Essays in Corporate Finance

by

Muhan HU

ORCID:0009-0008-2716-1660

A thesis submitted in partial fulfilment for the

degree of Doctor of Philosophy

in the

Department of Finance

Faculty of Business and Economics

THE UNIVERSITY OF MELBOURNE

June 2023

Abstract

This thesis explores the determinants and effects of corporate innovation and technology spillovers, emphasizing their roles in productivity and economic gains. It contains a literature review and two essays examining different aspects of finance at its intersection with innovation.

The first essay investigates how competition affects the economic value of innovation, the primary incentive for corporate R&D investments. I measure the economic value of innovation based on the changes in patenting firms' stock market value around patent issuance dates, following Kogan et al. (2017). The economic value of innovation is higher in industries with a low level of competition. Within an industry, firms at the technological frontier or those with relatively high pricing power enjoy higher economic returns from patents. I use a quasi-natural experimental design to compare the value of patents issued immediately before and after competition-altering events. Using horizontal M&A announcements as anti-competitive events, I show that an expected decrease in market competition leads to increased patent value.

In the second essay, Lyndon Moore and I study technology diffusion mechanisms using a unique historical setting: the introduction of the cyanide method of gold extraction on the Witwatersrand goldfields in the late 19th and early 20th centuries. Mines managed by the same mining house were more likely to adopt the new process, which increased extraction rates by around 50%. Cyanide technology was continually improved after its introduction. We find evidence of "technological know-how" spillovers from engineers and mine managers but not white-collar employees. Geographical spillovers are extremely localized in nature.

Declaration of Authorship

I, Muhan HU, declare that:

- The thesis comprises only my original work towards the PhD in Finance except where indicated in the preface;
- due acknowledgement has been made in the text to all other material used; and
- the thesis is fewer than the maximum word limit in length, exclusive of tables, maps, bibliographies and appendices as approved by the Research Higher Degrees Committee.

Signed: Muhan Hu

Date: May 2023

Preface

Chapters 2 and 3 comprise exclusively of my own original unpublished work.

Chapter 4 is co-authored with Professor Lyndon Moore and is in revision following peer review by *Management Science*.

All co-authorships are in accordance with the Graduate Research Training Policy of the University of Melbourne.

I acknowledge fundings from the Australian Government Research Training Program, the Faculty of Business and Economics Doctoral Program Scholarship, and the Faculty of Business and Economics Graduate Research Abroad Travelling Scholarship.

Acknowledgements

I would like to express my heartfelt gratitude to my supervisory committee members, Hae Won (Henny) Jung, Lyndon Moore, Garry Twite (Principal), and Ekaterina Volkova, for their invaluable guidance, support, and encouragement throughout my PhD journey. Their expertise, insights, and patience have been instrumental in shaping my research and helping me to achieve my academic goals. The completion of this thesis would not have been possible without them.

I am also appreciative to all the faculty members in the Department of Finance at the University of Melbourne, especially James Brugler, Andrea Lu, and Patrick Kelly, for their valuable contributions to this thesis. A special thanks to Bruce Grundy, Spencer Martin, and Garry Twite for their contributions toward the PhD program.

Thank you to my fellow PhD students, Sheng Huang, Chunxiao Lu, Haoyi Luo, Linxiang Ma, Fanqi Meng, Mengchuan Wang, Bochen Wu, Lin Wu, Shuying Wu, Zijin Xu, Yuyang Zhang, and Yifan Zhou. Their intellectual discussions and emotional support have been a source of inspiration and comfort.

I am grateful to the professional staff in the Department of Finance for their kindness, support, and assistance throughout my PhD candidature.

I would also like to express my sincere gratitude to Neal Galpin, Dalida Kadyrzhanova, and Martine Mariotti for their valuable collaborations that enriched my research and broadened my perspectives.

I thank the Department of Finance, Faculty of Business and Economics, University of Melbourne, for their generous financial support and provision of resources over the last five years. I also want to thank my cat, Gululu Mahu, whose endearing silliness has been my constant source of comfort and joy. Finally, I would like to thank my husband and best friend, Linxiang Ma, without whom I would not have started and been able to complete my PhD journey.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis focuses on the impact of corporate innovation and technology spillovers on economic gains. It investigates the influence of product market competition on the economic benefits derived from technological innovation, as well as the various mechanisms of technology transfer that contribute to productivity improvement.

This thesis begins with a literature review summarizing key studies in the fields of competition and innovation. The relationship between competition and innovation has long been a central focus in economic analysis, stemming from Joseph Schumpeter's concept of Creative Destruction (Schumpeter, 1942). Successful innovation is usually a source of temporary market power, which erode incumbent firms' economic rent and market position. Previous research has examined how competition influences firms' innovation outcomes, often measured by patent or citation counts (Aghion et al., 2005; Hashmi, 2013). While those measures capture the scientific contribution of innovation, which is socially beneficial, it differs from the economic gain firms can internalize from innovation, which is the incentive for firms to make further R&D investments (Abrams et al., 2020; Kogan et al., 2017). However, it remains unclear how competition affects the primary motivation for firms to engage in innovative activities.

The second part of the literature review focuses on knowledge spillovers, specifically the impact of external knowledge capital on firms' productivity and performance. One line of research examines the spillover effects of R&D investment and successful innovation (Bloom et al., 2013b; Jaffe et al., 1993; Matray, 2021; Thompson and Fox-kean, 2005; Zacchia, 2020) while the other examine the transfer of the knowledge of "know-how" (Conley and Udry, 2010; David, 1990; Foster and Rosenzweig, 1995, 2010). I discuss the "reflection problem" in identifying the causality of knowledge spillovers and review novel identification strategies proposed by recent studies. Additionally, it explores different mechanisms of technology transfer, such as geographical proximity, individual connections, and organizational and ownership structures.

The first essay examines the impact of product market competition on the economic value of patented innovation, which serves as the primary incentive for firms' innovative activities. Existing literature suggests an inverted-U relationship between industry innovation outcome, as measured by patent count and citation count, and product market competition (Aghion et al., 2005). However, utilizing a newly developed measure of innovation value and a novel quasi-natural experimental design, this paper demonstrates that intense competition leads to a lower economic value of innovation that the inventing firm can internalize.

This study employs a stock market return-based measure of innovation value, following Kogan et al. (2017), which captures the expected incremental cash flows associated with granted patents. Through the quasi-natural experiment, a comparison is made between the value of patents issued immediately before and after industry-wide events that alter competition. The findings indicate that the patent value increases by approximately 4.8% immediately after horizontal merger announcements, with a stronger effect observed for mergers expected to be more anti-competitive. An opposite effect is found when using an alternative event that is expected to intensify product market competition. Beyond academic research, the results in this paper shed light on the potential benefit of giving firms some extent of market power.

The second essay investigates how competing mechanisms of technology spillovers improved productivity in a single-product, fixed-price setting: the introduction of the cyanide method of gold extraction in South Africa in the 1890s. My coauthor and I explore efficiency gains through three potential knowledge transfer mechanisms: organizational, personnel, and geographical networks. We hand-collect mine-level monthly productivity data from archival sources. We find knowledge spillovers at both the extensive and intensive margins. Mines managed by the same "mining house" (large shareholders of individual mining companies) were more likely to adopt the cyanide process. The adoption of cyanide increased the gold recovery rate by more than 50%, which allowed marginally profitable mines to survive and new mines to open.

After the initial adoption of cyanide, we document continuous productivity improvements as the most effective use of this new technology was refined. The technological know-how was transmitted among mines in the same mining house via the appointment of common managers and engineers. In the second half of our sample, we observe that mining houses often appointed a consulting engineer for all mines under their control. The engineer was likely responsible for transmitting productivity improvements from one mine to another. In contrast, administrative personnel such as directors and secretaries show no evidence of aiding technology transmission. Mines learned from their geographical neighbours, but the spillovers were very localized. Mines had to be within a couple of kilometres to gain from their neighbours' knowledge. We contribute to the knowledge spillover literature by identifying the channels through which productivity gains from technology diffusion are transmitted.

Chapter 2

Literature Review

The literature on innovation, examined from various perspectives such as economics, management, and finance, is extensive. Recent survey papers in financial economics have covered aspects including how financing (Hall and Lerner, 2010; Kerr and Nanda, 2015), managerial incentive (Ederer and Manso, 2011), and institutions (He and Tian, 2020) affect innovation. Additionally, Kogan and Papanikolaou (2019) discuss the implications of corporate innovation on asset pricing. In this chapter, I review relevant recent research on innovation, focusing on how product market competition influences firms' innovative activities and the diffusion of technological knowledge among individuals and institutions.

2.1 Competition and Innovation

Product market competition and innovation are important topics in the field of economics, and their relationship has been the subject of many studies. Product market competition refers to the degree of rivalry among firms in a particular market, while innovation refers to the process of introducing new products, services, or processes into the market. This section aims to provide an overview of the existing research on the relationship between product market competition and innovation.

Since the pioneering work of Schumpeter (1942), there have been extensive discussions on the relationship between product market competition and innovation. Under the Schumpeterian view, perfect market competition discourages innovative activities because the low market share and the high threat of obsolescence would decrease the amount of innovation rent. In the absence of entry barrier, immediate imitation by competitors erodes the inventor's post-innovation market share and pricing power (Scherer, 1967). Additionally, the innovation rent only last until the next innovation arrives, making the incumbent knowledge obsolete (Aghion and Howitt, 1992). Intensive competition increases the chances of a new innovation happening sooner (obsolescence rate) (Gu, 2016), which, in turn, limits the monopoly rent of innovation. Scherer (1967) also suggests an innovator need to cultivate the new market and win consumers from its competitors over time. This penetration process process might also be curbed by competitive pressure. Thus, the payoff of innovation is higher in the absence of intensive competition.

Arrow (1962) is the first who formally analyzes this relationship under the assumption that a single firm makes discrete choices about innovation under perfect competition and monopoly. Although innovation increases firms' profit under both market structures, the pre-innovation rent under perfect competition, zero, is much lower than the pre-innovation rent for a monopoly firm. Therefore, the net gain from innovation is higher in a competitive market relative to a monopoly market.

These two competing arguments, despite their simplicity, provoked a discussion about the critical relationship between market structure and innovation throughout the next few decades. Loury (1979) provides an equilibrium model which accounts for the technological uncertainty and market uncertainty. His model confirms the Schumpeterian hypothesis that more competition is not necessarily socially optimal as it discourages firms from investing in innovation. Dasgupta and Stiglitz (1980) consider the endogenous nature of market structure and innovation and provides an analytical framework with multiple firms and complete information. Their results are consistent with the Schumpeterian belief that innovative activities are negatively associated with market competition. In addition, Futia (1980) provides a dynamic model to characterize the long-run equilibrium of the Schumpeterian Creative Destruction process and show that the market competition and technological progress are jointly governed by the ease of entry and the extent of innovative opportunities.

Although theoretical works are in favour of the Schumpeterian hypothesis, empirical evidence remains mixed. Kamien and Schwartz (1975) survey early empirical works and find that the relationship between creative activities and firm size/market power/diversification remains inconclusive. For example, industry concentration could either spur, be neutral to, or harm firms' innovative output. Aghion et al. (2005) are the first to show an inverted-U relationship between the industry competition and firms' innovation rate empirically and theoretically: firms' innovation first increase then decrease with competition intensity. There are two driving forces in this relationship. The first the well-documented Schumpeterian effect where competition reduces the economic rent of innovation; the second is the "escape competition" effect, where firms' incentive to innovate comes from the low competitive profit. While the Schumpeterian effect explains the decreasing part of the competition-innovation relationship, the "escape competition" effect drives the increasing part. The relative strength of the two forces depends on the difference in technology advancement levels among firms in a sector.

However, a follow-up paper using US data only finds a mildly negative relationship between competition and innovation (Hashmi, 2013). To reconcile this disparity in empirical evidence, the author proposes a partial equilibrium model to overcome several limitations in the general equilibrium model in Aghion et al. (2005). Solving the model numerically, the author finds three possible relationships between innovation intensity and competition: decreasing, inverted-U, or increasing, depending on each industry's technological gap. Hashmi (2013) argues that the inconsistency is driven by the fact that industry compositions are different in two countries in terms of technological advancement.

The works mentioned above on product market competition and innovation extensively focus on the amount of firm innovative input (for example, R&D intensity) and output (for example, patent counts, patent citations). However, theoretical and empirical evidence about the relationship between market competition and innovation value is lacking. Moreover, the relationship between market competition and innovation rent often serves as an assumption in the aforementioned theoretical works. Although this assumption is pivotal to model implications, researchers made different assumptions in those theoretical works, and there is no consensus about which assumption is more appropriate. This inconsistency calls for a thorough examination of the relationship between and product market competition and innovation value. Chapter 3 presents causal evidence for a negative relationship between competition and the private return to innovation using US data. This finding supports Hashmi (2013); innovation output is higher in less competitive market as the pecuniary incentive for innovation is higher.

Innovation is valuable in many different ways. An invention is valuable as it contributes to scientific progress even it does not generate profit for firms. The right associated with a patented innovation also creates value as it enables the inventor to enjoy the temporary monopolistic power of the invention under legal protection. Moreover, there is also a signalling value of patents, especially for start-ups (Hall and Harhoff, 2012). Trajtenberg (1990) argues that patents vary significantly in their importance or value. Since then, the number of patent citations is commonly used in empirical works as a proxy for patent value. Hall et al. (2005) explore the relationship between market value and the number of patents' citation and find that a one-unit increase in the citation-to-patent ratio is associated with a 3% increase in firm market value. Nevertheless, this relationship diminishes considerably over time (Kuhn et al., 2020). Harhoff et al. (1999) find that highly cited patents are more likely to be renewed and are associated with higher survey-based economic value estimates.

Nonetheless, recent work by Abrams et al. (2020) exploit proprietary patent licensing revenue data and find an inverted-U shape relationship between the value of patents and the number of citations, where patent value first increase then decrease in the number of citations. This hump-shaped relationship suggests that the number of citations might not be a good proxy for patents' economic value. Kogan et al. (2017) distinguish between patents' scientific value and economic value. The scientific value is measured by the number of a patent's forward citations, capturing a patent's importance in the scientific world. The economic value is based on stock market value fluctuations around the patent issuance dates, reflecting the expected value of discounted incremental future cash flows a firm obtained from a patent. Chapter 3 also shows that exogenous changes in product market competition affect the economic gain of existing patent but not the scientific value.

Pakes (1985) is the first who link the stock market reaction on patent announcement day to the economic value of patents. Under a simplified assumption, the author argues that the stock market reaction should partially reflect the expected discount present value of the patenting event. A similar approach is adopted by Austin (1993) who shows that the economic value of patents is associated with various attributes of the patent. It is noteworthy that the stock market reaction only captures the relative value of innovation and provides a lower bound for patents' actual value because the market already priced the probability of a successful patent application when a patent is being filed. After a patent issuance announcement, the market revises its initial assessment. Results in Chapter 3 is robust to incorporating both patent filing and issuance date market reactions when accessing patent value.

Later work tries to link the product market competition to the firms' market value of innovation. Blundell et al. (1999) show that the incremental firm market value per innovation stock is higher for dominant firms, in terms of sales volume, within an industry. Greenhalgh and Rogers (2006) investigate the stock market return to innovation in six sectors. By regress firms' market value on R&D expense scaled by total asset for each sector, they show that the coefficients on R&D intensity vary across industries and sectors with the lowest competition measure also have the highest innovation return coefficient and vice versa. These two studies hint at the potential relationship between competition and patents' value, yet, none of them directly tested this relationship and addressed possible endogeneity problems.

Im et al. (2015) examine how annual stock excess returns' sensitivity to innovation differs with competition intensity. They use the number of citations received from patents a firm applied in the previous year to measure innovation. The results imply that the innovation-return premium coefficient is the highest under moderate competition and declines when the market becomes more/less competitive. Nonetheless, Gu (2016) formally studies the R&D premium and market competition using portfolio sorting and shows that the positive relationship between research intensity and monthly the high-minus-low R&D portfolio return only exists in industries under high competitive pressure. These two papers investigate how stock excess return on innovation varies with competition and reach different conclusions. One potential explanation is that Gu (2016) looks at the stock's R&D premium while Im et al. (2015) focus on annual excess return related to realized innovation, and the former is riskier compared to the latter.

Previous research has examined the connection between stock market value, innovation, and market competition. However, no study has directly explored the relationship between market competition and the economic value of innovation. Therefore, Chapter 3 presents a comprehensive study on the impact of product market competition on patent value while holding other factors constant.

