Essays on the Direction of Technical Change

by

Bernardo Dionisi

Department of Business Administration
Duke University

Date: _____

Approved:

_____
Ashish Arora, Supervisor

_____
Sharique Hasan

_____
Colleen Cunningham

_____
Elia Ferracuti

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Business Administration
in the Graduate School of
Duke University

2023

ABSTRACT

Essays on the Direction of Technical Change

by

Bernardo Dionisi

Department of Business Administration
Duke University

Date: _____

Approved:

_____
Ashish Arora, Supervisor

_____
Sharique Hasan

_____
Colleen Cunningham

_____
Elia Ferracuti

An abstract of a dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Business Administration
in the Graduate School of
Duke University

2023

Abstract

The rate and direction of inventive activity are central to firms' competitiveness as well as economic growth. A fundamental question arises from the fact that innovation increasingly relies on public scientific findings, many of which are freely published and accessible. How can firms capture private value from publicly available knowledge? Chapter 2 investigates the concept of first-mover advantage in utilizing public science. It finds that being first to apply cutting-edge public science enables broader patents, but requires active internal R&D to recognize opportunities early. Chapter 3 surveys the extensive literature employing quasi-experimental techniques to quantify forces shaping the direction of innovation. It contributes software tailored for innovation data to assist future research. Chapter 4 leverages FDA data to empirically analyze post-acquisition innovation trajectories in medical devices, revealing the slowed evolution of acquired technologies.

Overall, this dissertation elucidates how firms derive competitive advantage from public science, surveys techniques for quantifying innovation's direction, and empirically examines merger impacts on medical technology advancement. These studies contribute novel data, methods, and insights to support researchers as well as firms, policymakers, and other stakeholders with an interest in the evolution of technical change.

Contents

List of Tables

List of Figures

Acknowledgements

I would like to express my gratitude to my advisor, Ashish Arora, for his guidance and support throughout these past several years. His knowledge, patience, and wisdom were a stable reference during easy and difficult times along the PhD path.

I am very grateful to my committee members Sharique Hasan, Colleen Cunningham, and Elia Ferracuti who have been a reassuring and positive force, and whose scholarship I look up to. Many thanks go also to the Fuqua Strategy group, which welcomed me and treated me as a close member.

A big thank you to the many past and current PhD fellows, within Fuqua and across the Duke community, for insightful and inspiring conversations and for the fun moments.

Finally, I would like to thank my brother Filippo, and my wife Alessia, for their support and constant love.

Chapter 1

Introduction

Understanding the drivers of technological innovation and advancement is a fundamental concern in the fields of economics, business strategy, and public policy. A large body of research emphasizes the importance of the rate and direction of technical change as key factors influencing economic growth and competitiveness (Rosenberg, 1969). Recent studies highlight the crucial role that science plays in enabling technical progress across many industries, as scientific knowledge provides fundamental inputs that spur new inventions and technologies (Arora et al., 2018, 2019; Fleming & Sorenson, 2004). However, a critical question arises regarding how actors, especially firms, are able to capture private value from public science to generate innovative outputs.

Chapter 2, co-authored with Ashish Arora and Sharon Belenzon, examines the relationship between scientific knowledge, much of which is publicly available, and the ability of firms to derive private value and competitive advantage. Chapter 2 notes that technical progress increasingly relies on scientific advances, many of which are published and accessible public knowledge. However, we question how firms can create proprietary value from knowledge inputs that are freely available to all competitors. The study finds that the average private returns from utilizing public science are relatively small, especially in crowded technical fields where many firms can access the same knowledge. However, private value generated is higher when a firm is the first to use a particular scientific breakthrough, as early movers can secure broader patents compared to later users, in technical areas where competition is high. To consistently recognize relevant scientific advances first requires active participation in scientific research itself, which raises firm familiarity with innovations before competitors. The study shows that corporate participation in basic science strongly predicts the first use of new public knowledge. This provides evidence that firms who invest in internal R&D are not just more likely to utilize public science, but are better positioned to extract private value by being among the earliest appliers of new public knowledge. Overall, the paper argues that firms can turn universally accessible scientific inputs into privately valuable resources

by moving first, especially in crowded areas.

Chapter 3, focusing on the direction of inventive activity, provides a survey of the literature on the economics of innovation with a focus on papers that employ difference-in-differences and event studies using innovation datasets, particularly patent-related data. I notice that many papers in the innovation domain structure the data following common patterns, thus I delineate these patterns, with an emphasis on the unit of analysis. I categorize the units or entities of analysis into patent classes, patents, inventors, and firms. I believe this codification can help researchers, especially early-stage researchers, in understanding how innovation scholars provide insights into the forces that shape the rate and direction of technical change. Furthermore, given the increasing availability of openly available datasets in the domain of innovation studies [1], we are then able to design software tailored around these datasets[2]; thus, along with a survey, and categorization of the literature, I contribute software that facilitates the analysis described.

Chapter 4 leverages the concepts and empirical design discussed in Chapter 3, along with a unique dataset I compiled from 510k submissions to the FDA. I investigate post M&A follow-on innovation, with a focus on the medical device industry. The institutional setting of the medical device industry allows me to track product continuation outcomes by leveraging the predicate device tree that I built. FDA guidelines require firms to document and demonstrate substantial equivalence of their class 2 medical device applications, and I use these applications to map medical device products over generations. Leveraging this newly constructed dataset and the observation that mergers and acquisitions are on the rise in the evolving American innovation ecosystem, I investigate the post-acquisition follow-on innovation of medical technologies. The findings reveal slowed evolution of acquired devices, though acquisitions seem to highlight quality devices ready for further advancement by external sources. This empirical research illuminates the complex interplay between acquisitions and technological progression in the medical device sector.

[1] Examples of which are PatentsView (USPTO, 2019), which indexes the universe of USPTO patents post 1976, as well as Reliance On Science, a newly developed database linking patents and scientific papers (Marx & Fuegi, 2019)

[2] As described in Chapter 3 the functionality is (at the time of writing) based on the listed sources: PatentsView, RelianceOnScience as well as data made available by the FDA under openFDA.

Chapter 2

First-mover advantage and the private value of public science

"... knowledge is regarded by economists as being 'on the shelf' and costlessly available to all comers once it has been produced. But ... it frequently requires a substantial research capability to understand, interpret and to appraise knowledge that has been placed upon the shelf ... . The cost of maintaining this capability is high, because it is likely to require a cadre of in-house scientists who can do these things. And, in order to maintain such a cadre, the firm must be willing to let them perform basic research. The most effective way to remain effectively plugged in to the scientific network is to be a participant in the research process. ...basic research may be thought of as a ticket of admission to an information network" (Rosenberg, 1990: 170).

2.1 Introduction

Technical progress and scientific advance are closely interlinked (Cohen et al., 2002; Mansfield, 1995). In many, if not most, sectors of the economy, new technologies are increasingly scientific in nature (Fleming et al., 2019; Narin et al., 1997). The percentage of utility patents that cite science has increased from approximately 6% in 1980 to 30% in 2015, and the average number of citations to science per patent has increased from 0.1 to 4.4 over the same period.[1] Much of this scientific knowledge is published and available for all to use. However, the resource-based view of the firm argues that creating private value requires idiosyncratic resources that are not reproducible by others and cannot be easily transferred (Barney, 1986). A challenge for firms, therefore, is how to create value from public science, which is of uncertain relevance and freely available to all.

In this paper, we show that one way for firms to capture value from public science is to use it before the competition, i.e., to be a first-mover (Lieberman & Montgomery, 1998). To be a successful first-mover, the firm has to participate in science to evaluate which scientific advances are relevant. Simply put, firms' engagement in basic research is a source of a "first-mover advantage,"

[1] Authors' calculations use front-page citations, with a confidence score equal or above 7, from the "Reliance on Science" dataset (Marx & Fuegi, 2019).

because participation in basic research increases a firm's ability to recognize and apply relevant extramural findings before others (Cohen & Levinthal, 1989; Rosenberg, 1990).

We quantify the average private value of inventions that use science, how the private value differs between early and late users, and how it depends on the number of potential users. Consistent with competition eroding the returns from the use of a common resource, the greater the number of potential users, the greater are the relative rewards of being an early user. Second, we explore a specific mechanism for potential first-mover advantage, namely the scope of patent protection obtained by the first- movers.[2] Third, we explore how being a first-mover is related to participation in scientific research by the firm, thus connecting firm-level enablers and environmental dynamics shaping first-mover benefits (Suarez & Lanzolla, 2007).

We make two contributions to the innovation and strategy literature. We theorize that firms can turn an input available to all into a privately valuable resource by moving first, thereby preempting rivals. This first-mover advantage is more salient in more crowded niches. Further, we provide evidence for the widely theorized, but under-documented, idea that firms that produce public research are also more likely to be among its *early* users. Thus, firms that invest in internal research are not just more likely to use science but are also better able to extract private benefits from using public science. Simply put, one source of returns to corporations from investing in scientific research is the early identification and use of relevant scientific discoveries. Our study focuses on the rival uses of science, wherein the use of science by one firm reduces its private value to another. However, firms may benefit from science in other ways that competition is less likely to erode, such as a reduction in the cost of invention.

On the methodological front, we leverage newly developed data (Kogan et al., 2017a) on the

[2]Bikard and Marx, 2020 provide an insightful example on how Amgen secured crucial intellectual property being the first to build on the breakthrough purification of erythropoietin (EPO) by Professor Goldwasser at the University of Chicago. In a paper published in 1977, Goldwasser showed that EPO stimulated production of red blood cells. Using the purified samples of EPO provided by Professor Goldwasser, Amgen scientists could learn about the structure of EPO and identify the relevant genes. This knowledge enabled Amgen to produce EPO using recombinant DNA techniques, ahead of its rivals, and obtain the vital patent rights that underpinned Amgen's subsequent commercial success.

value of patents to measure the private economic return to the specific invention, distinct from the technical quality of the patent. In so doing, we respond to the call by Lieberman and Montgomery, 2013 for forward-looking, rather than retrospective, measures of the benefits of being a first-mover, and for measures of economic return rather than market-share or survival (Lieberman & Montgomery, 1988). However, our study does not extend to entry into the product market but is limited to the pre-entry stage of invention.

We next discuss our findings and contribution in the context of the different literature that we draw upon and to which we contribute.

2.2 Prior literature and theoretical considerations

Though the two are related, it is helpful to distinguish between using science from engaging in science. Our paper is about the private value of using science, but we also show that engaging in science enables a firm to be an early user of science. The two main strands of literature we draw upon deal with the value of science and first-mover advantage. In addition, we draw from the large literature on how firms acquire absorptive capacity to access external knowledge (in our case, be early users of science).

2.2.1 Value of science

Scientific knowledge can benefit companies in a variety of ways. Scientific advances may directly lead to innovations (e.g., new drugs,) and innovations may sometimes arise as indirect outputs of scientific research (e.g., laser). As the literature on spillovers makes clear some of the ways in which firms benefit from science are rival-in-use. For instance, Arora et al., 2021 show that the private value to corporations of producing scientific research is lower when rivals cite the science in their patents: Innovations introduced by rivals compete with those of the focal firm, especially if they draw upon the same scientific discovery. Arguably there are other uses less susceptible to rivalry, at least in the short term. For instance, scientific knowledge can benefit the firm by improving R&D efficiency (Griliches, 1985), by guiding technical search toward more fruitful pastures (Evenson

& Kislev, 1976; Fleming & Sorenson, 2004; Gambardella, 1995), and by increasing absorptive capacity (Leten et al., 2022). By so doing, science can effectively reduce the cost of invention. However, the reduction in cost is not affected by whether others use the same scientific knowledge, as long as the other uses do not result in inventions that compete with those of the focal firm.[3] That is, the value generated by these other uses of science should not diminish with competition, and first-use therefore should not be valuable.

Notice also that competing inventions need not be from rivals in the product market if the rivals can acquire inventions from universities and startups. What matters is that the private value of an invention using a scientific discovery is lower if rivals get access to inventions that use the same discovery. In some cases, the presence of competition may even enhance the value derived from these sources, for instance as more organizations compete within an industry, the rate of knowledge spillovers may increase (Arora et al., 2021).

In this paper, we show that the private value of inventions that rely upon science is affected by the number of potential rival users and that first-movers enjoy an advantage over latecomers in terms of appropriating private value from using science. Engaging in scientific research enables a firm to be a first-mover in using science, by increasing its absorptive capability (Cohen & Levinthal, 1989; Rosenberg, 1990). Company scientists can help identify promising new inventions, engage with relevant outside researchers, and help assimilate and adapt outside technology. Being a first mover is privately valuable. However, there are other benefits of engaging in scientific activities which are clearly non-rival. Engaging in scientific research enhances the reputation of the firm and certifies the quality of its research to prospective investors, employees, government agencies, and sophisticated customers (D. Hicks, 1995). Also, to the extent that allowing employees to publish helps firms recruit more talented researchers, participating in the process of advancing science can be a profitable strategy for some firms (Roach & Sauermann, 2010; Stern, 2004).

A recent literature has studied the relationship between the use of science and patent value. A key difference with our work is that we distinguish measures of private value from measures

[3](Colen et al., 2023) document that science-based pharmaceutical inventions are more likely to likely to be developed in-house and more likely to succeed.

of patent quality or social value. Fleming and Sorenson, 2004 find that patents citing scientific prior art receive more follow-on citations, which is a measure of quality or social value. Our paper confirms that finding, but also shows that the mean private value of using science is low. Poege et al., 2019 show a positive relationship between the quality of the publications cited by a patent and the technical quality of the patent. We too find that patents citing articles published in high-impact-factor journals have higher private economic value relative to other patents. Krieger et al., 2022 provide evidence that patents closer to scientific publications tend to have higher private value than patents that are farther from science.[4] We too use the measure of monetary value of a patent developed by Kogan et al., 2017a. Different from Krieger et al., 2022, our objective is to understand how companies can create private value from public science. Accordingly, we take a within-firm approach to analyze how competition erodes the private value of using science and how first-use mitigates the effect of competition.

2.2.2 First-mover advantage

An influential literature has studied the conditions under which first-movers have an advantage, the possible means by which first-movers gain an advantage, and why some firms are more likely to move first e.g., Fosfuri et al., 2013; Lieberman and Montgomery, 1988, 1998. Whereas the first-mover literature has focused on product markets, we focus on first-movers in using a scientific discovery. This has some natural implications for how first-movers gain an advantage.[5] In our context, the first-mover, by being the first to patent, may gain a degree of market exclusivity, and in particular may obtain broader patents. Needless to say, these different sources of advantage are not mutually exclusive. A patent may itself provide the patenting firm with the time to develop the

[4]They define closeness to science following Ahmadpoor and Jones, 2017, so that, for instance, a patent has a distance of 1 from a publication if it cites the publication. It has a distance of 2 from a scientific publication if the patent cites a patent that in turn cites a scientific paper.

[5]Lieberman and Montgomery, 1988 identify three mechanisms whereby first-movers gain advantage: technology leadership, locking in buyers, and preemption of key inputs, or winning a patent race (Fosfuri et al., 2013).

complementary asset while rivals are trying to invent around.

Lieberman and Montgomery, 1998 argue that economic returns are the appropriate measure of first-mover advantage rather than market share or survival, the focus of much of the literature. They argue that ''*First-mover advantages exist when the pioneering firm earns positive present value of profits as the consequence of its early entry (i.e. positive profits net of those attributable to more general types of firm proficiency). A serious problem confronting those engaged in empirical work is the fact that disaggregated profit data are seldom obtainable.*" An advantage of our study is the availability of a proxy for the expected monetary value tied to the use of scientific findings – a measure of stock market returns associated with the patent (Kogan et al., 2017a). In addition to being a measure of anticipated future economic return, albeit with some caveats, it is also specific to the invention rather than an aggregate firm-level measure.[6] In other words, though limited to inventions that are patented, this is a much finer-grained measure of first-mover advantage in the use of science, which allows us to control for unobserved firm-specific effects. Thus, one can ask how economic returns from early versus late use of science differ within the same firm.[7]

Further, our measure allows us to explore how the returns are realized, which has been a challenge for the first-mover literature. We find that science-based patents are of higher technical quality, as reflected by citations received. Importantly, patents that are the *first* to use a scientific discovery also tend to be broader in scope compared to patents that are later users. That is, first-movers are able to gain broader patent protection.[8] Another strand of the literature has examined the capabilities and resources that allow firms to move first. Starting with Mitchell, 1989, the literature has examined a variety of such capabilities, including marketing capabilities (Mitchell, 1991) and specialized assets (Tripsas, 1997), technical capabilities (Klepper, 2002), and manufacturing (Arora & Ceccagnoli, 2006). We complement this literature by pointing to absorptive capacity as a source

[6]Indeed, Lieberman and Montgomery, 2013 did suggest using stock-market capitalization.

[7]Poletti et al., 2008 also use a stock market event-study approach with product market announcements.

[8]Firms also have other means of capturing value. For instance, firms with complementary assets, such as specialized machinery, skilled labor, or unique distribution channels, can enhance the value derived from scientific knowledge. For instance, Hartmann and Henkel, 2020 show that in Artificial Intelligence (AI), computational power and large datasets are key complements to public scientific knowledge.

of a first-mover advantage, especially in competitive niches.

In summary, although we draw upon the literature on first-mover advantage, there are important differences in context and measures. We use a forward-looking measure of economic return, rather than market share or survival. Our measure is tied to the invention itself rather than aggregate profitability of the firm as a whole. On the other hand, we do not observe entry into a market, only the timing of the use of science in invention. Further, we are confined to publicly traded firms.

Investment in research and absorptive capacity

The vast literature on absorptive capacity (Cohen & Levinthal, 1989; Rosenberg, 1990) emphasizes the importance of corporate participation in research for accessing outside knowledge (Berchicci, 2013; Fabrizio, 2009; Sauermann, 2010). While this literature does not typically discuss first-mover advantage directly, it suggests that the firms that invest in basic research are better able to benefit from scientific discoveries that are ostensibly available to all. Baruffaldi and Poege, 2020 show that firms that participate in a conference are more likely to cite a paper presented in that conference compared to papers at comparable conferences. We complement their findings by showing that firms are more likely to be early users when they participate in conferences and use scientific findings presented at the conference the firm attended. Our findings also resonate with the literature that has stressed the importance of pre-entry capabilities (Helfat & Lieberman, 2002) and specifically for direction of innovation (Helfat, 1994).

By combining the absorptive capacity perspective with the first-mover advantage literature, we show that the benefits of absorptive capacity are related to competition. Specifically, we show that being a first-mover is relatively more advantageous when there are more potential competitors, and by implication, the private benefit from absorptive capacity are greater in the face of more competiton. In so doing, we connect the absorptive capacity literature to the strategic competition literature. This is timely in view of the growing division of innovative labor in the American innovation ecosystem (Arora et al., 2018), whereby universities specialize in upstream research and corporations specialize in downstream development. Coupled with increasing reliance on science in invention (Fleming et al., 2019; Marx & Fuegi, 2019), this division means that while using public

science is becoming more important over time, profiting from it is getting harder as firms' scientific capabilities deteriorate.

2.2.3 A simple conceptual framework

Once a firm uses a scientific discovery in its inventions and discloses it by citing it in its patents, both the discovery and its relevance to invention become clearer to others, including firms that may lack absorptive capacity. In other words, once a scientific discovery has been successfully applied by others, even firms without significant absorptive capacity capabilities should be able to use it. On the other hand, only firms investing in science should be first-movers because they are more likely to be aware of the discovery and evaluate its relevance[9]. The first-mover has more time to build or acquire the required inputs for commercialization. However, it can also use the patent itself to carve out a broader exclusive zone, thereby appropriating a larger share of the rents from its invention. This type of preemption of rivals is more valuable the greater the number of rivals. Followers can more cheaply access scientific knowledge shown to be relevant by the first-mover, but must work around the invention claimed by the first-mover. Followers avoid having to invest in absorptive capacity, but may face a higher cost of working around the first-mover's patent.

The history of statin development from Gambardella, 1992 illustrates this idea. After scientific advance had demonstrated that high cholesterol levels were related to heart disease, firms such as Bristol Myers, developed Colestyramine, a compound that removed bile acids from the body, thereby reducing some cholesterol. However, this was far less effective than reducing the production of cholesterol itself. In the 1970s, Brown and Goldstein from the University of Texas elucidated how the body synthesized cholesterol. Thereafter, the search began for molecules that would block one of the 30 steps in the synthesis of cholesterol (Endo, 2010). Merck was the first to isolate lovastatin, the first effective compound in humans. Merck was first in part because its scientists were familiar with mevalonic acid, a key intermediate in the cholesterol pathway (Gambardella,

[9]The growing quantity of published research suggests that knowing where to look and what to use may be crucial (Jinha, 2010). For instance, Bikard and Marx, 2020 show that scientific findings produced close to patent inventors are more likely to be cited in the respective patents.

1992). Following the successful commercial debut of lovastatin (Mevacor) in 1987, other companies entered, but Merck maintained its leadership in the space by introducing simvastatin (Zocor). Parke-Davis (eventually acquired by Pfizer) was able to chemically synthesize a more potent statin, atorvastatin (Lipitor), by relying on molecular modeling of lovastatin. This example illustrates that internal scientific capability is useful for early users of science. However, once the applicability of a scientific discovery is demonstrated, followers need much lower levels of internal capability to apply the science, perhaps bringing to bear other capabilities (Gambardella, 1995).[10]

Competition conditions the private value of using public science and, hence, also conditions the first-mover advantage. With more competition, the number of potential rivals that could use a scientific advance is higher. With more competing uses of the same discovery, the average private returns would be lower. However, though average private returns may be lower, they would differ across firms. Those that manage to establish an advantageous position by obtaining broader patents or gaining control over some scarce complementary resource would have higher returns. In other words, competition would lower average private returns from using public science, but also make it relatively more valuable to be a first-mover.

In the empirical analysis, we investigate whether (i) firms with absorptive capacity are more likely to be first-movers and, (ii) especially for more relevant science. Moreover, we examine whether (iii) first-movers are more likely to obtain broader patents, (iv) especially when citing relevant science.

2.3 Data

We construct a patent-level dataset that includes patents assigned to U.S.-based firms granted between 1980 and 2010. We combine three main datasets: (i) patent value data Kogan et al., 2017a (ii) the DISCERN dataset, linking patents and firms Arora et al., 2017, and (iii) the "Reliance on

[10]For instance, the synthesis of atorvastatin required expertise in chiral chemical synthesis at large scale, whereas lovastatin was extracted from a fungus (Roth, 2002).

Science" dataset, which collects the universe of patent citations to scientific publications Marx and Fuegi, 2019. In addition, we use financial information on firms from CRSP/Compustat, scientific articles from Microsoft Academic (Sinha et al., 2015), and patent-level information from Google Patents. To build our sample of patents we start with patents granted since 1980 and are included in Kogan et al., 2017a. We restrict this sample to patents granted to firms from DISCERN, resulting in 853,382 patents issued by the U.S. Patent and Trademark Office (USPTO) from 1980 to 2010 to 3,767 publicly traded, U.S.-headquartered companies.

2.3.1 Measures

Patent value. (ξ_i) is a forward-looking measure of the monetary value, in 1982 millions of U.S. dollars, of a patent (Kogan et al., 2017a). Each patent value is estimated as a function of the idiosyncratic stock market returns of the company to which the patent is issued in a three-day window around the time the patent is granted, the ex-ante probability of a patent being granted, and a firm's market capitalization at the time the patent is issued.

The average patent value in our sample is $13.9 million in (1982 prices), while the median patent value is $6.4 million. Values vary across industries, ranging from $23.3 million in "Life Sciences", $15 million in "Chemical", "ICT" with an average of $14.7 million, to "Electronics" with an average value of $11 million.[11]

There are notable benefits and drawbacks to this measure. First, ξ_i is only available for public firms. Further, patents are typically issued once a week, so that all patents granted to a firm in a week are assigned the same value. As with all event-study measures, there is also an important question about the time window used to measure the abnormal return. The shorter the window, the more tightly the focal event is linked to the stock price movement on the day, but the less likely it will be that the market price incorporates all relevant information. Finally, as with event-studies, one has to make assumptions about what investors anticipated. In particular, the measure used assumes that investors had the same prior probability that the patents would be granted. That is, the

[11] Categories are from the NBER patent classifications. We winsorize patent values at the 1 and 99 percentiles. The results are not sensitive to winsorization.

ex-ante probability of a patent being granted is assumed to be the same across patents. However, we find similar results when we aggregate up to the firm-week level.[12] This suggests that the bunching of patent issuance on the Tuesday of a week is not biasing our results.

In addition, there may be important private benefits from using science that are not reflected in the value of the patent.[13] For instance, if science reduces the cost of invention, this increases profits. However, the cost of invention is sunk when the patent issues and therefore is not reflected in the value of the patent. These benefits to the use of science are non-rival – they are not sensitive to competition.[14] Perhaps most relevant for our paper, these benefits of using science should not depend on being a first-mover. Using science to reduce invention costs typically involves avoiding dead-ends and choosing more promising lines of attack. It is difficult for rivals to preempt a firm from using science in this way. In other words, our measure seems well-suited for investigating the rival uses of science – where not being able to exclude others from using science reduces the payoff to the focal firm – and how competition conditions the value of being a first-mover in this case.

This patent-level measure allows us to compare the value of inventions that cite science with inventions by the same firm that do not cite science. This is important because first-mover advantage is often measured empirically, using profit, market share, or business survival; these are backward-looking measures, aggregated to the firm level, and subject to a variety of selection biases.[15] Additionally, what matters for our results is not the absolute values, but whether the *relative* patent values are meaningfully measured. Indeed, we find similar results when using a different measure of private value, namely whether the patent is maintained to its full term. This measure, though less fine-grained, reflects the judgment of managers rather than that of investors, and does not require assumptions about whether the probability a patent application is successful.

Patent renewals As an alternative measure of patent economic value, we construct a variable indi-

[12] Results available on request.

[13] We are grateful to the reviewers for highlighting this issue.

[14] Arguably, if all firms could reduce invention costs, in the long run this also ought to reduce the gross value of invention as more firms introduce competing inventions.

