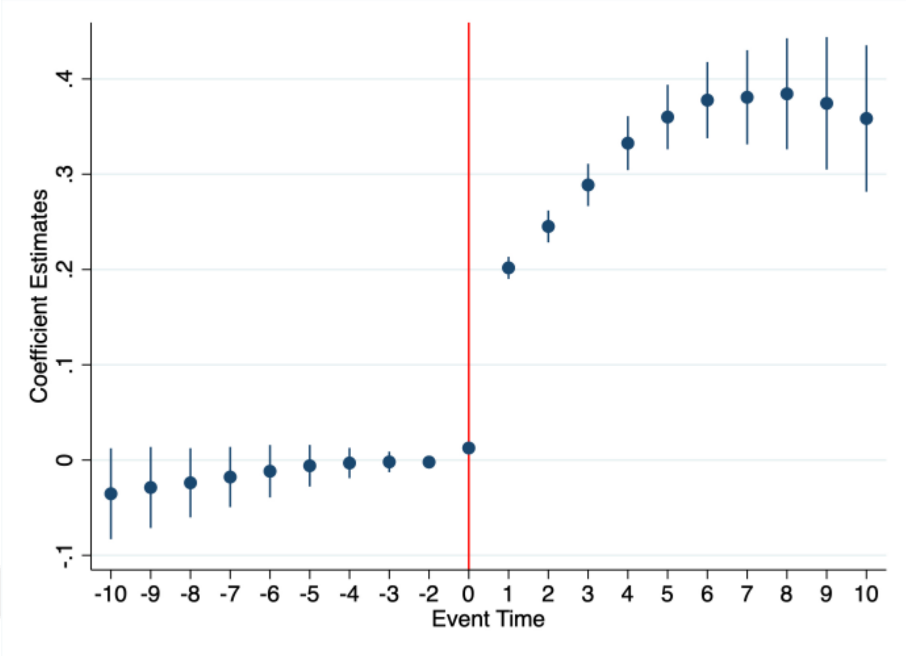

Figure 6.2: Event Study for Checking the Parallel Trends with No Self Citation as the Dependent Variable

6.2 Testing the Second Hypothesis

6.2.1 Regression Results

The results presented in Section 6.1 verify the learning-by-hiring hypothesis and demonstrate that inventor mobility facilitates knowledge transfers from the new hire to the employing firm. This finding provides the foundation for examining whether the America Invents Act (AIA) has influenced the extent to which mobility of inventors translates into knowledge spillovers. Table 6.3 presents the DiDiD estimation results designed to assess this effect, while Table 6.4 reports the 95% confidence intervals for the baseline specification to support the interpretation of the findings. Since my interest is in the effect of the AIA on knowledge transfers through inventor mobility I use No Self Citation as the dependent variable in all specifications. Similar to the previous analysis, I explore various fixed effect structures across columns (1) to (4). However, my main model is the specification in column (1), which includes matched-pair fixed effects, while the remaining specifications are considered for robustness.

Before turning to the parameter of interest, I first highlight several noteworthy findings. The main specification in column (1) finds no statistical evidence that the coefficient on Mobile differs from zero. Although this is not the case in column (2), the estimated coefficient is substantially smaller and much closer to zero than in the corresponding DiD specification. Similarly, unlike in the DiD model, the estimated coefficient on the interaction term Mobile×Post remains remarkably stable across specifications, even after the inclusion of inventor fixed effects. These results may suggest that the DiDiD specification better addresses inventor-level unobserved heterogeneity than its DiD counterpart, perhaps due to the higher-order differencing structure. However, it is difficult to determine whether the results reflect improved control for unobserved heterogeneity, and one should be cautious in drawing definitive conclusions.

Another noteworthy result is, the coefficient estimate on Move2013×Mobile is statistically indistinguishable from zero, suggesting that pre-move differences between the treatment and control groups are not statistically different across the pre- and post-2013 periods. Finally, the statistically significant and positive estimate for the coefficient on the interaction term Move2013×Post indicates that the differential in the outcome for the control group between the pre- and post-move periods is statistically different across the pre- and post-AIA periods, and is larger in the latter.