2.2 Technology Transfer and Knowledge Spillovers

This literature review section aims to examine previous studies on technology transfer and knowledge spillovers. According to Griliches (1979), knowledge spillovers refer to the impact of external knowledge capital on a firm's productivity. Spillovers can affect productivity through production or research and development processes. Griliches proposes that a firm's production function should incorporate its in-house knowledge capital and the outside knowledge capital available to it. Since then, a significant body of literature on knowledge spillovers has been developed (see Griliches (1992) for a review of early papers).

One strand of literature examines R&D spillovers, specifically the impact of one firm's R&D investment on the innovation of neighboring firms. Neighboring firms can be identified based on product market, technological, or spatial proximity. Bernstein and Nadiri (1989) test intra-industry R&D spillovers. Using the summation of the knowledge capital of all other firms within the same Standard Industrial Classification (SIC) 2-digit industry to measure R&D spillovers, they show that intra-industry R&D spillovers are associated with firms' cost reduction, production structure, capital accumulation rate, and R&D investment.

To test knowledge spillovers among technologically connected firms, Jaffe (1986) develops a pairwise technological proximity measure based on the similarity of patent distribution across patent classes. This measure is used to determine the weight of R&D expenditure from technology neighbors in calculating the technology spillover pool. Empirical findings indicate that firms with better access to external knowledge tend to receive more patents and achieve higher dollar returns, given their R&D expenditure.

Additionally, evidence suggests that knowledge spillovers are more likely to occur among firms in close geographical proximity. Jaffe et al. (1993) employ patent citation data to trace innovation spillovers between citing and cited firms. They calculate the fraction of citing patents within the same geographical area as the corresponding patents and compare it with a fraction computed using a sample of matched non-citing patents. Results demonstrate that non-self citing patents are more likely to be located in the same geographical area compared to matched non-citing patents. Furthermore, this difference diminishes as the definition of the geographical area becomes broader.

In addition to the knowledge of developing new technology, knowing the best use of new technology is equally important. The experience-based learning of know-how and its diffusion is critical to adopting new technology and its performance realization (David, 1990). A thinner but fruitful strand of literature examines the spillovers of "know-how". The idea that "...the profitability of any new technology grows over time as knowledge accumulates" (Foster and Rosenzweig, 1995) motivates many economists to explore how know-how spillovers enter production functions. For example, Foster and Rosenzweig (1995) develop a target-input model to incorporate learning-by-doing and learning from neighbours. When adopting new technology, a farmer updates his prior belief about the optimal input use according to information he obtained from his own experience and his neighbours' experience. Profitability increases with both learning by doing and learning from neighbours, with this positive feedback diminishing over time. The authors test the model's implications using the adoption of high-yielding seed varieties (HYVs) in the Indian agricultural sector, employing panel survey data to estimate the model and find results consistent with the model's predictions.

The knowledge spillover and social learning literature faces identification challenges, as highlighted by Manski (1993). There is a "reflection" problem: an outside observer without knowledge of optics and humans cannot understand whether (1) a person causes its mirror image to move, (2) a mirror image causes a person to move, or (3) they move simultaneously. Thus, without convincing identification, the inference that one reference group affects another's behaviour is akin to claiming that a mirror image causes a person to move.

Economists have employed various econometric techniques to address this challenge. One approach is to control for factors that could lead to simultaneous changes in productivity. Conley and Udry (2010) tackle the issue of learning from social networks by accounting for unobserved factors that may impact a farm's production. They study Ghana's agricultural sector in the 1990s, specifically the introduction of pineapple cultivation among local farms. The knowledge of fertilizer use is critical to the harvest of pineapple but new to the local farmers. The optimal input use level needs to be gradually acquired from both farmers' and peers' cultivation experience. They identify a farm's information neighbours and geographical neighbours, then use a farm's deviation from its geographical neighbours' fertilizer use to account for unobserved factors that might cause changes in fertilizer use. Conditional on this deviation, they find farmers adjust the fertilizer use intensity toward that of their information neighbours who experienced a sharp increase in outputs in the previous period. Nonetheless, similar results are not found in falsification tests in which farmers cultivated other traditional crops as fertilizer use for those crops is established knowledge.

Juhász et al. (2023) present empirical evidence of spatial know-how spillovers in the adoption of mechanized cotton spinning during the First Industrial Revolution in France. Using hand-collected plant-level archive data, they show that geographical distance to a high-productivity plant is negatively correlated with a cotton mill's productivity. To address the concern of potential confounding factors related to plant location, they account for variables such as access to waterpower, proximity to coal, and the extent of forest cover. Furthermore, the authors conduct falsification tests on the metallurgy and paper milling sectors, which have relatively standardized technologies during the sample period, finding no similar results.

An alternative approach to address the reflection problem is to employ a valid instrument for the behaviour of reference groups. Matray (2021) examines innovation spillovers from US-listed firms to private firms and uses the staggered adoption of Business Combination (BC) laws in each state as an exogenous shock to public firms' innovation activity. The rationale is that BC laws reduce the takeover threat faced by listed firms, leading to decreased innovation, but they should be unrelated to other determinants of private firms' R&D. The author finds that the number of patents filed by private firms decreases as the number of public firms' research labs affected by BC law increases within the same commuting zone. Furthermore, this negative effect diminishes when the affected labs are located in more distant areas.

Other similar examples include instrumenting reference firms' R&D using changes in federal and state tax credits (Bloom et al., 2013a), instrumenting connected firms' R&D using other sufficiently distant firms in the network space (Zacchia, 2020), and the relocation decisions of connected neighbours (Chu et al., 2019).

Conducting randomized experiments is another method to tackle this endogeneity problem, and it is commonly used to study knowledge diffusion of new financial and healthy technologies. The decision to adopt an invention requires knowledge about the advantages of using that technology. Adoption decisions related to various innovations have been investigated using this method, including Tax-Deferred Account retirement plan (Duflo and Saez, 2003), deworming pills (Kremer and Miguel, 2007), high-quality mosquito nets (Dupas, 2009), and menstrual cups (Oster and Thornton, 2012). These studies utilize different experimental designs in a small group and consistently find that an agent's decision to adopt new technology is influenced by whether the individuals they are connected to understand the advantage of adoption.

Another challenge in identifying R&D spillovers is distinguishing between technology/knowledge spillovers and the effect of product market rivalry (Bloom et al., 2013b). Combining these two types of spillovers can contaminate the estimated positive technology spillover effect. To address this issue, Bloom et al. (2013b) develop an analytical framework and empirically test model predictions. They measure spillovers from both the technology space and the product market space and employ the change in federal and state tax credits as an instrument for firms' R&D. Theoretical and empirical findings indicate a positive externality of technology neighbours' R&D on firm market value, patent output, and productivity. In contrast, product-market spillovers negatively affect firm value and show no significant impact on patenting activity and productivity. Overall, technological spillovers dominate the effect of product market rivalry.

The extant spillovers studies suffer from many of the issues outlined in the survey article by Syverson (2011). Most firms produce multiple products, and these products can be aggregated in various ways. Second, (p. 330) "even detailed producer microdata do not typically contain measures of output quantities". Third, measurement issues also arise for inputs – how to combine employees into a single measure of labor input and how to measure physical capital. Syverson cautions that any variation in unobserved input quality (p. 331) will show up as productivity.

Historical data can represent a unique opportunity to provide a detailed account of the technology diffusion. Nuvolari et al. (2023) examine between-country knowledge transfer. They analyze the technology transfer from Britain to France in the early stages of industrialization, using a comprehensive dataset of patents granted in France from 1791 to 1844. Leveraging unique aspects of French patent laws, they create various measures of patent quality. They show that patents originating from British inventors or French inventors with links to British inventors tended to have higher quality, revealing technology transfer from Britain to France.

Braguinsky et al. (2021) explore within-firm knowledge transfer mechanisms in Japanese cotton spinning industry from 1893 to 1914. With detailed product introduction data, they show that failed "push the boundary" trials help firm growth through successful implementation of accumulated knowledge to follow on horizontal product proliferation. Their findings support knowledge transfer mechanisms within the organization. Comin and Hobijn (2010) trace technology diffusion across 166 countries over the period from 1820 through 2003. They document that variations in technology adoption lag explain more than 25% of the differences in per capita income among their sample countries. These historical investigations enrich our understanding on technology diffusion.

Early studies on R&D spillovers assumed a connection between firms in the same product market or in close technological/geographical proximity. However, identifying the specific mechanisms through which knowledge is transmitted presents a challenge. The remainder of this literature review summarizes papers that investigate various channels affecting knowledge spillovers, such as personnel, organization, and geography.

Personnel

Technology transfer can also occur through the mobility of skilled workers, who carry tacit knowledge and skills when they switch jobs or move between countries. Babina and Howell (2022) show that employees in R&D-intensive firms are more likely to leave their old employer to start a new business. Schnitzer and Watzinger (2022) focus on the opposite direction, i.e., knowledge spillover from start-ups to mature firms. Both paper instrument firms' R&D expenditures use government R&D tax credits to address potential endogeneity problems. Similarly, Cao et al. (2022) documente the presence of "serial venture employees" that skilled labour from established entrepreneurial firms sequentially join other early-stage start-ups. Such inter-firm exchange of human capital predicts young start-ups' future performance. These three papers focus on "entrepreneurial" spillovers in start-up firms.

Worker mobility not only aids the spread of entrepreneurial ideas and management practice but also assists in transmitting technological knowledge. For example, Serafinelli (2019) documents local worker job hopping as an engine that transfers knowledge from high- to low-productivity firms. Song et al. (2003) study engineer mobility in the semiconductor industry and examine under what circumstances intra-industry knowledge spillovers are more likely to happen. They show that the likelihood of a hiring firm citing patents issued by the previous employer is associated with the hired workers' expertise. Stoyanov and Zubanov (2012) track worker movements among Danish manufacturing firms to investigate whether productivity can be gained by hiring labours from highly productive firms. Specifically, they test how firm productivity varies with the productivity gap between incumbent and new workers. They control for unobserved productivity shock using lagged productivity, capital investment, and its polynomial function. Results suggest a positive relationship between a firm's productivity gain and the average productivity gap between itself and incoming workers' previous employers. A similar approach is applied by Tambe and Hitt (2014), who investigate the diffusion of know-how of IT innovation. They use the conventional spillover measure, weighted IT investment intensity, to construct the IT know-how pool, where the weight equals the percentage of the focal firm's incoming IT workers from each other IT firm each year. Results suggest that firms' productivity increases with their IT knowledge pool.

These studies indicate that the flow of knowledge may not be reciprocal and symmetric but unidirectional. Specifically, firms that lag technologically tend to benefit from the positive externalities generated by firms at the technological frontier. Furthermore, while the findings on the positive impact of labour mobility on firms' productivity are compelling, caution should be exercised in interpreting the results. While they may suggest potential knowledge spillovers from high-productivity firms to low-productivity firms, other explanations are possible. It is unclear whether the productivity gains are due to the positive externalities generated by the job hoppers' previous employer or the skills possessed by the workers, which may be unrelated to their previous employer.

In addition to worker mobility, individuals' networks impact how knowledge flows. Networks enable individuals to interact with one another. The interaction could be face-to-face, such as in a conference or a bar, or virtual, such as exchanging emails or online meetings (Bloom et al., 2013b). Empirical evidence further confirms the critical role of networks in knowledge spillovers. Researchers have explored farmers' social network (Conley and Udry, 2010), inventors' past collaborations (Singh, 2005; Zacchia, 2020), co-authorship networks (Azoulay et al., 2010), and workplace peers (Waldinger, 2012), and demonstrates that social interaction facilitates the diffusion of technological knowledge, hence, innovation and productivity.

Regarding who matters for knowledge spillovers, most studies mentioned before focus on the rank-and-file employees without further distinguishing their roles. As for papers on technology spillover, scientists/inventors are arguably the most studied (see, for example, Azoulay et al., 2010; Breschi et al., 2017; Javorcik and Spatareanu, 2011; Kerr, 2008; Waldinger, 2012; Zacchia, 2020) because scientists are the primary force behind innovation. The available data on scientists/inventors' collaboration and output history also motivate those studies. For example, Zacchia (2020) constructs a dynamic research networking measure for each firm using its researchers' past co-patenting experience. The external knowledge pool of a firm is measured by its inventor-connection weighted external R&D capital. The author instruments connected firms' R&D using other sufficiently distant firms in the network space relative to the focal firm and shows firm performance and innovation in terms of patent counts increase with R&D spillovers through researchers.

However, whether firms' management and supervisory teams also contribute to integrating external knowledge to improve firm performance is less studied. Juhász et al. (2023) demonstrate that when a breakthrough technology is introduced, coordinating organizational structure to adapt to a new production process is critical for productivity improvement. Such change in organizational structure requires management participation and innovation. Thus, managers' ability to absorb information necessary for new technology adoption could also affect the value of new technology. Managers with superior access to information and knowledge about firm management practices also tend to perform better. Bloom et al. (2013a) study a sample of Indian firms using a randomized experiment and find that by accessing information about frontier management practice, managers increase firm productivity by an average of 17% in the first year. Yuan and Wen (2018) examine the influence of manager foreign experience on corporate innovation among a group of Chinese companies. Their findings indicate a positive impact, with senior managers having a more pronounced effect than junior managers. This difference suggests that, in addition to rank-and-file employees, top-level management can also promote innovation by applying the knowledge and technical skills they acquired in developed countries to developing economies.

The board of directors, which typically serves a supervisory or monitoring function, can also drive innovation within a company by leveraging the information and knowledge they have gathered from external sources. Research suggests that directors draw upon their expertise and experience to enhance the firm performance and value. For instance, foreign directors can apply their knowledge of their home country to assist companies in making better cross-border acquisition decisions (Masulis et al., 2012). Similarly, directors with foreign experience can employ their acquired foreign knowledge to benefit the businesses in their home country (Giannetti et al., 2015). Helmers et al. (2017) investigate the role of the board of directors in firm innovation. Using a regulation-induced exogenous change in board structure, they show that expanding directors' network size increases firms' innovation efforts (R&D) and outcomes (patenting). They also find evidence that the influence on innovation is brought about by the transmission of information through director interlocks. All papers suggest that, despite their primary supervisory role, directors integrate knowledge from external resources to enhance their current employers' performance.

Organization

Firms' organizational and ownership structures can play a critical role in facilitating technology transfer by providing various types of support. Collaboration among companies, especially in research and development, can facilitate technology transfer. Gomes-Casseres et al. (2006) study whether interfirm strategic alliance facilitate diffusion of technological knowledge using patent citation data. They find evidence that the probability of patent citation is higher between pairs of allied firms than non-allied pairs. Banal-Estañol et al. (2022) focus on firms' participation in research joint ventures (RJVs) and find that RJV involvement enables firms to enhance their absorptive capacity of technological spillovers and generate value. Recent studies explore the unique role of venture capitalists in technology transfer. Lindsey (2008) finds that firms that share a common venture capitalist are more likely to form strategic alliances. González-Uribe (2020) shows that venture capitalists can significantly increase their portfolio firms' returns to innovation by coordinating innovation resources. Hu (2023) presents evidence that venture capitalists can enhance labour mobility among start-up founders by creating an internal labour market among their portfolio firms. While he does not directly examine whether such a labour market aids technology diffusion, previous research has shown that increased worker mobility can facilitate knowledge spillovers. Therefore, it is plausible that such internal labour markets create a more conducive environment for knowledge spillovers.

Furthermore, theoretical frameworks are developed to capture knowledge spillovers through partial equity ownership (Ghosh and Morita, 2017) and common ownership (López and Vives, 2019), yet compelling empirical evidence is still lacking.

In addition to firms' organizational structure, external professional bodies could also serve as intermediaries which link the focal firm to external knowledge repositories. According to Wagner et al. (2014), professional service firms increase the likelihood of their clients drawing upon one another's knowledge pool.

Moreover, recent works in finance explore the role of business relationships and ownership structure in knowledge diffusion among large corporations. For instance, Chu et al. (2019) use the supplier-customer connection reported by US public firms to explore how geographical proximity affects knowledge transmission from downstream firms to upstream firms along the supply chain.

Geography

How geographical proximity between two entities affects the likelihood of technology spillovers between them? Jaffe et al. (1993) utilize patent citation information to trace innovation spillovers between citing and cited firms. They calculate the fraction of citing patents that are in the same geographical area as those cited patents and compare this number with the fraction based on a sample of matched non-citing patents. Results show that those (non-self) citing patents are more likely to be in the same geographical area than those matched non-citing patents. Moreover, this difference narrows when the definition of the geographical area is broader. Thompson and Fox-kean (2005) reassess the localization of knowledge spillovers using Jaffe et al. (1993) 's approach but with a more restricted control group selection criteria to control for "the existing geographic distribution of production", which would likely affect the probability of citation within an agglomeration. They find evidence for the localization of knowledge spillovers at the country level rather than at the statistical area and state levels.