[15] Measures derived from changes in equity prices have been used by Lieberman, 2005 to study first-mover advantage in the Internet sector.

cating whether the patent fees of the focal patent have been paid at the 12-year deadline. We use payments of maintenance fees as robustness over the stock market measure. Before the availability of KPSS, patent renewal data were used by scholars as a measure of private patent value (Schankerman & Pakes, 1986). If a patent holder determines that the costs of maintaining a patent outweigh the potential economic benefits, they would abandon the patent, allowing the underlying invention to enter the public domain. Patent fees can help ensure that valuable patents are maintained and protected, while less economically viable ones are allowed to expire.

Use of science. We follow the growing literature on the use of non-patent literature (NPL) cited by a patent as indicating that the patent used science. The "Reliance on Science" dataset developed by Marx and Fuegi, 2019 is a citation-level dataset that assigns to each patent an NPL citation string matched, with various degrees of confidence, to a scientific publication in Microsoft Academics.[16]

While patent citations, specifically patent-citations-to-patents, have been criticized as a measure of knowledge flows (Jaffe & Trajtenberg, 2002; Roach & Cohen, 2013), recent evidence suggests that patent citations to scientific papers suffer less from the drawbacks afflicting patent citations to patent. Specifically, patent citations to papers are plausible measures of inventor awareness and use of scientific findings.

Arora et al., 2017 validate patent citations to science as a measure of knowledge flows using responses of industrial R&D lab managers to the Carnegie Mellon Survey (Cohen et al., 2000). Their validation exercise indicates a positive relationship between citations to science made by a firm's patents and the firm's reported use of science in its R&D process. Bikard and Marx, 2020, in an attempt at understanding the process underlying patent citations to scientific literature, interview 21 inventors. In the interviews, the inventors confirmed their direct involvement in the addition of citations to science to their patents, implying that the citation reflected their awareness of the cited publication. Furthermore, in a large-scale analysis, Arora, Belenzon, and Suh, 2020 propose a novel approach to understanding the relationship between patent citations to science and academic literature. They show a positive relationship between technological classes citing science and the

[16]We use front-page NPL citations. We use the matched sample with a confidence score above 7. See Marx and Fuegi, 2019 for details.

average textual distance between the text of patents in the citing technological class and academic literature abstracts. Technological classes that are textually close to the academic literature have a higher propensity to cite scientific papers.

We create a binary indicator – *Science* – that is equal to one if the focal patent contains a front-page NPL citation linked to a scientific publication and zero otherwise. Approximately 26% of patents in our sample cite science. On average, a patent cites about two scientific publications, whereas 18% cite more than 1 publication. About 67% of the firms in our sample have at lease one patent that cites science.

We classify different types of citations to science using a common measure of journal quality, the Journal Impact Factor (JIF), as well as a new measure of the commercial relevance of the journal, suggested by Bikard and Marx, 2020: the Journal of Commercial Impact Factor (JCIF). The JIF and JCIF are calculated for each journal `J` and year `t` as the number of times the articles in journal `J` in years `t-1` and `t-2` were cited in year `t`, divided by the number of articles in journal `J` during years `t-1` and `t-2`. Whereas JIF employs citations from other papers, JCIF employs citations from patents. Approximately 20% of patents citing at least one scientific publication cite an article in a journal with a JIF in the top one percentile within a year, and 40% of patents that cite a scientific publication cite a paper from a journal with a JCIF in the top one percentile within a year.

First to use science To identify whether a patent was the first to cite a specific publication among all USPTO patents, we compare each cited publication with all the patents that cited the specific publication between 1901 and 2010. We classify the focal patent as the first to use a given publication if no other patent that cited the publication was granted before the focal patent. Specifically, we classify patent `T` as the first to use paper `S` if no other patent citing paper `S` was granted before `T`. Accordingly, *First Use of Science* is a binary variable equal to one if the focal patent is the first to cite a paper and zero otherwise. If a patent cites more than one publication, the first to use variable receives the value of one if at least one of these articles is cited for the first time. On average, a cited publication is referenced by approximately 3.8 patents, while the median cited publication is cited by two patents, with a range from one to 3,054 patents.

We build similar indices for citing scientific publications in different types of journals according to the above-described JIF and JCIF. We construct measures of first to use articles in top JIF and top JCIF by identifying whether a patent is the first to cite a paper that was published in a journal with a JIF/JCIF in the 99th percentile.

Relevant science Not all used discoveries are relevant for invention. Arguably, absorptive capacity should enable a firm to discern which discoveries will be relevant. Accordingly, we measure *relevant science* as publications that are highly cited by patents: papers that are in the top 95th percentile of patent-to-science citations. We construct the measure of *first to use relevant science* as a binary variable indicating whether a patent is the first to cite a relevant publication. As a robustness check, we also present results with papers that have been cited more than once, *Mult Tech Use*, as relevant science.

Corporate participation in science. We obtain corporate publication activity from Arora et al., 2021.[17] The corporate publication data allow us to identify scientific papers for which at least one of the authors is affiliated with U.S.-headquartered Compustat firm. We thus create our main measures of a firm engagement in science. Building on Rosenberg, 1990, who argues that active engagement in scientific research "buys" firms a ticket for admission into an information network, we propose that attending conferences in which relevant science is presented and being aware of relevant journals raises a firm's familiarity with applicable science, making it more likely that the firm is able to use the scientific advances before others.

Using corporate publications between 1980 and 2010, we identify conferences and journals with which a firm is "*familiar*" at each point in time. We construct two measures of familiarity: one based on whether the firm has published in a conference's proceeding, and one based on whether the firm has published in a journal. Specifically, we record an attendance by firm `F` of conference `C` if at least one employee of `F` is an author of a paper in the proceeding of conference `C`. Similarly, at

[17] Arora et al., 2021 match approximately 800,000 scientific in the "Science Citation Index" and "Conference Proceedings Citation Index - Science" of Web of Science, published between 1980 and 2015, to their sample of Compustat firms.

the journal level, we consider firm F to be familiar with a journal J if F has published in J. [18]

We use conference information from Microsoft Academics. There are about 3500 conferences linked to scientific articles, 46% of which are linked to articles cited in the front page of patents. In our sample, 11% of the patents that cite science, cite at least one proceeding from a conference in which the firm participated, while around 54% cite a publication that appeared in a journal the assignee of the patent had published in prior to the issuing of the focal patent.

We classify the focal patent P, filed by firm F and citing a scientific journal J, as familiar with J if an employee of F has authored a paper published in J prior to the filing of P. Similarly, at the conference attendance level, we indicate the focal patent P, filed by firm F and citing a paper published in the proceeding of conference C, as familiar with C if an employee of F has authored a paper presented at C.

Crowdedness of the technical space. We use the number of patenting firms as a proxy of technological competition in a market. Specifically, we count the number of publicly traded firms (PERMCOs) that are assigned patents in the same primary CPC, up to the group level, in the same year. When multiple firms develop similar innovations based on the same public scientific knowledge, the private rents for each firm are diminished. This is because the value of the innovation is shared across competitors, making it harder for any single firm to capture a significant portion of the market.

However, competition in the use of public science can also occur between firms that do not compete in the product market. These "tech-rivals" may, in a well-functioning market for technology (Arora & Gambardella, 2010), license or otherwise supply inventions to product market competitors of the focal firm. In some cases, they may embody their inventions into inputs they

[18] We use Microsoft Academics (MAG) Sinha et al., 2015 to identify the characteristics of each paper, such as the journal in which it was published, when it was published, what other papers it cited, and if it was presented at a conference. The *Corporate Pubs Data* from Arora et al., 2017 links Compustat firms to the Web of Science database (WoS). Since WoS and MAG use different identification codes for the same paper, we create a crosswalk between the WoS IDs in the *Corporate Pubs Data* and MAG IDs using TFIDF and fuzzy matching on papers' titles and authors, limiting comparisons within rolling windows of three years.

supply to product market rivals of the focal firm. By so doing, these "tech-rivals" also reduce the private returns to the use of the science by the focal firm.[19] By measuring competition as the number of firms patenting in the same technology class, we believe we can capture direct technological competition as well as the more indirect rivalry fostered by the market for technologies.[20]

Patent scope. A key contribution is to empirically demonstrate a source of first-mover advantage, namely broader patent scope. We use two measures to proxy for patent scope: the length of the shortest independent claim and the count of independent claims of the focal patent (Marco et al., 2019).[21] More independent claims, or fewer words in the shortest independent claim, correspond to a broader patent scope.

2.3.2 Non-parametric results

Our main sample and variables are at the patent level. The key variables for each patent are its dollar value to the firm, whether it cites science, and whether a citation to a specific scientific article first appeared in the focal patent. Table 2.1 presents descriptive statistics for our main variables. The median patent in our sample is valued at $6 million dollars (1982 prices). The median values of patents that cite science and that do not cite science are $7.3 million and $6 million, respectively.

The value of patents citing science vs. not citing science, however, differs along an important dimension: how many firms are patenting in the same technological space. Figure 2.1 shows that as the number of firms patenting in a technological class increases, the difference in value between patents citing science and patents not citing science steadily decreases.[22] This suggests that firms tend to create more value using scientific advances in less crowded technological fields, consistent with the idea that private returns to using science are low if rivals also have access to the same

[19] This kind of competition is not something that can be directly and easily observed.

[20] We would like to thank an anonymous reviewer for helpful comments on this point.

[21] Similar measures can also be used. For example, Kuhn and Thompson, 2019 propose the length of the first independent claim. The length of the shortest and the first independent claim, as well the average length, are highly correlated.

[22] The median patent is issued in a patent class populated by 20 firms.

science.

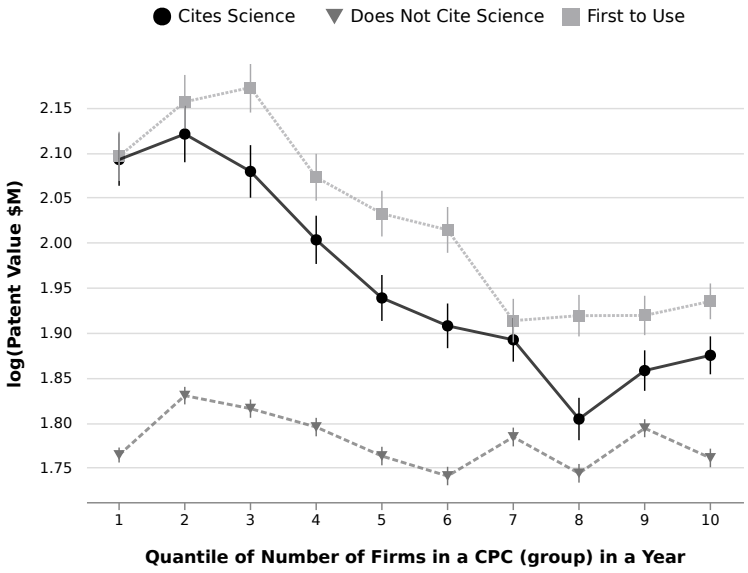

Figure 2.1: Use of Science in Crowded Spaces
This figure shows the average patent value for different bins of competition (Number of Firms Patenting in the 6-digit CPC). We group patents into ten quantiles based on the number of patenting firms in the same CPC (up to the group level); cutoffs are determined yearly. The values are presented for three different groups: (i) patents that do not cite science, (ii) patents that cite science for the first time, and (iii) patents that cite science, but not for the first time.

To explore first-mover advantage in the use of science, we split the sample of patents citing science into patents that are first to use a scientific article and patents that only cite scientific articles previously cited by other patents.[23] The difference between being first to use and not first to use tends to be larger in more crowded spaces. These trends suggest that the negative effect of competition is mitigated by being the first to use a particular research article. Note also that our measure of first use is absolute. It is possible that a follow-on use may, in fact, be the first use in the relevant market. If so, we are underestimating the private value of first use by mixing economically relevant first-use and follow-on use in a single category.

[23]If two patents are filed on the same day and they cite the same science, they are both first to use.

2.4 Econometric analysis

2.4.1 Patent value and use of science

We estimate the following specification for the i^{th} patent:

$$\ln(\xi_i) = \alpha_0 + \alpha_1 Science_i + Z_i'\gamma + \varepsilon_i \tag{2.1}$$

The value of the patent is represented by ξ_i and $Science_i$ is a dummy variable that takes the value 1 if the patent cites science and 0 otherwise. We control for factors Z_i that may influence the use of science and patent value. These include patent grant year fixed effects, because the use of scientific papers in the front page of patents varies over time; the logarithm of the firm's market capitalization on the day prior to the issuing of the patent; firm size; patent class fixed effects to account for variation of the use of science across different technologies; and firm fixed effects to control for the presence of unobservable firm characteristics that do not vary over time. We cluster the standard errors by firm to account for potential serial correlation in the use of science within firms.

Table 4.1 presents the estimation results for the relationship between patent value and the use of science.[24] The results show that patents that cite science are more valuable than patents that do not, but the difference is economically small. Column 1 presents estimates without controlling for firm fixed effects. The coefficient estimate of *Science Dummy* indicates that patents that cite scientific papers are about 4.1% more valuable than non-science-citing patents. When controlling for firm-fixed effects, this estimate is cut by about two-thirds (Column 2).

Columns 3-5 explore heterogeneity by the quality of the cited publication. Column 3 includes a dummy that equals 1 for cited publications that are published in journals with an impact factor (JIF) in the upper 1% of all journals in Microsoft Academics. This dummy variable receives the value of 0 for all other observations, so patents not citing science and patents citing science that is not in a top JIF are lumped into one category. The estimate indicates that patents citing a top JIF

[24] We also include results where we do not control for market capitalization in Table 2.9. The estimates are very similar to Table 4.1. In general, controlling for firm fixed effects and market caps reduces the science-premium from about 4% to about 1.5%.

publication are more valuable than patents citing lower-ranked publications. Similarly, we repeat the same exercise for patents citing science published in higher commercial impact factor journals (JCIF). Column 5 includes the two different quality dummies in a single specification. The value of using science is never bigger than 5% (summing all use of science-use dummies) with respect to the value of non-citing patents, even when the cited publication is published in a high-impact journal.

If the observed (small) difference in the economic value of citing scientific publications, vs. not, is driven by differences in the intrinsic quality of the two groups, then we would expect small differences in technical value as well (Hall et al., 2005; Nicholas, 2008). Columns 6-8 examine the relationship between science and technical importance, measured by forward citations. The estimate of the *Science Dummy* in Column 6 indicates that patents using science receive on average 20% more citations by follow-on inventions compared to patents not citing science. For the average patent in the sample, it translates into approximately two more citations. When we consider the use of high-quality science (science cited in journals with an impact factor on the upper 1%) the difference increases by about one more citation: citing science increases the number of forward citations by about 30%. Column 8 reports a linear probability model in which being in the top 1% of forward citations in a patent class-year (a “Home Run" patent) is related to the use of science. A patent that cites science has a 64% higher probability of being a “Home Run" when evaluated at the sample mean values.

Overall, the results from Table 4.1 indicate that, while patents that cite science are technically more important, the private economic value of using science is relatively low. Table 2.4 presents estimates by technology fields. The results indicate that using science in patented inventions appears to be most valuable in the “Life Sciences" and “Chemicals" categories, where the use of science is valued, respectively, at 3.3% and 2.1%. The estimates at the technology field level seem to confirm the basic pattern: the use of a resource that is available to many is not a source of significant private economic value.

2.4.2 Competition and first to use

If the private value of commercializing scientific research is *economically* small because science is non-excludable and thus can be used by all, we expect the value of using science to be lower in more crowded technical spaces. However, we would expect that in crowded niches, first movers would be more valuable. As discussed in the introduction, and as Figure 2.1 suggested, the average value of patents is smaller in crowded niches. However, the value of first-movers is higher than that of followers. Moreover, though both decline with competition, the difference between the two is higher in more crowded niches. That is, the value of being a first-mover is greater in more crowded spaces. Here we verify that those findings are robust to controlling for time and firm characteristics, including time-invariant firm effects. Accordingly, we estimate the following specification:

$$\begin{aligned}\ln(\xi_i) = \alpha_0 + \alpha_1 Science_i + \alpha_2 Science_i \times \ln(NFirms_i) + \alpha_3 \ln(NFirms_i) \\ + \alpha_4 First\ to\ Use_i \times \ln(NFirms_i) + Z_i'\gamma + \varepsilon_i\end{aligned} \tag{2.2}$$

We measure how crowded the niche is with the number of public firms that patent in the same CPCs (up to the group code) as the focal patent within the same year the focal patent is granted. For ease of exposition, we label firms operating in the same technology niche as likely technology competitors. Our interest is at α_2, where we expect $\hat{\alpha_2} < 0$. That is, the private value of commercializing science falls in the number of firms that have access to the publications cited by the focal patent. Table 2.5, Column 1 introduces our measures of competition: the natural log of 1 plus the number of different Compustat firms in the same technology class as the focal patent. The estimate of the coefficient on competition is negative, and the estimate on the dummy for science is consistent with our previous estimates (e.g., Table 4.1, Column 2) [25].

Column 2 adds an interaction between science and competition. As expected, $\hat{\alpha_2}$ is negative and statistically significant. The coefficient estimate on the science dummy jumps to close to 7%, indicating an increase in the private value of commercializing science in environments where fewer

[25] For completeness we also report estimates in which we do not control for market capitalization in Table 2.10. The estimates for the variables of interest are very similar to those reported in Table 2.5

firms patent compared to the baseline results in Table 4.1. The private value of using science falls with the number of patenting firms.

Column 3 introduces our main variable of interest: whether a patent is the first to cite a scientific paper. We test whether the negative effect of competition on private value can be mitigated by being a first-mover. We add an interaction term between first to use and competition. As expected, first-use mitigates the negative effect of competition on the private returns to using science. The coefficient estimate on the interaction term between science and competition is reduced by half when science is first cited by the focal firm (from -0.02 to -0.01). That is, as in the non-parametric analysis, we observe that being first to cite is associated with a reduction in the negative effect of competition on patent value by about 50%.

Columns 4-7 present estimates by technology field. These are broadly consistent, showing that first-use mitigates the effect of competition in all fields, albeit that the effects are not always precisely measured. The slight deviation of "Life Sciences" from this pattern may reflect the high reliance on science in bio-pharmaceuticals, but may also reflect the somewhat special role of patents in protecting pharmaceutical inventions. For instance, though Merck was the first-mover in statins, Parke-Davis (and, subsequently, Pfizer) was able to introduce a more effective product.

2.4.3 Patent Renewals

There are significant differences between stock market reactions and payment of renewal fees; Stock market reactions measure a forward-looking investor expectation of the monetary value of the patent, while payment of maintenance fees would more closely reflect the evolving managerial perception of the economic value of the patent.By employing maintenance fees as an alternative measure to proxy the economic value of patents, we confirm our main findings. Specifically, we re-estimate specification (2.2) using a binary variable as a dependent variable indicating whether the patent was renewed to full term (i.e., fees were paid to renew at the 4-, 8-, and 12-year marks).

Table 2.6 shows a very similar pattern as described in Section 2.4.1. The difference in the likelihood of renewal between patents citing science and patents not citing science is small. The difference is reduced further by an increase in the number of firms in the technical area (Column 2, Table 2.6). However, being first to use reduces the effects of number of firms for patents citing

science by half.

We provide these additional results as a validation of the overall pattern presented, that the private value of commercializing scientific research may be economically small because science is non-excludable and accessible to all, whereas the value of using science would be lower in more crowded technical spaces. Conversely, in crowded niches, first-mover advantage should be more valuable.

2.4.4 Rent dissipation or diminishing returns

One might also wonder if the higher value of first-mover patents merely reflect diminishing returns. That is, whereas our focus is on competition eroding rents and thereby lowering private value, the average quality of invention may also be lower when many firms mine the same seam of knowledge. The two possible explanations are not mutually exclusive. However, the evidence suggests that diminishing returns by themselves are unlikely to explain our results. Recall from columns 6 in table 4.1 that patents citing science receive nearly 20% more citations i.e., are of higher technical quality. In unreported regressions, we find that the quality premium associated with citing science does not meaningfully drop in more crowded niches, and neither is first-use associated with significantly higher quality.[26] In other words, our results are not the result of aggregate diminishing returns in a stream of innovation.

2.4.5 Patent Scope

We next investigate whether pioneering the use of scientific advances allows first-movers to secure broader patents. We estimate the following specification for the i^{th} patent:

$$Patent\ Scope_i = \beta_0 + \beta_1 First\ to\ Use_i + Z_i'\gamma + \varepsilon_i \tag{2.3}$$

Table 2.7 shows that patents that cite science have about 4% more independent claims (Column 1), but have about the same number of words in the shortest independent claims as patents

[26]Results are available upon request. We are grateful to an anonymous reviewer for highlighting this possibility.

that do not cite science (Column 3). However, first-use is associated with an additional 4% more independent claims (Column 1) and 6% fewer words in the shortest claim (Column 3). The gap between the patent scope of first-movers and followers is even greater when considering the use of relevant science. Table 2.7, Column 2, shows that the first-mover patents have about 7% more independent claims compared to follower patents. Similarly, Column 4 shows that the length of the shortest independent claim is about 9% shorter for first-mover patents relative to follower patents. In other words, one way in which being a first-mover in using science translates into private value is being able to get a broader patent, which can then discourage entry and provide a greater zone of exclusivity for the first-mover.

2.4.6 Scientific capabilities and first to use

The foregoing results naturally raise the question of which firms are more likely to be first-movers in using science. Obviously, firms can typically use their own science better than others. However, the amount of science produced outside the firm is significantly larger than the science produced internally. The theory of absorptive capacity (Cohen & Levinthal, 1989) predicts that corporations that produce science increase their ability to recognize and apply scientific knowledge originating outside the firm's boundaries, suggesting that corporate engagement in upstream research increases a firm's ability to recognize extramural knowledge.

Familiarity with science and first to use

Active engagement in scientific research may raise a firm's familiarity with relevant journals and conference proceedings. We explore whether this not only leads to citation to science, as Baruffaldi and Poege, 2020 show but, more specifically, whether it is related to a firm being the first to use the science. We estimate the following specification for the i^{th} patent:

$$\textit{First to Use}_i = \beta_0 + \beta_1 \textit{Familiarity with Science}_{if} + Z_i'\gamma + \varepsilon_i \tag{2.4}$$

We estimate this specification for the sample of patents citing science because both our dependent and main independent variables are constructed only for patents citing a scientific article.

First use of Science is a binary variable that equals 1 when the focal patent is the first to cite a scientific paper, and *Familiarity with Science* is a measure (at the patent level) indicating whether the focal patent cites a scientific paper published in a conference or in a journal where the firm has also published. We control for firm, year, and patent class fixed effects. If corporate participation in conferences and publications in certain journals increases a firm's absorptive capacity, we would expect $\beta_1 > 0$.

The results are reported in Table 2.8. Column 1 presents a validation test. Using a linear probability model, we find that a patent that cites internal science is also more likely to cite a paper that is being cited for the first time. The estimates from Column 1 imply that patents that make internal citations, i.e., a citation to papers published by an employee of the focal firm prior to the filing of the patent, have about a 6% higher probability to be first to use. Evaluated at the average of the distribution of *first to Use*, this translates into an increase of approximately 12%.

Columns 2-6 explore how first use is related to familiarity with the science cited. As previously defined, familiarity relates to whether the assignee of the patent has published in the journal cited by the focal patent or if the assignee has attended the conferences cited in the focal patent. Estimates in Column 2 show that if a firm is citing a paper presented at a conference it attended, then its patent is 11.5% more likely to be a first-user of science. Similarly, if a patent is citing a paper published in an outlet the firm has published in before the filing of the patent, then the patent is 5% more likely to be the first to cite a paper (Column 3). Hence, a patent that cites science from a familiar domain is also more likely to cite a paper that is being cited for the first time. Column 4 shows that familiarity through journals and familiarity through conference participation are largely independent in terms of their effects on first use. Columns 5 and 6 consider the first use of papers published in journals with a high journal impact factor (Column 5) and with a high commercial impact factor (Column 6); the estimates in Columns 5 and 6 are in line with the baseline results.

We assumed that absorptive capability implies that a firm is not only aware of scientific discoveries but is also able to assess the relevance of the discovery for use in its inventive activities better than other firms; we thus restrict our first-use measure to relevant science. In Columns 7 and 8 of Table 2.8, we assess whether engagement in the publication process increases the likelihood of

pioneering the use of papers (Column 7) that *will be* cited more than once – *First to Use Mult Tech Use* – or the use of papers (Column 8) that *will receive* enough citations to place them in the upper 5% of papers with respect to patent citations count – *First to Use Relevant Paper*. Evaluating our estimates at the average of the distribution of *First to Use - Mult Tech Use*, being familiar with the science used translates into an increase in the likelihood of being *First to Use - Mult Tech Use* by about 50%. Similarly, evaluated at the average of the distribution of *First to Use - Relevant Paper*, being familiar with the science used translates into an increase in the likelihood of being *First to Use - Relevant Paper* by about 80%.

2.5 Discussion and Conclusion

Our results suggest that firms that engage in basic research are more likely to be aware of relevant scientific advances. These firms are, therefore, also more likely to be better positioned to use relevant scientific advances sooner than others. The resulting inventions are likely to be more valuable, in part because the patents would be broader in scope. The benefits of being able to recognize and use scientific advances are more salient in more competitive technology fields. Taken together, investing in internal research is likely to help a firm identify and use publicly available knowledge before its competitors. In this way, the firm is able to create private value from a public good.

Preemption plays a critical role in capturing value from using science, as evidenced by the advantage conferred by the broader scope of patents available to first-movers. Insofar as first-movers are likely to be firms engaged in conducting and publishing scientific research, this suggests that intellectual property protection may complement public science. Though we focus on patents, other preemptive strategies could include cornering complementary assets, such as specialized machinery, skilled labor, or unique distribution channels, which can enhance the value derived from scientific knowledge. A case in point is Artificial Intelligence (AI). AI technologies often require large datasets and substantial computational power, which serve as key complements to public sci-

entific knowledge. Firms that can effectively utilize these complements alongside their scientific knowledge can further enhance their competitive advantage. As AI continues to reshape industries, the importance of preemptive strategies, such as securing broad patent protection and cornering complementary assets like large datasets and computing resources, may be increasingly critical for firms seeking to maximize the private value derived from public scientific knowledge.