Having discussed these secondary findings, I now turn to the coefficient on the triple interaction term Move2013×Mobile×Post, which is the parameter of primary interest. As discussed in Section 5.2, this term captures the differential effect of inventor mobility on knowledge transfers from the new hire across the pre- and post-AIA periods. In all specifications, from column (1) to column (4), the estimates of the DiDiD coefficient are statistically significant at the 1 percent level and around -0.38. This indicates that, on average, inventor mobility events that occurred in the post-AIA period result in 0.38 fewer additional citations per year from the destination firm to the new hire's prior inventions, compared to the pre-AIA period. Considering that the coefficient estimate on Mobile×Post, which reflects the additional annual citations attributable to inventor mobility in the pre-AIA period, is approximately 0.44 across all specifications, the estimated reduction of 0.38 citations is economically significant. Therefore, my results suggest that the AIA signifi-canlty reduced the extent to which inventor mobility facilitates knowledge transfers from the new hire to the employing firm, thereby supporting my second hypothesis.

Table 6.3: DiDiD Estimation Results for the Effect of AIA on Knowledge Transfers Through Inventor Mobility

	Dependent Variable: No Self Citation			
	(1)	(2)	(3)	(4)
Regression Model:	Pair FE	Firm FE	Inventor FE	Inv. & Firm FE
`Mobile`	0.009	0.009***	—	—
	(0.006)	(0.002)	—	—
`Post`	0.151***	0.155***	0.143***	0.145***
	(0.007)	(0.005)	(0.006)	(0.006)
`Move2013`	—	-0.072***	0.167***	-0.058
	—	(0.004)	(0.031)	(0.063)
`Move2013×Mobile`	-0.004	-0.004	—	—
	(0.007)	(0.003)	—	—
`Move2013×Post`	0.090***	0.097***	0.099***	0.098***
	(0.008)	(0.006)	(0.007)	(0.007)
`Mobile×Post`	0.449***	0.449***	0.433***	0.435***
	(0.028)	(0.029)	(0.028)	(0.028)
`Move2013×Mobile×Post`	-0.387***	-0.387***	-0.381***	-0.382***
	(0.033)	(0.033)	(0.032)	(0.032)
Number of Observations	13,413,723	13,413,723	13,413,723	13,413,723
R^2	0.081	0.013	0.270	0.275
F-Statistic	797.240	1146.720	857.900	874.970

Notes: Standard errors in parentheses. Robust standard errors clustered at the inventor level. Results are robust to clustering at the matched-pair level. All specifications include year and tenure fixed effects. Column (1) includes matched pair fixed effects. Column (2) includes firm fixed effects. Column (3) includes inventor fixed effects. Column (4) includes both inventor and firm fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6.4: Coefficient Estimates and 95% Confidence Intervals for DiDiD Model with Matched-Pair Fixed Effects

Variable	**Coefficient**	**Std. Error**	**95% Confidence Interval**
`Mobile`	0.009	0.006	[–0.0028, 0.0208]
`Post`	0.151	0.007	[0.1373, 0.1647]
`Move2013` × `Mobile`	–0.004	0.007	[–0.0177, 0.0097]
`Move2013` × `Post`	0.090	0.008	[0.0744, 0.1056]
`Mobile` × `Post`	0.449	0.028	[0.3942, 0.5038]
`Move2013` × `Mobile` × `Post`	–0.387	0.033	[–0.4517, –0.3223]
Fixed Effects	Matched Pair, Year, and Tenure		

Notes: Robust standard errors clustered at the inventor level. All estimates come from the specification using matched-pair, year, and tenure fixed effects. Dependent variable is *No Self Citation*.