The studies above proxy the closeness among firm pairs based on whether two firms are located within the same administrative/political unit, for example, county, state, and country. However, under this approach, two adjacent firms located on two sides of a borderline are not considered geographical neighbours despite the physical distance between them being at a minimal level. Hence, it is important to consider how physical distance affects technology diffusion to gain a more comprehensive understanding of spatial knowledge spillovers. Singh and Marx (2013) reexamine the localization of technology spillovers within administrative units following Jaffe et al. (1993) but account for spatial proximity. They find evidence that physical proximity impedes knowledge spillovers; nonetheless, even after controlling for geographical distance, both the state and country borders still appear to be an obstacle that impedes knowledge spillovers, as patent citation likelihood is significantly higher among pairs of patents invented in the same state.

In addition to the limitation of using administrative borders, using patent citations to measure knowledge flow also faced criticism. Although in the technology spillover literature, a citation dyad is typically viewed as an indication that the citing patent incorporates knowledge from the cited patent, the United States Patent and Trademark Office (USPTO) has noted that "The basic purpose for citing prior art in patent files is to inform the patent owner and the public in general that such patents or printed publications are in existence and should be considered when evaluating the validity of the patent claims.[1]"

Patent citations are often used as a sign of knowledge transfer from the cited invention to the invention citing it. Therefore, the localization of patent citations is taken as evidence that such knowledge transfer is also localized. The merit of using citation data to investigate knowledge spillovers is twofold. First, citations leave a paper trail by which knowledge flows may be measured and tracked.

[1] See: https://www.uspto.gov/web/offices/pac/mpep/mpep-2200.html

In addition, the US Patent and Trademark Office also provides information on the patent inventors' locations, making research on the geography of knowledge spillovers feasible. However, relying on citations to measure knowledge flows has limitations. Alcácer and Gittelman (2006) demonstrate that a significant proportion of citations (about 63%) are added by patent examiners rather than inventors, with around 40% of citing patents having all citations contributed by examiners. While inventor-added citations are more likely to capture genuine knowledge flow, examiner-added citations often reflect the inventor's lack of awareness (Jaffe et al., 2002). Consequently, studies that are based on all patent citations may introduce bias when estimating the localization effect of knowledge spillovers.

To address this issue, Thompson (2006) utilizes examiner-added citations as a benchmark for inventor-added citations. The assumption is that patent examiners do not have a geographical bias when adding additional citations to the patents under review, as they all work in Alexandria, Virginia. Thus, examiner-added citations can account for other geographical or industry unobservable factors that lead to the observed localization in patent citations. Leveraging identification strategies employed by Jaffe et al. (1993) and Thompson and Fox-kean (2005), he finds that inventor-added citations exhibit a higher degree of localization than examiner-added citations. Similarly, Singh and Marx (2013) adopt examiner-added citations as a benchmark in a regression framework and incorporate physical proximity in their analysis in addition to administrative borders. They find similar evidence supporting stronger localization in inventor-added citations compared to examiner-added citations.

However, even with access to recently publicized data distinguishing between citations added by inventors and examiners, using patent citations to establish localized knowledge spillovers still face challenges. Firstly, while inventors' and examiners' citations are not mutually exclusive, they cannot be separately identified using public data. Kuhn et al. (2020), using proprietary USPTO internal data, reveal that approximately 7% of examiner-added citations (as indicated in public USPTO data) are also included in the initial patent application submitted by the inventor(s). Thus, public data cannot perfectly identify all citations made by inventors, potentially concealing some knowledge flows behind examiner citations. Secondly, Barber and Diestre (2022) find evidence that inventors strategically avoid making specific citations to seek more lenient examiners and obtain more favourable outcomes for their applications. This strategic behaviour suggests that specific knowledge flows may intentionally be obscured. Lastly, Arora et al. (2018) challenge the assumption that the localization of patent citations implies localization in knowledge spillovers. They identify two categories of citations: one less likely to capture genuine knowledge flow and one more likely to do so. They find no systematic difference in localization between these two citation groups. This evidence raises doubts regarding previous research that relied on patent citation patterns to support the localization of knowledge spillovers.

Researchers also test the existence of knowledge spillover by examining whether a focal firm's productivity and innovation are affected by its geographical neighbours' knowledge. The following papers show that the positive effect of knowledge spillovers diminishes with physical distance. Matray (2021) examines the impact of an exogenous variation in public firms' R&D on local private firms' innovation outcomes. He shows that public firms' R&D has a robust positive externality on private firms in the same commuting zone. This spillover effect halves for firms in an adjacent commuting zone but does not affect firms in more distant areas. Similar results have been demonstrated by Autant-Bernard (2001) using French patent data. They find that knowledge spillovers are bounded within the French administrative department, with R&D from a close neighbour department not significantly affecting the focal department's innovation activities. The differences in findings could be due to the different sizes of the administrative units used in the studies, highlighting the importance of physical distance in the diffusion of technological knowledge. In addition, von Graevenitz et al. (2022) explore innovation in trademark tokens and find that distance hampers the diffusion of new ideas.

While these studies indicate that technology spreads more easily when people are located nearby, it remains unclear how close people need to be to learn from each other. According to Eriksson (2011), only the inflow of skills related to a plant's existing knowledge base from less than 50 kilometres away positively impacts the focal plant's performance. Meanwhile, Singh and Marx (2013) document that the probability of being cited drops by 36% when the distance between inventors doubles. These studies collectively suggest that the division of political units and physical distance play important roles in the diffusion of technological knowledge, with 50 kilometres being a potential minimum distance for knowledge spillovers.

While the knowledge spillover literature commonly finds technology diffusion is bounded by country borders, some evidence suggests that knowledge can also flow across country borders. Foreign direct investment (FDI) is one of the mechanisms through which knowledge can flow. For example, Abebe et al. (2022) investigate FDI in Ethiopia and find that, during the four years following the opening up of FDI plants, domestic plants located in treated districts (with FDI openings) experienced increases in total factor productivity and employment relative to untreated firms. Javorcik and Spatareanu (2011) study FDIs in Romania and find that the "brain gain" effect from FDI varies based on the origin of foreign investors. Nonetheless, FDI might also crowd out their local peers in less developed regions by bringing in competition with their superior technology. Bwalya (2006) studies FDI in Zambia and find the rise in FDI within a sector affects the productivity of domestic firms negatively, which could indicate the negative rivalry effects brought by FDI. Moreover, he finds knowledge spillovers occur vertically, wherein foreign companies in upstream sectors transmit knowledge to local firms in downstream sectors.

In addition to FDI, international trade and multinational enterprises (MNEs) also significantly facilitate international technology spillovers. Ayerst et al. (2023) utilize the input-output tables of country pairs to show that global trade enables importing countries to learn about new technological knowledge embodied in the imported goods, facilitating new technology to diffuse across countries. According to a study by Marino et al. (2020), inventors with ethnic backgrounds tend to integrate knowledge from their home country into the innovation process of MNEs. Moreover, ethnic scientists residing in other countries can also facilitate the flow of technological knowledge to their home countries. Kerr (2008) finds that an increase in the immigration of ethnic scientists in technological frontier countries, resulting from exogenous changes in US immigration quotas, can lead to a rise in manufacturing output in those scientists' countries of origin. Additionally, as mentioned earlier, foreign directors or directors with foreign experience can bring valuable foreign skills to their hosting firms (Giannetti et al., 2015; Masulis et al., 2012). These findings underscore the importance of scientists and labour mobility in the diffusion of technology across country borders.

Development in transportation infrastructure could also overcome the barrier in knowledge flows imposed by physical distance. Better connectedness among cities increases the likelihood that innovators access knowledge inputs from more distant neighbours, thus, facilitating the flow of local knowledge. Bahar et al. (2023) gathered direct flight data and employed a regression discontinuity design to show that a 10% increase in nonstop flights between two cities results in a 3.4% rise in citations and a 1.4% increase in collaborative patents between the departing city and the destination city. Similarly, Murata et al. (2014) examine the effect of transportation infrastructure on innovation by analyzing the introduction of interstate highways. They find that a 10 per cent increase in a region's stock of highways leads to a 1.7 per cent higher patent application in the following five years. Thus, enhancing transportation infrastructure can boost regional growth and promote technological development through a better flow of technological knowledge.

In the digital age, advanced information technology is changing the traditional ways knowledge transfer occurs, such as physical interaction. A study by Zheng and Wang (2020) find that the unexpected ban on Google in China limits Chinese inventors' ability to explore remote technological areas compared to a control group unaffected by the ban. However, inventors with well-connected collaboration networks can alleviate the observed adverse effects caused by the Google ban. Similarly, Wang et al. (2022) discovered that without access to the Google search engine, analysts affiliated with domestic brokers in China exhibited more optimistic views about events outside China than domestic events. This difference did not exist for analysts affiliated with brokers with foreign ownership, as they may access some uncensored foreign information. Thus, while Internet technology allows information to flow across borders, individual interaction still plays a significant role in disseminating technological knowledge.

To sum up, technology transfer and knowledge spillovers are essential concepts that promote innovation and economic growth. While there are various mechanisms and channels through which they can occur, their extent and nature can be influenced by factors such as worker mobility, geographical proximity, and even transportation infrastructure. Further research is needed to understand better the dynamics of technology transfer and knowledge spillovers and their implications.

Chapter 3

Competition and the Value of Innovation

3.1 Introduction

The United States has spent 2-3% of GDP on research and development (R&D) over the past few decades.[1] The economic gain a firm obtains from innovation is the primary motivation for corporate R&D investment but we know little about how different factors might influence this value. Investigating such impacts can help firms' management and policymakers to make better-informed decisions to promote economic growth. Researchers have long realized that competition is crucial in determining innovation outcomes.[2] Nonetheless, the competition-innovation value relationship is *a priori* unclear. The available empirical evidence is further limited because the commercial value of innovation is not directly observable, and competition and the value of innovation simultaneously affect each other.

I address this gap by studying the causal impact of product market competition on the economic value of innovation and find that competition negatively affects the value of innovation. I focus on patented innovation and measure the value

[1] Source: https://data.worldbank.org/indicator/GB.XPD.RSDV.GD.ZS?locations=US
[2] See, for example, Aghion et al. (2005, 2001); Aghion and Howitt (1992); Schumpeter (1942).

of a patent following Kogan et al. (2017). This measure is based on the patenting firm's cumulative abnormal return on its equity around a three-day patent issuance window and captures the present value of expected incremental future cash flows associated with the underlying innovation. To establish causality, I use a quasi-experimental design to compare the value of patents issued immediately after competition-altering events to the value of patents issued before those events. I use horizontal Merger and Acquisition (M&A) announcements as anti-competitive events and the United States granting Permanent Normal Trade Relations (PNTR) to China as an event that is expected to intensify competition. On average, patents issued immediately after horizontal merger announcements have a 4.8% higher market value than patents issued immediately before those announcements. Such a positive effect is more pronounced for deals that are expected to be more anti-competitive: 9.1% for "stealth mergers" with value just below the antitrust scrutiny threshold (see Kepler et al., 2021) and 11.1% for mergers in concentrated markets with high product similarity (see Fathollahi et al., 2022). On the contrary, patents' value is significantly lower after the US-China PNTR agreement than before the announcement.

Innovation grants a firm temporary pricing power to extract an economic rent. A firm's post-innovation rent will be higher than its pre-innovation rent. This incremental amount captures the firm's private economic gain from innovation and might be affected by competition. On the one hand, a high level of competition limits the profits a firm can obtain without innovation (Arrow, 1962). In the extreme case, firms earn zero profits under perfect competition. Low or zero economic rents encourage firms to innovate to "escape" from intense competition and earn higher profits. In addition, when product substitutability is high, a characteristic of intense competition, innovation can more easily help the inventor to attract market shares from its competitors (Raith, 2003), hence, increasing the value of innovation. On the other hand, a firm's ability to derive economic rents from innovation is also constrained by high competition intensity. For example, intense competition might increase the pace of follow-on innovation and lead to technological obsolescence (Gu, 2016; Ma, 2021). Therefore, the net impact of competition on the value of innovation remains unclear.

There is anecdotal evidence that companies engage in anti-competitive activities to enjoy a higher innovation rent. For example, the European Union Commission investigated a technology acquisition and five patent settlement agreements of Servier, a French pharmaceutical company, related to perindopril, Servier's main product.[3] It was established that these transactions blocked generic entry, which would have led to a significant drop in perindopril's price and a shift of market share from Servier to generic companies. By preempting competition, Servier prolonged its exclusivity over perindopril, maintained its market share, and hence enjoyed a higher economic rent from an established invention. Servier is not a single case; empirical evidence finds that pharmaceutical firms often acquire a target and 'kill' the target's innovative projects (Cunningham et al., 2021), suggesting that competition might reduce the expected profit firms gain from innovation.

In my baseline test, I measure product market competition based on the industry average price-cost margin, a measure of industry monopolistic power, following Nickell (1996) and Aghion et al. (2005). The higher the industry average price-cost margin, the lower the product market competition. I follow Kogan et al. (2017) to measure the economic value of a patent, which is a dollar amount based on the patenting firm's stock return around a 3-day patent grant window. If the stock market is efficient, then all future incremental profits from a patent should be priced by the market immediately after the patent is granted.[4] Thus, the change in stock market value around the patent issuance date, after adjusting for systematic factors, noise in the market and the ex-ante expectation of the patent being granted, should reflect the economic value of the patent.[5]

My baseline results show that product market competition is negatively associated with the economic value of patents conditioning on firm and patent characteristics, firm, patent granting year, and patent class fixed effects. A 1% increase

[3]Case AT.39612 — Perindopril (Servier)

[4]The USPTO published patent grants every Tuesday. Kogan et al. (2017, p. 674) show significant increase in trading volume around patent issuance day, suggesting that patent issuance imparts valuable information to the market. As prices adjust efficiently to publicly available firm information in the US market (Fama, 1970, 1991), unexpected stock returns around patent grants should reflect, among other things, unanticipated patenting decisions.

[5]Details in Appendix B.

in the competition level is associated with a 1% decrease in patent value.

I conduct split-sample analysis according to industries' technological gap, defined as the average distance between industry leaders' and followers 'productivity. Industry technological leaders experience greater private returns from innovation compared to followers when there is a significant technological gap between leaders and laggards. However, competition affects leaders' patent value more negatively compared to laggards' when the technological gap is insufficient to establish a clear leader position.

Moreover, within an industry, firms with relatively high markups benefit more from innovation whereas high market shares do not guarantee higher innovation rent. Additionally, the inverse relationship between competition and patent value is more pronounced for high-value patents compared to low-value ones. However, the impact of competition on patents with high scientific significance is not different from its effect on those with lower scientific contributions.

The relationship between competition and patent value is endogenous. For instance, breakthrough innovation typically generates substantial economic gains and might alter the current market structure, introducing a reverse causality problem. To address this issue, I apply a quasi-experimental design, in which the patent issuance date is plausibly random relative to the announcement date of the competition-altering events within a short event window. I consider horizontal merger announcements in the same industry as the patenting firm as anti-competitive events and the US-China PNTR event in October 2000 as an event that intensifies competition.[6,7] I leverage on "stealth mergers" and mergers in concentrated markets with high product similarity to explore the heterogeneous

[6] A horizontal merger is defined to be a takeover deal in which both the acquirer and the target operate in the same industry. Industries are classified using the 4-digit SIC code in the primary analysis. Robustness tests are performed using the Text-based Network Industry Classification (TNIC).

[7] In October 2000, US Congress passed legislation that granted Permanent Normal Trade Relations (PNTR) to China. Before this event, China had been subject to low NTR tariff rates since 1980, which was a temporary status that required annual renewals. The granting of PNTR to China prevented future tariff increases on Chinese imports.

anti-competitive effects of horizontal M&As. All those events alter market expectations about product market competition but are unlikely to directly affect the value of patents filed by non-merging firms in the past through other channels.[8]

For identification, I do not require M&A announcements and the US-China PNTR event to be exogenous. Instead, the assumption is whether a patent is granted before or after these events is random within a $\pm$ 5-week event window. This assumption is plausible because 1) the United States Patent and Trademark Office (USPTO) decides the exact date of patent grants, not firms; 2) the long patenting process makes the exact issuance date hard to anticipate or manipulate; 3) I find no evidence suggesting patent characteristics are systematically different in the pre-merger sample and the post-merger sample.