Large corporations have progressively withdrawn from scientific research over the last three decades (Arora et al., 2018). Along with the very substantial increase in publicly funded research, this implies that scientific advances from public institutions have become crucial inputs into corporate innovation. In turn, this has highlighted the tension analyzed in this paper, namely how to create private value from an input that is freely available to all, including competitors. Our study emphasizes the importance of acquiring absorptive capacity to move first in utilizing scientific knowledge. Firms that actively participate in scientific research are better positioned to be the first to use relevant science, as their involvement fosters familiarity with the latest advances. Maintaining such a capability is likely to require a cadre of in-house scientists, who can plug into the relevant flows of scientific knowledge. Arguably, an effective way to remain plugged into the scientific network is to be a participant in the research process.

Yet, because established firms appear to be withdrawing from scientific research, our findings also raise a potential concern about the health of the American innovation ecosystem. While the growing specialization of universities in upstream research and of firms in downstream development should make each sector more productive in its respective activity, it also means that the need to link these two activities together is also becoming more important. If by withdrawing from science corporations lose their access to the scientific community, using scientific advances produced by universities for downstream inventions would be harder. A central challenge for policymakers and business managers alike is how to connect public science to downstream invention in an ecosystem where firms themselves are not engaged in scientific research, and thus, are less able to understand and use science.

Table 2.1: Summary Statistics for Patent Level Data

	Mean	Std	1%	25%	50%	75%	99%	Count
Patent Value $M	13.95	25.26	0.23	2.76	6.37	14.54	123.71	853382
log(Patent Value $M)	1.83	1.30	-1.47	1.02	1.85	2.68	4.82	853382
Market Cap	41557.12	70177.70	25.69	2271.32	10627.07	43768.29	342794.96	853382
log(Market Cap)	9.08	2.16	3.25	7.73	9.27	10.69	12.74	853382
Science Dummy	0.26	0.44	0.00	0.00	0.00	1.00	1.00	853382
First Use Science Dummy	0.13	0.34	0.00	0.00	0.00	0.00	1.00	853382
Science Dummy (High JIF)	0.05	0.22	0.00	0.00	0.00	0.00	1.00	853382
Science Dummy (High JCIF)	0.11	0.31	0.00	0.00	0.00	0.00	1.00	853382
Fw Cites	4.26	7.23	0.00	1.00	2.00	5.00	33.00	853356
Home Run	0.02	0.14	0.00	0.00	0.00	0.00	1.00	853356
N Patenting Firms	33.16	32.69	1.00	8.00	21.00	50.00	122.00	853382
log(N Patenting Firms)	2.90	1.24	0.00	2.08	3.04	3.91	4.80	853382
N Indep Claims	3.00	2.35	1.00	2.00	3.00	4.00	12.00	843685
log(N Indep Claims)	0.88	0.65	0.00	0.69	1.10	1.39	2.48	843685
Words in Ind Claim	137.22	96.72	13.00	80.00	118.00	172.00	450.00	843474
log(Words in Ind Claim)	4.72	0.67	2.56	4.38	4.77	5.15	6.11	843474

Table 2.2: Summary Statistics for Patent Level Data [Cite Science Sample]

	Mean	Std	1%	25%	50%	75%	99%	Count
First Use Science Dummy	0.51	0.50	0	0	1	1	1	222018
Science Dummy (High JIF)	0.19	0.39	0	0	0	0	1	222018
Science Dummy (High JCIF)	0.41	0.49	0	0	0	1	1	222018
Self Citation	0.12	0.33	0	0	0	0	1	222018
Attended Conference	0.11	0.31	0	0	0	0	1	222018
Published in Journal	0.54	0.50	0	0	1	1	1	222018
First Use (High JIF)	0.06	0.24	0	0	0	0	1	222018
First Use (High JCIF)	0.14	0.35	0	0	0	0	1	222018
First Use (Mult Tech Use)	0.38	0.49	0	0	0	1	1	222018
First Use (Relevant Paper)	0.07	0.26	0	0	0	0	1	222018

Table 2.3: The Use of Science and the Patent Economic and Technical Value

Dependent Variables:	log(Patent Value $M)					log(Fw Cites + 1)		Home Run
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Science Dummy	0.0413	0.0159	0.0116	0.0078	0.0071	0.1991	0.1596	0.0127
	(0.0212)	(0.0075)	(0.0065)	(0.0048)	(0.0048)	(0.0067)	(0.0076)	(0.0009)
Science Dummy (High JIF)			0.0383		0.0304		0.0553	
			(0.0151)		(0.0118)		(0.0164)	
Science Dummy (High JCIF)				0.0233	0.0153		0.0958	
				(0.0126)	(0.0108)		(0.0124)	
log(Market Cap)	Yes	Yes	Yes	Yes	Yes			
Fixed-effects								
Grant Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Patent Class	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	853,382	853,382	853,382	853,382	853,382	853,356	853,356	853,356
R^2	0.56902	0.78767	0.78770	0.78769	0.78770	0.18769	0.18857	0.04260
Adjusted R^2	0.56869	0.78657	0.78659	0.78658	0.78660	0.18346	0.18435	0.03762

Clustered (Firm) standard-errors in parentheses

This table presents the association between the use of science and measures of monetary and technical value of the focal patent. The dataset used is a patent level data of patents granted between 1980 and 2010 to US-based public firms that can be found in the DISCERN (Arora et al 2020) data. *Patent Value $M* (in 1982 USDm) is the estimated patent monetary value derived from excess stock returns of the filing company the days after the issuing of the patent, sourced from Kogan et al. (2017). *Science Dummy* is a binary variable equal to one if the focal patent has a front page citation to a scientific publication, sourced from Marx and Fuegi (2019). *Market Cap* is the market capitalization of the focal firm the day before the patent was issued. *Science Dummy (High JIF)* is a binary variable equal to one if the focal patent cites a scientific publication published in a journal with a Journal Impact Factor (sourced from Marx and Fuegi (2019)) in the top 99th percentile within year. *Science Dummy (High JCIF)* is a binary variable equal to one if the focal patent cites a scientific publication published in a journal with a Journal of Commercial Impact Factor (sourced from Marx and Fuegi (2019)) in the top 99th percentile within the year. *Fw Cites + 1* is the number (plus one) of forward citations received by the focal patent within 5 years from its issuing. *Home Run* is a binary variable equal to one if the focal patent received a number of forward citations in the top 99th percentile among the patents granted in the same year and the same patent cpc (up to the subclass).

Table 2.4: Patent Value and Use of Science (Industry Level)

Dependent Variable:	log(Patent Value $M)				
NBER Industry		Chemical	Electronics	ICT	Life Sci
	(1)	(2)	(3)	(4)	(5)
Science Dummy	0.0159	0.0208	0.0084	-0.0021	0.0326
	(0.0075)	(0.0086)	(0.0094)	(0.0072)	(0.0109)
log(Market Cap)	Yes	Yes	Yes	Yes	Yes
Fixed-effects					
Grant Year	Yes	Yes	Yes	Yes	Yes
Patent Class	Yes	Yes	Yes	Yes	Yes
Firm	Yes	Yes	Yes	Yes	Yes
Observations	853,382	149,529	194,033	253,004	70,200
R^2	0.78767	0.79716	0.77889	0.77315	0.89993
Adjusted R^2	0.78657	0.79407	0.77617	0.77098	0.89758

Clustered (Firm) standard-errors in parentheses

This table presents the association between the use of science and a measure of monetary value of the focal patent across industries (defined using the NBER classification). The dataset used is a patent level data of patents granted between 1980 and 2010 to US-based public firms that can be found in the DISCERN (Arora et al 2020) data. *Patent Value $M* (in 1982 USDm) is the estimated patent monetary value derived from excess stock returns of the filing company the days after the issuing of the patent, sourced from Kogan et al. (2017). *Science Dummy* is a binary variable equal to one if the focal patent has a front page citation to a scientific publication, sourced from Marx and Fuegi (2019). *Market Cap* is the market capitalization of the focal firm the day before the patent was issued.

Table 2.5: Patent Value, Use of Science & Number of Patenting Firms

Dependent Variable:	log(Patent Value $M)						
NBER Industry				Chemical	Electronics	ICT	Life Sci
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Science Dummy	0.0161	0.0687	0.0758	0.0257	0.0400	0.0685	0.0273
	(0.0075)	(0.0221)	(0.0229)	(0.0195)	(0.0247)	(0.0330)	(0.0284)
log(N Patenting Firms)	-0.0441	-0.0397	-0.0399	-0.0112	-0.0228	-0.0201	-0.0042
	(0.0151)	(0.0139)	(0.0139)	(0.0044)	(0.0077)	(0.0180)	(0.0061)
Science Dummy × log(N Patenting Firms)		-0.0164	-0.0207	-0.0059	-0.0140	-0.0196	0.0001
		(0.0057)	(0.0057)	(0.0069)	(0.0065)	(0.0089)	(0.0091)
First Use Science Dummy			-0.0150	-0.0226	-0.0077	-0.0222	0.0062
			(0.0125)	(0.0201)	(0.0204)	(0.0321)	(0.0227)
log(N Patenting Firms) × First Use Science Dummy			0.0091	0.0169	0.0125	0.0077	0.0010
			(0.0037)	(0.0073)	(0.0070)	(0.0081)	(0.0073)
log(Market Cap)	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed-effects							
Grant Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Patent Class	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	853,382	853,382	853,382	149,529	194,033	253,004	70,200
R^2	0.78814	0.78818	0.78819	0.79723	0.77909	0.77328	0.89994
Adjusted R^2	0.78704	0.78708	0.78709	0.79414	0.77637	0.77111	0.89758

Clustered (Firm) standard-errors in parentheses

This table presents the association between the use of science and a measure of monetary value of the focal patent, the mediating aspect of competitions (the number of patenting firms in a CPC group) and the first use of science by the focal patent. Industries are defined using the NBER classification. The dataset used is a patent level data of patents granted between 1980 and 2010 to US-based public firms that can be found in the DISCERN (Arora et al 2020) data. *Patent Value $M* (in 1982 USDm) is the estimated patent monetary value derived from excess stock returns of the filing company the days after the issuing of the patent, sourced from Kogan et al. (2017). *Science Dummy* is a binary variable equal to one if the focal patent has a front page citation to a scientific publication, sourced from Marx and Fuegi (2019). *N Patenting Firms* is the number of public firms patenting in the same year and in the same CPC (up to the group level) as the focal patent. *Market Cap* is the market capitalization of the focal firm the day before the patent was issued. *First Use Science Dummy* is a binary variable equal to one if the focal patent is the first to cite, in its front page, a scientific publication.

Table 2.6: Patent Renewals, Use of Science & Number of Patenting Firms

Dependent Variable:	Fully Renewed		
	(1)	(2)	(3)
Science Dummy	0.0047	0.0293	0.0292
	(0.0046)	(0.0090)	(0.0108)
log(N Patenting Firms)		0.0044	0.0043
		(0.0030)	(0.0030)
Science Dummy × log(N Patenting Firms)		-0.0076	-0.0101
		(0.0028)	(0.0034)
First Use Science Dummy			-0.0016
			(0.0083)
log(N Patenting Firms) × First Use Science Dummy			0.0057
			(0.0028)
Fixed-effects			
Grant Year	Yes	Yes	Yes
Patent Class	Yes	Yes	Yes
Firm	Yes	Yes	Yes
Observations	824,745	824,745	824,745
R^2	0.17106	0.17112	0.17120
Adjusted R^2	0.16663	0.16669	0.16677

Clustered (Firm) standard-errors in parentheses

This table presents the association between the use of science and full term renewals of the focal patent, the mediating aspect of competitions (the number of patenting firms in a CPC group) and the first use of science by the focal patent. The dataset used is a patent level data of patents granted between 1980 and 2010 to US-based public firms that can be found in the DISCERN (Arora et al 2020) data. *Fully Renewed* is a binary variable equal to one if all the fees of the focal patent: 4, 8 and 12 years marks, have been paid. *Science Dummy* is a binary variable equal to one if the focal patent has a front page citation to a scientific publication, sourced from Marx and Fuegi (2019). *N Patenting Firms* is the number of public firms patenting in the same year and in the same CPC (up to the group level) as the focal patent. *First Use Science Dummy* is a binary variable equal to one if the focal patent is the first to cite, in its front page, a scientific publication.

Table 2.7: First to Use and Patent Scope

Dependent Variables:	log(N Indep Claims)		log(Words in Ind Claim)	
	(1)	(2)	(3)	(4)
Science Dummy	0.0425	0.0230	0.0059	-0.0092
	(0.0039)	(0.0041)	(0.0058)	(0.0050)
First Use Science Dummy	0.0436	0.0476	-0.0605	-0.0488
	(0.0050)	(0.0051)	(0.0069)	(0.0063)
Science Dummy (Relevant Paper)		0.0365		0.0268
		(0.0042)		(0.0069)
First Use (Relevant Paper)		0.0210		-0.0492
		(0.0066)		(0.0074)
Fixed-effects				
Firm	Yes	Yes	Yes	Yes
Grant Year	Yes	Yes	Yes	Yes
Patent Class	Yes	Yes	Yes	Yes
Observations	843,685	843,685	843,474	843,474
R^2	0.17254	0.17277	0.18084	0.18096
Adjusted R^2	0.16819	0.16842	0.17654	0.17665

Clustered (Firm) standard-errors in parentheses

This table presents the association between the use of science and measures of patent scope of the focal patent. The dataset used is a patent level data of patents granted between 1980 and 2010 to US-based public firms that can be found in the DISCERN (Arora et al 2020) data. *N Indep Claims* is the number of independent claims made by the focal patent. *Science Dummy* is a binary variable equal to one if the focal patent has a front page citation to a scientific publication, sourced from Marx and Fuegi (2019). *First Use Science Dummy* is a binary variable equal to one if the focal patent is the first to cite, in its front page, a scientific publication. *Science Dummy (Relevant Paper)* is a binary variable equal to one if the focal patent cited a paper that is in the upper 5% of patent citations to science (within year of paper publication). *First Use (Relevant Paper)* is a binary variable equal to one if the focal patent is the first to cite a paper that is in the upper 5% of patent citations to science (within year of paper publication). *Words in Ind Claim* is the number of words of the shortest independent claim made by the focal patent.

Table 2.8: Firms' familiarity with science and first to use

Dependent Variables:	First Use Science Dummy				First Use (High JIF)	First Use (High JCIF)	First Use (Mult Tech Use)	First Use (Relevant Paper)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Self Citation	0.0585			0.0218	0.0281	0.0605	0.0502	0.0367
	(0.0076)			(0.0090)	(0.0048)	(0.0081)	(0.0076)	(0.0028)
Attended Conference		0.1151		0.1041			0.1164	0.0318
		(0.0104)		(0.0115)			(0.0116)	(0.0046)
Published in Journal			0.0515	0.0438	0.0302	0.1188	0.0791	0.0238
			(0.0070)	(0.0079)	(0.0034)	(0.0075)	(0.0071)	(0.0020)
Fixed-effects								
Firm	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Grant Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Patent Class	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	222,018	222,018	222,018	222,018	222,018	222,018	222,018	222,018
R^2	0.09436	0.09721	0.09498	0.09889	0.21907	0.13935	0.07910	0.06760
Adjusted R^2	0.08157	0.08446	0.08220	0.08615	0.20804	0.12719	0.06609	0.05443

Clustered (Firm) standard-errors in parentheses

This table presents the association between the first use of science and measures of familiarity to science of the focal firm: whether the firm actively participates in the scientific community by publishing in journals and participating to conferences. The dataset used is a patent-level data, restricted to the sample of patents which cites science, of patents granted between 1980 and 2010 to US-based public firms that can be found in the DISCERN (Arora et al 2020) data. *First Use Science Dummy* is a binary variable equal to one if the focal patent is the first to cite, in its front page, a scientific publication. *Self Citation* is a binary variable equal to one if the focal patent cites a scientific publication authored by at least one employee of the patent assignee. *Attended Conference* is a binary variable equal to one if the focal patent cites a scientific publication published in a conference proceeding attended by at least one employee of the patent assignee. *Published in Journal* is a binary variable equal to one if the focal patent cites a scientific publication published in a jorunal in which the firm had published a paper prior to the filingof the focal patent. *First Use (High JIF)* is a binary variable equal to one if the focal patent is the first to cite a scientific publication published in a journal with a Journal Impact Factor (sourced from Marx and Fuegi (2019)) in the top 99th percentile within the year. *First Use (High JCIF)* is a binary variable equal to one if the focal patent is the first to cite a scientific publication published in a journal with a Journal of Commercial Impact Factor (sourced from Marx and Fuegi (2019)) in the top 99th percentile within the year. *First Use (Mult Tech Use)* is a binary variable equal to one if the focal patent is the first to cite a paper that is cited by at least two patents. *First Use (Relevant Paper)* is a binary variable equal to one if the focal patent is the first to cite a paper that is in the upper 5% of patent citations to science (within year of paper publication).

Table 2.9: The Use of Science and the Patent Economic Value [with and without market capitalization as a control]

Dependent Variable:	log(Patent Value $M)			
	(1)	(2)	(3)	(4)
Science Dummy	0.0379	0.0413	0.0216	0.0159
	(0.0329)	(0.0212)	(0.0083)	(0.0075)
log(Market Cap)		Yes		Yes
Fixed-effects				
Grant Year	Yes	Yes	Yes	Yes
Patent Class	Yes	Yes	Yes	Yes
Firm			Yes	Yes
Observations	853,382	853,382	853,382	853,382
R^2	0.16647	0.56902	0.70351	0.78767
Adjusted R^2	0.16584	0.56869	0.70196	0.78657

Clustered (Firm) standard-errors in parentheses

This table presents the association between the use of science and measures of monetary and technical value of the focal patent. *Patent Value $M* (in 1982 USDm) is the estimated patent monetary value derived from excess stock returns of the filing company the days after the issuing of the patent, sourced from Kogan et al. (2017). *Market Cap* is the market capitalization of the focal firm the day before the patent was issued. *Science Dummy* is a binary variable equal to one if the focal patent has a front page citation to a scientific publication, sourced from Marx and Fuegi (2019).

Table 2.10: Patent Value, Use of Science & Number of Patenting Firms [with and without market capitalization as a control]

Dependent Variable:	log(Patent Value $M)					
	(1)	(2)	(3)	(4)	(5)	(6)
Science Dummy	0.0382	0.0416	0.0217	0.0161	0.0782	0.0758
	(0.0328)	(0.0213)	(0.0083)	(0.0075)	(0.0283)	(0.0229)
log(N Patenting Firms)	-0.0336	-0.0275	-0.0396	-0.0441	-0.0358	-0.0399
	(0.0265)	(0.0194)	(0.0182)	(0.0151)	(0.0169)	(0.0139)
First Use Science Dummy					-0.0186	-0.0150
					(0.0147)	(0.0125)
Science Dummy × log(N Patenting Firms)					-0.0192	-0.0207
					(0.0073)	(0.0057)
log(N Patenting Firms) × First Use Science Dummy					0.0093	0.0091
					(0.0046)	(0.0037)
log(Market Cap)		Yes		Yes		Yes
Fixed-effects						
Grant Year	Yes	Yes	Yes	Yes	Yes	Yes
Patent Class	Yes	Yes	Yes	Yes	Yes	Yes
Firm			Yes	Yes	Yes	Yes
Observations	853,382	853,382	853,382	853,382	853,382	853,382
R^2	0.16677	0.56922	0.70389	0.78814	0.70393	0.78819
Adjusted R^2	0.16613	0.56889	0.70235	0.78704	0.70239	0.78709

Clustered (Firm) standard-errors in parentheses

This table presents the association between the use of science and a measure of monetary value of the focal patent, the mediating aspect of competitions (the number of patenting firms in a CPC group) and the first use of science by the focal patent *Patent Value $M* (in 1982 USDm) is the estimated patent monetary value derived from excess stock returns of the filing company the days after the issuing of the patent, sourced from Kogan et al. (2017). *Market Cap* is the market capitalization of the focal firm the day before the patent was issued. *Science Dummy* is a binary variable equal to one if the focal patent has a front page citation to a scientific publication, sourced from Marx and Fuegi (2019). *First Use Science Dummy* is a binary variable equal to one if the focal patent is the first to cite, in its front page, a scientific publication. *N Patenting Firms* is the number of public firms patenting in the same year and in the same CPC (up to the group level) as the focal patent.

Chapter 3

Differences in the Rate and Direction of Technical Change

"What forces, then, determine the directions in which a firm actually goes in exploring for new techniques? Since it cannot explore in all directions, what are the factors which induce it to strike out in a particular direction? Better yet, are there any factors at work which compel it to look in some directions rather than others?" (Rosenberg, 1969)

3.1 Introduction

In this paper, I propose a review of the literature in the fields of economics and management of innovation, with a particular focus on studies employing difference-in-differences (DiD) methodologies. This approach aims to establish shifts in the rate and direction of patenting activity following the occurrence of specific events or a series of analogous events. By examining these changes, this survey seeks to provide insights into procedures employed to uncover underlying mechanisms that drive innovation.

The literature in the economics and management of innovation encompasses numerous studies that utilize difference-in-differences to assess the impact of a diverse set of events on patenting activity. In the following sections, I will outline the work of multiple researchers who have investigated the influence of a broad range of events such as compulsory licensing (Baten et al., 2017; Moser & Voena, 2012), immigration (Moser et al., 2014), demand shocks (Galasso & Luo, 2021, 2022), corporate break-ups (Poege, 2022; Watzinger & Schnitzer, 2022), revisions in legal framework (Aydin Ozden & Khashabi, 2022; Kang, 2021), cartel formation or breakups (Kang & Lee, 2022), shifts in tax regimes (Lichter et al., 2021), changes in patents scope Lee, 2023, patent pledges (de Rassenfosse & Palangkaraya, 2023), litigation (Giebel, 2021) and hiring of inventors (Akcigit & Goldschlag, 2023b).

Building upon Joseph Schumpeter's early 20th-century insights, economists have recognized the crucial role that technical change plays in shaping economic growth and societal prosperity (J. Schumpeter, 1942). As a result, the long-term economic performance of individual companies, industries, and nations is increasingly seen as reliant on their ability to foster and implement technological advancements (Cohen, 2010). In light of recent debates, on whether are we facing a future of stagnation (Gordon, 2016), or whether these worries mirror historical concerns of the ending of economic growth (Mokyr et al., 2015), it is of central interest to underline that growth potential is often determined by the direction firms take in their innovative endeavors.

Innovation and technical change, as the cornerstone for numerous economic issues, hinge significantly on organizations' and firms' choices and the direction of their innovative pursuits. Industrial societies have developed the aptitude to address specific challenges, with firms leading the charge in driving innovation and shaping its trajectory. However, an important reminder is to recognize that firms typically address only a limited range of problems, despite possessing the capacity to tackle a broader array of issues. This realization suggests that a deeper understanding of the process of technological change can be achieved by examining how firms identify, prioritize, and approach problems in their innovative endeavors (Rosenberg, 1969). Environmental changes can be a driving force in identifying the factors that shape the rate and direction of inventive efforts; whereas technological transformation can take center stage in interpreting the dynamics of economic growth.

Event studies and difference-in-differences can be employed to analyze the impact of quasi-exogenous changes in the environment in which organizations operate on their innovation activities. For example, various studies focus on identifying and examining events that trigger changes in regulations, market conditions, or consumer preferences, which could potentially influence the direction and pace of innovation. This paper will label these types of events, *inducement events* [1]. We can define inducement events as events that spurred technological change[2]; in the absence of an

[1]Unlike Hick's lineage of induced theory of innovation, inducement events can comprise changes other than the change in the relative prices of factors of production, and leading to a broader set of inventions, not restricted to "economizing the use of a factor which has become relatively expensive" (J. Hicks, 1963).

[2]In surveying the innovation literature in section 2, we will see that studying the events that induced

inducement event, the inventive trajectory would have been different [3]. By utilizing data before and after the occurrence of an inducement event, for entities involved in the events and a comparable control group[4], researchers can assess the causal relationship between the event and the subsequent innovation activities. This enables us to determine whether the event has led to a change in research and development investments or shifts in the rate and direction of inventive focus. Furthermore, heterogeneity analysis can help uncover the underlying economic mechanisms that drive organizations to respond to these inducement events.

In the following sections, I will explore recent studies in the economics of innovation literature, presenting the economic rationale of the research, and highlighting the conflicting aspects that arise from economic theory in answering the questions to which various scholars propose an empirical answer. I will then explain, to the best of my understanding, the data structure used by the various authors in their analysis, the approach used to identify a treatment and a control group, and the investigated outcomes[5]. The overview of the different investigations will highlight the broad and diverse set of research questions scholars have set forth and answered using the procedures involving patent data and difference-in-differences/event study estimation; reviewing various applications will facilitate the understanding of how one can streamline this type of analysis to further knowledge of inducement events.

To aid future research, this paper introduces pyDRAD [6], an open-source software specifically

---

inventive activity will allow us to determine economic factors and the mechanisms through which the technological trajectory changed.

[3]As we will see in section 2, we are assuming that the trend in the outcome (the measure of technical change in use) for entities undergoing the event: treated patent classes, patents, firms, inventors would have been parallel to the trend in the outcome of the control group, had the treated entities not undergone an event.

[4]The following sections will provide a detailed explanation of the data involved and it will also introduce the datasets that are distributed along this paper.

[5]I will not review specific findings of the studies surveyed; please review the referenced articles for details on the results and their implications.

[6]pyDRAD, an acronym for “Differences in the Rate and Direction". py stands for Python: the language it is written in.

designed to facilitate difference-in-differences estimation and inference in the realm of innovation and technical change. pyDRAD leverages the increasing availability of open-source data on patenting to provide researchers with a user-friendly and efficient tool for analyzing the impact of various events on inventive activity. Specifically, pyDRAD is designed to (1) construct longitudinal datasets derived from raw innovation datasets, such as patent data files distributed by the USPTO, PatentsView or Google Patents; (2) simplify the process of selecting a treated and a never-treated group, incorporating options for selecting never-treated entities based on best practices; (3) provide access to a range of variables, frequently employed in the economics of innovation literature, which can be constructed using open-access I3-indexed datasets; (4) streamline the data preparation for difference-in-differences analysis.

The remainder of this paper is organized into the following two sections:

In Section 3.2, the paper surveys a portion of the literature on the economics of innovation, specifically articles employing DiD and event studies to investigate questions posed by scholars in this field and explores the distinctive analytical setups used by the various authors. The aim here is not only to broadly map the methodological terrain of innovation economics but also to elucidate why the questions raised are intellectually stimulating and significant.