6.2.2 Event Study

As shown by Olden and Men (2022), the DiDiD estimator can be computed as the difference between two DiD estimators. Although this is not a necessary condition, if the parallel trends assumption holds for both of the underlying DiD estimators, it will also hold for the triple-difference estimator. Based on this, to assess the validity of the parallel trends assumption for my DiDiD model, I examine whether it holds for each of the two underlying DiD specifications separately. To do so, I estimate the event study regression specified in Equation (6.3) separately for mobility events that occurred in the pre- and post-reform periods.

$$no\_self\_citation_{ijt} = \alpha_{ij} + \lambda_t + \sum_{\tau=-k}^{-2} \delta_\tau D_{ij\tau} + \sum_{\tau=0}^{m} \delta_\tau D_{ij\tau} + \varepsilon_{ijt} \tag{6.3}$$

The structure of the event study analysis is similar to that in Subsection 6.1.2. The main difference lies in the adjustments made to the event window for the post- and pre-AIA subsamples, due to data limitations. For the post-AIA sample, there are no observations extending 10 years after the mobility event, and the highest estimable lag is event time +8. Accordingly, the event window for this sample spans from 10 years before to 8 years after the mobility event. In contrast, for the pre-AIA sample, there are no observations 10 years prior to the mobility event, and the highest estimable lead is event time -8. As a result, the event window for the pre-AIA sample spans from 8 years before to 10 years after the mobility event.

I present the event study graphs for the pre- and post-2013 periods in Figures 6.3 and 6.4, respectively. The graphs show that the initial jump in citations at the event year is comparable across the two periods—rising from approximately zero to 0.18 in the pre-AIA period and to 0.22 in the post-AIA period. However, interestingly, the persistence of the mobility effect on knowledge transfers observed in the pre-reform period appears to diminish substantially following the AIA. An immediate explanation for this pattern could be that inventors switch firms more frequently in the post-AIA period, thereby reducing the opportunity for their knowledge to diffuse within any single employing firm. If that were the case, we would expect the average number of moves per mobile inventor to be higher in the post-AIA period compared to the pre-AIA period. However, unless there was a sharp decline in the labor supply of mobile inventors, Figure 4.2 suggests the opposite: it shows that the total number of mobility events drops markedly after 2013. Nevertheless, for a more complete picture, it is necessary to compute the average employment duration in the pre- and post-2013 periods.

More importantly, although the parallel trends assumption appears to hold for the post-2013 period, it does not hold for the pre-reform DiD specification, as evidenced by lead coefficients that are significantly different from zero. Citations from the destination firm to mobile inventors begin to rise gradually in the years leading up to the mobility event, relative to their matched immobile counterparts. The presence of these pre-treatment trends in the pre-2013 period calls for caution when interpreting the triple-difference estimate in Table 6.3 as causal.