In the quasi-experimental regression, patents' value is, on average, 4.8% higher immediately after horizontal merger announcements than before those announcements. In other words, an expected decrease in competition positively contributes to patent value. This positive effect is more pronounced when proposed mergers just avoid antitrust scrutiny or are in concentrated markets with high product similarity. On the contrary, patents issued immediately after the US-China PNTR event have lower economic value than those granted before.

I perform additional analysis and find no evidence that the significant differences in patent value between the pre-event group and the post-event group are driven by (1) non-random assignment into the two groups, (2) patents in the pre-event group driving horizontal M&A announcements, (3) patenting firms' stock market reactions due to horizontal merger announcements, (4) changes in firms' risk due to horizontal merger announcements, and (5) changes in other unobservable factors due to M&A announcements.

This paper is related to a growing literature that studies firm R&D investments

[8] Cunningham et al. (2021); Fathollahi et al. (2022); Hankir et al. (2011); Kepler et al. (2021) find evidence that horizontal takeovers are anti-competitive and Autor et al. (2020); Pierce and Schott (2016) suggest that tariff cuts or prevention of tariff increases intensify competition.

and patenting outputs.[9] One stream of the studies focuses on the determinants of firms' innovative activities. For example, studies have shown that institutional investors (Aghion et al., 2013; Brav et al., 2018; González-Uribe, 2020), managerial incentives (Ederer and Manso, 2013; Manso, 2011), knowledge spillovers (Arora et al., 2021; Bloom et al., 2013b; Matray, 2021), takeovers (threat) (Atanassov, 2013; Phillips and Zhdanov, 2012), the firm's boundary (Seru, 2014), covenant violations (Chava and Roberts, 2008), creditor rights (Mann, 2018), foreign competition (Autor et al., 2020; Bloom et al., 2016), analyst coverage (He and Tian, 2013), and stock liquidity (Fang et al., 2014) play important roles in corporate R&D investment and patenting outcomes.

Another stream of this literature examines how R&D investments affect firm stock returns and market value. For instance, firms' R&D capital stock is positively associated with their stock market value (Blundell et al., 1999). R&D investment efficiency (Hirshleifer et al., 2013) and R&D ability (Cohen et al., 2013) can predict a positive future stock return. A few studies investigate the interaction between competition and R&D investments in determining stock market value and stock returns. Greenhalgh and Rogers (2006) show that the positive correlation between market value and R&D is larger in sectors with a high sectoral profit persistence, an indicator of lack of competition. Gu (2016) studies the interaction between the R&D premium and market competition and shows that the R&D premium is higher in low market concentration portfolios than in high market concentration portfolios. The author argues that this difference is driven by the higher risk in R&D investments in more competitive industries than in less competitive industries.[10]

I complement prior studies by directly examining how product market competition affects the monetary gains shareholders can enjoy from successful R&D investments. Greenhalgh and Rogers (2006) and Gu (2016) analyze the unconditional effect of competition and R&D on asset prices. By contrast, I focus on the effect of competition on the value of R&D outputs conditional on successful

[9] See Ederer and Manso (2011); He and Tian (2020); Kerr and Nanda (2015) for detailed surveys.

[10] R&D premium is the significant positive return on the high-minus-low R&D-intensity portfolio.

R&D. Besides having a different focus; I aim to establish causality. To my knowledge, this is the first study that identifies the causal impact of competition on the value of innovation. The finding in this paper can be applied in the capital budgeting process when a firm is making investment decisions. In addition, the results also have implications for the valuation of patent acquisition and patents that are pledged as collateral for loans.

I also contribute to the competition and innovation literature. Both theoretical and empirical works find mixed evidence regarding the relationship between competition and R&D effort and outcomes. Aghion et al. (2005) find an inverted-U relationship between citation-weighted patents and industry competition using UK data, while Hashmi (2013) finds a negative relationship using US data.[11] Bloom et al. (2016) show that European firms increase their patenting activities when facing more intensive foreign competition from China. In contrast, Autor et al. (2020) find that foreign competition from China reduces US firms' R&D expenditure and patent production. Theoretical works have inconsistent implications about how competition changes firms' innovation rent (Aghion et al., 2005; Aghion and Howitt, 1992; Futia, 1980; Loury, 1979).[12] I complement previous work by providing direct causal evidence that competition reduces the economic value of innovation. Prior research focuses on the incidence of innovation rather than on the economic returns of innovation. I explore the monetary gain of innovation, a novel yet critical aspect of corporate innovation outcomes that provides motivation and a decision rule for firms' R&D investments.

This study also relates to the literature on the distinction between the number of patent citations and patent value. Trajtenberg (1990) suggests that patents vary significantly in their importance, and the number of citations a patent receives has become a popular measure of innovation outcome since then. Nonetheless, recent studies find that the relationship between the number of citations and patent economic value is not monotonic (Abrams et al., 2020; Kogan et al., 2017).

[11] The inverted-U relationship has competition on the horizontal axis.

[12] For example, Loury (1979) and Futia (1980) assume the intensity of competition has no impact on the incremental cashflow rewarded to successful innovation; Aghion and Howitt (1992) assume more competition reduces the expected present value of monopolistic rent from innovation; Aghion et al. (2005) assume competition could either increase or decrease the incremental benefit of innovation conditional industry characteristics.

My findings highlight the difference between citations received and patent value. I document that, all else equal, competition only affects the economic value of innovation but not the number of citations. This result suggests researchers need to be cautious when using citation-related measures to measure innovation outcomes.

This paper also has implications beyond academic research. Literature has documented a rise in the aggregate market power and industry concentration (De Loecker et al., 2020; Grullon et al., 2019) in the US economy over the past few decades. I show that this trend will have a positive impact on innovation rent. Furthermore, antitrust laws aim to limit monopoly power and foster competition. The government can, to some extent, influence the level of competition in the economy.[13] The results in this paper shed light on the potential effect of modifying industry competition levels.

3.2 Data

The sample consists of all patents issued to US public firms from 1976 to 2020 that have accounting information in Compustat and stock return information in the Center for Research in Security Prices (CRSP).[14] The USPTO online patent dataset provides detailed information on patent filing dates, grant dates, patent classifications, and citation history. The patent identification number in the USPTO data is matched to CRSP permanent security identification number using the patent-firm match data provided by Kogan et al. (2017) (KPSS hereafter).

The USPTO/KPSS data is then matched to Compustat to obtain firm accounting data. The industry classification is obtained from the Compustat historical segment data tape, and it is defined using the Standard Industrial Classification (SIC) 4-digit industry classification in which a firm generates the highest sales.

[13] For example, the New Deal policies under the Roosevelt administration abandoned antitrust prosecutions and approved cartel-led price-reflation efforts in the hope of reinvigorating the collapsing economy during the Great Depression (Phillips Sawyer, 2019).

[14] The analysis is restricted to patents granted after January 1, 1976, as the US Patent and Trademark Office (USPTO) publishes information on those patents in machine-readable form.

The sample excludes firms with missing or negative assets and/or sales and firms in industries classified as financial (SIC 6000-6999), utility (SIC 4900-4999), and miscellaneous (SIC codes ending in 9).[15] Firm characteristics are measured at the fiscal year-end prior to the patent issue date for each patent. Stock return data come from CRSP. The text-based industry classification (TNIC) data used in the robustness test is obtained from the Hoberg-Phillips Data Library.

I obtain Merger and Acquisition (M&A) data from 1979 to 2019 from Thomson Reuters SDC Platinum and restrict the sample to deals in which the acquirer and the target are both US public firms with identifiable GVKEYs. Only completed mergers in which the acquirer has 100% control over the target after the transaction are considered in the sample.[16] The M&A data contain announcement dates of each deal, which are used as the event dates in the identification strategy. M&A deals are classified into horizontal mergers and non-horizontal mergers. A horizontal merger is a deal in which both the acquirer and the target firms are in the same four-digit SIC industry, and non-horizontal mergers are all others. I also adopt alternative definitions of industries in the robustness section.

The outcome variable is the economic value of patents, obtained from the KPSS patent dataset, reflecting *"the present value of monopoly rents associated with that patent"* (Kogan et al., 2017). This measure is a dollar value based on the stock market cumulative abnormal return (CAR) around a 3-day patent issuance window. Furthermore, KPSS adjust the 3-day CAR for (1) the noise in return, (2) the under-expectation that arises from ex-ante probability assessment of a successful patent application, and (3) firms' market capitalization.

In addition, when multiple patents are issued to a firm on a day, the patent value of each patent equals the estimated dollar amount divided by the number of patents issued on a firm day. Acknowledging the noise introduced by this simple division, I use a sample of single patents issued on a firm day in the

[15]The reasons for excluding miscellaneous is that industry competition is not well defined for miscellaneous industries (Clarke, 1989).

[16]The qualitative results are the same when including withdrawn mergers in the sample. The sample excludes withdrawn cases as I can not identify whether those are partial acquisitions or complete acquisitions.

main analysis and show that the results are consistent when employing the entire patent sample. Appendix B discusses the estimation procedure and assumptions; a complete description of the data construction can be found in KPSS. I adopt KPSS's patent value estimates in the primary analysis and discuss and relax their assumptions in Appendix B.

I measure competition based on firms' price-cost margin following Aghion et al. (2005). The price-cost margin (or Lerner index) captures a firm's ability to charge a price above its marginal cost, reflecting its market power. The marginal cost is not directly observable; thus, the firms' markup is used to proxy for the price-cost margin. Specifically, the markup up is defined as a firm's operating profit over sales:

$$Markup_{i,t} = \frac{Operating\ Profit_{i,t}}{Sales_{i,t}}$$

where $Operating\ Profit_{i,t}$ is sales minus cost of goods sold, selling, general and administrative expense, and depreciation. The market competition is measured by one minus industry sale weighted average markup:

$$Competition_{s,t} = 1 - \sum_{i \in s} \frac{Sales_{i,t}}{\sum_{j \in s} Sales_{j,t}} Markup_{i,t}$$

where s is a subscript for the four-digit SIC industry.

This profit margin-based competition measure is more appropriate than the commonly used concentration measure, the Herfindahl-Hirschman Index (HHI), to address the research question. First, this markup-based measure is closely connected to firms' monopolistic pricing power. Second, the HHI has been criticized as it is sensitive to the definition of a product/geographical market (De Loecker et al., 2020) and such measure based on public firm data suffers from measurement error as it ignores all private firms in the market (Ali et al., 2008). In contrast, the profit margin-based competition measure incorporates information from private firms as the price is an equilibrium outcome of all firms competing in an industry. The HHI based on census data for manufacturing firms is used as an alternative explanatory variable as a robustness check.

Table 3.1 presents the summary statistics of 1.5 million unique patents, 41,872 unique patenting firm-years, and 10,208 unique industry-years observations for all patents issued to US public firms between 1976 to 2020. I also report the summary statistics for a sub-sample of patents that contains single patents issued on a firm day, a total of 225,917 patents, which is about 15% of the all patent sample. All dollar amounts are shown in 1996 US dollars using the inflation index.[17] Variable definitions are listed in Table 3.A1. All continuous variables are winsorized at 1% and 99% levels in each year to avoid extreme values biasing the result.

Panel A shows that the average patent value is $23.79 million, with a standard deviation of $46.11 million. An average patent makes/receives 19/18 citations and has a review time of 2.74 years.[18] Compared with the entire sample, the single patent sample consists of patents with a higher average economic value and receives more citations made by following patents.

In panel B, the average patenting firm has $4.2 billion in total assets, 19% book leverage, -4% profitability, a market-to-book ratio of 1.98, 9% R&D intensity, and an 8% markup. In the full sample, the median firm does not issue multiple patents on a firm day. The median of firms in the entire sample and the single patent sample are comparable, and the average firms in the two samples are also comparable based on accounting ratios like profitability, market-to-book, R&D intensity, leverage, and markup. Nonetheless, on average, patenting firms in the single patent sample are smaller in size. This discrepancy can be explained by the tendency for larger firms to file more patents than smaller firms; hence, large firms are more likely to get multiple patent grants on the same day.

[17] The inflation adjustment is based on Organization for Economic Co-operation and Development, Consumer Price Index: Total All Items for the United States [CPALTT01USA661S], retrieved from FRED, Federal Reserve Bank of St. Louis.

[18] I report the raw citation data as updated on January 18, 2022 by the USPTO. In the empirical analysis, I adjust the number of citations a patent makes and receives following Lerner and Seru (2022) to address to data truncation issue of using US patent data.

Panel C presents industry-level data, where industries are classified based on the 4-digit SIC code. The average competition is 0.9 with a standard deviation of 0.06, and the average HHI is 0.45 with a standard deviation of 0.27. On average, an industry has 15 firms. Similarly, the industry characteristics are comparable in the two samples, indicating that focusing on single-issuance patents does not bias our sample toward a specific set of industries.

3.3 Empirical methodology

I first run a baseline regression to explore the association between market competition and the economic value of patents, then apply a quasi-experimental design to establish the causal relationship. All variables are defined in Table 3.A1.

3.3.1 Baseline Regression

The baseline OLS regression is as follows:

$$Log(Patent\ Value_{i,j,s,t}) = \alpha_i + \beta Log(Competition_{s,t}) + \gamma X_{i,j,s,t} + \delta_t + \omega_j + \varepsilon_{i,j,s,t} \tag{3.1}$$

where i, j, s, t are subscripts for firm, patent, industry, and patent granting year; $Patent\ value_{i,j,s,t}$ is the economic value of patent j, I use the log patent value to avoid extreme patent values on the right tail driving the result. The results are quantitatively similar without log transformation. $Competition_{s,t}$ is the variable of interest, which is measured as one minus sales-weighted industry average markup, $X_{i,j,s,t}$ is a vector of controls including log firm size; leverage, profitability, market-to-book ratio, research intensity, the adjusted number of citations a patent makes, the adjusted number of citations a patent receives, the patent examination time, and the number of patents issued to a firm on a day.[19] α_i, δ_t, and ω_j are firm, patent granting year, and patent class fixed

[19]A second-degree polynomial terms in $Adj.\ Forward\ Citations$ and $Multiple\ Issuance(N)$ are included in all regressions that control for patent characteristics, but only the first-degree terms are reported for brevity.

effects, respectively. $\varepsilon_{i,j,s,t}$ is the error term. Standard errors are adjusted for non-independence of observations within firms and granting years.

A significant estimate of β in baseline regression suggests some correlation between product market competition and the economic value of patents but does not indicate causality. Endogeneity problems could exist. For example, unobservable industry trends might drive the industry to be more competitive and reduce the economic return of innovation simultaneously, introducing omitted variable bias. A breakthrough innovation with significant economic value might also alter the product market competition intensity, leading to a reverse causality problem.

3.3.2 Quasi-natural Experiment

To establish causality, I propose a quasi-natural experiment to compare the value of patents that are granted immediately before versus after competition-altering events. The identification relies on the expected competition-altering effects of selected events, and the random assignment of patents into the two groups, which will be discussed in more detail shortly.

In the baseline analysis, I consider horizontal M&A announcements as events that are expected to lessen market competition, and the 4-digit SIC industry to be the event industry. The figure below illustrates the details of the empirical design. M&A announcement dates occur at $t = 0$. Only patents issued to the event industry within the [-35, 35] day window relative to the merger announcement dates are included.[20] Next, patents issued to either the acquirer or the target firm are excluded as the merging firm might experience some fundamental changes that could affect their patent value. Furthermore, patents issued within the [-7, 7] event window are excluded to address the concern that M&A activities might affect peer stock returns (Fee and Thomas, 2004; Shahrur, 2005). As a result,

[20]All patents are filed at least 35 days prior to their issuance dates in the regression sample.

rivals' patents issued within the [-35, -8] window fall in the pre-merger group and within the [8, 35] window fall in the post-merger group.[21]

Timeline of the quasi-natural experimental design

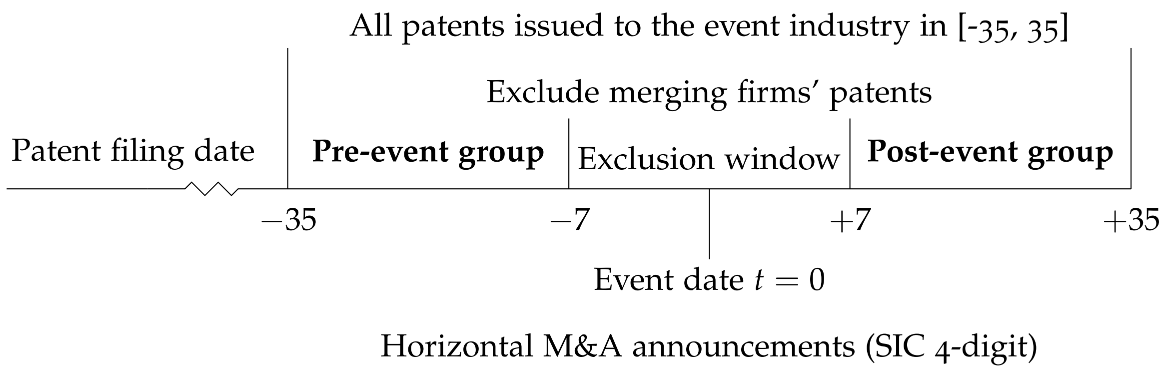

The identification strategy relies on the expected anti-competitive effect of horizontal merger announcements and the randomness of patent issuance dates within a short window. Horizontal M&As combine two firms within the same industry, leading to a more concentrated and less competitive market (see, for example Cunningham et al., 2021; Fathollahi et al., 2022; Hankir et al., 2011; Kepler et al., 2021).[22] Although M&A activities can affect firms' innovation efforts and outcomes and vice versa (Bena and Li, 2014; Phillips and Zhdanov, 2012; Seru, 2014), the identification strategy only requires that whether a patent is granted before or after merger announcements within a 2-month window is random.