In Section 3, the paper proposes an overview of pyDRAD—an open-source software crafted to facilitate difference-in-differences estimation and inference within the domain of innovation and technical change. This section explores the software's key features, capabilities, and the practical implications it holds for future research.

3.2 Rate and Direction of Technical Change

Over the years many scholars have uncovered patterns and relationships in patent data [7]. One of the earliest attempts in the explanations of different rates of inventions across industries is (Merton,

[7]Patent statistics *"are available; they are by definition related to inventiveness, and they are based on what appears to be an objective and only slowly changing standard. No wonder that the idea that something interesting might be learned from such data tends to be rediscovered in each generation." (Griliches, 1998)*

1935). Merton's discussion of *"the determination of the factors involved in the focusing and shifting of inventive interest, i.e., the variations in the amount of invention within given industries."* [8] His work pioneered the interest and research using patent statistics [9] to uncover shifts in inventive activities across sectors and the factors influencing them.

In the next paragraphs, I will focus on describing setups that use patent data. In order, I will first define the *outcomes* that are of academic and policy interest when working with patent statistics, then I will list and highlight the core components of much of the empirical research in the economics of innovation domain that employ difference-in-differences with patent data. The core components are the elements discussed: mainly the choice of *entities* that are "treated" and entities "never-treated" (when relevant), as well as the *events e* under investigation[10]. For each choice of unit of analysis, that is the type of entity selected to construct a panel dataset, I will survey a number of research articles related to that unit of analysis.

We will use patent counts within a defined time period (typically counts in a year) to measure the rate and direction of inventive activity [11]. However, not all patents are created equal. They differ across various observable or harder-to-observe dimensions. Some patents may rely on scientific knowledge (Marx & Fuegi, 2020), they may be covering product or process inventions (Ganglmair et al., 2022); they may differ in their private economic value (Kogan et al., 2017b) or social usefulness (Griliches et al., 1986; Jaffe et al., 1993); they can be more novel (Arts et al.,

[8] "Fluctuations in the Rate of Industrial Invention" (Merton, 1935), page 1

[9] While (Merton, 1935) focuses on patent statistics it advocates for the study of patent statistics in conjunction with case studies in the history of technology in order to explore comparative rates of technological change.

[10] We will see that the interest underlying the study of an event *e* lies in the economic factors that the event represents.

[11] The use of patent data also comes with significant limitations, most notably that not all innovations are patented. Firstly, not every invention can satisfy the patent eligibility requirements defined by the USPTO, which stipulate that the invention must be unique, non-obvious, and have a commercial purpose. Secondly, the choice to patent is a strategic one for the inventor, who may instead choose to depend on confidentiality or other forms of securing ownership rights. (Hall et al., 2001)

2021), broader (Marco et al., 2019), or more impactful (Kelly et al., 2021). The nature of the relationships among these characteristics is complex (Higham et al., 2021) and subject to countless studies. Many scholars have provided access to calculated or estimated measures of the different patents' aspects, allowing for a fine representation of the type of invention embedded in the patent documents, which provides a platform for characterizing the patent counts across multiple classes (CPCs)

$$\underbrace{WPC_t(\alpha,\delta)}_{\substack{\text{Weighted Patent}\\ \text{Count}}} = \sum_{i=1}^{n_t}(1^{\delta}+A_i^{\alpha}) = \underbrace{n_t^{\delta}}_{\substack{\text{Number of}\\ \text{patents}}} + \underbrace{\sum_{i=1}^{n_t}A_i^{\alpha}}_{\substack{\text{Number of}\\ \text{patents}\\ \text{with attribute A}}} \tag{3.1}$$

where $\alpha > 0$ and $\delta \in \{0,1\}$ [12]. This formulation allows us to adjust the emphasis placed on the attribute A in the calculation of the *Weighted Patent Count*.

Allowing a more flexible definition of patent counts we can better characterize the nature of inventive activity. These properties could be interpreted to embody distinct aspects of innovation and technology development, thus offering a rich perspective on technical trajectories. A weighing scheme often used in the literature is by citation, that is A is `citation received by the patent` (Trajtenberg, 1990), where the number of citations received (forward citations) is considered a measure of the technical or social value of the patent. Other aspects could relate to the `novelty` of the invention that has been measured with the text of the patent (Arts et al., 2021) or leveraging patent classification schemes (Fleming, 2001). An attribute (A) need not be bound to aspects of the specific invention, but it could be, for example, related to the organization producing

[12] In the equation above, $WPC_t(\alpha)$ represents the *Weighted Patent Count* at time t, a metric that takes into account various characteristics of patents.On the right side of the equation, the sum $\sum_{i=1}^{n_t}(A_i^{\alpha}+1^{\delta})$ consists of two parts: the total number of patents, denoted by n_t^{δ}, and the number of patents with a specific attribute A, denoted by $\sum_{i=1}^{n_t}A_i^{\alpha}$. In the sum $\sum_{i=1}^{n_t}(A_i^{\alpha}+1^{\delta})$, each patent i from 1 to n_t (total number of patents at time t) is assigned a weight of $(A_i^{\alpha}+1^{\delta})$. Here, A_i^{α} represents the weight for patent i having the attribute A, and the exponent α is a parameter greater than 0. If a patent has the attribute A, it contributes A_i^{α} to the sum. On the other hand, every patent contributes 1^{δ} to the sum. In this equation, δ is a binary parameter that takes the values of 0 or 1. In the variables used in the empirical sections, when constructing weighted patent counts I will use $\alpha = 1$ and $\delta = 0$.

the focal invention. For example, if we set A to `assigned to small entity`, we can investigate how certain events/treatments induce technical production by small entities.

So far we have defined counts within a timeframe t, we need to add a second dimension that specifies the unit of analysis. The unit of analysis will determine the type of research questions that can be answered with a focus on the rate and direction of technical change. We need to have an entity and a time to set up a longitudinal dataset. The entities that are often used in this domain are *technology classes, patents, firms, inventors* [13]. In the following sections, I will describe in detail different setups which use *technology classes, patents, firms, inventors* as entities; the aim of the discussion is to provide a general (enough) framework that would allow the reader to design a difference-in-differences analysis in this domain. A central aspect of this design is the selection of the entities that are untreated, which is necessary when the event/treatment is fixed for all the treated entities [14] but not essential when the treatment is staggered [15]. I will provide a survey of papers that use setups with different entity classes to answer questions related to the factors shaping outcomes related to patenting. In surveying the articles I will focus on the economic motivations underlying the different research questions and I will try to describe [16] the empirical setup used by the authors. [17]

[13]The 4 listed classes of entities are not an exhaustive list of the set of entities that can be and have been used, in the study of patent data. Other relevant classes, for example, are defined by geography (Ballesteros, 2021). The 4 classes that I will focus on (tech, patent, firms, inventors) are the most used in the economics of innovation literature, however, selecting the entity class to study falls into the creativity of the researcher and the research question.

[14]Refer to Table 1.1 for the case in which the treated entities undergo an event fixed in time and a number of entities never undergo an event/treatment.

[15]Refer to Table 1.3 for the case in which all entities eventually undergo an event/treatment.

[16]To the best of my understanding

[17]Reviewing and discussing the findings of the different papers is beyond the scope of this survey. Please refer to the cited papers to study the empirical findings.

3.2.1 Entities

Technology Class

A natural starting point to analyze different rates of patenting and the direction of development efforts is the patent classification system. A patent classification is a systematic codification that offers a structured approach to categorizing inventions according to their technological field or specific features. This method enables the efficient organization and retrieval of patent information, helping patent examiners, researchers, and other stakeholders to navigate through the vast collection of patents [18]. By using a consistent and comprehensive classification system, such as the CPC, inventions can be effectively grouped, compared, and analyzed, facilitating the tracking of technological trends.

The organization of the CPC classification makes it more suitable, with respect to USPC, for studies of patenting changes across areas. Given that the CPC is harmonized across jurisdictions one can more easily compare patenting activity across countries.

Several studies employ USPC (Baten et al., 2017; Moser & Voena, 2012; Moser et al., 2014) or CPC (Galasso & Luo, 2021, 2022; Watzinger & Schnitzer, 2022) in combination with a difference in difference estimation strategy to study a diverse set of questions related to the economics and management of inventive efforts.

Moser and Voena, 2012 articulate a tension in the ability, of domestic inventors, to freely use foreign inventions. This scenario reflects policy interventions regarding compulsory licensing, which refers to the practice of allowing third parties to use patented inventions without the consent of the patent holder, usually in exchange for a fee. They argue that the availability of foreign inventions in a domestic market could reduce incentives to research and develop alternative technologies through a price mechanism; on the other hand, the paper argues that compulsory licenses may increase learning and incentives for complementary inventions. In order to propose evidence for the latter mechanism Moser and Voena, 2012 they exploit the Trading with the Enemy

[18] In addition to serving as a searchable collection of patents with similar subject matter, patent classifications are used as tools for finding patents during patentability searches and for aiding in the designation of patent applications examiners for examination purposes.

Act (TWEA). A 1918, amendment of TWEA enabled the seizure of patents owned by the enemy nations. Consequently, by 1919, an organized initiative was underway to license German-owned patents to American businesses. Moser and Voena, 2012 measured the effects of compulsory licensing, comparing shifts in the quantity of US patents across 7,248 USPC subclasses from 1875 to 1939 (in organic chemistry), 336 of which are considered treated, where treatment is defined as whether a patent in the class was involved in one of the 727 compulsory licenses.

Baten et al., 2017 mirrors Moser and Voena, 2012 taking the opposite perspective; they ask what is the effect of compulsory licensing of foreign patents on the incentive to innovate by foreign firms. The authors motivate this question by underlying that economic theory can support two outcomes: on one hand, compulsory licensing could potentially discourage innovation in the long run by diminishing the perceived effectiveness of patents, which in turn reduces the anticipated returns on investments in research and development. This could discourage firms from committing resources to inventive activities. Conversely, models of competition and innovation suggest that compulsory licensing might stimulate innovation by intensifying the threat of competition. When faced with the possibility of new entrants or rivals gaining access to their patented technologies, market leaders may be motivated to invest more in R&D to maintain or expand their competitive edge.

Continuing with a setup centered on USPC, but focusing on the role of individuals in the US aggregate inventive output, Moser et al., 2014 suggests there is an intriguing tension concerning the impact of German Jewish scientists who fled Nazi Germany on US innovation. Some historical recollections propose that these émigrés significantly transformed and boosted innovation in the United States. However, alternative narratives suggest that their contributions may have been limited due to administrative obstacles and prevailing anti-semitism, which during the economically challenging 1930s, was a prevalent reality not only in American universities but also in other sectors of the US economy (Sachar, 1993). This tension highlights the complex interplay between the potential positive impact of the German Jewish émigrés on US innovation and the social and institutional barriers they faced, which could have limited their ability to contribute to scientific advancements and technological progress.

Shifting the focus to a CPC level analysis, Watzinger and Schnitzer, 2022 ask a timely and consequential question: what is the effect of the breakup of a large corporate conglomerate on innovation [19]? The breakup of large dominant companies is among the potential solutions in antitrust cases that investigate unlawful actions such as exclusionary practices (Fumagalli et al., 2018). Practices, by dominant companies, that may hinder innovation have recently been under scrutiny (Cunningham et al., 2021) creating momentum for empirical evidence of the consequences, on the innovation ecosystem, of both firm actions as well as potential remedies, such as corporate break-ups. To provide evidence on some of the consequences of large conglomerate break-ups Watzinger and Schnitzer, 2022 study US patenting activity around the 1984 break-up of Bell System [20]. To study patenting arount the break-up the authors construct a panel where the entities are CPC groups; the treated entities are CPC groups in which Bell Labs had at least 5 patents prior to 1984 and the untreated entities are the CPC groups that are hierarchically connected by a CPC subclass to the treated group.

The hierarchy defined by the Cooperative Patent Classification scheme, allows researchers to take a fine-grained approach by defining treated and untreated entities within the same narrow CPC group. After identifying treated entities [21] researchers can define an untreated group using the CPC dot hierarchy. Galasso and Luo, 2021 exploit CPC subgroups to study how shifts in market demand, driven by consumer risk perception, influence the rate and nature of firms' innovative reactions. The overall effects that risk perception may have on firms' inventive efforts is ex-ante uncertain, as the authors set forth: heightened risk perception may steer innovation toward technologies that mitigate risk: what the authors refer to as risk mitigating technologies. On the contrary, an amplified perception of risk may dampen the demand for that product class, which could in turn curb re-

[19] Where innovation is measured by patenting

[20] Bell System was a vertically integrated telecommunications giant, and, at the time, one of the largest companies in the US

[21] Identifying treated groups can be done in different ways depending on the event of interest. Methods often rely on selecting objects that are at different "levels" in the data, for example by identifying patents that are involved in an event by using the text of the patent or other relevant information such as the assignee and then using the patent classification system to define what groups are treated.

search and development investments; additionally, as Rosenberg, 1976 suggested, in order for a demand-driven mechanism to function in technological advancements, the necessary technological capabilities must be present. Galasso and Luo, 2021 take advantage of a significant incident involving radiation overdoses that occurred in Los Angeles, specifically at Cedars-Sinai Medical Center. Over an 18-month period from 2008 to 2009, it was discovered that more than 200 patients were exposed to excessive radiation during CT brain perfusion scans, a type of imaging test used to diagnose strokes. In 2010 the incident made national news shifting the public perception of the devices involved. To study shifts in inventive efforts around that date the authors build a balanced panel, of CPC subgroups in which the treated entities are a set of CPC subgroups that the authors closely determine [22] to be related to radiation mitigation technologies. The untreated entities are CPC subgroups which are indented under the CPC main group of the treated subgroups.

The same authors (Galasso & Luo, 2022), using a similar approach, still in the medical device industry, take advantage of the bankruptcy of the leading producer of TMJ implants, to study how downstream inventions are affected by product liability shocks. Similarly to their previous paper, they highlight a tension with regard to the creation of risk-mitigating technologies: on one side, heightened liability could dampen innovation incentives by suppressing the demand or increasing the costs for technologies linked to higher risk. Conversely, it might boost demand for technologies that mitigate risk and reduce the probability of injuries, increasing the profitability of safer product designs. The empirical setup follows closely Galasso and Luo, 2021 but the technological classification used is the USPC. The authors build a panel, ranging between 1985 and 1995, of USPC sublcasses and they categorize to be treated if they "contain" a patent related to implants. [23]

Firms

Literature has shown substantial interest in estimates of firms' direct efforts in inventive activity, often with regard to competitive behavior. Early empirical models of competition, a central aspect

[22] By carefully reviewing the descriptions of the CPC subgroup, additionally including the main group as treated

[23] They classify patents to be related to implants by using the patent text

of economic activity, have shown higher social welfare in competitive environments where prices are lower and production more efficient. With regard to increases in social benefits, the provision of innovation has been demonstrated to be a driver of economic growth, increasing overall returns. Scholars have thus displayed broad interest in the factors that may lead to the underprovision of inventive activity, for instance, a notable observation supporting the underprovision of inventive activity, in both theoretical and empirical models of research and development (R&D) investments, is the significant disparity between social returns, facilitated by knowledge spillovers, and private returns (Arora et al., 2021; Bloom et al., 2013; Griliches, 1992). The interrelationship between competition and innovation has thus been central since early discussions on whether the presence of market power is beneficial to inventive activity (J. Schumpeter, 1942) or whether a more competitive environment would create the correct incentives to spur innovation (Arrow, 1972). Empirical findings have so far pointed to an inverted U relationship (Aghion et al., 2005) which would indicate that a moderate amount of market power may lead to greater innovation.

Recent attempts to provide evidence in this debate have used difference-in-differences with the focal entity of interest firms and their patenting activity. Kang, 2021 takes advantage of the information provided by the US Department of Justice (DOJ) to [24] provide differences in estimates of patenting activity around the *formation* and *breakup* of price-fixing cartels thus trying to establish a relationship between price competition and patenting intensity.

Patents are not the only mechanism firms utilize to protect their intellectual property. In industries in which patenting is not the main protection instrument, secrecy and other means are fundamental for sustaining companies' competitive advantage (Cohen et al., 2000). Firms can however shift their reliance on patenting vs secrecy, as they often do based on modifications in the legal landscape surrounding the protection of trade secrets (Contigiani et al., 2018).

Non-competes are a legal instrument firms employ for protection from departing employees who have access to valuable information. Building on the evidence of the significant risk that employee turnover poses in terms of information leakages, (Kang & Lee, 2022) investigate how this particular threat influences a company's utilization of patents or secrecy for safeguarding its

[24] DOJ antitrust case filings

proprietary information.

Kang and Lee, 2022 propose that companies adapt their knowledge protection strategies based on the perceived risk of employee attrition. In order to ascertain a connection between the threat of employee turnover and strategic choices related to patenting, the researchers utilize a pivotal 1998 court decision (Application Group, Inc. v. Hunter Group, Inc., 61 Cal. App. 4th 881) that altered the risk of employee turnover for companies outside California. They build a balanced longitudinal dataset, spanning from 1994 to 2002, of firms and assign firms to a *treated group*: firms headquartered in states, in the US, with high non-compete enforceability and an *undreated group*: firms headquartered in states with low non-compete enforceability. They then compare changes in patent filings before and after 1998 between firms in high-enforceability states and low-enforceability states with the idea that firms in high-enforceability states would shift their mean of intellectual property protection to patenting given that they would not be legally protected by non competes if essential employees would move to California after 1998. Legal system changes represent a significant factor of economic relevance in innovation studies due to their potential for exploitation. The nature of such changes often has far-reaching implications on the landscape of innovation, directly affecting how firms protect their intellectual property rights, how they strategize their R&D, and how they position themselves in competitive markets. Shifts in legal stipulations governing patents, trade secrets, or other intellectual property protections, for instance, can create strategic opportunities or challenges for firms. Organizations may take advantage of these transitions in legal frameworks to fortify their market positions, drive technological advancements, and maximize economic benefits. Concurrently, these changes can stimulate broader sectoral or societal shifts, thereby forming a crucial area of interest in the literature of innovation economics. The dynamism and economic implications of these legal shifts underscore the imperative for an intricate understanding of law and its interplay with innovation processes.

Aydin Ozden and Khashabi, 2022 exploit a case heard by the Supreme Court of the United States in 2006, eBay Inc. v. MercExchange, L.L.C. [25], a decision that significantly changed the standard for issuing permanent injunctions in patent infringement cases. By setting up a firm-

[25]MercExchange owned a patent that it claimed eBay's "Buy it Now" feature was infringing. The district court found that eBay did infringe on the patent but declined to issue a permanent injunction against eBay, a

year panel dataset Aydin Ozden and Khashabi, 2022 estimate the difference between licensing rates [26] before and after, eBay v. MercExchange, where the treated firms are US-headquartered firms and European firms are never-treated. They find that the implications of these legal measures bear significant weight on the licensing strategies employed by firms. The potential imposition of an injunction could bolster the bargaining leverage of an innovator during licensing discussions, thereby enhancing the probability of reaching an agreement. Conversely, the prospect of accruing damages, such as continuous royalties, may not prove as potent in encouraging preemptive licensing arrangements or in dissuading potential infringements.

Patents

A number of scholars have used variation within patents, and across time, to investigate changes in patent assignments, patent licensing, or follow-on innovations, due to factors such as litigation (Giebel, 2021), patent pledges (de Rassenfosse & Palangkaraya, 2023) or changes unexpected in

---

decision which was overturned by the Court of Appeals for the Federal Circuit (which specializes in patent law) on the ground that permanent injunctions should be issued once infringement and validity have been adjudged.

However, the Supreme Court reversed the decision of the Federal Circuit, ruling that the decision to grant or deny permanent injunctions in patent infringement cases should not be automatic or general, but should instead be subject to the discretion of the courts. According to the court, a four-factor test must be satisfied before a court may issue a permanent injunction: (1) That the plaintiff has suffered an irreparable injury; (2) That remedies available at law, such as monetary damages, are inadequate to compensate for that injury; (3) That considering the balance of hardships between the plaintiff and defendant, a remedy in equity is warranted; (4) That the public interest would not be disserved by a permanent injunction. The decision effectively made it harder for patent holders to obtain permanent injunctions against infringers because they now need to satisfy this four-factor test rather than simply proving infringement. This was particularly significant for "non-practicing entities" or "patent trolls" (entities that own patents and generate revenue through licensing and litigation, rather than manufacturing products or supplying services based on the patents they own), as courts may be less likely to find that they suffer irreparable harm from infringement or that a remedy in equity is warranted.

[26]The number of licensing agreements the focal firm is involved with over the patent stocks of the firm.

patent scope (Lee, 2023).

By refining the scope of analysis from a general patent class (Cooperative Patent Classification, CPC) down to a specific individual patent, the stakeholders who stand to benefit from such analyses also shift. This transition in focus offers nuanced insights that could prove invaluable for different entities, primarily firms that might be directly impacted by the outcomes of these studies.

A broader level of patent analysis usually delivers results that are of interest for policy-making entities. Such evaluations typically involve the macroscopic examination of patent classes, enabling the generation of high-level insights about trends, and consequently, enabling the formulation of robust policies that promote innovation, protect intellectual property rights, and drive economic growth. On the other hand, a granular analysis focusing on an individual patent can provide insights that are of practical significance to firms, especially those involved in the development and exploitation of the specific technology under investigation. Understanding how various economic factors may mold the trajectory of a particular invention would be an asset for these companies. For instance, the study of shifts in follow-on innovation, i.e., innovation that directly builds upon a pre-existing piece of technology, can be enlightening. Changes in this dimension due to exogenous factors might directly impact the future development and application of a specific technology. These shifts could be induced by a plethora of factors such as regulatory changes, technological advancements, market demands, or even geopolitical considerations.

Three examples of this type and level of analysis are: Giebel, 2021, de Rassenfosse and Palangkaraya, 2023 and Lee, 2023. Giebel, 2021 analyzes how patent enforcement through litigation affects subsequent innovation. The author uses patent litigation cases in the U.S. as the "events" in a difference-in-differences analysis. The data consists of utility patents filed with the USPTO, which is merged with patent litigation data to identify litigated vs non-litigated patents. To construct the panel data, the author matches each litigated patent with a similar non-litigated patent based on characteristics like technology class, application year, etc. This may account for potential selection bias in which patents get litigated. The outcome variable is forward citations received in a given year, a common proxy for follow-on innovation. Using this difference-in-differences approach, Giebel, 2021 finds that subsequent innovation, as measured by forward citations, increases

after a patent litigation case is filed. The effect is largest during the litigation period but stays positive in the years after the case closes. This may suggest enforcement encourages follow-on innovation, perhaps by signaling value and reducing information asymmetry. However, the author finds the new patents cite the litigated patent heavily, implying more incremental than radical innovation.

de Rassenfosse and Palangkaraya, 2023 examines a different legal mechanism: patent pledges, which are public commitments by patent holders to limit enforcement. Patent pledges have become increasingly popular, but their effects on innovation have been an open question. Large companies like Google, IBM, Tesla have pledged patents, as well as universities and other institutions. de Rassenfosse and Palangkaraya, 2023 investigate how such patent pledges influence downstream innovation, leveraging a novel dataset of over 1,200 pledged patents from 2005-2017. Understanding the impacts of patent pledges is important for several reasons. Pledges can theoretically facilitate technology diffusion by reducing barriers to accessing the patented invention, which may accelerate follow-on innovation as others can more easily build on it. Analyzing pledges also provides insights into the role of intellectual property and voluntary licensing approaches in shaping downstream innovation outcomes. To estimate the causal effect of patent pledging, the authors implement a difference-in-differences methodology comparing pledged patents to matched non-pledged controls. The control group is carefully matched on characteristics like technology class, application year, quality metrics, and textual similarity using the full patent documents. This helps address potential endogeneity issues in which patents are selected for pledging. The outcome variable is forward citations received in a given year, capturing follow-on innovation building directly on the patented technology. The empirical strategy compares how forward citations change after pledging for the treated patents versus the matched control group. The authors show patent pledging substantially increases subsequent forward citations - by over 50% in the 5 years after pledging based on the estimates. This effect is particularly strong for high-quality patents. However, pledges related to software patents have a more limited impact, perhaps due to more restrictions. Overall, the findings suggest publicly committing to open access does accelerate follow-on innovation.

In contrast to the prior two studies which considered forward citation as the main outcome vari-

able, (Lee, 2023) analyzes the relationship between patent scope and licensing propensity. Prior studies have shown mixed results on how patent scope affects licensing, and have been limited to specific industries or settings. This paper combines a broad set of publicly reported licensing deals with a novel methodology to systematically study the patent scope-licensing linkage. The analysis leverages USPTO patent data to construct a panel of patents over time. The key independent variable is an exogenous shock to patent scope based on the discovery of similar prior art after the priority date. Using a difference-in-differences approach, the paper compares how licensing propensity changes for treated patents versus matched controls. The results show decreasing patent scope substantially reduces licensing propensity of inventions. Furthermore, the negative effect is stronger for high-quality, science-based, novel, and small inventor patents.

Across the papers reviewed, constructing panel data at the patent-year level provides important analytical advantages for studying the impacts of patent policies and events on innovation outcomes. The patent-year structure allows the analyses to leverage within-patent variation over time, such as before and after a litigation event or policy change. This helps establish causality and avoids bias from time-invariant differences between treatment and control patents. Combining patent data with information on specific events or policy reforms is crucial to generate exogenous shocks to patent strength, scope, or enforcement. The resulting patent-year panel data enables difference-in-differences designs that compare how outcomes change over time for affected patents versus unaffected controls. This quasi-experimental approach facilitated by the panel structure underpins the causal inferences on the impacts of patent policies made across the literature.

Inventors

The rate of inventive activity explained by organizations' R&D dollars can be broken down into multiple factors one of which is the utilization of human capital inputs to the innovation. The relationships between inventors-related aspects and innovations have long been central to economics and management. Recent interest around the relationship between human capital, talent (Akcigit et al., 2020), parental resources and childhood environments (Bell et al., 2019) and innovation led to large-scale and accurate disambiguation of inventors in open access datasets such as PatentsView

(Monath et al., n.d.) as well as in Census data (Akcigit & Goldschlag, 2023a).

Access to disambiguated inventor data allows researchers to include inventors among the classes of entities that can be utilized in difference-in-differences analysis to investigate factors that shape inventors' choices and outcomes related to the rate of their efforts and the direction of their pursuits.