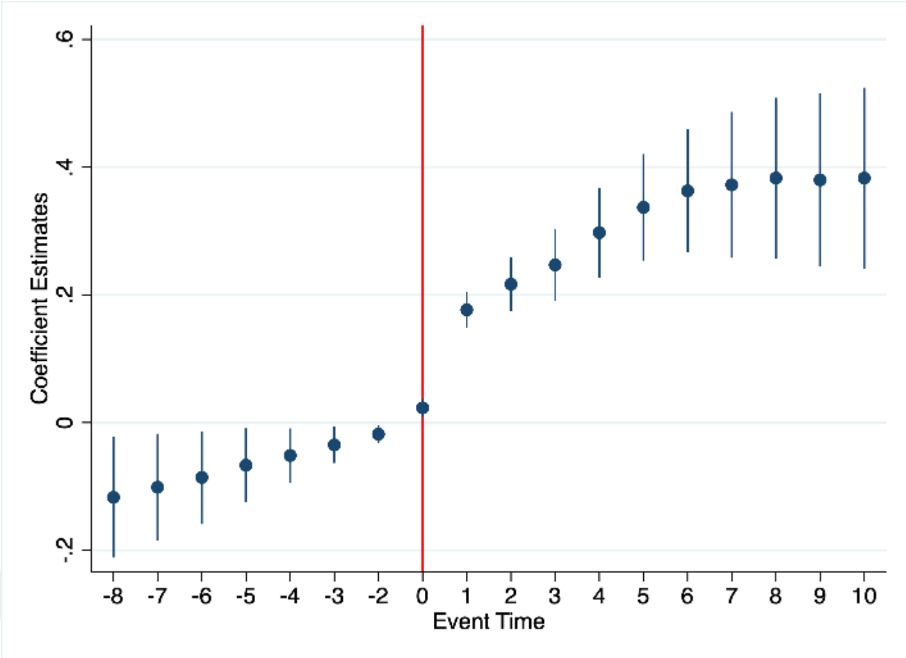

Figure 6.3: Event Study Graph for Mobility Instances in the Pre-2013 Period

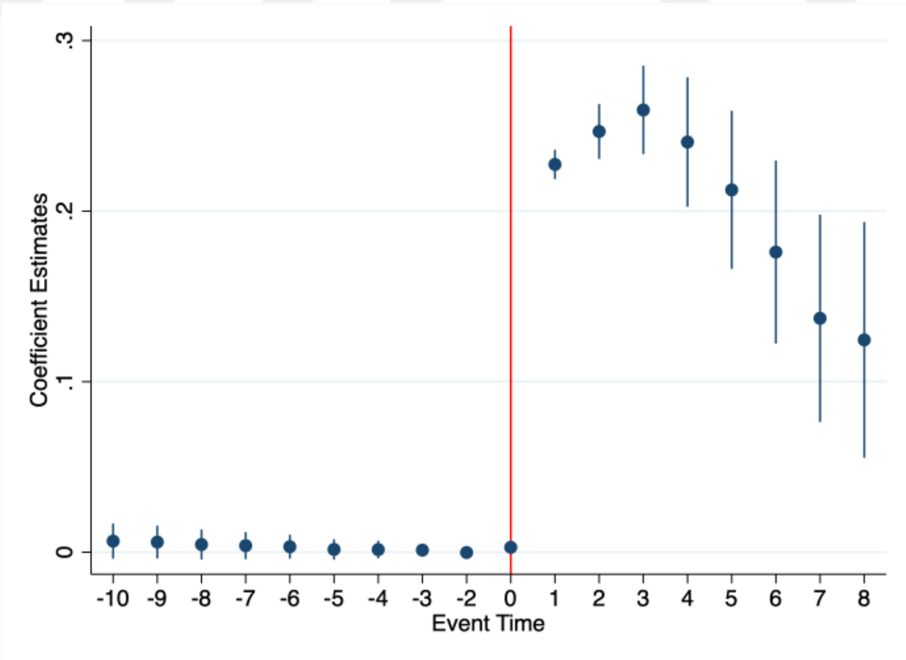

Figure 6.4: Event Study Graph for Mobility Instances in the Post-2013 Period

Chapter 7

SUMMARY AND CONCLUDING REMARKS

On September 16, 2011, President Barack Obama signed into law the Leahy-Smith America Invents Act, marking the most significant change to the U.S. patent system since the Patent Act of 1952. The central provision of the reform was the transition from the century-old first-to-invent patent awarding system to a first-to-file rule. Proponents of the reform argued that this shift would simplify the patenting process, lower administrative costs for the USPTO, and most importantly, encourage inventors to file applications more quickly. This, in turn, was expected to lead to the earlier disclosure of new technological advancements and accelerate follow-on innovation. While prior research has examined whether the policy achieved its goals by analyzing firms' R&D intensity choices and patenting behavior, its impact on the labor market for inventors remains largely unexplored, leaving the overall evaluation of the reform incomplete. This thesis addresses that gap by focusing on mobile inventors—those who switch firms—and examines the effect of the AIA on the extent to which inventor mobility facilitates knowledge spillovers.