The random assignment assumption is plausible for several reasons. First, the USPTO, instead of patenting firms, decides whether and when to grant a patent, with the granting decision purely depending on the novelty of the application.

[21] Because USPTO announces patent grants every Tuesday, the $\pm$5-week inclusion window is used instead of the $\pm$ 1-month to ensure that four patent issuance days are included in both the pre-event and post-event sample regardless of the dates of merger announcements. The results are robust to alternative event windows.

[22] Not all mergers have anti-competitive intentions. For instance, mergers in declining industries may primarily seek consolidation for resource reallocation (Anand and Singh, 1997; Dutz, 1989). Mergers might also come in waves, driven by economic and industry shocks (Gort, 1969; Mitchell and Mulherinb, 1996) or market timing (Rhodes-Kropf and Viswanathan, 2004; Shleifer and Vishny, 2003). Despite various ex-ante motivations, the anticipated ex-post impact may still be anti-competitive due to the more concentrated market.

Hence, firm and industry activities should not affect the granting and timing decision. In addition, the average patent examination time, i.e. the interval between the patent application and grant date, is 2.7 years, with a standard deviation of 1.6 years. Therefore, it would be difficult to anticipate whether a patent will be issued in the pre-event or post-event window. Furthermore, a merger decision based on rivals' patent grants within this pre-event window is unlikely because M&A negotiation is costly and typically takes more than 35 days, my window length.[23] Furthermore, I perform balance tests to verify the random assignment assumption in the empirical section and find no evidence of systematic differences between the characteristics of patents and patenting firms in these two groups or change in M&A propensity immediately after peer patent grants.

I then compare the value of patents in the pre and post-event group conditioning on patent and firm characteristics, and firm, year, patent class fixed effects:

$$Log(Patent\ Value_{i,j,s,t}) = \alpha_i + \beta Post\ Merger_{s,t} + \gamma X_{i,j,s,t} + \delta_t + \omega_j + \varepsilon_{i,j,s,t} \quad (3.2)$$

where $Post\ merger_{j,s,t}$ is a dummy that equals one if patents are issued in the [8, 35] event window and 0 if patents are issued in the [-35, -8] window. The same set of controls $X_{i,t}$ and fixed effects as in the baseline regression are included. The coefficient of interest, β, intends to capture the change in patent value due to horizontal merger announcements.

[23]The identification strategy does not require exogeneity of the M&As with respect to all industry innovation. Two firms may merge as a result of their rival's innovation. Nonetheless, as long as patents in the pre-announcement group do not drive the takeover, there is no reverse causality concern. I argue that this reverse causality is unlikely because M&A negotiation processes are so costly and time-consuming that a planned M&A conditional on rivals' patenting outcomes or announcing a M&A as an immediate response to innovation are both unlikely.

3.4 Results

3.4.1 Graphical Evidence

Figure 3.1 displays the unconditional relationship between patent value and the number of citations received (top panel), patent value and product market competition (middle panel), and the number of citations received and product market competition for all patents in the single patent sub-sample. Each figure displays binned scattered dots representing the average y-value for each x-value.

First, the top panel reveals that the relationship between patent citation and patent economic value is not monotonically increasing. In line with Abrams et al. (2020), there is a U-shaped relation between patent citation and economic value. This pattern implies that patent value decreases with the number of citations in the low citation region, but increases with the number of citations in the high citation region. Furthermore, the figure demonstrates a highly skewed distribution of the number of patent citations, with only a small proportion of patents receiving a substantial volume of citations.

Next, the middle panel depicts a downward-sloping relationship between patent value and the intensity of product market competition faced by the patenting firm. This graph indicates that, on average, patent value tends to be higher in industries with a low level of product market competition. Nonetheless, the bottom panel suggests that the number of citations does not monotonically decrease in competition. Instead, there is an inverted-U-shaped relationship between patent value and patent citations. These findings collectively highlight the significance of distinguishing between patent economic value and scientific value (as often proxied by the number of citations a patent received) when examining the relationship between competition and innovation.

Figure 3.2 provides visual evidence of the negative relationship between competition and patent value over time. The top panel depicts the average patent economic value and the average product market competition over the sample period. The average patent value was relatively low from 1976 to 1990, during which the average competition remained at a relatively high level. The patent economic value spiked around 2000, 2008 and 2020, shortly after the average competition level reached local minimum levels.

The bottom panel plots the average patent value over time for industries in the top and bottom terciles of competition. Patents granted to firms in low-competition industries have a higher average value than those in high-competition industries throughout the sample period. These visual evidences imply a negative relationship between market competition and innovation rent.

It is likely that patent value spikes when the aggregate stock market value is high and remains at a low level when the stock market crashes. I plot the average patent value against the total US stock market value in Figure 3.3. Although both aggregate stock market value and average patent value have increased substantially from 1975 to 2020, they are not always positively correlated. This figure suggests that the fluctuation in the economic value of patents is not purely driven by changes in stock market value.

3.4.2 Baseline Regression

Table 3.2 presents the baseline results on competition and patent value using the specification in Equation 3.1. In Columns (1) to (3), estimation is based on the single patents sub-sample. Without conditioning on firm and patent characteristics, Column (1) suggests that a one percent increase in competition is associated with a 2.3 percent decrease in patent value. As indicated by the high adjusted R-squared, firm, year, patent class fixed effects alone explain almost 85% of the variation in patent value.[24] After controlling for firm and patent characteristics, the economic magnitude of the estimated coefficient of interest is more than halved. The coefficient estimate in Column (3) indicates that patent value is approximated 1% higher in industries with 1% lower competition intensity.

[24] Table 3.A2 shows how the inclusion of different fixed effect combinations affects the estimation results. According to the R^2 reported in that table, product market competition alone explains 3% of the variation in patent value, and firm fixed effects explain almost 80% of the variation in patent value.

In Column (4), the estimation includes all sample patents. To alleviate noise in patent value estimates caused by multiple issuances on a day, I control for the quadratic form of the number of patents issued to a firm on a day. The result based on the entire sample shows a stronger economic magnitude but is noisier compared to Columns (1) to (3). In sum, the results in Table 3.2 suggest that high pricing power allows firms to enjoy a greater economic benefit of innovation.

Aghion et al. (2005) provide evidence of a non-monotonic relationship between competition and innovation (measured by industry average citation-weighted patent counts), suggesting that the relationship between competition and patent value may also exhibit nonlinearity. The top panel in Table 3.A3 reports the regression results that include the square of market competition (without log transformation) to allow for the nonlinearity in the competition-patent value relationship. The coefficient estimates suggest that patent value is higher in high competition industry only when competition intensity is below 0.725 (global maximum at $x = -7.438/(2 \times -5.127) = 0.725$) and that patent value decreases in competition when competition intensity is above 0.725. However, it is worth noting that only a small proportion (3%) of the observations in the regression sample have competition intensities below 0.725, and the one percentile patent faces a competition intensity of 0.7. As a result, the positive correlation in the low-competition region is very weak, if exists.

To further investigate the nonlinearity in the competition-patent value relationship, I examine the relationship across quartiles of competition intensity within the single patent sample. Panel B of Table 3.A3 presents the coefficient estimates of the regressions. Across all subsamples, there is still a statistically significant negative correlation between patent value and competition. Although the economic magnitude of the coefficient is smaller in the bottom competition quantile (Column (1)) compared to other subsamples, the coefficient is still negative and significant, which is consistent with our prior findings.

3.4.3 Cross-sectional Results

I now investigate the heterogeneous effects of competition on patents' value. First, I examine how this relationship varies with the industry technology gap and firms' position on the industry technological spectrum. According to Aghion et al. (2005), industries can be classified into neck-to-neck (NN) industries, in which all firms have similar technologies, and leader-laggard (LL) industries, in which technology leaders have more advanced technology than laggards. High competition leads to higher economic value of patents than in low competition industries for firms in an NN industry. On the contrary, high competition leads to higher economic value of patents for leader firms but lower economic value of patents for laggard firms in an LL industry (details see Appendix B).

I use firms' total factor productivity (TFP), defined as the exponential residual of the estimated Cobb-Douglas production function, to identify the technological hierarchy within an industry. A high value of TFP means that the firm is closer to the technology frontier compared to its peers. The average distance between firms' TFP and the highest TFP within an industry is used to proxy the industry technology gap following Aghion et al. (2005).

Industries are split into terciles based on the industry technology gap, where the bottom tercile represents NN industries, and the top tercile represents LL industries. Firms with the highest 10% TFP in each industry are classified as technology leaders. I extend the baseline regression by including an indicator for industry technology leaders and its interaction term with the competition:

$$\begin{aligned} Patent\ value_{i,j,s,t} = &\alpha_i + \beta_1 Competition_{s,t} + \beta_2 Leader_{i,s,t} + \beta_3 Competition_{s,t} \times Leader_{i,s,t} \\ &\gamma X_{i,j,s,t} + \delta_t + \omega_j + \varepsilon_{i,j,s,t} \end{aligned} \tag{3.3}$$

where $Leader_{i,t}$ is an indicator for industry technology leaders, $Competition_{s,t} \times Leader_{i,j,t}$ is an interaction term between industry leaders and competition. For

NN industries (bottom technology gap tercile), β_1 is expected to be positive and significant. For LL industries (top technology gap tercile), the leader and laggard firms are expected to have different slope coefficients. Thus, β_1 is expected to be negative (laggard firms), and $\beta_1 + \beta_3$ is expected to be positive (leader firms).

Column (1) in Table 3.3 suggests that competition has a negative impact on the value of innovation for both leaders and laggards, which is consistent with previous results. In addition, patents issued to firms at the technological frontier achieve 0.6 percent higher economic returns than those issued to technology laggards. Nevertheless, competition does not have a statistically significant differential impact on leaders. Columns (2) to (4) present the split sample results. Consistently with the baseline results, the coefficient estimates are negative and statistically significant in all columns. Industry technology leaders enjoy higher private economic returns from innovation only when the industry technology gap is sufficiently large, while technology leaders in neck-neck industries do not receive extra benefits from their innovation. This finding can be interpreted as firms are clustered at a relatively small range of the technology spectrum in the NN industry; thus, the leader position does not necessarily mean a firm has substantial technological advantage to enjoy a higher economic gain from innovation.

Concerning the coefficient estimates of the interaction term of competition and industry leader, I find that industry leaders are more negatively affected by competition in the NN industries but less negatively (albeit insignificant) affected by competition in LL industries. In NN industries, competition would make it harder for technology leaders to widen the industry technological gap so that the industry can transit from NN to LL, thus, detrimental to patent value. In contrast, technology leaders in LL industries would not be more adversely affect by competition owing to its more prominent technological advantages.

The findings in this table do not support the hypothesis made in Aghion et al. (2005). One explanation is that industries in the estimation sample cluster at a small range in the technological gap spectrum. However, an unreported histogram shows that the distribution of the technological gap resembles a normal distribution. Thus, these results are not driven by the lack of variation in the industry technological gap. Another possible explanation is that Aghion et al. (2005) model oligopoly industries, while the estimation sample is based on US industries whose majority are not oligopolies.

Table 3.4 reports additional cross-sectional analysis based on Equation 3.3, with the market technology leader be replaced by other characteristics of interest. Columns (1) and (2) test how within industry variation in market power and market share and their interaction with competition affect patent value, respectively. If firms in industries with relatively high monopolistic power gain more from innovation, we would also expect firms with high market power within an industry to earn a higher innovation rent. The results indicate that, under the same industry competition intensity, firms with above-median markups have higher patent values, while above-median market shares do not correspond to a higher patent value. Nonetheless, there is no evidence that firms within the same industry would be affected by competition differently.

In Columns (3) and (4), I test whether competition affects patent value differently for economical or scientifically more valuable patents. Patent value is more sensitive to competition when it is economically more valuable, but not when it is scientifically more meaningful. Column (3) suggests that a one per cent increase in competition is associated with a 1.48 per cent decrease in patent value for those patents with above-median economic value, which magnitude is twice as negative as the coefficient estimate of the below-median economic value patents (-0.75). This finding is consistent with prior studies that find core patents are more closely related to firm value, and the negative impact of obsolescence is more pronounced for those core patents (Akcigit et al., 2016; Ma, 2021). In contrast, the relationship between competition and patent value does not exhibit a statistically significant difference for patents with above-median citation counts compared to those with below-median citations received.

3.4.4 Identification: Quasi-natural Experiment

Quasi-experiment regression analyses are conducted for both the single patent sample and the entire sample. Summary statistics of these two samples are reported in Table 3.A4. On average, the quasi-experiment sample consists of more valuable patents, larger patenting firms, and industries with lower competition intensity compared to the baseline sample. A more detailed comparison of sample distribution across different industry classifications, patent classifications, and time between the baseline sample and quasi-natural experiment sample can be found in Appendix C.

Table 3.5 provides results of the quasi-natural experimental design specified in Equation 3.2. Consistent with the baseline results, the estimates of the coefficient of interest are positive and statistically significant across all columns. Column (3) indicates that patent economic value increased by approximately 4.8 per cent immediately after horizontal merger announcements. The economic magnitude of the coefficient of interest in the single patent sample and the full sample are also comparable, although the full sample has larger standard errors.

Table 3.6 reports balance tests to validate the identifying assumption that patents are randomly assigned to the pre-event and post-event samples. Panel A presents the summary statistics for the pre-event and post-event samples. Overall, there is no statistically significant difference in patent and firm characteristics between the two samples.

Since the patent value estimates rely on stock price movement around patent issuance dates, stock market fluctuations, which might be affected by horizontal merger announcements, could also drive the observed results. Hence, I also compare the stock market capitalization for all sample firms, and the 3-day CARs for all rivals of the merging firms. The difference in market capitalization between the two groups is indistinguishable from zero. To further alleviate the concern that increases in patenting firms' market value due to merger announcements drive the results, I control for the market capitalization used in the patent value estimates as a robustness check; results are consistent and will be discussed in the next section.

The 3-day CAR calculation is consistent with the method of Kogan et al. (2017).[25] The first row shows the average 3-day CARs for each day within the [-35, -8] and [8, 35] announcement window for all rivals. The second row shows the average 3-day CARs on each Tuesday only (the USPTO issues patents every Tuesday). Although the post-event group has significantly lower 3-day CARs in both rows, this difference would drive the results in Table 3.5 to the opposite direction. Therefore, observed increases in patent values after merger announcements are unlikely to be caused by stock market reactions to the announcements.

In the Panel B of Table 3.6, I regress patent characteristics on the *Post-merger* dummy conditioning on all control variables and fixed effects included in Table 3.5. If the random assignment assumption is valid, patent characteristics unrelated to market competition should not differ significantly between the two groups. Consistent with the assumption, Columns (1) to (4) suggest that the average number of citations made and received, the patent review time, and the number of multiple grants in the pre-event and post-event groups are not statistically different. This balance test includes all patents issued within the pre-event and post-event windows instead of only single patents because patent characteristics under examination are not be affected by multiple issuances on a firm day. In unreported results, I find the results in Columns (1) to (3) still hold for the single patent sample.