For instance Akcigit and Goldschlag, 2023b, exploit Census data to show changes in earnings and patent application after inventors get hired by incumbent firms.

Brown and Roche, 2022 investigate how reduced patent enforceability affects the mobility of inventor-employees. The specific event is the Supreme Court's eBay v. MercExchange ruling in 2006, which substantially decreased the use of permanent injunctions in patent infringement cases. This reduced the threat of injunctions against infringing firms, lowering the risks and penalties for patent infringement. By decreasing patent enforceability, the ruling potentially made it easier for competitors to hire inventors and use patented technologies. To analyze the impact on inventor mobility, the paper constructs a panel dataset tracking over 50,000 early career inventors before and after the eBay ruling. Using a difference-in-differences approach, it compares how an inventor's likelihood of departing their job changes based on the firm's reliance on intellectual property, as measured by patents per inventor in the field. Firms with more patents per inventor provide a treatment group more affected by the reduced enforceability. The results show inventors at high patent-reliant firms have increased mobility after the ruling. By leveraging the shock of the eBay ruling and constructing inventor-level panel data, the paper provides causal evidence that weakening patent enforceability expands mobility options for valuable inventors.

3.3 pyDRAD

In this section, I will describe the pyDRAD software, trying to abstract from the specific implementation details as much as possible in order to provide the reader with an overview that is independent of the execution specifics. pyDRAD stands for "python Differences in the Rate And Direction." As the name suggests, pyDRAD is written in Python and allows users to quickly download and organize datasets from different sources related to innovation. Currently, pyDRAD supports importing

data from Reliance on Science, PatentsView, and openFDA - each containing various datasets that can be accessed openly from the respective websites. However, pyDRAD takes this a step further by providing an interface to pull the data directly from the sources and organize it locally on the user's machine in a standardized way for fast querying. The raw files are automatically converted to Apache Parquet format, which is optimized for nested data and provides performance improvements for large datasets.

After downloading, the data is organized into folders. The user can then call relevant functions in pyDRAD to quickly load the data into memory or scan it lazily. Beyond simple data access, pyDRAD also contains predefined pipelines built around common datasets used in the economics of innovation literature. The main purpose of a data pipeline is to automate the repetitive steps needed to make raw data ready for downstream analysis. Pipelines should improve efficiency, reduce errors, and provide consistency in data processing. They allow the researcher teams to focus on high-value tasks like analysis rather than manual data wrangling.

pyDRAD contains predefined pipeline functions that preprocess common datasets for tasks frequently required in research projects. These pipelines encapsulate multi-step data processing logic to clean, transform, and join data from sources like PatentsView and Reliance on Science. By using the pipelines, researchers can avoid having to write boilerplate data-wrangling code themselves for every new project. Instead, they can focus on higher-level analysis and modeling, enabled by pyDRAD's reusable data processing functions. Some examples include: joining together datasets from different sources and creating new derived datasets, such as calculating forward citations, weighted citations, novelty scores, etc. Thus, pyDRAD goes beyond simply accessing raw data from public sources. It provides an easy way to download, organize, and preprocess datasets frequently used in innovation economics.

Beyond common preprocessing steps, pyDRAD allows for pipelines that are designed around econometric tasks, for example creating datasets that are structured as panel (longitudinal) datasets. As we have previously observed, structuring the data as a panel is a typical setup in the literature of the economics of innovation, especially if the aim is estimating treatment effects, for example using difference-in-differences. Starting from the raw data, pyDRAD allows the user, by stating which

columns to use as "entity" and "time", to preset the steps that one would need to take in order to convert a dataset into panel format. These steps often involve aggregations and joins of different datasets; these steps can be declared before execution which means that the pipeline would return a plan rather than data already structured as a panel. The benefit of returning a plan is that execution is deferred and the user can change the configuration that was used to preset the panel-transformation. Configurations may include the period frequency at which the panel is constructed, the entities to be included in the final data[27], definitions of variables though aggregations at the entity or entity-time level, or joins with other plans (as long as the entity units are comparable). By allowing the configuration of panel datasets through declarative plans rather than imperative code, pyDRAD enables flexibility and optimizations in constructing panel data. Users can iteratively refine the panel structure without needing to re-execute previous steps. Planning the panel structure also allows querying larger-than-memory datasets by isolating the entities needed. Overall, pyDRAD's panel data capabilities provide advanced pipelines to facilitate econometric workflows[28].

---

[27]By enabling the user to preconfigure the entities to be included prior to execution, the manipulation of the data is done only on a subset of entities which eases memory requirements and enables the user to interact with larger than memory datasets with ease.

[28]I named the python object that allows panel data workflows DataPanel; this object automates multiple steps involved in creating panel datasets and it can integrate with different data sources, as long as the data has two dimensions that represent entities and times. To the best of my knowledge, the functionality of DataPanel as I described in the text is novel and unique: I do not know of any other software in any language that has methods that can create a plan to structure data as a longitudinal dataset with deferred execution. Of course one could write this functionality in any software that allows for lazy (deferred) evaluation, but the focus on panel data planning is unique to DataPanel, I think making it a contribution for both practitioners as well as developers. The key innovation is the DataPanel object, which automates complex workflows to convert data into panel structure. DataPanel allows configuring the panel structure through high-level declarations rather than imperative code. This includes specifying the entity and time dimensions, aggregation granularity, entities to include and joins with other data sources. Execution of the panel data plan is deferred, enabling iterative refinement of the structure without re-running. DataPanel also facilitates analysis on large datasets by isolating the subset of entities needed before execution. The ability to plan and optimize panel data workflows is a novel capability not found in other software. By leveraging deferred execution and op-

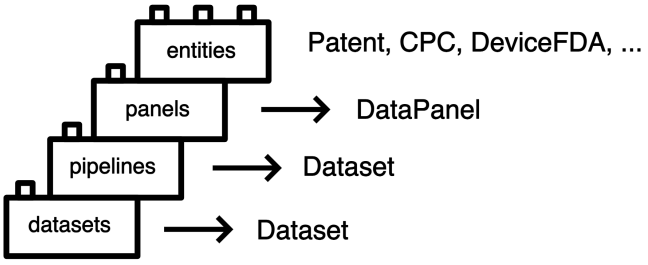

Figure 3.1: This figure shows the structure of the software. The rectangles represent subpackages, and they are arranged in a visual way to represent that the user can interact and plug in both data as well as functionality at each step

Figure 3.1 provides a zoomed-out picture of the current organization and design choices of the software. There are four sub-packages which build on each other. The baseline sub-package named "datasets" contains functions to access the raw dataset as they are distributed by the source, may it be PatentsView USPTO, 2019, Reliance on Science, or openFDA. [29] The functions in the "datasets" and in the "pipelines" sub-packages all return [30] a "Dataset" object [31]: a "Dataset" object is a wrapper around the DataFrame [32] which extends its functionality with methods to store and display the metadata associated with the data. I believe that the storage and display of metadata

timizations, DataPanel provides both usability and performance benefits for econometric analysis compared to constructing panel data manually. DataPanel is included in the package named 'differences' (which I authored and currently maintain); 'differences' is a Python package is built on top of Polars, and it facilitates data manipulation for econometric analysis dependency of pyDRAD and it currently contains functionality for inference and estimation of difference-in-differences models. pyDRAD depends on 'differences', meaning that some of the classes and functions that are defined in pyDRAD directly call classes and functions from 'differences'.

[29]These are the three sources included in pyDRAD at the time of writing. Once pydrad.datasets is imported into the namespace, one can call a function named after the source: datasets.patents_view(), datasets.reliance_on_science() or datasets.open_fda(). Given that these sources contain multiple different datasets, the first argument of each function would need to contain the name of the requested dataset.

[30]The arrow on the right of the boxes represents a functional return sign

[31]Specifically a Python Class.

[32]Specifically a Polars DataFrame or a LazyFrame.

greatly facilitates and releases the mental burden for the researcher when dealing with large and multifaceted datasets. The metadata for the raw data (in the "datasets" sub-package) is retrieved from source. On the other hand, "pipelines" that manipulate the raw data would need to either store preconfigured metadata or user-added metadata. The sub-package named "panels", as mentioned previously contains specialized pipelines that store a plan to convert datasets returned in "datasets" or "pipelines" into longitudinal data. These are necessary steps that a researcher often is faced with when needing to conduct econometric analysis, especially when concerned about over dynamics. The object returned by the functions in "panels" is not a Dataset in the form of a longitudinal panel, it is a DataPanel which is a wrapper for a DataFrame that contains a plan to build the variables relevant for a particular panel dataset. The plan is stored in DataPanel and the user, with a handle on the DataPanel object has the ability to materialize that plan into a DataFrame that has a longitudinal form (with entity and time identifiers). The particular settings to construct the panel, for example, period frequency or entities to consider, can be easily changed at runtime. Finally, the "entities" sub-package contains a set of Python Classes designed to interact with the datasets, pipelines, or panels through the lens of a particular entity/unit type. Let's take the CPC object, it is designed to interact with the data with a focus on patent classes. One functionality could be retrieving the count of patents by CPC, or the count of patents by CPC-year. Another functionality could be creating a set of treated and control CPCs, a common task in innovation studies as we have seen in Section 3.2.1. The CPC object contains access to the hierarchy structure of the CPC and can facilitate the construction of treated and control classes; for example, starting from a single CPC can define connected CPCs that can be considered treated according to the researcher's assumptions. Figures 3.2 to 3.5 show possible options of such functionality. Other objects in the "entities" subpackage would have similar functionality but with references to other units/entities.

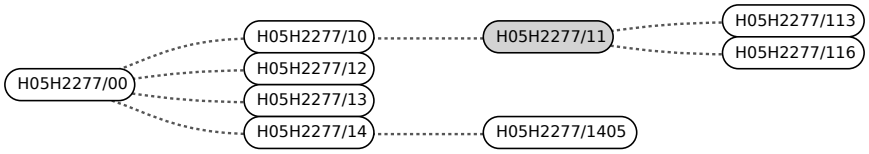

Figure 3.2: This figure shows the setup in which only the selected CPC (H05H 2277/11) is considered treated.

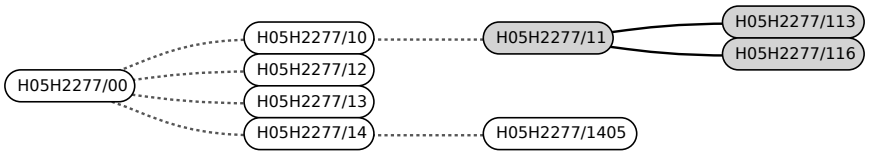

Figure 3.3: This figure shows the setup in which the selected CPC (H05H 2277/11) is considered treated along with its descendants.

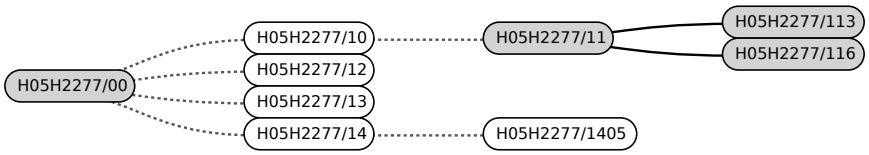

Figure 3.4: This figure shows the setup in which the selected CPC (H05H 2277/11) is considered treated along with its descendants and the root of H05H 2277/11' subtree.

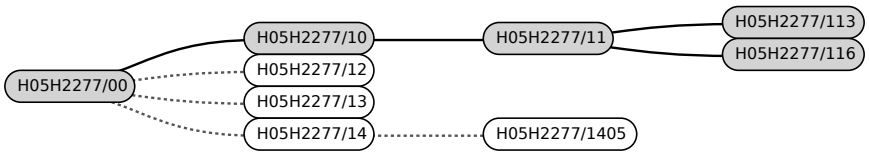

Figure 3.5: This figure shows the setup in which the selected CPC (H05H 2277/11) as well as all the descendents of H05H 2277/11' subtree.

3.4 Conclusion

In this paper, I first tried to give an overview of the rich academic literature employing difference-in-differences estimation in the economics of innovation, its practical applications, and the resulting nuances that arise from this analytical practice. The aim of this review has been to inform researchers how quasi-exogenous changes, or inducement events, can induce shifts in innovative activities and affect the rate and direction of technical change.

Second, I introduced pyDRAD, an advancement in the form of open-source software, specifically designed to aid difference-in-differences estimation in the realm of innovation and technical change, with the hope that the availability of this software will aid future research in this field, potentially simplifying the analytical process and broadening the accessibility of robust investigative methodologies.

Chapter 4

Predicating Acquisitions. Post M&A Follow-on Innovation in the Medical Device Industry

4.1 Introduction

The last several decades have witnessed significant shifts in the American innovation ecosystem, characterized by a growing division of innovative efforts between large corporate entities and small startups, often backed by university research; these changes led to a growth of external technology sourcing at the expense of internal research (Arora, Belenzon, Patacconi, & Suh, 2020). This trend is further evident in the rise of VC-backed startup exits via mergers and acquisitions, surpassing the previous dominance of initial public offerings (Gao et al., 2013).

Prompted by the surge in M&A activity, researchers have highlighted a rise of oligopoly power (Ederer & Pellegrino, 2023). While the motivations for M&A activity are complex, two broad economic drivers often underlie a firm's decision to acquire another company: greater efficiency (possibly due to complementary capabilities) or increased market power. Acquisitions driven by complementary capabilities allow the acquiring firm to more efficiently develop and commercialize the target's products than the target could alone. The combined capabilities may generate efficiencies in R&D, production, and distribution, ultimately benefiting consumers through lower prices or improved products. Careful analysis is needed to verify that efficiency gains are the primary motivation rather than excuses for increasing market dominance. Alternatively, M&A deals may aim to increase market power by absorbing competitors. This consolidated market power enables firms to raise prices and restrict output, directly harming consumer welfare. Distinguishing between these motivations is a key challenge for antitrust enforcement. While research has extensively examined the price effects of mergers [1], recent work has began unraveling efficiency implications (Demirer &

[1] A large literature has shown the effect of acquisitions on prices in various industries and segments (Prager & Hannan, 1998; Town, 2001; Vita & Sacher, 2001)

Karaduman, 2022), but more research on the trajectory of a target's products after acquisition could further illuminate product dynamics post-acquisition.

The focus of this paper is to measure the extent to which acquiring firms continue the development of acquired, but readily commercializable [2], products. This is an open empirical question that builds on the opposing views of why acquisition happens; on one hand, synergy and cost advantage considerations pose that the acquiring firm may be best positioned in the further development of technical products (Ahuja & Katila, 2001), especially in the presence of complementary assets (Teece, 1986). On the other hand, there are considerations that would lead to terminating product trajectories rooted in multiple factors. These include duplication of research efforts (Cassiman et al., 2005), integration challenges (Paruchuri et al., 2006), lack of fit with the acquirer's existing portfolio (Cloodt et al., 2006), or an intent to eliminate future competition and counteract the "gale of creative destruction" (J. A. Schumpeter, 2013) by absorbing smaller companies (Cunningham et al., 2021).

The recent theory, strikingly termed "killer acquisitions" (Cunningham et al., 2021), has gained momentum in policy discussions; according to the theory, established firms (often dominant or incumbent firms in an industry) acquire innovative startups or smaller competitors not necessarily for the potential of the smaller firm's products or technologies, but to preemptively shut down or curtail the acquired firm's projects that could have become future competition. Essentially, the acquisition is made to "kill" the potential threat (Cunningham et al., 2021). Given the potentially anti-competitive nature of "killer acquisitions", they have drawn attention from antitrust authorities and policymakers. It is a challenge for regulators to distinguish between legitimate acquisitions and those made with the intent to stifle competition, especially since the acquired technology or project might be in the early stages and its future success uncertain.

Notably, during the 2020 shortage [3] of ventilators, essential for addressing the COVID-19 pandemic caused by SARS-CoV-2, some argued that the 2012 acquisition of Newport Medical Instru-

[2]The development stage of the products under consideration is a key difference from this paper and the theory & evidence provided by Cunningham et al., 2021 which studies products that are in early stage development.

[3]Coronavirus US ventilator shortage, NYT 2020-03-29

ments by Covidien for USD 108 million acted as an anti-competitive move. This acquisition is believed to have potentially slowed or even halted the advancement of newer, more affordable ventilators needed for various respiratory viruses, including COVID-19 (OECD, 2020). Such claims initiated debates about whether the Newport Medical Instruments/Covidien merger serves as empirical evidence supporting Cunningham et al., 2021 theory. Others suggest considering factors like the relatively minor deal value, the probable existence of other competitors, and the potential lack of market overlap between Newport's and Covidien's ventilators, which might argue against antitrust intervention [4].

Beyond deliberate anti-competitive actions, the post-acquisition management of internal technology development and associated continuation or termination decisions pose several challenges, many of which I would argue are parallel to managing an increasing asset portfolio. Under the assumption that managers aim to maximize the expected value of technology in their continuation or termination choices, when the choice set increases, the average probability of continuation may decline, for instance in an economy with N firms, where each firm i offers a set of products M_i, we can think of the probability firm i developments/continues product z as $p_{iz} = \frac{p_0}{1+\beta M_i}$, where p_0 is the baseline propensity for firm i to develop product j when it has no existing products and β captures the diminishing propensity for product development as a firm's portfolio grows. Now if firm i acquires firm j, the probability product j will be developed, for the merged entity, abstracting away possible multiplicative factors that could be created by synergies or complementary assets, or dampening factors derived from post-merger complexities, would be $p_{iz} = \frac{p_0}{1+\beta(M_i+M_j)}$.

In addition to the issues pertaining to the post-acquisition management of multiple assets, decision-makers face relevant frictions stemming from information asymmetry, agency conflicts, and cultural mismatches between the acquiring and target entities. These challenges can skew optimal decision-making, potentially resulting in the inefficient allocation of resources.

The Newport Medical Instruments/Covidien ventilator case underscores the critical need for

[4]Todo: explain further in this footnote the points made by The Covidien/Newport Merger: Killer Acquisition or Just a Killer Story? and page 16 "OECD (2020), Start-ups, Killer Acquisitions and Merger Control"

in-depth examination of the innovation trajectory landscape within the medical device sector. The complex interplay of strategic acquisitions, particularly those that could stifle advancements or potential innovation, has profound implications for not only industry dynamics but also public health and our ability to innovate quickly in times of need.

A recent manifestation of regulatory vigilance in this sector was the FTC's close scrutiny of CooperCompanies proposed acquisition of Cook Medical's reproductive health portfolio, which produces medical hardware for in vitro fertilization and obstetrics/gynecology. The federal oversight led to CooperCompanies' decision to withdraw from the acquisition, underscoring the regulators' commitment to preserving competitive vigor in healthcare markets. The director of the FTC's Bureau of Competition emphasized this stance, stating, "The FTC is committed to protecting patients from higher costs and preserving the incentive to innovate. This deal termination protects competition and is a win for patients." [5]

In this study, I gather and analyze a novel and comprehensive dataset by combining multiple sources [6] to study companies' development trajectories of medical device technology after an

[5]FTC Statement Regarding Termination of CooperCompanies' Attempted Acquisition of Cook Medical's Reproductive Health Business.

The preservation of innovation trajectories is at the center of antitrust enforcement operations, in the speech, titled "Competition: The Mother of Invention", Margrethe Vestager, EU's competition commissioner, said "when we look at high-tech mergers, we do not just look at whether they might raise prices. We also assess whether they could be bad for innovation. Last year, we looked at a merger between the drug company Pfizer and its rival, Hospira. We only approved the deal after Pfizer agreed to sell the European rights to an arthritis drug it was developing. One concern was that Hospira already had a competing drug on the market, and we thought Pfizer might stop work on its own drug if the deal went ahead as planned. Which would have meant less of the innovation that we depend on as patients" (as reported in Haucap et al., 2019)

[6]Most if not all of these sources do not "speak" to eachother: combining them requires careful effort in standardization of the sources to be able to join them. For example in order to join company names between sources, I not only employed fuzzy matching techniques, these would lead to multiple matches that to the computer (and often to the human eye) would be correct, I also geocoded the universe of FDA applicants as well as corporate locations from the different corporate datasets used to filter names matched on text by restricting matches to various radii of geographical distance.

acquisition.

Firstly, I take advantage of the openFDA database to get detailed information on medical devices, including their IDs, specialties, and the dates they were approved. The openFDA database helps me build a picture of the medical device landscape from 2002 to 2020. I then use the FDA's 510k database to programmatically download summaries[7] of device applications, which, as I will describe in detail in the following sections, allows me to track how these devices evolve over generations. I then obtain data on acquired companies from SDC, Zephyr, and Crunchbase and I meticulously match firm names to FDA applicants to create a crosswalk between the FDA database and M&A data.

I employ a difference-in-differences framework to study the post-acquisition likelihood of technology continuation. This allows me to estimate the probability a medical device is used as a *predicate* by subsequent submissions to the FDA.

In the sections ahead, I'll delve into the literature that discusses the challenges companies face after an acquisition, focusing on technology development. Next, I'll discuss the data sources and a novel dataset on predicate relationships between devices as well as applicants' corporate acquisitions. While researchers in public health have recently started to use the information embedded in the *predicate* tree, it is still a fresh perspective in the field of strategic management. This unique dataset offers new insights into how acquisitions impact technology trajectories. I will explain the methods used to build the dataset and conduct my analyses.

The findings reveal that devices acquired tend to have a slower evolution into the next generation, with fewer being used as predicates post-acquisition. This observation is prominent when setting never-acquired devices as a baseline or when considering devices that haven't been acquired yet, while excluding the never-acquired ones. However, this doesn't suggest that acquirers are not

[7]In practice the documentation included in the 510k summaries is stored in multi page PDFs that companies need to attach to their applications and that are linked on the FDA portal. In order to obtain the PDFs I wrote a program that would pull the information from the portal. I retrieved summaries and additional documentation for over 65.000 applications approved between 1980 and 2022; prior to 2002 the summaries are not available in the same percentage of the total approved devices as post 2002, thus I will restrict my analysis to post 2002, see Figure 4.1 for more details.

developing the acquired devices. Rather, it indicates a shift in the pace of internal development compared to non-acquired devices. A plausible explanation could be acquirers concentrating on a specific subset of the acquired devices, mirroring typical product portfolio challenges. In addition, some evidence shows a surge in external development for acquired devices when juxtaposed with never-acquired ones. This insinuates that acquisitions might act as a quality beacon, signifying which devices are prime for further enhancement. Lastly, the acquirers' behavior post-acquisition is compelling. With the data set's inherent staggered acquisition timelines, it's evident that acquirers lean towards intensified development after an acquisition, likely catalyzed by internal knowledge dissemination.

This paper does not aim to test the "killer acquisition" theory, as there are marked differences between this empirical context and the theory and context in Cunningham et al., 2021. In the pharmaceutical context studied by Cunningham et al., 2021, acquisitions take place when the target firm's innovative project is still in its developmental phase. As a result, further development is both necessary and costly, and the eventual success of the project remains uncertain. Conversely, in the context of the medical device sector examined in this paper, the devices/products in question have already received FDA approval and are ready for commercialization. While development in this scenario pertains to the expansion of an existing product line, it is still accompanied by significant uncertainties. On the technical front, for example, devices might pose potential adverse events necessitating refinements. Similarly, uncertainties persist on the commercial front. However, the extent of technical uncertainty associated with early drug development, as explored by (Cunningham et al., 2021), is considerably more pronounced than that seen with FDA-approved 510k applications. Furthermore, the authors define acquisitions based on the degree of technical overlap between the acquiring and target firms. In contrast, this study deliberately sidesteps that dimension to lay out fundamental patterns of product continuations. I then introduce additional examinations to discern if product continuations stem from in-house development initiatives or external sources.

4.2 Related literature

Mergers and acquisitions (M&A) are undertaken for a variety of strategic reasons, but two broad economic rationales often underpin one firm's decision to acquire another: seeking greater efficiency or increasing market power. Acquisitions motivated primarily by efficiency gains involve consolidating complementary capabilities that allow the acquiring firm to more effectively develop and commercialize products compared to the target alone (Andrade et al., 2001). For example, combining the R&D strengths of the acquirer with the manufacturing expertise of the target can generate synergies that enhance innovation capacities and accelerate time-to-market. Bena and Li, 2014 finds that both acquirers and targets tend to be active innovators, but with different characteristics - acquirers often have large patent portfolios but lower R&D expenses, while targets frequently have high R&D expenses but slowing growth in patentable innovations. This complementary innovation profile supports merger motivations of synergies. Further, the probability of a merger increases with technological overlap between firms, as evidenced by proximity of patent portfolios, shared knowledge bases, and mutual patent citations. This provides firm-level evidence for mergers driven by desires to unlock synergistic gains. The combined capabilities and resulting innovations may also allow for efficiencies in production through economies of scale and scope, as well as improved distribution leverage utilizing existing networks and relationships. Ultimately, consumers stand to benefit from M&A activity driven by efficiency gains through higher quality products, lower prices, and/or greater variety. However, claims of efficiency motivations must be carefully scrutinized. Acquirers may offer efficiency arguments as pretext for deals that are primarily driven by motivations to increase market power and dominance Baker and Bresnahan, 1985. In-depth analysis is required to verify that true complementary capabilities exist between acquirer and target, that material efficiencies are likely to result, and that these efficiencies are pass-through to consumers in the form of lower costs or improved products. Antitrust authorities weigh these factors closely in reviewing proposed M&A deals.