Before examining the effect of the AIA, I first empirically test whether inventor mobility facilitates knowledge transfers, a prerequisite channel I refer to as the Learning-by-Hiring Hypothesis. To test it, I identify mobile inventors using USPTO patent data and find approximately 920,000 instances of mobility between 2007 and 2018. To address the differences in observable characteristics between mobile and immobile inventors, I employ a Coarsened Exact Matching (CEM) procedure and construct a control sample of non-mobile inventors. Then, using a difference-in-differences (DiD) approach, I show that inventor mobility leads to an increase of 0.46 citations per year from the destination firm to the pre-move inventions of the new hire. This suggests that hiring firms increase their use of new recruits' pre-move inventions after employing them. However, the impact of inventor mobility extends beyond direct use: once I exclude self-citations by the mobile inventor, my most conservative DiD estimate implies that inventor mobility leads to an increase of 0.274 citations per year from other inventors at the destination firm to the inventions of the new hire. This suggests that, following the move, mobile inventors' new colleagues at the destination firm increase their use of the new hire's pre-move knowledge, providing evidence of knowledge transfers through inventor mobility and confirming my first hypothesis. Finally, to examine the dynamics of knowledge transfers through mobility, I conduct an event study. I document that the effect of mobility on citation flows gradually increases over eight consecutive years after the inventor moves to the destination firm. Moreover, even ten years after the move, the effect remains significant, suggesting that the tacit knowledge of the new hire diffuses throughout the hiring firm and becomes part of its broader knowledge base.

Once I verify that inventor mobility leads to knowledge spillovers, I turn to testing the effect of the AIA on the extent to which mobility facilitates knowledge transfers. In particular, I argue that the America Invents Act decreased knowledge transfers through inventor mobility, and I refer to this hypothesis as the AIA Effect. To test this hypothesis, I employ a Difference-in-Difference-in-Differences (DiDiD) design, which effectively compares the causal effect of inventor mobility on knowledge transfers before and after the transition to first-to-file rule. My estimation results indicate that inventor mobility events occurring in the post-AIA period result in 0.38 fewer additional citations per year from the destination firm to the new hire's prior inventions, compared to the pre-AIA period. Considering that the same DiDiD model predicts an increase of 0.44 citations due to mobility for the pre-AIA period, the estimated post-AIA reduction of 0.38 citations is economically significant. These findings suggest that the AIA significantly reduced the extent to which inventor mobility facilitates knowledge transfers from the new hire to the employing firm.

My findings have several implications for policymakers. The results on the Learning-by-Hiring hypothesis suggest that there are benefits from inventor mobility in the form of knowledge spillovers. Therefore, any policy that restricts inventor mobility could have a negative effect on technological progress by making it more difficult and costly for firms to access external knowledge. On the other hand, knowing that inventor turnover spreads knowledge—potentially to competitors—firms may become more cautious about investing in human capital in the first place. Therefore, optimal policy on restricting inventor mobility depends on which effect dominates the other.

The validity of the AIA Effect, on the other hand, points to an unintended consequence that runs counter to the reform's objectives: As I document, the AIA weakened an important channel -inventor mobility- through which new ideas and knowledge spread across the innovation landscape. This may have detrimental effects on U.S. innovative activity by making it more costly and difficult for firms to access external knowledge. Nevertheless, caution is warranted when interpreting these findings in a policy context. While the results show that the AIA reduced knowledge transfers through inventor mobility, this may not be the primary channel through which the reform affects innovation. That is, rather than reduced knowledge transfers harming aggregate inventive activity, it is possible that the AIA first led to a decline in firms' innovative activity through a different primary mechanism, which then resulted in fewer knowledge transfers via inventor mobility. Clarifying this ambiguity requires future research to empirically identify and test the mechanisms explaining the observed reduction in knowledge spillovers.

BIBLIOGRAPHY

Abrams, D. and Wagner, R. (2013). Poisoning the next apple? the america invents act and individual inventors. *Stanford Law Review*, 65:517–563.

Almeida, P. and Kogut, B. (1999). Localization of knowledge and the mobility of engineers in regional networks. *Management Science*, 45:905–917.