Not all horizontal M&As have the same extent of anti-competitive effect. Next, I examine the heterogeneous effects of horizontal merger announcement events on rival patent value, specifically, whether deals expected to be more anti-competitive would lead to a stronger post-announcement on patent value improvement. First, I focus on "stealth mergers" that fall just below regulatory

[25]The idiosyncratic return is defined as the stock's return minus the market return: $\tilde{r_{i,t}} = r_{i,t} - r_{m,t}$. The daily idiosyncratic return is then compounded to get the 3-day CAR: $CAR_{i,[t,t+2]} = (1 + \tilde{r}_{i,t}) \times (1 + \tilde{r}_{i,t+1}) \times (1 + \tilde{r}_{i,t+2}) - 1$.

bodies' antitrust scrutiny threshold. In 1976, the Hart-Scott-Rodino (HSR) Antitrust Improvements Act specified a premerger notification threshold that mandates the merging firms to notify the Department of Justice (DOJ) and Federal Trade Commission (FTC) of their merger intent if the deal exceeds this specified value, which is renewed annually after 2000. Kepler et al. (2021) show that some firms deliberately manipulate the merger deal value to avoid antitrust scrutiny, ultimately reducing competition. Furthermore, Cunningham et al. (2021) find that acquisitions that fall below the antitrust scrutiny threshold are more likely to preempt competition by terminating the target firm's research projects, known as "kill acquisitions".

The results for stealth mergers are reported in Columns (1) and (2) of Table 3.7. Patents issued immediately after the announcement of just-below-threshold deals exhibit an 8 per cent higher increase in their value compared to patents issued after deals just above the threshold. It is important to note that while the coefficient estimate of $PostxStealthMerger$ in the single patent sample is not statistically significant at the 10% level, the joint significance of the coefficients on $Post$ and $PostxStealthMerger$ is significant at 5% level ($F = 19.52$, $Prob > F = 0.048$). This result suggests that patent value is 9.10% higher immediately after stealth merger announcements, doubling the average effect of 4.8% reported in Table 3.5. In the all patent sample, the results show that patent value following stealth mergers is 12 per cent higher compared to non-stealth mergers, which is statistically significant at the 5% level. The joint significance of the coefficients on $Post$ and $PostxStealthMerger$ is significant at 5% level ($F = 20.31$, $F = 0.046$). This finding further supports the notion that stealth mergers have a more pronounced impact on increasing patent value.

The second set of heterogenous tests focuses on horizontal mergers in industries characterized by high concentration and industry product similarity. Both theoretical and empirical evidence suggests that those deals are expected to be more effective in enhancing the market power of incumbent firms (Farrell and Shapiro, 1990; Fathollahi et al., 2022; Perry and Porter, 1985). I utilize the 10-K text-based measure of industry product similarity (IPS) constructed by Fathollahi et al. (2022) and the census-based industry concentration measure compiled by Keil (2017) to identify high concentration and high IPS industries. In order to quantify the differential impact of mergers on patent value in industries characterized by high concentration and high IPS, I regress patent value on a post-merger dummy variable, a dummy variable indicating high (above-median) concentration and IPS industries, and their interaction term.

Column (3) and (4) of Table 3.7 suggest that horizontal merger announcements in high-concentration and high-IPS industries lead to an 11% increase in peer firms' patent value in the single patent sample. This effect is 8.3% higher than the average effect observed in M&As in the remaining benchmark industries. Moreover, patents issued to high-concentration and high-IPS industries do not show a statistically significant difference in economic value compared to other patents without merger announcements. These findings suggest that horizontal merger announcements that are more likely to have substantial anti-competitive effects result in greater improvements in patent value.

Table 3.8 replicates the quasi-natural experiment results using alternative events that are expected to impact competition differently. First, I use non-horizontal merger announcements, which involve the combination of two firms from different industries, as placebo events to rule out this alternative explanation that M&A announcements might deliver other information that causes patents in the post-event groups to have a higher economic value. For example, a M&A can create synergy for the combined entity, which can further spill over to their industry peers or signal an increased likelihood of industry peers being future targets. In either case, the patent value might increase due to these confounding effects of M&A rather than the expected change in product market competition. While non-horizontal mergers could still generate substantial synergy or signal future targets, their impact on market competitiveness is less pronounced than a horizontal merger. By employing non-horizontal mergers as placebo events, any observed effects on patent value due to horizontal merger announcements can plausibly be attributed to the change in competition rather than other confounding factors.

Panel A reports the regression results by using the acquirer's industry and the target's industry as the event industry. Columns (1) and (2) demonstrate that while the coefficient estimates of interest remain positive, they are not statistically significant at any conventional level. Similarly, Columns (3) and (4) reveal no significant difference in patent value before and after non-horizontal merger announcements in target industries. This lack of significance helps alleviate concerns that changes in patent value following merger deals are driven by factors other than alterations in product market competition.

Panel B displays results using horizontal merger withdrawal announcements and the grant of US-China PNTR as two alternative events that are expected to intensify product market competition. In Columns (1) and (2), I focus on a subsample of withdrawn horizontal mergers and use the withdrawal announcement as event date zero. The coefficient estimates of the *Post* dummy are negative, despite being insignificant. Using merger withdrawals as "pro-competition" events might have limitations that introduce noise and lead to this insignificant result. First, the withdrawal sample might include some partial acquisition; thus, the reversed effect would be milder than complete acquisition if the deal is partial. Second, some withdrawals are due to a change of the acquirer or the target; withdrawing one deal could mean initiating another. Thus, the anti-competitive effect is not reversed by the withdrawal announcement. Last, there might be the anticipation effect of deal withdrawal. Despite noises in the estimation, the negative sign still conveys that the market adjusts the patent value downwards following merger withdrawal announcements.

In Columns (3) and (4), I consider the US grant of Permanent Normal Trade Relations (PNTR) to China in October 2000 as a positive shock to market competition for US firms following Pierce and Schott (2016). PNTR increases the competition faced by US firms as this agreement eliminates the potential increase in tariffs on Chinese imports from annual renewals of the tariff rate. The results correctly indicate that patents issued within the [+8, +35] day window relative to the agreement had been signed have lower economic value compared to those granted before the event. The finding is consistent with Autor et al. (2020) who find that rising import exposure from China increases competitive pressure and lowers US firm R&D spending and patent production in industries with higher import competition. These alternative events further confirm that changes in product market competition affect the economic value of patents.

3.4.5 Robustness

3.4.5.1 Robustness to Baseline Regression

In Table 3.9 and Table 3.10, I examine the robustness of the baseline and quasi-experimental regressions. Panel A in Table 3.9 presents the baseline result using alternative measures of competition: (1) one minus the equally weighted 4-digit industry average markup, (2) one minus the sales-weighted 3-digit SIC industry average markup, and (3) the 4-digit SIC Herfindahl–Hirschman Index (HHI) based on census data. The results using markup-based competition measures are consistent with the baseline results. There is also a positive correlation between industry concentration and patent value; however, this relationship is not statistically significant. The insignificant coefficient might be due to the sticky nature of the concentration measure, as the US census concentration measure is only updated every five years. Therefore, it lacks variation to identify a significant relationship.

Panel B in Table 3.9 presents the baseline regression results with the adjusted number of citations as the dependent variables, including (1) the adjusted number of total citations received, (2) the adjusted number of self-citations received, and (3) the adjusted number of nonself-citations received.[26] The coefficient estimates indicate that the number of citations received by a patent does not vary significantly with product market competition.

[26]A citation is classified as self-citation if the two patents in the citing-cited patent pair have a common assignee.

Panel C in Table 3.9 shows firm-level baseline regression results with three proxies for firm innovation: (1) R&D intensity, (2) aggregate patent economic value in a firm-year, and (3) aggregate patent scientific value in a firm-year. Although R&D intensity is not a direct measure of patent value, it reflects firms' willingness to invest in innovation. Conditioning on all firm controls and firm and year fixed effects, a high level of competition is negatively associated with firms' R&D intensity and the aggregate economic value of their patent portfolios that are issued in the following year. Nonetheless, there is no evidence suggesting competition significantly correlates with the scientific importance of patents. The insignificant result in Column (3) alleviates the concern that unobservables simultaneously drive competition and firm innovation activities and highlights the importance of distinguishing patent economic and scientific value.

Panel D in Table 3.9 reports baseline results with additional control variables or fixed effects. The market-based estimates of patent value are sensitive to stock market attributes that are not related to the economic value of patents. Hence, in Column (1), I include stock market characteristics that could affect the patent value estimates: (1) the patenting firms' market capitalization (million $) on the day before patent issuance,[27] (2) firm beta on patent issuance dates, and (3) firm idiosyncratic volatility.[28] The result in Column (1) is similar to the baseline results.

Furthermore, prior studies on innovation acknowledge that the economic importance of innovations varies significantly and suggest some characteristics of innovations, such as "generality" and "originality" of a patent (e.g. Hall et al., 2005). Panel D Column (2) presents the estimation results conditional on patent generality and originality, and the coefficient estimate of interest remains negative and significant. Last, some unobservable technological trends might simultaneously affect the value of innovation and industry competition. Thus, I include patent class-year fixed effects in the baseline regression to test the within technology class and year variation in patent value caused by competition. Patents issued

[27] This variable is the market capitalization that enters patent value estimation formula.

[28] Beta and idiosyncratic volatility are estimated based on the market model using daily stock return and 252 trading day estimation window with a minimum number of observations of 126.

to firms in industries with lower competition intensity enjoy a higher economic return than patents issued in the same year and patent class but to firms facing more intensive competition.

3.4.5.2 Robustness to Identification Regression

In Table 3.10, I examine the robustness of the quasi-natural experiment results. The primary analyses focus on patents issued in the [-35, -8] and [8, 35] windows to ensure randomness in the post-merger group assignment, and the return does not contain information from the merger event. One might suspect whether the result is sensitive to the event window's choice. Thus, I replicate the quasi-experiment results with narrower or wider event windows. Panel A shows that the main results are little changed if the event inclusion window is defined as 21, 28, 42, or 56 days around the merger announcements.

Next, I include stock market controls, patent controls, M&A deal controls, and patent class-year fixed effects in the quasi-natural experiment regression. Panel B presents the results with these additional controls. The main finding, which indicates an increase in patent value by approximately 4-5 percent following horizontal merger announcements, holds across all specifications. This consistency reinforces the robustness of the observed effect.

Table 3.A5 reports more robustness tests for the identification strategy. One potential concern arises from consecutive merger events announced within a 70-day gap. In such cases, a patent may be included in multiple events, potentially biasing the results. To address this issue, Columns (1) and (2) of the table present regression results based on subsamples of patents exclusively included in either the first or last consecutive announcement. The regression results continue to support the main finding of this paper, suggesting a positive and significant effect of horizontal merger announcements on patent value.

Columns (3) and (4) include announcements of horizontal merger deals that were eventually withdrawn. It is worth noting that those withdrawn deals might contain partial mergers and mergers that are ex-ante unlikely to go through, thus, biasing the results downwards. This robustness test finds a positive change in patent value after horizontal merger announcements.

To further validate the results, industry classifications are defined according to the Text-based Network Industry Classifications (TNIC) (Hoberg and Phillips, 2010, 2016). The TNIC provides a similarity score ranging from zero to one for each firm-peer pair, reflecting their proximity in product markets, with a score of one indicating the highest similarity. In this analysis, mergers are classified as high-similarity if the bidder and the target have an above-median similarity score (0.05). Rival firms are defined as those with an above-median similarity score with both the bidder and the target. Columns (5) and (6) present the regression results based on these TNIC industry classifications. The results remain consistent and robust when using these TNIC-based industry classifications. The findings support the main conclusion that horizontal merger announcements have a positive and significant impact on patent value. These additional robustness tests confirm the validity of the observed relationship positive impact of anti-competitive events on patent value.

3.4.5.3 Innovation and M&A

There may remain concerns about the identification strategy, especially reverse causality. Bena and Li (2014) document that firms' R&D intensity and patent portfolio would affect the likelihood of being a target, and firm pairs' technological overlap can affect the merger transaction incidence. Furthermore, Chen et al. (2020) show that a firm's propensity to acquire another firm significantly increases if its competitors were awarded the *R&D 100 Award* in the previous three years. The identification assumption will be flawed if two firms announce mergers as a strategic response to rivals' patents which are issued within the pre-announcement window. Acknowledging that rivals' innovation can affect the probability of a merger deal, however, I argue that merger announcements are unlikely to be caused by patents issued in days t-35 to t-8.

This reverse causality is unlikely for several reasons. Firstly, M&A is a lengthy process involving multiple stages, such as target screening, valuation, engaging investment bankers, due diligence, and negotiation. These activities typically take more than 35 days before the announcement of the deal to the public (Eckbo, 2014).[29] Therefore, M&A events are unlikely to be influenced by patenting events within the [-35, -8] window. Additionally, if a patent is important enough to trigger a merger deal, it is expected to hold substantial economic value. As my findings suggest an increase in patent value after horizontal merger announcements, valuable patents in the pre-merger group would introduce attenuation bias to my results.

3.5 Conclusion

In conclusion, this study addresses the gap in our understanding of how product market competition influences the economic value of innovation. By examining the causal impact of competition on the value of patents, this research provides valuable insights for firms' management and policymakers in making informed decisions to promote economic growth.

The findings of this study indicate that competition has a negative effect on the economic value of innovation. The analysis reveals that patents issued immediately after horizontal merger announcements have a significantly higher market value compared to patents issued before such announcements, while other patent characteristics show no significant difference. This positive effect is more pronounced for deals expected to be more anti-competitive. On the other hand, patents' value decreases significantly after the US-China Permanent Normal Trade Relations (PNTR) agreement, which is expected to intensify competition.

These findings have important implications for firms' investment decisions, the valuation of patent acquisitions, and the use of patents as collateral for loans. By

[29]For example, the prospectus typically takes one to three months to write, and the target shareholder vote is often scheduled three to six months after the initial merger proposal is signed.

understanding the impact of competition on the value of innovation, firms can better assess the potential returns and risks associated with their R&D investments.

Policymakers can also consider these insights to design effective policies that balance innovation and competition. While antitrust laws intend to promote competition, they may disincentivize firm from investing in innovation because the prospect of innovation rents is low. Therefore, there should be an optimal balance, where the market is competitive enough to encourage efficiency and welfare in the present, while leaving enough monopolistic rent to incentivize firms to invest in research and development for the future. To strike this balance, policymakers can encourage innovation by providing incentives, such as research grants and tax credit, while also prevent inefficient anti-competitive practices.

It is important to acknowledge the limitations of this study. The analysis focuses on patented innovation, and the findings may not fully capture the value of non-patented forms of innovation. Future research could explore this limitation and further investigate the mechanisms through which competition affects the value of innovation.

Table 3.1: Summary Statistics

The estimation sample consists of 1,505,856 patents issued to US public firms over the period 1976 through 2020, among which 225,917 are single patents granted to a firm on a day. Panel A to C report the patent-level, firm-year, and industry-level characteristics of the all patents sample and the single patents sub-sample. The table reports means, standard deviations, and medians of the main variables employed in the empirical analysis. Variables in levels are measured in 1996 dollars. All variables are defined in Table 3.A1.