While mergers and acquisitions offer the potential for significant synergies, the empirical evidence paints a more nuanced picture, with many deals failing to achieve the desired goals and benefits. Information gaps between the acquirer and target represent a salient challenge. Acquirers

often lack comprehensive insights into the true upside or inherent limitations of the technologies and innovation capacities they are acquiring. Without full transparency, they may prematurely terminate development projects that still hold promise under new leadership or, conversely, double down on initiatives that have constrained viability in new markets or under new strategies (Cassiman et al., 2005). Navigating integration complexities further complicates matters. Combining two formerly distinct entities involves more than just consolidation of physical assets and IP. A successful merger requires re-calibrating teams, systems, and cultures in a unified direction (Larsson & Finkelstein, 1999). This human dimension of integration is precarious; employees from the acquired firm may feel marginalized if their expertise is discounted or contributions go unrecognized post-acquisition. The loss of key talent integral to the target's innovation capacities disrupts integration momentum (Weber, 1996). Leadership from both sides must invest in transparent communication and cross-cultural team building to retain human capital through uncertain transitions. The "not invented here" syndrome compounds matters as acquirer employees resist externally sourced capabilities (Katz & Allen, 1982). Agency issues stemming from divergent management incentives also challenge merger integration (Jensen, 1986). Short-term performance metrics and personal biases may skew the resource allocation decisions of executives in ways misaligned with the strategic rationale underlying the deal (Hitt et al., 2001). Effective integration requires judicious distribution of finite resources across the expanded portfolio to nurture promising projects without over-investing in peripherals. While mergers promise synergies, realizing these potential benefits requires overcoming information barriers, thoughtful cultural integration, and a realignment of management incentives towards strategic – rather than short-term – goals. While internal integration factors play a major role, the broader innovation ecosystem also influences post-acquisition technology trajectories. In the medical device industry specifically, the regulatory regime creates an open system where any firm can leverage the approvals of "predicate devices" - existing products previously cleared by the FDA - as the basis for new submissions. This means that, theoretically, external firms remain able to build upon and extend the technological lineage of acquired products, even without the participation of the new owners. However, in practice, acquisitions can erect barriers for external players through strengthened intellectual property protections. The scale and resources of acquir-

ing firms allow them to fortify the IP defenses around acquired technologies. While this shields the competitive interests of the acquirer, it may also intimidate external innovators.

Focusing on innovative outcomes, scholars have often paid attention to post-acquisition efforts by the targets on technology development, as a measure of technology acquisition success. Companies in industries where R&D is pivotal are progressively turning to acquisitions to secure advanced technological expertise. The trend of purchasing smaller, technologically-focused entities has gained traction to bolster internal innovation, particularly in times characterized by rapid and consistent technological advancements (Sarkar et al., 2006; Bower, 2001; Chaudhuri & Tabrizi, 1999). The driving force behind these acquisitions is not a lack of capacity for internal innovation but the aim to bypass uncertainty, often time-related, associated with internal innovation. For instance, Cisco often opts for technology acquisition over internal development when the latter would take more than six months (Aguilera & Dencker, 2004; Sears, 2018).

More closely related to innovation outcomes, Joshi and Nerkar, 2011 investigate how strategic alliances shape innovation trajectories. They find that in the optical disc industry, a unique form of R&D collaboration, termed "patent pools", leads to a notable decline in both the volume and quality of patents by both licensors and licensees. This suggests a potential paradox wherein R&D collaboration, rather than facilitating innovation, might serve as an impediment to it, stifling systemic innovation amongst participant firms.

While many studies focus on post-acquisition performance at an aggregate level—examining stock market reactions or assessing innovation through aggregated patent counts—the evidence at the product level regarding continuation or termination decisions remains limited. One of the challenges in obtaining such data lies in the ability to clearly define the developmental trajectory of a product, requiring the linkage of products in a chronological sequence of technological advancements. However, the medical device industry, which is the focus of this study, offers a unique opportunity to analyze product-level post-acquisition technological trajectories. Class 2 medical devices are mandated to demonstrate substantial equivalence to previously marketed products approved by the FDA. By leveraging the substantial equivalence linkages, I can construct a tree representing product generations and their associated decisions of continuation or termination for each Class 2

medical device approved by the FDA [8].

4.3 The medical device industry

Medical devices, ranging from simple tools like stethoscopes to complex machinery such as MRI machines and cardiac ventricular assist devices, are intrinsic to modern medical practices. The industry drives significant direct and indirect economic value: with around 6,500 firms in the US, an output of $380 billion, a direct impact on about 519,000 jobs, and significant research & development investments, its technology has contributed to the reduction in the number of hospitalization days by 59% over the 1980-2010 period, reducing the number of work absences with relevant economic spillovers across the economy [9].

The industry's advancements also have direct health benefits. As medical devices evolve, they have the potential to offer more effective treatments, less invasive procedures, reduced hospitalization times, and improved quality of life for patients. Furthermore, the technical linkage of this industry with sectors like information technology, materials science, and biomedical research makes its development an indicator of broader technological and scientific progress.

The medical device industry is a critical sector that affects the well-being of millions of individuals worldwide, serving as a nexus of technological innovation, public health, and economic growth. With rising healthcare costs, an aging population, and emerging health threats, understanding the innovation dynamics of the medical device industry has become of central interest to scholars and policymakers. This paper delves into the economic dynamics of this industry, particularly in the United States, by examining the innovation patterns of firms after corporate acquisitions.

4.3.1 Regulatory system

A medical device is, "an instrument, apparatus, implement, machine, contrivance, implant, in vitro reagent, or other similar or related article, including a component part, or accessory..." (Section

[8] See Section 4.3.1 for more details.

[9] AdvaMed: Medical Device Industry Facts, Job Creation

201(h)(1) of the Food, Drug, and Cosmetic Act)[10].

In the United States, there is a regulatory framework in place to ensure the safety and effectiveness of medical devices. The Food and Drug Administration (FDA) stands as the primary authority for the oversight and approval of these devices. Based on the potential risks associated with their use, devices are categorized into three classes:

Class I Devices: These are considered low-risk devices. They are mainly subject to basic controls such as proper labeling. An example might be a stethoscope. In many cases, manufacturers of these devices are not required to notify the FDA before marketing them.

Class II Devices: These are medium-risk devices. Often, they require a premarket notification, known as 510(k), to be submitted to the FDA before they can be sold. During this process, manufacturers need to show that their device is essentially equivalent in terms of safety and effectiveness to a device that is already on the market. An example of a Class II device might be an electric wheelchair.

Class III Devices: This category includes high-risk devices. These are devices that play a crucial role in sustaining or supporting human life, are extremely important in preventing human health impairment, or have a considerable chance of causing an unreasonable risk of illness or injury. Devices in this class have to undergo a more stringent Pre-Market Approval (PMA) process. During this process, manufacturers are required to provide clinical data to prove the safety and effectiveness of their product. A heart pacemaker is an example of a Class III device.

In the United States, the primary pathway for introducing most medical devices to the market is through the 510(k) regulatory process. During the 510(k) procedure, also known as the Premarket Notification, manufacturers are required to demonstrate that their proposed device is "substantially equivalent" to a previously approved device, known as a "predicate" device [11]. Here, "substantially equivalent" implies that the new device is as safe and effective as the predicate one, even if there

[10]How to Determine if Your Product is a Medical Device

[11]The term "predicate device" refers to a device that has already been legally marketed for a similar purpose. Essentially, this earlier device serves as a reference point for the FDA when assessing the safety and effectiveness of the new device.

are differences in design or materials. This process is typically faster than a PMA, often taking just a few months up to a year. From a cost perspective, the 510(k) is generally more affordable. This is largely because it often does not necessitate the extensive clinical trials that the PMA process requires.

For example, if a company wants to introduce a new type of blood pressure monitor, they would look for a similar blood pressure monitor that has already been approved and is on the market. They would then need to provide evidence to the FDA that their new monitor performs at least as well as this existing one, despite any differences that might exist between the two. This method ensures that technological advancements in the medical field continue while preserving the health and safety of the patients using them.

On the other hand, the PMA process is the most rigorous regulatory pathway for medical devices within the FDA's purview. This process is mandatory for Class III devices, which might include those that support or sustain human life, are critical in preventing health impairment, or have the potential to pose an unreasonable risk of illness or injury. Unlike the 510(k), the PMA process requires manufacturers to provide robust clinical trial data, proving both the safety and effectiveness of the device. Given the in-depth nature of this evaluation, the PMA can span several years. Not surprisingly, it is also more expensive, with costs reflecting the in-depth clinical trials, FDA application fees, and other associated expenses.

In summary, the choice between the 510(k) and PMA pathways is not based on manufacturer preference but hinges on the nature, risk profile, and classification of the medical device. The 510(k) is faster and more affordable but only suitable for devices with a similar counterpart already in the market. In contrast, the PMA, while more extensive and costlier, provides a comprehensive evaluation designed for high-risk devices or those introducing novel uses or principles.

4.4 Empirical Analysis

4.4.1 Data

The empirical analysis is based on three main sources of data:

- openFDA: Serving as the foundation of my investigation, openFDA provides pivotal administrative data pertinent to the 510(k) clearance process. Within its vast repository, I extract critical information such as device IDs, associated medical specialties, and the pertinent dates of approval. The breadth and depth of this data allow us to meticulously construct a main panel dataset, capturing the nuanced landscape of the medical device arena from 2002 through 2020.

- FDA 510(k) Database: Complementing the administrative insights from openFDA, the FDA's 510(k) database is a veritable treasure trove of summary PDFs corresponding to each device application. By downloading and analyzing these summaries, we are not only privy to the technical details and claims of each device but also discern the complex predicate relationships that underpin them. This information is crucial, enabling us to map out the generational changes and evolutionary paths that medical devices undergo over time.

- Corporate Acquisition Databases – SDC, Zephyr, and Crunchbase: While the aforementioned databases facilitate an understanding of the medical device landscape from a technological vantage, the exploration of the research at hand would be incomplete without insights into the realm of corporate acquisitions. To this end, I turned to three premier databases – SDC, Zephyr, and Crunchbase. Collectively, these databases provide a comprehensive view of acquisition events, allowing us to ascertain the timing of acquisitions.

510k Summaries and Predicate Devices

The 510(k) database, maintained by the U.S. Food and Drug Administration (FDA), stands as an invaluable repository for scholars and professionals engaged in the medical device domain. Within this resource, researchers can access detailed approval histories for medical devices approved in the United States. Such rich data facilitates in-depth academic exploration into the economics of innovation, offering insights into specific applications by medical/technological areas, current and past firms, and manufacturers' information such as their name and location, in addition to a rich set of administrative data that is provided according to regulatory directives. Additionally, the database serves as an indispensable tool for identifying and studying predicate devices—those pre-existing

and approved entities utilized as benchmarks for evaluating new device submissions.

Users of the FDA's 510(k) database can find details about predicate devices in summary PDFs available in the online portal. The platform's search capability aids in the retrieval process by allowing keyword inputs, specific device names, or even the 510(k) number when known. After executing the search and identifying the desired device, there's usually an option to download its "Summary" in PDF format. Upon accessing this document, it provides comprehensive information about the device and its associated FDA clearance procedures.

However, while the database is a valuable resource, extracting information from the PDF summaries at scale presents its own set of challenges. Firstly, PDFs are not inherently structured for bulk data extraction. Unlike databases with tables or spreadsheets, the information in a PDF is often locked in a static format, making automated extraction difficult. This means that data extraction often requires specialized software or manual effort. Moreover, summaries can vary in layout, design, and the amount of detail they provide, depending on the device and the period of submission. As a result, the extraction process might not be consistent across different summaries, leading to gaps or inconsistencies in the data collected. Furthermore, these PDFs might contain embedded images, graphs, or diagrams, representing data in a visual format that cannot be easily translated into plain text. This necessitates additional processing or manual interpretation, further complicating the extraction process.

To extract detailed information from the summaries in the FDA's 510(k) database, I undertook a systematic approach. My initial step was to download all available summaries, resulting in a collection of approximately 70,000 multi-page PDFs. Given the heterogeneity of these documents, I categorized them into two primary types: “readable PDFs" and scanned PDFs. The term “readable PDFs" refers to those documents that are text-based and typically correspond to more recent submissions. These PDFs inherently contain text that can be programmatically parsed with relative ease. Utilizing specific software tools, I was able to extract the contained information directly from these readable documents. However, a significant portion of the PDFs presented a more intricate challenge: they were scans of original documents. Scanned PDFs, unlike their readable counterparts, are essentially image-based representations of text, rendering them impervious to conventional text

extraction techniques. To transform these scanned images into extractable text, I employed Optical Character Recognition (OCR) [12]. Once the scanned PDFs were processed using Tesseract, I successfully converted the content into textual form. This transformation allowed me to systematically extract the necessary details in a programmatic manner, similar to how I handled the readable PDFs.

Diving into the content of the Summary, the initial section you would encounter is the 510(k) number. This is an exclusive identifier, tailored for each unique submission. Following this, the document introduces the device's commercial name and details about the applicant, which could be the manufacturer or the distributor responsible for the 510(k) notification.

A particularly informative section of the Summary is dedicated to the predicate device or devices. Here, the new device's comparison points are laid bare. It offers insights into the name of the predicate device and its 510(k) number. With this data in hand, users can initiate a fresh search to delve deeper into the predicate's specifics. The Summary does not stop there. It furnishes a meticulous depiction of the device, shedding light on its intended use, the nuances of its design, and the materials from which it is crafted. The indications for use are also elucidated, portraying the specific circumstances or patient demographics that the device is tailored for. As the Summary unfolds, a comprehensive discussion on substantial equivalence emerges. This narrative elaborates on how the new device stands in relation to its predicate, emphasizing aspects of safety and effectiveness. Moreover, it houses a section dedicated to performance data, encapsulating results from various avenues like bench testing, animal studies, and perhaps clinical trials. Drawing toward its conclusion, the summary document rounds off with a section on the device's labeling. This encompasses a spectrum of data including, but not limited to, instructions for use, cautionary notes, contraindications, and other promotional materials.

In essence, the Summaries available in the 510(k) database are a trove of information, streamlining the process of understanding devices, grasping their clearance basis, and pinpointing potential predicate devices for subsequent submissions.

[12]OCR, or Optical Character Recognition, is a sophisticated computational technique that converts images of typed, handwritten, or printed text into machine-encoded text. This process essentially transforms the image-based content of scanned documents into a format that is programmable and thus more easily parsed. For my OCR needs, I turned to https://github.com/tesseract-ocr/tesseract

The use of predicate device linkages offers a nuanced perspective on the evolution and direction of medical device innovations. It also provides keen insights into the strategic choices made by firms.

Each device emerging on the market does not spring forth in isolation; it often draws upon the innovations and designs of its predecessors. By studying predicate device linkages, researchers can track the progression and iteration of devices over time. This not only paints a picture of technological evolution but also provides valuable insights into firm strategies. When a firm chooses to base its new device on a particular predicate device, it reveals a deliberate choice about where to invest and which technological pathways to pursue. Understanding these choices is crucial. It reveals patterns of innovation, highlighting areas of burgeoning technological advancements and areas that may be stagnating or becoming obsolete.

In many settings, patent citations have been the go-to method for tracking technological trajectories. Patent citations, a cornerstone of patent data and innovation analysis, serve multiple pivotal roles. Beyond merely acknowledging prior knowledge, these citations fundamentally demarcate the boundaries of a patent's novelty. They define the scope of the new invention by indicating how it differentiates from prior patented works, thereby setting the grounds for its protection. Over time, researchers have exploited these citations as proxy indicators for knowledge flows and spillovers, capturing how inventions influence subsequent technological developments.

Yet, while patent citations possess significant value, they come with inherent challenges when utilized as trackers of technological trajectories. For one, the sheer volume of citations accompanying a single patent can convolute its lineage. Many of these citations might not directly correlate to the core essence of the new invention but are rather added to bolster its distinctiveness. Furthermore, patent examiners frequently append additional citations during the review process to ensure comprehensive coverage, further complicating the original intent and direction of the citing patent. Consequently, while patent citations might indicate a certain direction of innovation, they can often be noisy indicators, clouded with extraneous references that might not directly bear upon the patented innovation's core essence.

In juxtaposition, predicate device linkages in the medical device sector offer a more streamlined

approach. When a new medical device references a predicate, it is making a direct and tangible connection, signaling that the novel device incorporates, improves upon, or shares functionalities with the cited predecessor. This lineage, while reminiscent of patent citations, is often more explicit in its intention, devoid of the 'noise' often inherent in the multifaceted world of patent references.

In essence, while both patent citations and predicate device linkages aim to trace technological pathways and choices, the latter often presents a clearer, more direct map of practical advancements and strategic decisions within the medical device industry. Integrating this understanding augments our broader exploration of tracking predicates, illuminating its potential as a tool for deciphering the evolutionary arcs of medical technology.

Surprisingly, the vast potential of analyzing predicate device linkages has remained largely underutilized. Much of the research focus has been public health-centric, with studies geared towards understanding the health implications and safety of medical devices. While these are undeniably vital areas of study, the realm of business strategy and competitive dynamics has been somewhat overlooked in this context.

This paper takes a pioneering step in this direction. By leveraging the predicate graph [13], I dive into the nuances of post-acquisition dynamics and I offer a fresh lens through which to view the innovation landscape of the medical device industry. In essence, this paper goes beyond the conventional use of the predicate device data, unveiling its potential to inform strategic management and competition in this vital sector.

---

[13]The predicate graph is a tree composed of devices and links between a device and its predicates.

[14]Detailed description of 4.1: [LEFT] The figure on the left reports the number of 510k applications approved between 2002 and 2020: the years that are relevant in the analysis. The overall height of the bars indicates the total number of devices approved each year. The shaded part of the bar reports the number of devices approved in the focal year for which I was able to retrieve predicate information. Predicate information may be missing for three main reasons, (1) either the application is a de-novo application, thus it is not required to provide substantial equivalence to a previously approved device; (2) the summary documents associated with the device have not been distributed by the firm, these summaries can be requested with a FOIA request but that would need to be done case by case. These include summary documents for which information was intentionally redacted. (3) My algorithm did not identify the predicate information. This can

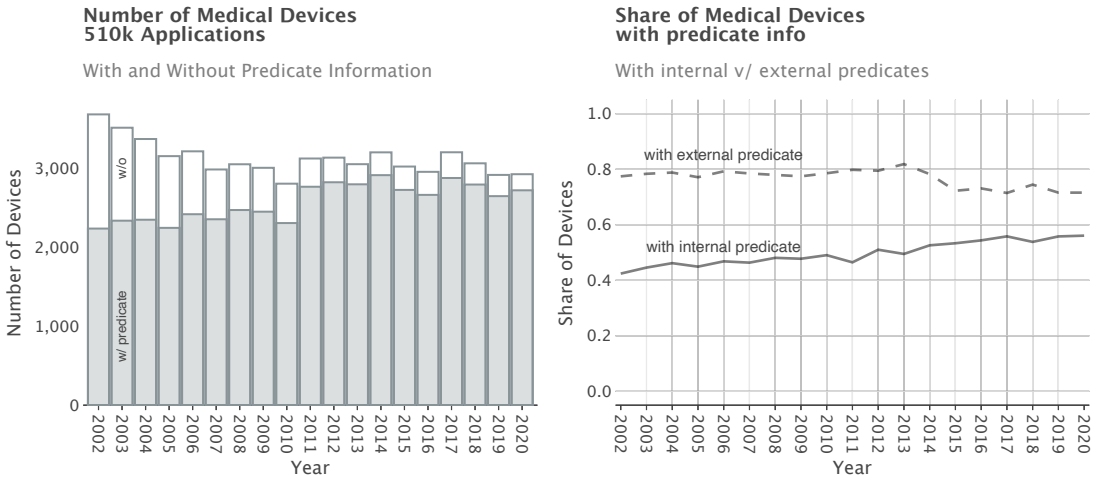

Figure 4.1: Number and share of devices over time. [14]

---

be the case for some scanned documents with peculiar formatting, or most often when the predicate information does not include an identifier for the 510k application, in these cases the predicate could be identified by using the name of the device, but often the name reported does not easily match the actual name in a standardized way. [RIGHT] The figure on the right reports, conditionally on having identified a predicate for the device, what is the share of devices in each year that claim substantial equivalence with an internal device (solid line) and what is the share of devices in each year that claim substantial equivalence with an external device (dashed line). [Internal] A device claims substantial equivalence to an "internal predicate" if the predicate device was in turn assigned to the applicant of the focal (claiming) device. Internal use is defined as used by the original applicant of the focal device or by the acquirer of the applicant of the focal device if the claim happens post-acquisition. [External] A device claims substantial equivalence to an "external predicate" if the predicate device was in turn assigned to an applicant that was not the applicant of the focal (claiming) device. External use may include the acquirer of the original applicant of the claims if the predicate relationship falls in a pre-acquisition year. When a 510(k) submission must show that the new device is substantially equivalent in intended use, technological characteristics, and performance to a legally marketed predicate device, it can demonstrate it by using one or more predicate devices. Using multiple predicates may help establish equivalence in different aspects of the new device. Thus a device using more than one predicate, device may be counted both in the share of devices using an internal predicate and the share using an external predicate, if these predicates have been assigned to multiple different applicants.

Mergers and Acquisitions

I then complement the medical device information, the predicate graph and the FDA applicants with a comprehensive list of acquisitions [15] between 1990 and 2021. To do so I retrieve data from three different sources: SDC, Zephyr, and Crunchbase. A share of the acquisitions is covered by all three sources but they complement providing a larger set of deals [16]. In the analysis, I will have to restrict to medical devices approved after 2002, given that the available predicate information "starts" in 2002, but covering acquisitions prior to 2002 allows me to drop devices that have already experienced acquisitions. Once the acquisition deals data is retrieved I use the main (and common to the three databases) variables: target name, acquired name, and acquisition date [17]. I then fuzzy match the names [18] of the target to the FDA applicants to identify which 510k applications were part of at least one M&A transaction [19]. Noticing that the fuzzy match led to many false positives, I geocode [20] the companies in both sets of data. On the FDA data side, I extract the names of the

[15] I employ the term "acquisition" to denote one corporation's takeover of another complete corporation or a segment of an active corporation. Most of the acquisitions taken into consideration could be characterized as "horizontal acquisitions", i.e. acquisitions within the same industry/sector.

[16] If we look at earlier years, prior to 2000, SDC has a better coverage.

[17] Filtering out the rumored acquisition, and acquisitions that do not take full control of the company: minority share acquisitions or share repurchases.

[18] The integration of M&A datasets and FDA applicant data poses multiple challenges. Besides misspellings, abbreviations, and changes in company names there are frequent acquisitions of business lines which are often difficult to distinguish from acquisitions of corporate acquisitions. For example in 2001, Gyrus Group PLC acquired the ENT (ear, nose and throat) business from Smith & Nephew, the target of this acquisition is recorded in the SDC dataset as "Smith & Nephew PLC-ENT" which is matched to the corporate entity "Smith & Nephew" when fuzzy matching with the FDA applicant dataset. I deal with these cases by geocoding each applicant as described in the text. So if a division has a separate address than the HQ then I am able to discern it from full corporate acquisitions.

[19] I am unable to identify non-corporate acquisitions: transactions of assets that do not involve shares acquisitions, that is I do not know if a company is selling a bundle of devices to another firm without selling shares of the company.

[20] I identify the latitude and longitude associated with the company names

applicants [21] and their address [22]from the FDA portal as well as the PDFs summaries [23]. I then standardize the addresses and use the Google Map API to locate the applicants' (and correspondents') geographical positions. Then I repeat this process from the targets in the acquisition databases, using the address provided by SDC, Zephyr, and Crunchbase. With the latitude and longitude of both sets of data, I calculate the geographical distance between the applicant and the string-matched target name, to then filter the names using a combination of string similarity and geographical distance [24].

4.4.2 Analysis

First I construct a medical-device, year panel of class 2 medical devices approved between 2002 and 2020 [25]. The panel is pre-processed according to best practices that have emerged in recent econometric literature (Callaway & Sant'Anna, 2021).

I drop all of the devices that have been acquired before the start of their time-series. This approach is preferred since, under the assumption of parallel trends, these already-treated units do not provide meaningful insight into estimating the impact of treatment. Even though TWFE regressions and in general event study regressions can mechanically operate with these observations, including them might introduce unfavorable comparisons (Goodman-Bacon, 2021). This step leads

[21]There are over 40 thousand unique string names, that when cleaned and standardized using regular expressions tailored to corporate identifiers reduce to about 25 thousand names. The addresses are harder to standardize in a way to reduce their number substantially.

[22]The string that identifies the location of the applicant or correspondent.

[23]The FDA portal reports the address of the correspondent for the focal application. The correspondent may be internal to the firm or it often is an external consultant hired by the firm to interact with the FDA, this is often the case for smaller companies that rely on regulatory consultants to navigate the 510k submission process. The summary PDFs on the other hand report both the address of the correspondents as well as the address of the applicants.

[24]I Identify good combinations of textual similarity scores and graphical distances heuristically.

[25]I restrict to medical devices approved after 2002 because predicate data is not as available in years prior to 2002.

to a panel of 510k applications approved between 2002 and 2020 with a set of devices that are acquired between 2002 and 2022 (when the acquisition data ends) and a set of devices that have never been acquired [26]. Acquired and never-acquired constitute the full sample in Table 4.1. Additionally, I will conduct the analysis by dropping the never-acquired set of devices. In a classic DID setting in which each entity is treated at the same time-period, one needs a never-treated group to be able to take 3 differences, however, when studying corporate acquisitions, the treatment, the M&A event, is staggered over time: different devices may be treated at different times, which allows us to use not-yet-treated units as the sole control group [27]. By conducting the difference-in-differences analysis in the sample with no never-acquired devices I will rely more heavily on the estimation procedure suggested by (Callaway & Sant'Anna, 2021). With treatment effect heterogeneity [28], in a staggered treatment setting TWFE regression can lead to biased estimates, especially in the absence of a never-treated group (Goodman-Bacon, 2021; Sun & Abraham, 2021). The estimation procedure suggested by (Callaway & Sant'Anna, 2021) simply splits the data sample into all possible group-times, where the group is the cohort of treatment (year of acquisition) and the time is the panel-year, and estimates a simple 2x2 difference in difference. The result is a group-time average treatment effect which can be then aggregated to conduct an event study or to obtain an overall treatment effect.

By conducting the analysis without a never-treated group I am able to add an additional dependent variable which is not available for never-treated entities. I can measure whether a device is used as a predicate by the acquirer of that same device. I thus measure use before and after the

[26] To the best of my knowledge, excluding possible measurement error in the matching process since for example a firm may change the name over time or the database I have used Zephyr, SDC, and Cruchbase do not include the acquisition.

[27] Not-yet-treated units can and are often used as part of the control group even in the presence of a never-treated group, but when there is no never-treated group then they constitute the entire control group.

[28] In this setting it is reasonable to expect heterogeneity across cohorts. As technology accumulates, companies have an ever greater set of devices to choose from, the likelihood a device is used as a predicate may change over time given the increasing number of devices available, which would lead to acquirers having more options to choose from the later for acquisitions that happened at different periods of time.

acquisition took place, this exercise can identify on average for acquired devices the timing of use by the acquiring company.