Babina, T., Bernstein, A., and Mezzanotti, F. (2020). Crisis innovation. Working Paper 27851, National Bureau of Economic Research.

Boucher, P. M. (2012). Recent developments in us patent law. *Physics Today*, 65(1):27–32.

Brown, B. L. (2020). The bell versus gray telephone dispute: Resolving a 144-year-old controversy [scanning our past]. *Proceedings of the IEEE*, 108(11):2083–2096.

Halac, M., Kartik, N., and Liu, Q. (2017). Contests for experimentation. *Journal of Political Economy*, 125(5):pp. 1523–1569.

Hoisl, K. (2007). Tracing mobile inventors—the causality between inventor mobility and inventor productivity. *Research Policy*, 36(5):619–636.

Hounshell, D. A. (1981). Two paths to the telephone. *Scientific American*, 244(1):156–163.

Huang, R., Li, L., Lu, L. Y., and Wu, H. (2021). The Impact of the Leahy-Smith America Invents Act on Firms' R&D Disclosure. *European Accounting Review*, 30(5):1067–1104.

Iacus, S., King, G., and Porro, G. (2008). Matching for causal inference without balance checking. *Universit degli Studi di Milano, UNIMI - Research Papers in Economics, Business, and Statistics*, 20.

Lemley, M. A. and Chien, C. V. (2003). Are the u.s. patent priority rules really necessary? *Hastings Law Journal*, 54:1299–1332. UC Berkeley Public Law Research Paper No. 124.

Lerner, J., Speen, A., and Leamon, A. (2015). The leahy-smith america invents act: A preliminary examination of its impact on small businesses. Technical report, Bella Research Group, Salem, MA. Prepared for the U.S. Small Business Administration, Office of Advocacy. Contract number SBAHQ-14-Q-0011.

Lo, S. and Sutthiphisal, D. (2009). Does it Matter Who Has the Right to Patent: First-to-invent or First-to-file? Lessons From Canada. NBER Working Papers 14926, National Bureau of Economic Research, Inc.

Mansfield, E. (1988). Industrial r&d in japan and the united states: A comparative study. *The American Economic Review*, 78(2):223–228.

Miller, D. L. (2023). An introductory guide to event study models. *The Journal of Economic Perspectives*, 37(2):pp. 203–230.

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

Miyagiwa, K. and Ohno, Y. (2015). International harmonization of the patent-issuing rules. *International Journal of Industrial Organization*, 39:81–89.

Mueller, W. F. (1962). The Origins of the Basic Inventions Underlying Du Pont's Major Product and Process Innovations, 1920 to 1950. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, NBER Chapters, pages 323–358. National Bureau of Economic Research, Inc.

Olden, A. and Men, J. (2022). The triple difference estimator. *The Econometrics Journal*, 25.

Rantanen, J. (2012). Peripheral disclosure. *University of Pittsburgh Law Review*, 74(1):1–52. U Iowa Legal Studies Research Paper No. 11-33.

Rosenkopf, L. and Almeida, P. (2003). Overcoming local search through alliances and mobility. *Management Science*, 49(6):751–766.

Scotchmer, S. and Green, J. (1990). Novelty and Disclosure in Patent Law. *RAND Journal of Economics*, 21(1):131–146.

Singh, J. and Agrawal, A. (2011). Recruiting for Ideas: How Firms Exploit the Prior Inventions of New Hires. *Management Science*, 57(1):129–150.

Song, J., Almeida, P., and Wu, G. (2003). Learning-by-hiring: When is mobility more likely to facilitate interfirm knowledge transfer? *Management Science*, 49(4):351–365.

Varshney, M. (2023). Learning-by-hiring: How do rival firms learn from focal firm's hiring. *Research Policy*, 52(2):104664.

Yilmaz, M. M., Akdoğu, E., and Simsir, S. A. (2024). Post-aia innovation acquisition: A race to the patent office or a race to acquire patents?