	All Patents Sample				Single Patents Sample			
Variable	N	Mean	SD	Median	N	Mean	SD	Median
Panel A: Patent-level								
Patent Value ($mln)	1,505,856	23.79	46.11	10.33	225,917	29.12	64.43	9.16
Backward Citations	1,505,856	19.50	92.89	7.00	225,917	21.11	51.97	9.00
Forward Citations	1,505,856	18.26	54.92	5.00	225,917	27.14	70.59	9.00
Examination Time (Year)	1,505,856	2.74	1.62	2.34	225,917	2.54	1.52	2.14
Panel B: Firm-level								
Assets ($bln)	41,872	4.18	19.35	0.33	40,922	3.11	12.78	0.31
Sales ($bln)	41,872	3.53	13.61	0.32	40,922	2.92	11.91	0.31
ROA	41,872	−0.04	0.31	0.04	40,922	−0.04	0.31	0.04
Market Cap. ($bln)	41,872	4.67	21.23	0.37	40,922	3.42	13.24	0.36
Market-to-Book	41,872	1.98	2.43	1.27	40,922	1.98	2.44	1.26
R&D Intensity	41,872	0.09	0.15	0.04	40,922	0.09	0.15	0.04
Book Leverage	41,872	0.19	0.16	0.17	40,922	0.19	0.16	0.17
Markup	41,872	0.08	0.08	0.07	40,922	0.08	0.08	0.07
Multiple Issuance (N)	41,872	1.59	4.52	1.00	40,922	1.00	0.00	1.00
Panel C: Industry-level								
HHI	10,208	0.45	0.27	0.38	10,122	0.44	0.27	0.37
Competition	10,208	0.90	0.06	0.91	10,122	0.90	0.06	0.91
#Firms	10,208	14.78	28.24	7.00	10,122	14.86	28.34	7.00
Technology Gap	8,187	0.47	0.11	0.49	8,144	0.47	0.11	0.49

Table 3.2: Product Market Competition and Patent Value

This table presents the results of regressing log patent value on log product market competition, firm controls, patent controls, and firm, year and patent class fixed effects. In Columns (1)-(3), estimation is based on single patent grants on a firm-day. Column (4) includes all patent grants in the sample. Firm and industry characteristics are measured at the end of fiscal years prior to patent grants. A second-degree polynomial terms in $Adj.ForwardCitations$ and $MultipleIssuance(N)$ are included in the regressions, but only the first term is shown for brevity. All variables are defined in Table 3.A1. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	Single Patent Grants			All Patent Grants
	(1)	(2)	(3)	(4)
Log(Competition)	-2.340***	-0.945***	-0.946***	-1.421***
	(0.340)	(0.183)	(0.183)	(0.454)
Log(Total Assets)		0.641***	0.642***	0.414***
		(0.012)	(0.012)	(0.044)
Net Income (Loss)		0.180***	0.181***	0.257**
		(0.051)	(0.051)	(0.126)
Market-to-Book		0.159***	0.159***	0.224***
		(0.012)	(0.012)	(0.019)
Book Leverage		-0.754***	-0.753***	-0.765***
		(0.081)	(0.081)	(0.197)
R&D Intensity		0.342***	0.343***	0.004
		(0.094)	(0.094)	(0.335)
Adj. Forward Citations			0.004***	0.003***
			(0.001)	(0.001)
Adj. Backward Citations			0.000	-0.000
			(0.001)	(0.001)
Examination Time (Year)			0.001	-0.000
			(0.001)	(0.002)
Multiple Issuance (N)				-0.026***
				(0.007)
Observations	225,917	225,917	225,917	1,505,856
Adjusted R-squared	0.874	0.925	0.925	0.815
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes

Table 3.3: Industry Technology Gap and Technological Leaders

This table performs the baseline regression analysis allowing the interaction of industry competition with industry technological leaders. Column (1) is estimated based on all single patent grants with non-missing industry technology gap measures. In column (2) to (4), patents are split into terciles based on the issuing firms' industry technology gap following Aghion et al. (2005), and the estimation is based on patents issued to firms in industries with a low/moderate/high level of technology gap, respectively. Firm and industry characteristics are measured at the end of fiscal years prior to patent grants. All variables are defined in Table 3.A1. All columns include firm, granting year, and patent class fixed effects. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	Industry Technology Gap			
	(1) All	(2) Low	(3) Moderate	(4) High
Technology Leader	0.063***	-0.030	0.096**	0.137***
	(0.021)	(0.036)	(0.036)	(0.040)
Competition	-0.877***	-0.887***	-1.207***	-0.774**
	(0.198)	(0.282)	(0.255)	(0.307)
Technology Leader * Competition	-0.139	-0.844***	0.092	0.554
	(0.150)	(0.279)	(0.218)	(0.343)
Observations	176,223	59,331	61,098	54,860
Adjusted R-squared	0.934	0.942	0.941	0.947
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes

Table 3.4: Cross-sectional Variation

This table performs the baseline regression analysis allowing the interaction of industry competition with HighPower, a dummy variable that equals one if the patenting firm's markup is above the industry-year median (Column (1)), HighShare, a dummy variable that equals one if the patenting firm's market share is above the industry-year median (Column (2)), HighValue, a dummy variable that equals one if the patent's economic value is above the industry-year median (Column (3)), and HighCitation, a dummy variable that equals one if the patent's number of forward citations is above the industry-year median (Column (4)). The estimation is based on all single patent grants on a firm-day. Firm and industry characteristics are measured at the end of fiscal years prior to patent grants. All variables are defined in Table 3.A1. All columns include firm, granting year, and patent class fixed effects. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	Firm Characteristic		Patent Characteristic	
	(1) Market Power	(2) Market Share	(3) Patent Value	(4) Patent Citation
Competition	-0.999***	-0.870***	-0.747***	-0.923***
	(0.161)	(0.156)	(0.241)	(0.186)
HighPower	0.159***			
	(0.026)			
HighPower * Competition	0.127			
	(0.244)			
HighShare		0.003		
		(0.032)		
HighShare * Competition		-0.096		
		(0.193)		
HighValue			0.268***	
			(0.023)	
HighValue * Competition			-0.727***	
			(0.169)	
HighCitation				-0.005
				(0.004)
HighCitation * Competition				-0.050
				(0.036)
Observations	225,917	225,917	225,917	225,917
Adjusted R-squared	0.926	0.925	0.933	0.925
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes

Table 3.5: Quasi-experimental Regression

This table presents the results of the quasi-experimental regression in which I compare patent value between patents issued immediately before and after horizontal merger announcements in the same industry as the patenting firms. $Post_{j,s,t}$ is a dummy equals one if patent j is issued in the [8, 35] day window after a horizontal merger announcements and zero if it is issued in the [-35, -8] window. Patents issued to the bidder and the target are excluded from this analysis. Firm and industry characteristics are measured at the end of fiscal years prior to patent grants. All variables are defined in Table 3.A1. All columns include firm, granting year, and patent class fixed effects. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	Single Patent Grants			All Patent Grants
	(1)	(2)	(3)	(4)
Post-merger	0.050***	0.048***	0.048***	0.036***
	(0.010)	(0.009)	(0.010)	(0.009)
Log(Total Assets)		0.645***	0.644***	0.448***
		(0.032)	(0.032)	(0.067)
Net Income (Loss)		0.147***	0.149***	0.225***
		(0.052)	(0.051)	(0.071)
Market-to-Book		0.095***	0.094***	0.105***
		(0.019)	(0.019)	(0.029)
Book Leverage		-0.597***	-0.591***	-0.417***
		(0.182)	(0.180)	(0.086)
R&D Intensity		0.589***	0.592***	0.215
		(0.115)	(0.118)	(0.313)
Adj. Forward Citations			0.011***	0.003***
			(0.002)	(0.001)
Adj. Backward Citations			0.005***	0.001
			(0.001)	(0.002)
Examination Time (Year)			-0.002	0.001
			(0.006)	(0.002)
Multiple Issuance (N)				-0.092***
				(0.012)
Observations	12,929	12,929	12,929	110,399
Adjusted R-squared	0.867	0.914	0.915	0.893
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes

Table 3.6: Balance Test

This table performs balance tests of the quasi-experimental sample. Panel A displays baseline differences of the pre-announcement and post-announcement sample for key variables employed in the analysis and stock market characteristics that could potentially affect the patent value estimates. This panel reports means and standard deviations of the pre- and post-announcement samples, and differences and standard error of the differences between these two samples. Stock market characteristics are reported for all non-merging firms in the industry as the merging firms instead of estimation sample firms. The 3-day CARs are estimated for each trading day in the [-35, -8] and [8, 35] windows around merger announcements. Betas and idiosyncratic volatility are estimated based on market model using 252 trading day daily returns before and after the event, respectively. Panel B shows the quasi-experimental regression results with the outcome variable being the adjusted number of backward citations (Column (1)), the adjusted number of forward citations (Column (2)), patent examination time (Column (3)), and the number of patent grants on a firm-day (Column (4)). The estimation in Panel B is based on all patent grants. All variables are defined in Table 3.A1. Columns (1) - (3) include firm, granting year, and patent class fixed effects, and Column (4) includes firm and granting year fixed effects. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

Panel A: Summary Statistics	Pre-announcement		Post-announcement		Difference	
	Mean	SD	Mean	SD	Pre-Post	StdError
Adj. Backward Citations	1.200	2.765	1.185	2.716	0.015	0.016
Adj. Forward Citations	1.191	2.901	1.184	2.987	0.006	0.017
Examination Time	3.132	1.673	3.119	1.657	0.013	0.010
Firm characteristics						
Log(Total Assets)	6.487	2.042	6.482	2.050	0.005	0.036
Net Income (Loss)	-.059	0.338	-.067	0.377	0.007	0.006
Market-to-Book	2.989	3.548	2.919	3.239	0.069	0.060
Book Leverage	0.122	0.156	0.121	0.157	0.000	0.002
R&D Intensity	0.160	0.163	0.159	0.165	0.000	0.002
Multiple Issuance (N)	2.772	5.859	2.791	6.003	-.019	0.105
MarketCap	26.552	66.317	26.958	67.775	-.405	0.867
Stock market characteristics						
CAR (%)	0.121	7.668	0.001	7.706	0.120	0.008
CAR (%, Tue)	0.118	7.788	0.001	7.755	0.117	0.018
Beta	0.979	1.068	0.981	0.799	-.001	0.005
Ivol	0.041	0.019	0.042	0.019	-.000	0.000

Table 3.6: Balance Test (Continued)

Panel B: Regression Analysis	(1) #Backward Citations	(2) #Forward Citations	(3) Examination Time	(4) #Patent Grants
Post-merger	-0.003	-0.002	-0.001	0.007
	(0.003)	(0.003)	(0.001)	(0.075)
Log(Total Assets)	0.033***	-0.027***	0.027***	1.035***
	(0.010)	(0.009)	(0.009)	(0.279)
Net Income (Loss)	-0.019	-0.004	-0.004	0.423***
	(0.015)	(0.023)	(0.021)	(0.131)
Market-to-Book	-0.002	0.010***	-0.004***	-0.298**
	(0.001)	(0.002)	(0.001)	(0.114)
Book Leverage	-0.133***	-0.053	-0.064**	0.361
	(0.031)	(0.042)	(0.030)	(1.136)
R&D Intensity	-0.179***	0.077	-0.004	3.771***
	(0.052)	(0.059)	(0.059)	(1.204)
Multiple Issuance (N)	0.002**	-0.004***	-0.000	
	(0.001)	(0.001)	(0.000)	
Observations	110,399	110,399	110,399	13,072
Adjusted R-squared	0.159	0.150	0.253	0.663
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	No

Table 3.7: Heterogenous Effect of Mergers on Patent Value

This table performs the quasi-experimental regression analysis for horizontal merger deals with different characteristics. In columns (1) and (2), I focus on deals with a deal value that are 10% below and above the US antitrust scrutiny threshold, where *Stealth Mergers* is an indicator for those deals with values that fall just below the threshold, following Kepler et al. (2021). In columns (3) and (4), I identify merger deals in highly concentrated industries with high product similarity. The dummy variable *High concentration and IPS* takes a value of one if the event industry has an above-median concentration and above-median concentration industry product similarity. The estimation is based on all single patent grants on a firm-day. Firm and industry characteristics are measured at the end of fiscal years prior to patent grants. All variables are defined in Table 3.A1. All columns include firm, granting year, and patent class fixed effects. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	Stealth Merger		High Concentration and IPS	
	(1) Single	(2) All	(3) Single	(4) All
Post-merger	0.005	-0.075	0.028**	0.029***
	(0.052)	(0.028)	(0.012)	(0.008)
Post x Stealth Merger	0.086	0.121**		
	(0.032)	(0.017)		
Post x High Concentration and IPS			0.083***	0.140***
			(0.025)	(0.024)
High Concentration and IPS			0.156	0.000
			(0.103)	(0.086)
Observations	422	4,056	7,160	70,215
Adjusted R-squared	0.967	0.947	0.914	0.894
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes

Table 3.8: Alternative Shocks on Product Market Competition

The table performs the quasi-experimental regression using non-horizontal merger announcements (Panel A), horizontal merger withdrawal announcements (Panel B, Column (1)-(2)), and US–China Relations Act of 2000 (Panel B, Column (3)-(4)) as alternative events that affect product market competition. Non-horizontal mergers are those M&A deals in which the acquirer and the target are not in the same SIC 4-digit industry. The event industry is the acquirer's industry in Column (1) and (2), and the target firm's industry in Column (3) and (4). Horizontal merger withdrawal announcements use the withdrawal announcement dates as event dates and compare patent value immediately before and after those announcements. The US–China Relations Act of 2000 granted Permanent Normal Trade Relations (PNTR) to China on 10th October 2000. Estimation is based on patents issued within the [-35, -8] and [8, 35] event windows. Firm and industry characteristics are measured at the end of fiscal years prior to patent grants. All variables are defined in Table 3.A1. All columns include firm, granting year, and patent class fixed effects. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

Panel A: Non-horizontal Merger				
	Acquirer Industry		Target Industry	
	(1) Single	(2) All	(3) Single	(4) All
Post-event	0.016	0.042	0.001	-0.002
	(0.028)	(0.031)	(0.015)	(0.009)
Observations	7,867	46,768	6,585	37,274
Adjusted R-squared	0.939	0.909	0.937	0.899
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes
Panel B: Other Events				
	Deal Withdrawal		US-China PNTR	
	(1) Single	(2) All	(3) Single	(4) All
Post-event	-0.030	-0.026	-0.183***	-0.216***
	(0.030)	(0.020)	(0.064)	(0.053)
Observations	1,240	13,767	726	4,257
Adjusted R-squared	0.950	0.923	0.784	0.779
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes

Table 3.9: Robustness: Baseline Regression

This table presents robustness tests of the baseline regression in Table 2 Column (3) using alternative definitions of product market competition (Panel A), citation-based measures as outcome variables (Panel B),firm-level innovation outcomes (Panel C), and additional control variables (Panel D). All variables are defined in Table 3.A1. All columns include firm, granting year, and patent class fixed effects. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

Panel A: Competition Measures	(1)	(2)	(3)
	Equally Weighted Competition	SIC-3 Competition	SIC-4 HHI
Log(Competition)	-1.133***	-0.915***	0.024
	(0.224)	(0.238)	(0.026)
Observations	225,917	225,917	121,096
Adjusted R-squared	0.925	0.925	0.919
Panel B: Citation Regression	(1)	(2)	(3)
	Citation Received	Self Citation	Non-self Citation
Log(Competition)	-0.005	-0.145	0.031
	(0.056)	(0.098)	(0.044)
Observations	225,917	225,917	225,917
Adjusted R-squared	0.203	0.228	0.180
Panel C: Firm-level Outcomes	(1)	(2)	(3)
	R&D Intensity	Patent Economic Value	Patent Scientific Value
Log(Competition)	-0.030*	-0.659*	0.385
	(0.016)	(0.372)	(0.318)
Observations	40,669	32,529	32,529
Adjusted R-squared	0.819	0.899	0.709
Panel D: Additional Controls	(1)	(2)	(3)
	Stock Market Char.	Patent Controls	Class-Year FE
Log(Competition)	-1.037***	-1.092***	-0.951***
	(0.179)	(0.177)	(0.171)
Market Cap. (t-1)	0.002		
	(0.001)		
Market Beta (t-1)	0.218***		
	(0.029)		
IVol (t-1)	4.306***		
	(1.053)		
Patent generality		0.007	
		(0.005)	
Patent originality		-0.003	
		(0.007)	
Observations	206,071	194,902	225,589
Adjusted R-squared	0.929	0.925	0.930

Table 3.10: Robustness: Quasi-experimental Regression

This table presents robustness tests of the quasi-experimental regression in Table 5 Column (3) using alternative event windows (Panel A), and conditioning on additional control variables (Panel B). All variables are defined in Table 3.A1. All columns include firm, granting year, and patent class fixed effects. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

Panel A: Event Window	(1)	(2)	(3)	(4)
	± 21 Days	± 28 Days	± 42 Days	± 56 Days
Post-merger	0.036***	0.042***	0.047***	0.046***
	(0.011)	(0.011)	(0.011)	(0.009)
Observations	6,209	9,589	16,217	19,614
Adjusted R-squared	0.914	0.915	0.915	0.916
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes

Panel B: Additional Controls	(1)	(2)	(3)	(4)
	Stock Market Controls	Patent Controls	Deal Controls	Class-Year FE
Post-merger	0.047***	0.053***	0.049***	0.044***
	(0.012)	(0.012)	(0.010)	(0.008)
Market Cap. (t-1)	0.004			
	(0.002)			
Market Beta (t-1)	0.283***			
	(0.050)			
IVol (t-1)	2.463			
	(2.282)			
Patent Generality		-0.011		
		(0.025)		
Patent Originality		0.067**		
		(0.032)		
Deal Value			0.000	
			(0.000)	
Combined Share			-0.174	
			(0.123)	
Observations	11,727	11,096	12,855	12,611
Adjusted R-squared	0.919	0.911	0.914	0.924
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes

Figure 3.1: Patent Value, Citation, and Product Market Competition

This figure displays the unconditional relationship between patent value and the number of citations received (top panel), patent value and product market competition (middle panel), the number of citations received and product market competition (bottom panel) using binned scatter plots with quadratic fit lines. The sample includes all patents issued to US public firms between 1976 to 2020. All variables are winsorized at 1% and 99% levels.