Table 4.1: Predicate Usage

	Full Sample				Acquired		Never Acquired	
	Mean	Std	Min	Max	Mean	Std	Mean	Std
Used as Predicate	0.138	0.345	0	1	0.153	0.36	0.136	0.342
Used as Predicate [external]	0.087	0.282	0	1	0.092	0.289	0.086	0.281
Used as Predicate [internal]	0.06	0.238	0	1	0.073	0.26	0.058	0.234

Used as Predicate is a binary variable indicating whether the focal device (the 510k application included in the panel) was used as a predicate device by any other device approved in the panel-year. *Used as Predicate [internal]* is a binary variable indicating whether the focal device was used as a predicate device by a device belonging to the original applicant of the focal device or to the acquirer of the original applicant, approved in the panel-year (if it's used by the acquirer, only post-acquisition use is considered internal). *Used as Predicate [external]* is the opposite of *Used as Predicate [internal]*. Note that a device may have internal and external predicate use since during the years considered in the study a company may choose to use multiple predicates.

I then employ two main approaches to studying the effect of acquisition on subsequent technology development, which I measure as a binary variable indicating whether a device is used as a predicate by the panel-level device.

First I use a liner probability model with a two-way fixed effects specification, as follows:

$$\underbrace{y_{it}}_{\substack{\text{Device used as}\\ \text{a predicate}\\ \text{in a year}}} = \underbrace{\sum_{d=-D, d\neq -1}^{D} 1(t-e_i=d)\beta_d}_{\substack{\text{Dynamic effects}\\ \text{of firm acquisition}}} + \underbrace{\alpha_t}_{\substack{\text{Year}\\ \text{effects}}} + \underbrace{\gamma_i}_{\substack{\text{Device}\\ \text{effects}}} + \varepsilon_{it} \tag{4.1}$$

however under staggered treatment, and treatment effects heterogeneity, recent literature has shown that the two-way fixed effect regression may lead to biased estimates, the bias is further

exacerbated with the absence of a never-treated group [29].

ATTgt

I complement the two-way fixed effects results with computing group-time average treatment effects on the treated (ATTgt) following the procedure and framework suggested by Callaway and Sant'Anna, 2021 [30]. This estimation procedure consists in splitting the multi-period [31] panel structure into 2x2 panels: panels composed of 2 time periods and 2 groups a treated and an untreated group. This is done for each cohort (time of treatment), in my case there are as many cohorts as calendar times: there is at least one acquisition every year under analysis. As mentioned, to mitigate concerns of intrinsic differences between acquired v/ non acquired devices, I will also take into consideration the subset of devices that are eventually acquired, excluding never-acquired devices. In this case, when constructing the control group for each group time, I keep only the devices that have not yet been acquired, but that will be acquired in subsequent years; for example if we take the panel subset consisting of the years 2003 and 2004, for the cohort 2002, the estimate will tell us the difference in outcome between 2004 and 2003 after 1 year entities have been treated, minus the difference in outcome between 2004 and 2003 for medical devices that have not yet been acquired, for example, a subset of those would be devices that are acquired in 2005. After all ATT estimates have

[29] While I do observe devices that are never acquired, I will conduct this exercise both by including and excluding the never-acquired group. When dropping the never-treated group I will take advantage of the staggered nature of the treatment: devices are acquired at different times, to compare acquired devices to not-yet-acquired devices (but which will eventually get acquired). This exercise should mitigate the concern that acquired devices are intrinsically different than non-acquired devices.

[30] In order to estimate the group-time ATT and its aggregations I rely on a software that I wrote (https://github.com/bernardodionisi/differences). This software (as of October 2023) covers all the functionality (and more) as the R library distributed by Callaway and Sant'Anna, 2021 with the advantage of working more efficiently with large datasets thanks to parallelization of the estimation. differences is tested against R 'did' (Callaway & Sant'Anna, 2021) to ensure the numerical equivalence of the results.

[31] In my case the time periods go from 2002 to 2020. The time-series for each device starts at the time of approval, thus, an approved device will enter the panel at the time of its approval. I then truncate the time-series for the device 10 years after its approval. Results are similar for any different choice of time-series length, up to keeping all the years since approval to 2020.

been computed for each group-time, Callaway and Sant'Anna, 2021 suggest weighted aggregation schemes that will take into consideration the share of devices in each group, these aggregations are designed to remove the bias that the canonical two-way fixed-effects regression introduces under staggered treatment timing and cohort heterogeneity. In this study I will consider event time aggregations, what would typically be considered an event study as well as overall aggregations, that will return an aggregated estimate over the sample period. For completeness and robustness the overall aggregations that I conduct and report include overall aggregation of group-time treatment effects as well as overall aggregations of other aggregation schemes.

4.4.3 Results

The estimation results are reported in the figures from Figure 4.2 to Figure 4.6. In order, initially, I provide estimates for the two-way fixed effects specification as defined in Equation 4.1 where the estimates are relative to one event period prior to the acquisition; this can be called 'universal' base period. Estimates are computed with the -1 reference, this is in contrast to a time 'varying' base period, which I use in the group time estimation of the average treatment effects in Figure 4.4 to 4.6. When using the time-varying base period I will shift the reference period forward for each group time, the reference period will vary for each cohort and calendar time. Doing so I aim to provide evidence that the different procedures lead to similar estimates.

In Figure 4.2 the estimation is conducted on the full sample of an unbalanced panel dataset of medical devices, the sample includes both devices that underwent an acquisition as well as never acquired devices; [32] the dependent variable is a binary variable indicating whether the focal device (the 510k application included in the panel) was used as a predicate device by any other device

[32] I will drop the never acquired devices when applying the framework brought forward by Callaway and Sant'Anna, 2021. The two-way fixed effects estimates are reported for completeness, however recent literature Goodman-Bacon, 2021 has shown bias in these types of regressions in the presence of multiple treatment cohort and treatment effects heterogeneity, this bias is especially evident when the number of never-treated entities is low compared to treated ones; thus I avoid reporting TWFE estimates with no never-treated entities.

approved in the panel-year. Thus we can interpret the estimates as the percentage point change in the likelihood a device is used as a predicate compared to the year prior to the acquisition: one year after the acquisition, that likelihood decreases by 2 percentage points, which is a 15% decrease if compared to the overall likelihood a device is used as a predicate in a given year. The dynamics are consistent with the acquired devices being on average less developed since we do not see use in the 10 years after the acquisition. Empirically the years prior to the acquisition seem to follow a parallel trend where the acquired devices follow a similar trend as the control group, still relative to -1.

In Figure 4.3 I estimate the same model but now I construct variables that indicate internal use or external use, in the left panel of Figure 4.3 I use a binary variable indicating whether the focal device was used as a predicate device by another device approved in the panel-year and assigned to the applicant of the focal device (internal use). Internal use is defined as used by the original applicant of the focal device or by the acquirer of the applicant of the focal device if the panel-year is post-acquisition. The dependent variable in the right panel is a binary variable indicating whether the focal device was used as a predicate device by another device approved in the panel-year and assigned to an applicant other than the applicant of the focal device (external use). External use includes use by the acquirer of the applicant of the focal device use if the panel-year is pre-acquisition. Figure 4.3 indicates that the overall drop in the continuation of the trajectory, into new product generation, as measured by the predicate relationship is driven by a drop in internal use: acquired devices experience less internal use than the control group. One year after the acquisition the likelihood a device is used as a predicate drops 5 percentage points relative to the year prior to the acquisition. In the right panel, conversely, we observe that there is an uptake in external use after the acquisition, I attribute this to a signaling effect on the quality of the devices: prior to the acquisition, there is uncertainty on the viability that a device can have if used as a predicate [33]. Even a 1.5 percentage point increase relative to the -1 baseline translates in a large change given the average likelihood of external use is 0.08%. Following the TWFE estimates I then turn to the

[33]Choosing the wrong predicate may be costly for the firm, given it may cause delays in the approval process and further investigations on the safety and effectiveness of the product.

estimating group-time ATTs using the framework brought forward by Callaway and Sant'Anna, 2021 that I described above. Given the overall patterns are consistent with the main findings, even when excluding never-acquired devices from the analysis I would turn the attention to the bottom-left panel of Figure 4.6, I now provide estimates using as dependent variable a binary indicator for whether the devices is used as a predicate by the acquirer of the firm that was the original applicant, this variable is defined only for acquired devices, thus the sample excludes never-acquired devices. The estimates seem to suggest that some devices are being developed by the acquired and the development happens post-acquisition.

4.5 Conclusion

Understanding post-acquisition product trajectories is of central interest when studying the direction of innovation. Uncovering these patterns is of increasing interest over the past few decades, given that the division of innovative labor between large corporations and startups has expanded, with a clear shift toward external technology sourcing. By leveraging the unique institutional setting that is the medical device industry, this study offers product-level evidence on post-acquisition continuation decisions, closely examining the evolution of technology trajectories.

Using difference-in-differences I provide product level evidence of the development of medical devices post-acquisition. Devices that have been acquired tend to progress more sluggishly into subsequent generations. However, this finding does not imply a stalling of developmental activities but rather suggests a re-calibration in the pace of internal progression in contrast to those not acquired. An essential takeaway, and likely focus of future research on this domain, is the potential re-calibration of focus by acquirers onto specific subsets of acquired devices—a reflection of the classical challenges seen in product portfolio management. Furthermore, there is an observable inclination for heightened external development of these acquired devices, insinuating their latent potential for refinement and progression.

The case of Newport Medical Instruments and Covidien, among others, accentuates the urgency and significance of comprehensively understanding this innovation landscape. Not only do these

dynamics possess ramifications for industry operations, but they have far-reaching consequences for public health, especially in exigencies like the COVID-19 pandemic. The multifaceted challenges tied to post-acquisition decision-making—ranging from information disparities to potential cultural misalignments—underscore the intricacies involved in managing innovation.

Contrary to the “killer acquisition" narrative, this study sheds light on the nuances specific to the medical device sector, emphasizing the divergence between early-stage innovation and development post-approval by the FDA. While the study doesn’t set out to directly affirm or contest the "killer acquisition" theory, it aims to contribute valuable perspectives on the patterns of product continuations with some potential discussion on the underlying determinants.

In summary, the medical device industry stands out as a critical focal point in the broader innovation landscape. As antitrust authorities and policymakers cast an increasingly vigilant eye over M&A activities, understanding these dynamics becomes paramount. With the challenges and opportunities the sector presents, future research endeavors are needed to discern the mechanism underlying the findings of this study.

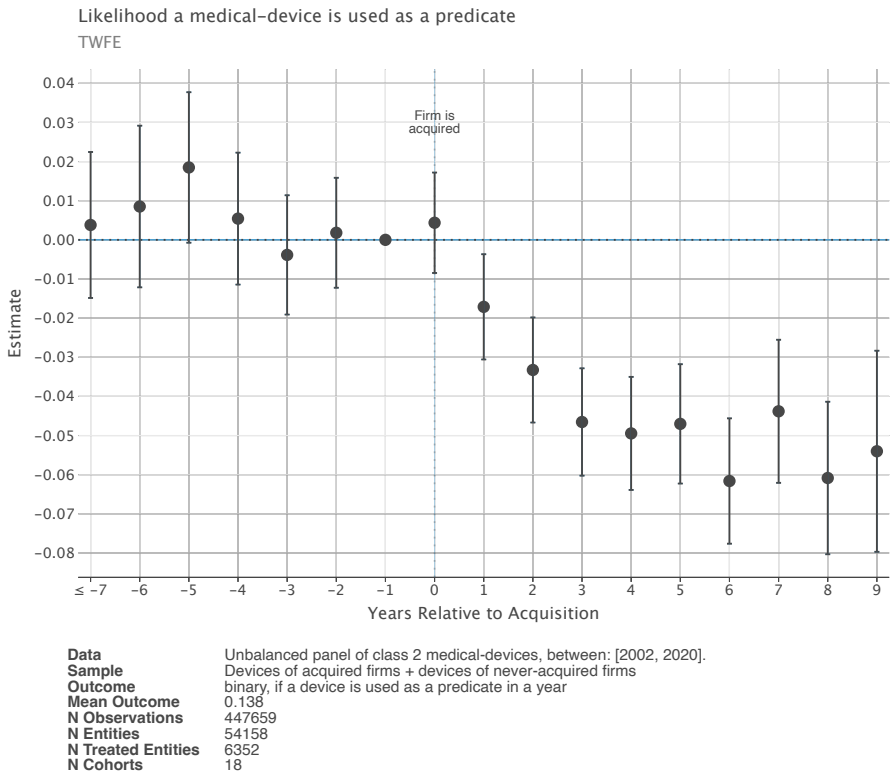

Figure 4.2: This figure reports the estimates of the two-way fixed-effects regression as specified in equation 4.1. The dataset used is an unbalanced panel of class 2 medical devices (510k approvals) where each device enters the panel the year of approval and exits the panel 10 years after the approval year. The sample includes acquired and never-acquired patents. The dependent variable is a binary indicator for whether the focal device was used as a predicate device by any other device approved in the panel-year.

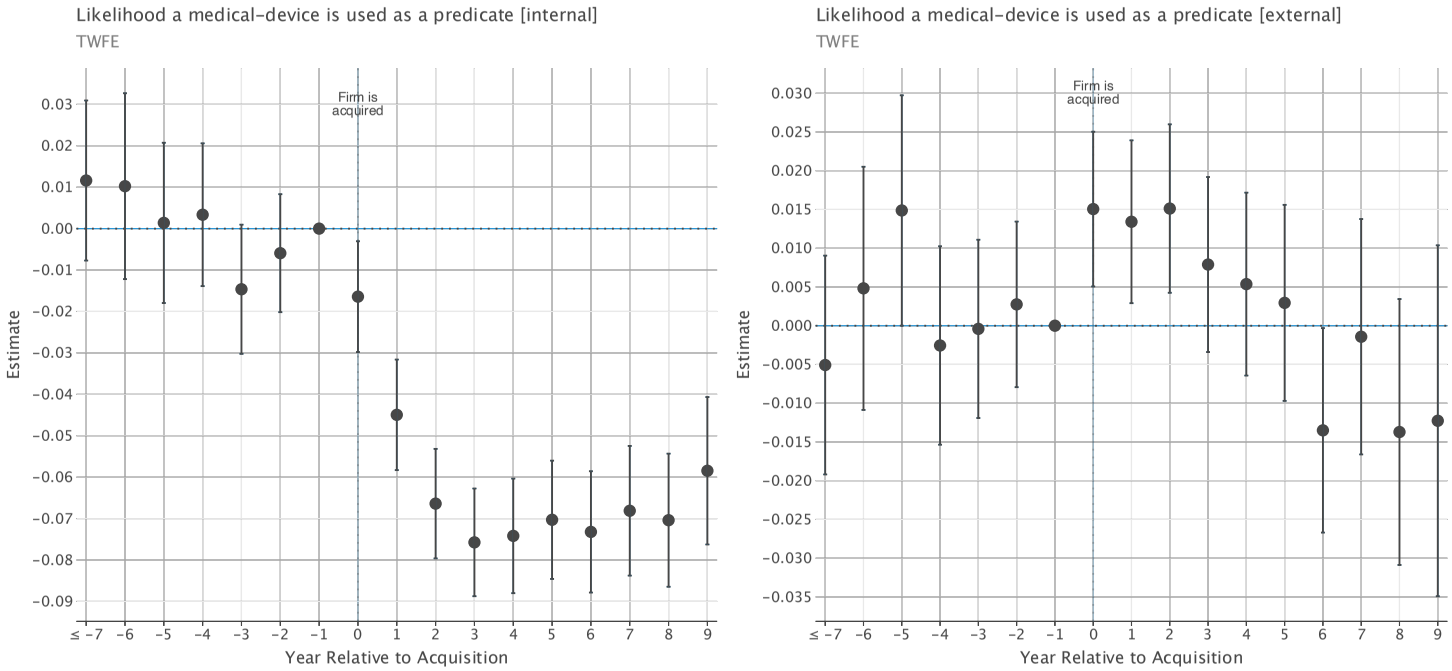

Data	Unbalanced panel of class 2 medical-devices, between: [2002, 2020].
Sample	Devices of acquired firms + devices of never-acquired firms.
Outcome [internal]	binary, if a device is used as a predicate [internal] in a year; internal use is by the original applicant or by the acquirer (after the acquisition).
Mean Outcome [internal]	0.0727
Outcome [external]	binary, if a device is used as a predicate [external] in a year; external use is by firms other then the applicant or the acquirer (after the acquisition).
Mean Outcome [external]	0.0871
N Observations	447659
N Entities	54158
N Treated Entities	6352
N Cohorts	18
N Times	19

Figure 4.3: Estimates of the two-way fixed-effects regression] [34]

[34]Detailed description of Figure 4.3. The figure reports the estimates of the two-way fixed-effects regression as specified in equation 4.1 for two different dependent variables: the dependent variable in the chart on the LEFT is a binary variable indicating whether the focal device was used as a predicate device by another device approved in the panel-year and assigned to the applicant of the focal device (internal use). Internal use is defined as used by the original applicant of the focal device or by the acquirer of the applicant of the focal device if the panel-year is post-acquisition. The dependent variable in the chart on the RIGHT is a binary variable indicating whether the focal device was used as a predicate device by another device approved in the panel-year and assigned to an applicant other than the applicant of of the focal device (external use). External use includes use by the acquirer of the applicant of the focal device use if the panel-year is pre-acquisition. The dataset used is an unbalanced panel of class 2 medical devices (510k approvals) where each device enters the panel the year of approval and exits the panel 10 years after the approval year. The sample includes acquired and never-acquired patents.

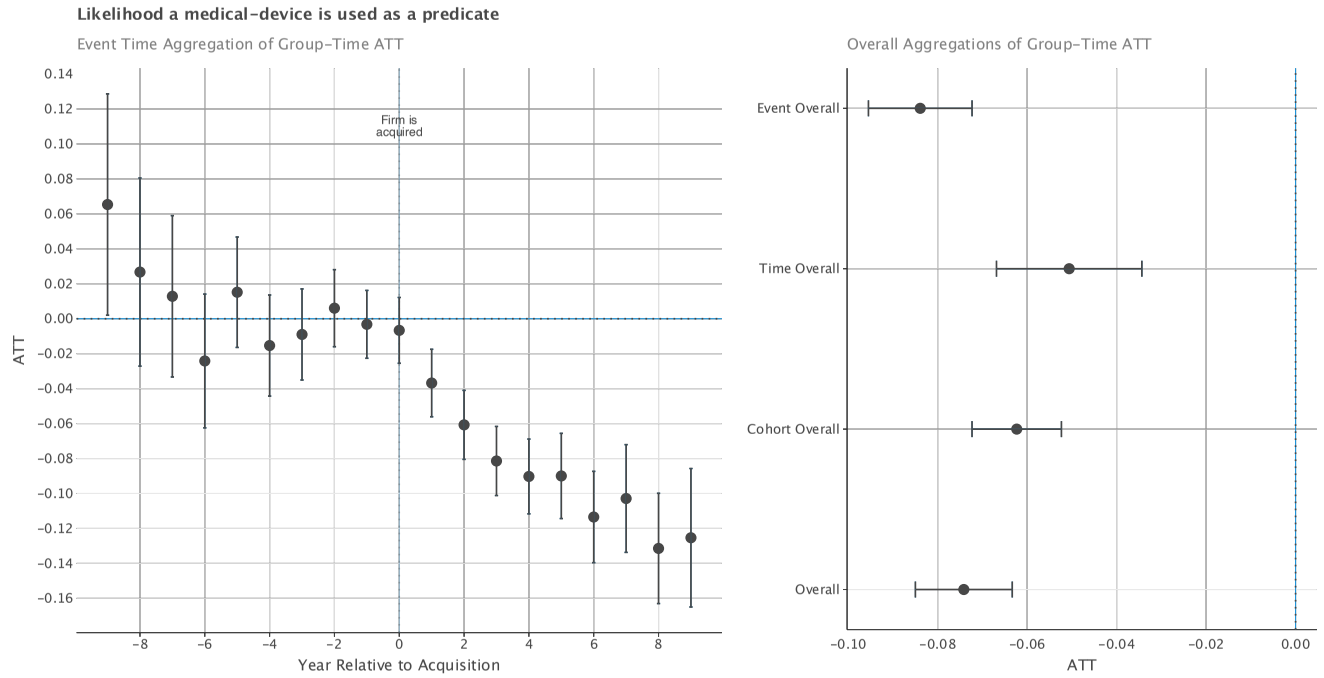

Data	Unbalanced panel of class 2 medical-devices approved between 2002 and 2020
Sample	Devices of acquired firms + devices of never-acquired firms
Dependent Variable	binary, if a device is used as a predicate in a year
Base Period	Varying
Control Group	Not Yet Treated
SE	Bootstrapped. Clustered at the device level
N Observations	447659
N Entities	54158
N Treated Entities	6352
N Cohorts	18
N Times	19

Figure 4.4: Estimates of the event-study aggregation of group-time ATT [35]

[35]Detailed description of Figure 4.4: The figure reports the estimates of the event-study aggregation of group-time ATT as described in the ATTgt description (4.4.2). The sample used includes both acquired as well as never-acquired devices. The control group is set to include not yet acquired devices (alongside never-acquired devices), that is, for the calculation of the group-time ATT, when constructing each 2x2, the control devices include both devices that will never be acquired (up to 2022), and devices that have not yet been acquired at the focal time considered in the group-time split.[LEFT] The left chart reports event time aggregation of the group-time ATTs. [RIGHT] The chart on the right reports overall average treatment effects (on the acquired) calculated at different levels of aggregation. The "Overall" estimates calculate a collective ATT when aggregating all the ATTgt together. The "Cohort Overall" would first aggregate by cohort and then calculate an overall effect for each cohort. The "Time Overall" would first aggregate by calendar time and then calculate an overall effect for each time. The "Event Overall" would first aggregate by event-time and then calculate an overall effect for each event-time.

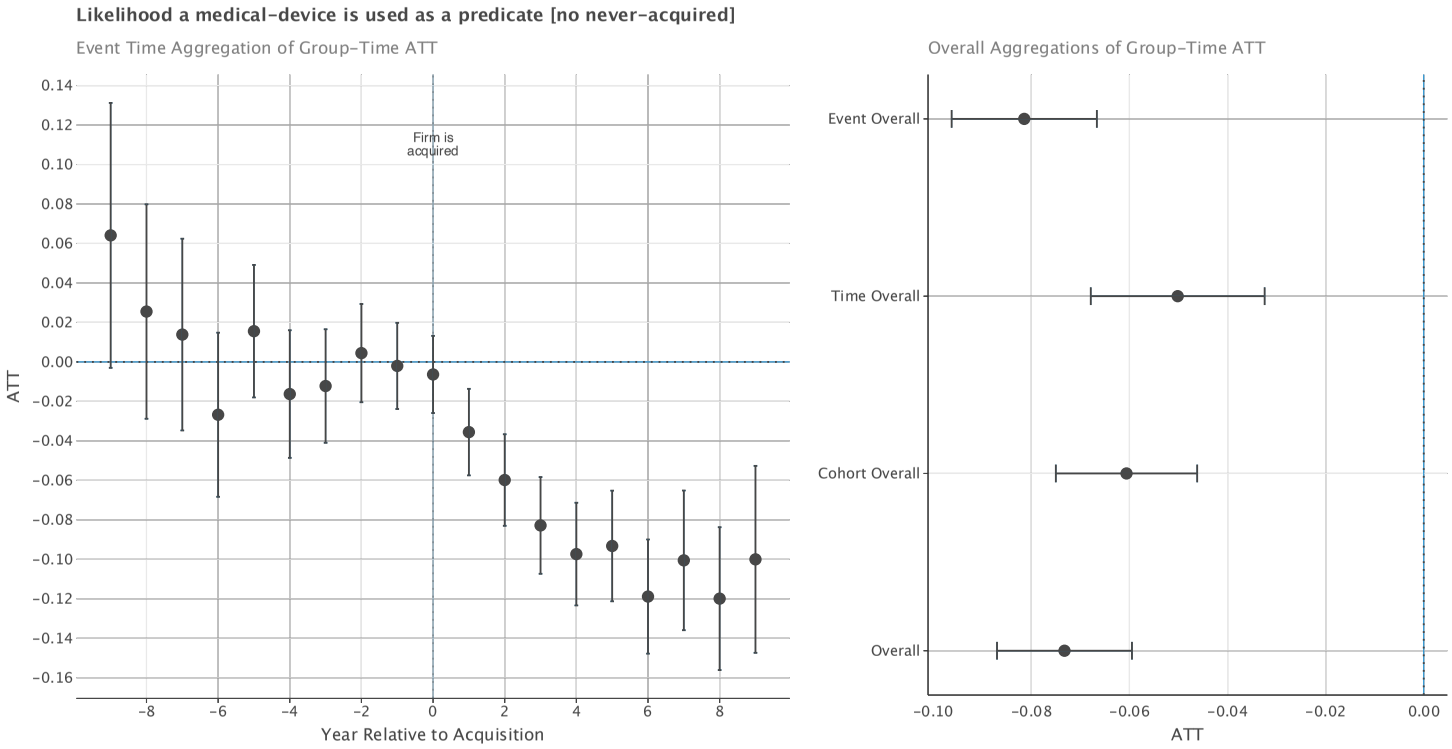

Data	Unbalanced panel of class 2 medical-devices approved between 2002 and 2020
Sample	Devices of acquired firms + devices of eventually-acquired firms [no never-treated]
Dependent Variable	binary, if a device is used as a predicate in a year
Base Period	Varying
Control Group	Not Yet Treated
SE	Bootstrapped. Clustered at the device level
N Observations	75500
N Entities	7680
N Treated Entities	6352
N Cohorts	18
N Times	19

Figure 4.5: Estimates of the event-study aggregation of group-time ATT (excluding never-acquired devices) [36]

[36]Detailed description of Figure 4.5: This figure reports the estimates of the event-study aggregation of group-time ATT as described in the ATTgt description (4.4.2). The sample used includes only medical devices that will eventually be acquired [excluding never-acquired devices]. The control group is thus composed of not yet acquired devices, that is, for the calculation of the group-time ATT, when constructing each 2x2, the control devices include devices that have not yet been acquired at the focal time considered in the group-time split.[LEFT] The left chart reports event time aggregation of the group-time ATTs. [RIGHT] The chart on the right reports overall average treatment effects (on the acquired) calculated at different levels of aggregation. The "Overall" estimates calculate a collective ATT when aggregating all the ATTgt together. The "Cohort Overall" would first aggregate by cohort and then calculate an overall effect for each cohort. The "Time Overall" would first aggregate by calendar time and then calculate an overall effect for each time. The "Event Overall" would first aggregate by event-time and then calculate an overall effect for each event-time.

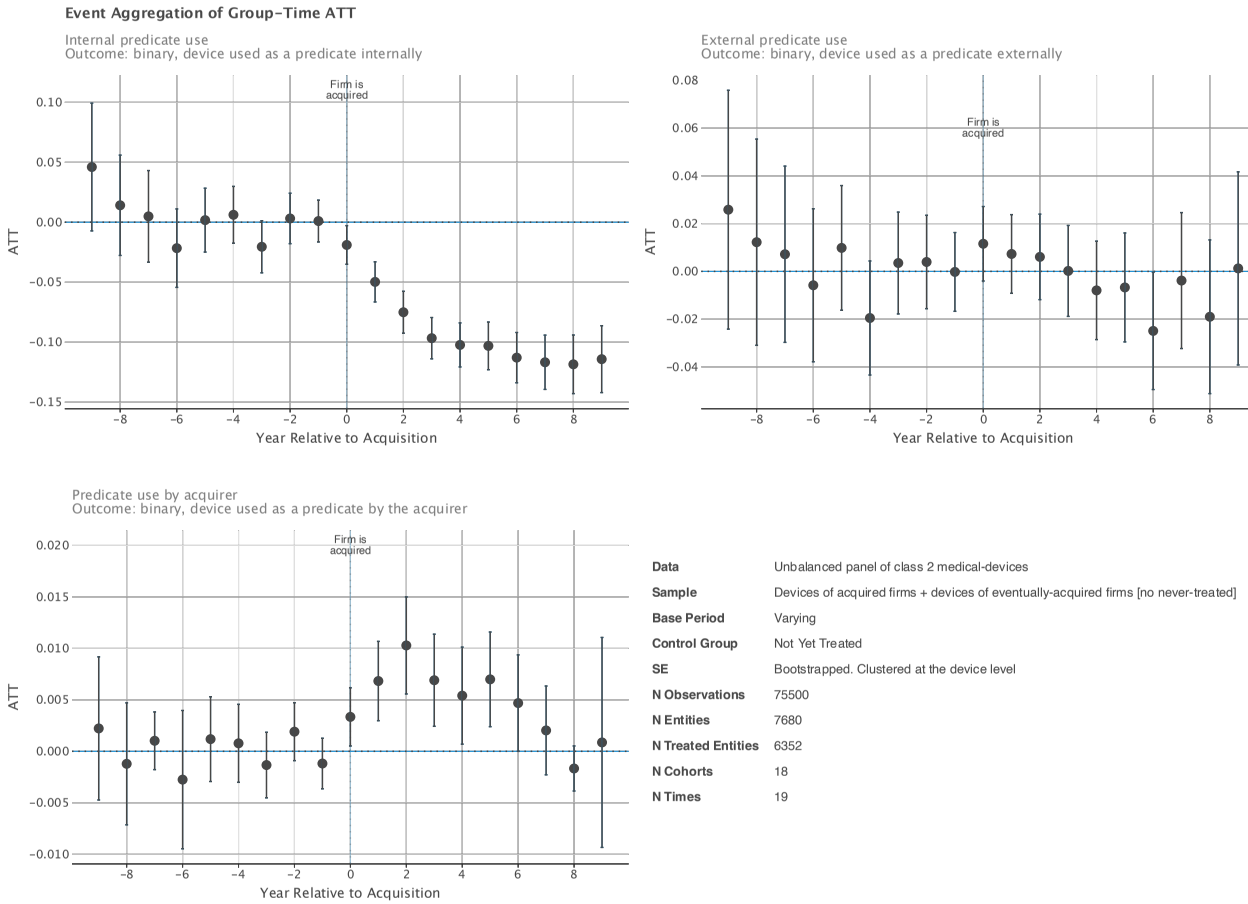

Figure 4.6: Event-study aggregation of group-time ATT: internal, external and by acquirer (excluding never-acquired devices) [37]

[37]Detailed description of Figure 4.6. The figure reports the estimates of the event-study aggregation of group-time ATT as described in the ATTgt description (4.4.2). The sample used includes only medical devices that will eventually be acquired [excluding never-acquired devices]. The control group is thus composed of not yet acquired devices, that is, for the calculation of the group-time ATT, when constructing each 2x2, the control devices include devices that have not yet been acquired at the focal time considered in the group-time split. [TOP-LEFT] The top-left chart reports event time aggregation of the group-time ATTs when using as the dependent variable an indicator for use of predicates that are internal to the applicant: internal use is defined as used by the original applicant of the focal device or by the acquirer of the applicant of the focal device if the panel-year is post-acquisition. [TOP-RIGHT] The top-right chart reports event time aggregation of the group-time ATTs when using as the dependent variable an indicator for the use of predicates that are external to the applicant: external use is a binary variable indicating whether the focal device was used as a predicate device by another device approved in the panel-year and assigned to an applicant other than the applicant of the focal device (external use). External use includes use by the acquirer of the applicant of the focal device use if the panel-year is pre-acquisition. [BOTTOM-LEFT] The bottom-left chart reports event time aggregation of the group-time ATTs when using as the dependent variable an indicator of whether the device is used as a predicate by the acquiring company: the company that acquired the applicant of the focal device. (Note that this dependent variable can be identified only by restricting the sample to devices that are going to be eventually acquired [excluding never-acquired devices.]

Chapter 5

Conclusions

The study of innovation and technical change has become a vital area of economics research, providing insight on the dynamics driving inventive activity and technological progress. As discussed in this dissertation, difference-in-differences and event study analysis have emerged as prominent empirical approaches for investigating these relationships. However, many open questions remain regarding the factors that facilitate or hinder advancements. Ongoing scholarly inquiry using rigorous methodologies is critical for deepening our understanding and informing effective policies that promote innovation and growth.

Recent research highlights the increasing role of mergers and acquisitions in the innovation ecosystem. An examination of the medical device sector reveals a deceleration in the post-acquisition evolution of acquired technologies, though acquisitions may serve as quality signals. This sheds new light on the complex interplay between acquisitions and technological progression.

Additionally, an analysis of firms' use of public scientific knowledge finds limited private value in fields with extensive competition. However, returns are higher for early adopters who can secure broader patents. Corporate participation in public science strongly predicts priority access, underscoring the value of scientific engagement.

The overall offers a framework for assessing existing literature at the nexus of economics and technical change. As the innovation ecosystem evolves, persistent empirical work leveraging robust methodological tools will be critical for comprehensively mapping the drivers of inventive activity. Considerable efforts remain to translate these insights into policies that effectively foster economic growth.

Bibliography

Aghion, P., Bloom, N., Blundell, R., Griffith, R., & Howitt, P. (2005). Competition and innovation: An inverted-U relationship. *The Quarterly Journal of Economics*, *120*(2), 701–728.

Aguilera, R. V., & Dencker, J. C. (2004). The role of human resource management in cross-border mergers and acquisitions. *The International Journal of Human Resource Management*, *15*(8), 1355–1370.

Ahmadpoor, M., & Jones, B. F. (2017). The dual frontier: Patented inventions and prior scientific advance. *Science*, *357*(6351), 583–587.

Ahuja, G., & Katila, R. (2001). Technological acquisitions and the innovation performance of acquiring firms: A longitudinal study. *Strategic management journal*, *22*(3), 197–220.

Akcigit, U., & Goldschlag, N. (2023a). *Measuring the characteristics and employment dynamics of us inventors* (tech. rep.). National Bureau of Economic Research.

Akcigit, U., & Goldschlag, N. (2023b). *Where have all the" creative talents" gone? employment dynamics of us inventors* (tech. rep.). National Bureau of Economic Research.

Akcigit, U., Pearce, J. G., & Prato, M. (2020). *Tapping into talent: Coupling education and innovation policies for economic growth* (tech. rep.). National Bureau of Economic Research.

Andrade, G., Mitchell, M., & Stafford, E. (2001). New evidence and perspectives on mergers. *Journal of economic perspectives*, *15*(2), 103–120.

Arora, A., Belenzon, S., & Patacconi, A. (2018). The decline of science in corporate R&D. *Strategic Management Journal*, *39*(1), 3–32
doi: 10.1002/smj.2693.

Arora, A., Belenzon, S., Patacconi, A., & Suh, J. (2019). *The Changing Structure of American Innovation: Some Cautionary Remarks for Economic Growth*.

Arora, A., Belenzon, S., Patacconi, A., & Suh, J. (2020). The changing structure of american innovation: Some cautionary remarks for economic growth. *Innovation Policy and the Economy*, *20*(1), 39–93.

Arora, A., Belenzon, S., & Sheer, L. (2017). *Back to basics: Why do firms invest in research?* (Tech. rep.). National Bureau of Economic Research.

Arora, A., Belenzon, S., & Sheer, L. (2021). Knowledge spillovers and corporate investment in scientific research. *American Economic Review*, *111*(3), 871–98.

Arora, A., Belenzon, S., & Suh, J. (2020). Science and the market for technology.

Arora, A., & Ceccagnoli, M. (2006). Patent protection, complementary assets, and firms' incentives for technology licensing. *Management Science*, *52*(2), 293–308.

Arora, A., & Gambardella, A. (2010). The market for technology. *Handbook of the Economics of Innovation*, *1*, 641–678.

Arrow, K. J. (1972). *Economic welfare and the allocation of resources for invention*. Springer.

Arts, S., Hou, J., & Gomez, J. C. (2021). Natural language processing to identify the creation and impact of new technologies in patent text: Code, data, and new measures. *Research Policy*, *50*(2), 104144.

Aydin Ozden, S., & Khashabi, P. (2022). Patent remedies and technology licensing: Evidence from a supreme court decision. *Strategic Management Journal*.

Baker, J. B., & Bresnahan, T. F. (1985). The gains from merger or collusion in product-differentiated industries. *The Journal of Industrial Economics*, *33*(4), 427–444.

Ballesteros, L. (2021). Can natural disasters affect innovation? evidence from hurricane katrina. *Evidence from Hurricane Katrina (December 13, 2021)*.

Barney, J. B. (1986). Strategic factor markets: Expectations, luck, and business strategy. *Management science*, *32*(10), 1231–1241.

Baruffaldi, S., & Poege, F. (2020). A firm scientific community: Industry participation and knowledge diffusion. *Max Planck Institute for Innovation & Competition Research Paper*, (20-10).

Baten, J., Bianchi, N., & Moser, P. (2017). Compulsory licensing and innovation–historical evidence from german patents after wwi. *Journal of Development Economics*, *126*, 231–242.

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

Bena, J., & Li, K. (2014). Corporate innovations and mergers and acquisitions. *The Journal of Finance*, *69*(5), 1923–1960.

Berchicci, L. (2013). Towards an open r&d system: Internal r&d investment, external knowledge acquisition and innovative performance. *Research Policy*, *42*(1), 117–127.

Bikard, M., & Marx, M. (2020). Bridging academia and industry: How geographic hubs connect university science and corporate technology. *Management Science*, *66*(8), 3425–3443.

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

Brown, D. J., & Roche, M. (2022). The ties that no longer bind: Inventor mobility and patent litigation. *Academy of Management Proceedings*, *2022*(1), 11502.

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

Cassiman, B., Colombo, M. G., Garrone, P., & Veugelers, R. (2005). The impact of m&a on the r&d process: An empirical analysis of the role of technological-and market-relatedness. *Research policy*, *34*(2), 195–220.

Cloodt, M., Hagedoorn, J., & Van Kranenburg, H. (2006). Mergers and acquisitions: Their effect on the innovative performance of companies in high-tech industries. *Research policy*, *35*(5), 642–654.

Cohen, W. M., & Levinthal, D. A. (1989). Innovation and learning: The two faces of R & D. *Economic Journal*, *99*(397), 569–596.

Cohen, W. M. (2010). Fifty years of empirical studies of innovative activity and performance. *Handbook of the Economics of Innovation*, *1*, 129–213.

Cohen, W. M., Nelson, R. R., & Walsh, J. P. (2000). *Protecting their intellectual assets: Appropriability conditions and why us manufacturing firms patent (or not)* (tech. rep.). National Bureau of Economic Research.

Cohen, W. M., Nelson, R. R., & Walsh, J. P. (2002). Links and impacts: The influence of public research on industrial r&d. *Management science*, *48*(1), 1–23.

Colen, L., Belderbos, R., Kelchtermans, S., & Leten, B. (2023). Many are called, few are chosen: The role of science in drug development decisions. *The Journal of Technology Transfer*, 1–26.

Contigiani, A., Hsu, D. H., & Barankay, I. (2018). Trade secrets and innovation: Evidence from the "inevitable disclosure" doctrine. *Strategic Management Journal*, *39*(11), 2921–2942.

Cunningham, C., Ederer, F., & Ma, S. (2021). Killer acquisitions. *Journal of Political Economy*, *129*(3), 649–702.

Demirer, M., & Karaduman, O. (2022). *Do mergers and acquisitions improve efficiency: Evidence from power plants* (tech. rep.). Working paper.

de Rassenfosse, G., & Palangkaraya, A. (2023). Do patent pledges accelerate innovation? *Research Policy*, *52*(5), 104745.

Ederer, F., & Pellegrino, B. (2023). The great start-up sellout and the rise of oligopoly. *AEA Papers and Proceedings*, *113*, 274–278.

Endo, A. (2010). A historical perspective on the discovery of statins. *Proceedings of the Japan Academy, Series B*, *86*(5), 484–493.

Evenson, R. E., & Kislev, Y. (1976). A stochastic model of applied research. *Journal of Political Economy*, *84*(2), 265–281.

Fabrizio, K. R. (2009). Absorptive capacity and the search for innovation. *Research Policy*, *38*(2), 255–267.

Fleming, L. (2001). Recombinant uncertainty in technological search. *Management science*, *47*(1), 117–132.

Fleming, L., Greene, H., Li, G., Marx, M., & Yao, D. (2019). Government-funded research increasingly fuels innovation. *Science*, *364*(6446), 1139–1141.

Fleming, L., & Sorenson, O. (2004). Science as a map in technological search. *Strategic management journal*, *25*(8-9), 909–928.

Fosfuri, A., Lanzolla, G., & Suarez, F. F. (2013). Entry-timing strategies: The road ahead. *Long Range Planning*, *46*(4-5), 300–311.

Fumagalli, C., Motta, M., & Calcagno, C. (2018). *Exclusionary practices: The economics of monopolisation and abuse of dominance*. Cambridge University Press.

Galasso, A., & Luo, H. (2021). Risk-mitigating technologies: The case of radiation diagnostic devices. *Management Science*, *67*(5), 3022–3040.

Galasso, A., & Luo, H. (2022). When does product liability risk chill innovation? evidence from medical implants. *American Economic Journal: Economic Policy*, *14*(2), 366–401.

Gambardella, A. (1992). Competitive advantages from in-house scientific research: The us pharmaceutical industry in the 1980s. *Research Policy*, *21*(5), 391–407.

Gambardella, A. (1995). *Science and innovation: The us pharmaceutical industry during the 1980s*. Cambridge University Press.

Ganglmair, B., Robinson, W. K., & Seeligson, M. (2022). The rise of process claims: Evidence from a century of us patents. *ZEW-Centre for European Economic Research Discussion Paper*, (22-011).

Gao, X., Ritter, J. R., & Zhu, Z. (2013). Where have all the ipos gone? *Journal of Financial and Quantitative Analysis*, *48*(6), 1663–1692.

Giebel, M. (2021). Patent enforcement and subsequent innovation. *Available at SSRN 3992810*.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254–277.

Gordon, R. J. (2016). The rise and fall of american growth. In *The rise and fall of american growth*. Princeton University Press.

Griliches, Z. (1992). The search for r&d spillovers' scandinavian journal of economics, 94. *S29-S47*.

Griliches, Z. (1985). *Productivity, r&d, and basic research at the firm level in the 1970s* (tech. rep.). National Bureau of Economic Research.

Griliches, Z. (1998). Patent statistics as economic indicators: A survey. In *R&d and productivity: The econometric evidence* (pp. 287–343). University of Chicago Press.

Griliches, Z., Pakes, A., & Hall, B. H. (1986). The value of patents as indicators of inventive activity.

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

Hall, B. H., Jaffe, A. B., & Trajtenberg, M. (2001). *The NBER patent citation data file: Lessons, insights and methodological tools*.

Hartmann, P., & Henkel, J. (2020). The rise of corporate science in ai: Data as a strategic resource. *Academy of Management Discoveries*, *6*(3), 359–381.

Haucap, J., Rasch, A., & Stiebale, J. (2019). How mergers affect innovation: Theory and evidence. *International Journal of Industrial Organization*, *63*, 283–325.

Helfat, C. E. (1994). Evolutionary trajectories in petroleum firm r&d. *Management Science*, *40*(12), 1720–1747.

Helfat, C. E., & Lieberman, M. B. (2002). The birth of capabilities: Market entry and the importance of pre-history. *Industrial and Corporate Change*, *11*(4), 725–760.

Hicks, D. (1995). Published papers, tacit competencies and corporate management of the public/private character of knowledge. *Industrial and Corporate Change*, *4*(2), 401–424.

Hicks, J. (1963). *The theory of wages*. Springer.

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

Hitt, M. A., Harrison, J. S., & Ireland, R. D. (2001). *Mergers & acquisitions: A guide to creating value for stakeholders*. Oxford University Press.

Jaffe, A. B., & Trajtenberg, M. (2002). *Patents, citations, and innovations: A window on the knowledge economy*. MIT press.

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

Jensen, M. C. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *The American economic review*, *76*(2), 323–329.

Jinha, A. E. (2010). Article 50 million: An estimate of the number of scholarly articles in existence. *Learned Publishing*, *23*(3), 258–263.

Joshi, A. M., & Nerkar, A. (2011). When do strategic alliances inhibit innovation by firms? evidence from patent pools in the global optical disc industry. *Strategic Management Journal*, *32*(11), 1139–1160.

Kang, H. (2021). How does price competition affect innovation? evidence from us antitrust cases. *Evidence from US Antitrust Cases (October 10, 2021). USC Marshall School of Business Research Paper*.

Kang, H., & Lee, W. (2022). How innovating firms manage knowledge leakage: A natural experiment on the threat of worker departure. *Strategic Management Journal*, *43*(10), 1961–1982.

Katz, R., & Allen, T. J. (1982). Investigating the not invented here (nih) syndrome: A look at the performance, tenure, and communication patterns of 50 r & d project groups. *R&d Management*, *12*(1), 7–20.

Kelly, B., Papanikolaou, D., Seru, A., & Taddy, M. (2021). Measuring technological innovation over the long run. *American Economic Review: Insights*, *3*(3), 303–320.

Klepper, S. (2002). The capabilities of new firms and the evolution of the us automobile industry. *Industrial and corporate change*, *11*(4), 645–666.

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

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

Krieger, J. L., Schnitzer, M., & Watzinger, M. (2022). *Standing on the shoulders of science* (tech. rep.). Harvard Business School Working Paper.

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

Larsson, R., & Finkelstein, S. (1999). Integrating strategic, organizational, and human resource perspectives on mergers and acquisitions: A case survey of synergy realization. *Organization science*, *10*(1), 1–26.

Lee, H. (2023). The heterogeneous effects of patent scope on licensing propensity. *Research Policy*, *52*(3), 104696.

Leten, B., Kelchtermans, S., & Belderbos, R. (2022). How does basic research improve innovation performance in the world's major pharmaceutical firms? *Industry and Innovation*, *29*(3), 396–424.

Lichter, A., Löffler, M., Isphording, I. E., Nguyen, T.-V., Poege, F., & Siegloch, S. (2021). Profit taxation, r&d spending, and innovation. *ZEW-Centre for European Economic Research Discussion Paper*, (21-080).

Lieberman, M. B. (2005). Did first-mover advantage survive the dot-com crash. *Unpublished working paper, UCLA*.

Lieberman, M. B., & Montgomery, D. B. (1988). First-mover advantages. *Strategic Management Journal*, *9*(S1), 41–58.

Lieberman, M. B., & Montgomery, D. B. (1998). First-mover (dis) advantages: Retrospective and link with the resource-based view. *Strategic management journal*, *19*(12), 1111–1125.

Lieberman, M. B., & Montgomery, D. B. (2013). Conundra and progress: Research on entry order and performance. *Long Range Planning*, *46*(4-5), 312–324.

Mansfield, E. (1995). Academic research underlying industrial innovations: Sources, characteristics, and financing. *Review of Economics and Statistics*, *77*(1), 55–65.

Marco, A. C., Sarnoff, J. D., & Charles, A. (2019). Patent claims and patent scope. *Research Policy*, *48*(9).

Marx, M., & Fuegi, A. (2019). Reliance on science in patenting. *Available at SSRN*.

Marx, M., & Fuegi, A. (2020). *Reliance on science by inventors: Hybrid extraction of in-text patent-to-article citations* (tech. rep.). National Bureau of Economic Research.

Merton, R. K. (1935). Fluctuations in the rate of industrial invention. *The Quarterly Journal of Economics*, *49*(3), 454–474.

Mitchell, W. (1989). Whether and when? probability and timing of incumbents' entry into emerging industrial subfields. *Administrative Science Quarterly*, 208–230.

Mitchell, W. (1991). Dual clocks: Entry order influences on incumbent and newcomer market share and survival when specialized assets retain their value. *Strategic Management Journal*, *12*(2), 85–100.

Mokyr, J., Vickers, C., & Ziebarth, N. L. (2015). The history of technological anxiety and the future of economic growth: Is this time different? *Journal of economic perspectives*, *29*(3), 31–50.

Monath, N., Madhavan, S., DiPietro, C., McCallum, A., & Jones, C. (n.d.). Disambiguating patent inventors, assignees, and their locations in patentsview.

Moser, P., & Voena, A. (2012). Compulsory licensing: Evidence from the trading with the enemy act. *American Economic Review*, *102*(1), 396–427.

Moser, P., Voena, A., & Waldinger, F. (2014). German jewish émigrés and us invention. *American Economic Review*, *104*(10), 3222–3255.

Narin, F., Hamilton, K. S., & Olivastro, D. (1997). The increasing linkage between us technology and public science. *Research policy*, *26*(3), 317–330.

Nicholas, T. (2008). Does innovation cause stock market runups? evidence from the great crash. *American Economic Review*, *98*(4), 1370–96.

OECD. (2020). Oecd (2020), start-ups, killer acquisitions and merger control.

Paruchuri, S., Nerkar, A., & Hambrick, D. C. (2006). Acquisition integration and productivity losses in the technical core: Disruption of inventors in acquired companies. *Organization science*, *17*(5), 545–562.

Poege, F. (2022). Competition and innovation: The breakup of ig farben. *Available at SSRN*.

Poege, F., Harhoff, D., Gaessler, F., & Baruffaldi, S. (2019). Science quality and the value of inventions. *Science advances*, *5*(12), eaay7323.

Poletti, M., LING, H., & ENGELLAND, B. (2008). An empirical study of launch order valuation based upon stock market reaction. *Marketing Management Journal*, *18*(1).

Prager, R. A., & Hannan, T. H. (1998). Do substantial horizontal mergers generate significant price effects? evidence from the banking industry. *The Journal of Industrial Economics*, *46*(4), 433–452.

Roach, M., & Cohen, W. M. (2013). Lens or prism? patent citations as a measure of knowledge flows from public research. *Management Science*, *59*(2), 504–525.

Roach, M., & Sauermann, H. (2010). A taste for science? phd scientists' academic orientation and self-selection into research careers in industry. *Research policy*, *39*(3), 422–434.

Rosenberg, N. (1990). Why do firms do basic research (with their own money)? *Research policy*, *19*(2), 165–174.

Rosenberg, N. (1969). The direction of technological change: Inducement mechanisms and focusing devices. *Economic development and cultural change*, *18*(1, Part 1), 1–24.

Rosenberg, N. (1976). Science, invention and economic growth, ch. 15 in: Perspectives in technology, by n. rosenberg.

Roth, B. D. (2002). 1 the discovery and development of atorvastatin, a potent novel hypolipidemic agent. In *Progress in medicinal chemistry* (pp. 1–22, Vol. 40). Elsevier.

Sachar, H. M. (1993). *A history of the jews in america*. Vintage.

Sauermann, W., Henry Cohen. (2010). What makes them tick? employee motives and firm innovation. *Management Science*, *56*, 2134–53.

Schankerman, M., & Pakes, A. (1986). Estimates of the value of patent rights in european countries during the post-1950 period. *The Economic Journal*, *96*(384), 1052–1076.

Schumpeter, J. (1942). 142 capitalism, socialism and democracy. In *What are the questions? and other essays* (pp. 142–143). Routledge.

Schumpeter, J. A. (2013). *Capitalism, socialism and democracy*. routledge.

Sears, J. B. (2018). Post-acquisition integrative versus independent innovation: A story of dueling success factors. *Research Policy*, *47*(9), 1688–1699.

Sinha, A., Shen, Z., Song, Y., Ma, H., Eide, D., Hsu, B.-j. P., & Wang, K. (2015). An overview of microsoft academic service (mas) and applications. *Proceedings of the 24th international conference on world wide web*, 243–246.

Stern, S. (2004). Do scientists pay to be scientists? *Management Science*, *50*(6), 835–853.

Suarez, F. F., & Lanzolla, G. (2007). The role of environmental dynamics in building a first mover advantage theory. *Academy of Management Review*, *32*(2), 377–392.

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175–199.

Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research policy*, *15*(6), 285–305.

Town, R. (2001). The welfare impact of hmo mergers. *Journal of Health Economics*, *20*(6), 967–990.

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

Tripsas, M. (1997). Unraveling the process of creative destruction: Complementary assets and incumbent survival in the typesetter industry. *Strategic Management Journal*, *18*(S1), 119–142.

USPTO. (2019). PatentsView.

Vita, M. G., & Sacher, S. (2001). The competitive effects of not-for-profit hospital mergers: A case study. *The Journal of Industrial Economics*, *49*(1), 63–84.

Watzinger, M., & Schnitzer, M. (2022). The breakup of the bell system and its impact on us innovation.

Weber, Y. (1996). Corporate cultural fit and performance in mergers and acquisitions. *Human relations*, *49*(9), 1181–1202.