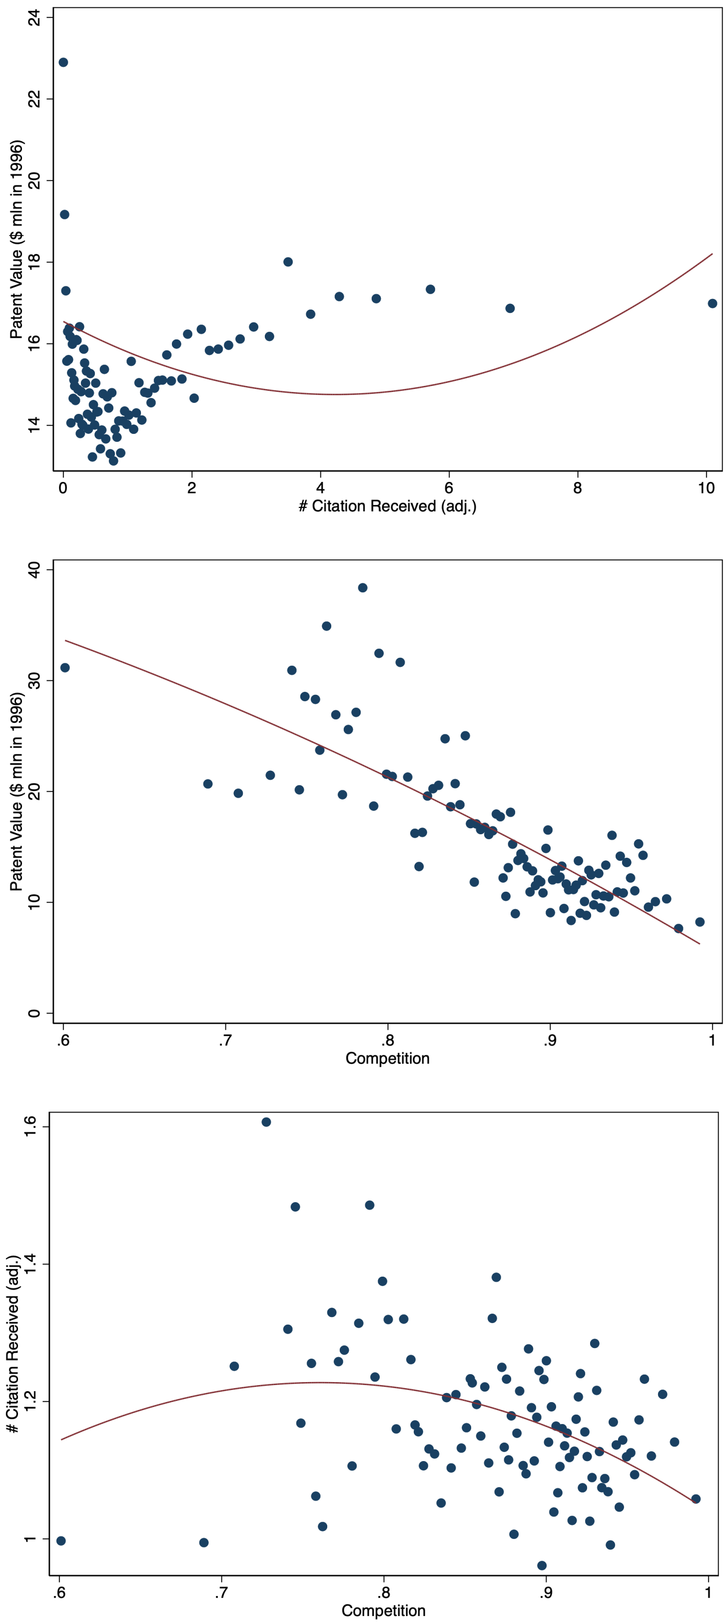

Figure 3.2: Competition and Patent Value over Time

This figure displays the relationship between patent value and product market competition over time. The sample includes all patents issued to US public firms between 1976 to 2020. All variables are winsorized at 1% and 99% levels.

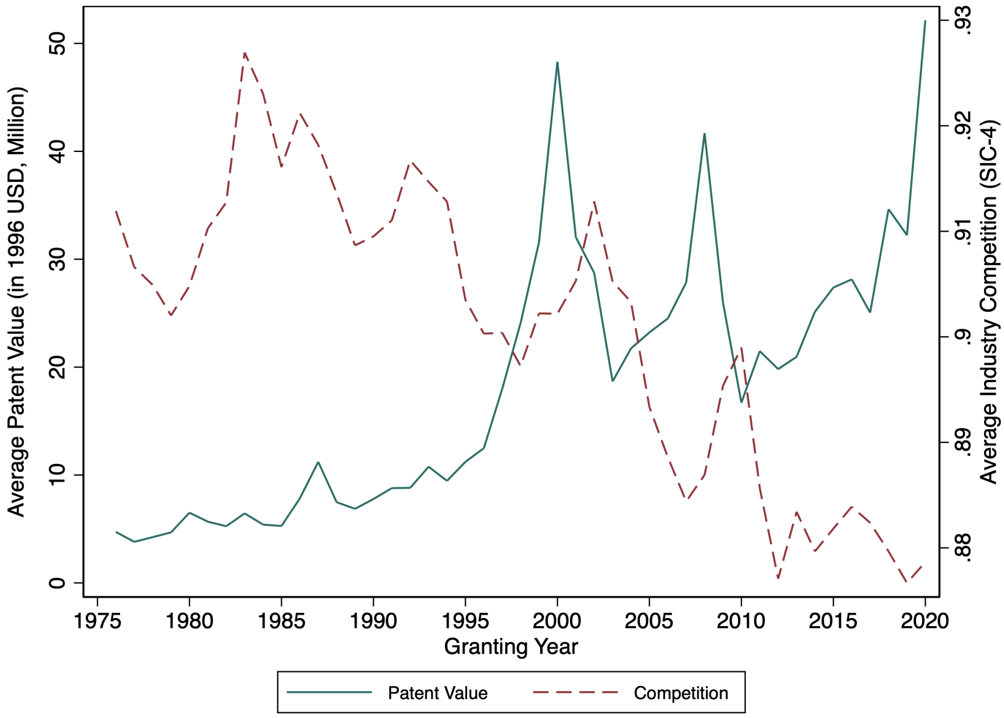

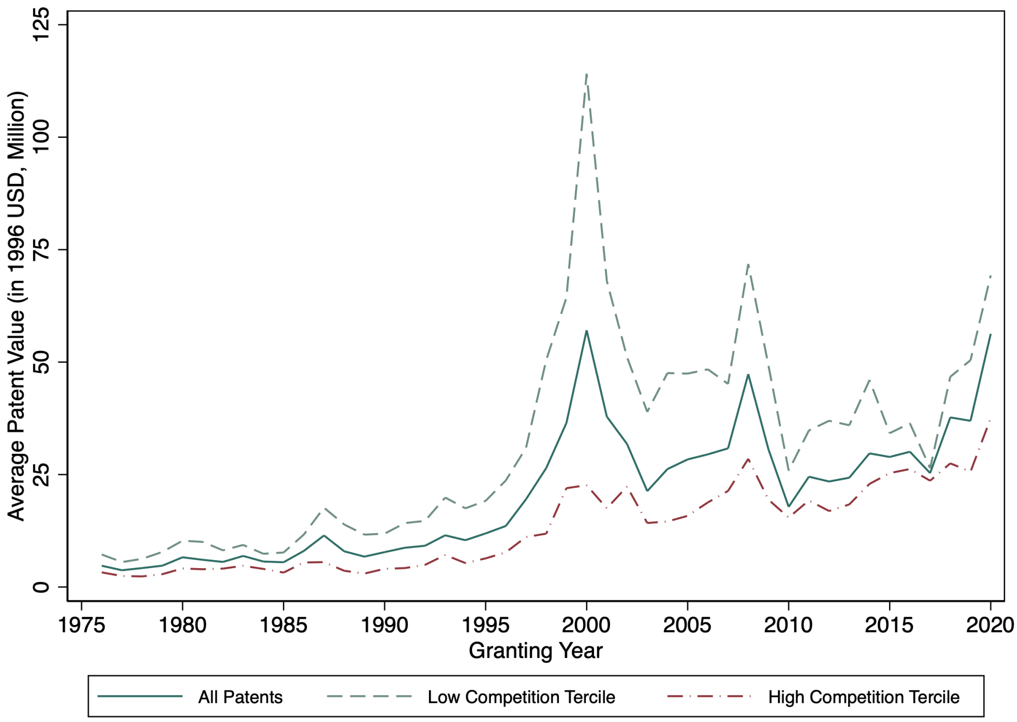

Figure 3.3: Patent Value and Stock Market Capitalization

This figure displays the relationship between patent value and stock market capitalization over time. The sample includes all patents issued to US public firms between 1976 to 2020. All variables are winsorized at 1% and 99% levels.

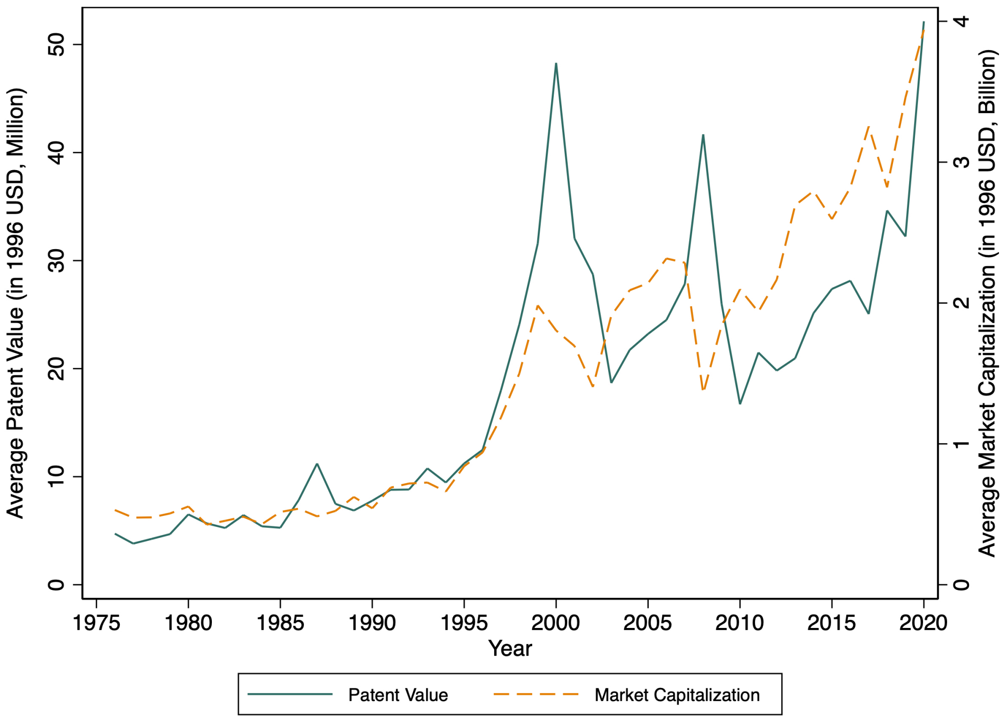

Table 3.A1: Variable Definition

This table provides definitions of variables used in our analysis along with the relevant data sources in square brackets.

Variable	Definition
Patent Class	Three-digit Cooperative Patent Classification (CPC), including 125 patent classes. [USPTO]
Patent Value	Stock market return based estimates of patent value in Kogan et al. (2017) (Million) [Kogan et al. (2017)]
Adj. Backward Citations	The number of backward citation a patent made to prior patents, adjusted by CPC 3-digit patent class and filing year average. [USPTO]
Adj. Forward Citations	The number of forward citation a patent received from other patents, adjusted by CPC 3-digit patent class and filing year average. [USPTO]
Examination Time	The number of years between a patent's filing and grant date. [USPTO]
Multiple Issuance	The number of patents issued to a firm on a day. [USPTO/Kogan et al. (2017)]
Assets	Total Assets in 2000 US dollars (Million) [Compustat]
Sales	Total Sales in 2000 US dollars (Million) [Compustat]
Profitability	Net Income / Total Assets [Compustat]
Market-to-Book	(Market Value + Total debt + Preferred Stock - Deferred Taxes and Investment Tax Credit) / Total Assets [Compustat]
R&D Intensity	R&D Expense / Total Assets [Compustat]
Book Leverage	Total Debt / Total Assets [Compustat]
Markup	(Sales - Cost of Goods Sold - Selling, General and Administrative Expense - Depreciation) / Sales [Compustat]
Multiple Issuance	The number of patents issued by a firm in a year [Compustat/USPTO]
Market Capitalization	Number of Shares Outstanding * Share Price [CRSP]
Patent Economic Value (Firm)	Aggregate patents' economic value of all patents granted to a firm in a year in 2000 US dollars (Million) [Kogan et al. (2017)]
Patent Scientific Value (Firm)	Aggregate patents' scientific value of all patents granted to a firm in a year [Kogan et al. (2017)]
Competition	1- Industry sales-weighted average markup [Compustat]
Technology Gap	$\frac{1}{N_{s,t}} \sum_{i=1}^{N} \frac{TFP_{highest s,t} - TFP_{i,t}}{TFP_{highest s,t} - TFP_{lowest s,t}}$, where N is the number of firms in a industry, $TFP_{i,t}$, $TFP_{highest s,t}$, and $TFP_{lowest s,t}$ are the TFP of firm i, the highest firm TFP in a industry s, and the lowest firm TFP in the industry s, respectively. [Imrohoroğlu and Tüzel (2014)]
HHI	Industry Herfindahl–Hirschman Index based on concentration data provided by U.S. Census Bureau at the U.S. Department of Commerce. [Keil (2017)]
IPS	An industry product similarity measure based on firm 10K fillings and ranges from zero to one with a higher value indicating a greater similarity. [Fathollahi et al. (2022)]
No. of firms	The number of US publicly listed firms in an industry [Compustat]
TNIC Similarity	Firms' pair-wise product market similarity score from the Text-based Network Industry Classification (TNIC) (calibrated to be as granular as three-digit SIC codes) [Hoberg and Phillips Data Library]

Table 3.A2: Baseline Regression with Alternative Specifications

This table presents the results of regressing log patent value on log product market competition and different combinations of firm controls, patent controls, and firm, year and patent class fixed effects. Estimation is based on single patent grants on a firm-day. Firm and industry characteristics are measured at the end of fiscal years prior to patent grants. All variables are defined in Table 3.A1. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	Patent Value					
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Competition)	-3.232***	-4.118***	-1.910*	-3.279***	-2.340***	-0.946***
	(1.052)	(0.542)	(1.080)	(0.896)	(0.340)	(0.183)
Observations	226,972	226,038	226,972	226,852	225,917	225,917
Adjusted R-squared	0.030	0.815	0.082	0.069	0.874	0.925
Firm FE	No	Yes	No	No	Yes	Yes
Year FE	No	No	Yes	No	Yes	Yes
Patent Class FE	No	No	No	Yes	Yes	Yes
Controls	No	No	No	No	No	Yes

Table 3.A3: Baseline Regression with Nonlinearity

This table presents the baseline regression allowing for a non-linear relationship between patent value and product market competition. Panel A presents the regression results which regress patent value on the competition variable and its squared form (both variables are not log transformed to avoid perfect multicollinearity). Panel B displays the baseline regression results for four subsamples based on competition quartiles. The estimation is based on single patent grants on a firm day. Firm and industry characteristics are measured at the end of fiscal years prior to patent grants. All variables are defined in Table 3.A1. All columns include firm, granting year, and patent class fixed effects Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

Panel A: Non-linearity	Single Patent Grants			All Patent Grants
	(1) Single	(2) Single	(3) Single	(4) All
Competition	7.031*	7.437***	7.438***	2.409
	(3.855)	(1.508)	(1.506)	(6.731)
Competition Squared	-5.888**	-5.126***	-5.127***	-2.492
	(2.265)	(0.917)	(0.916)	(4.002)
Observations	225,917	225,917	225,917	1,505,856
R-squared	0.877	0.927	0.927	0.816
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes

Panel B: Competition Quantile	Competition Quantile			
	(1) Low	(2) ⟵ Competition	(3) ⟶	(4) High
Log(Competition)	-0.398*	-1.392***	-3.447***	-2.307***
	(0.203)	(0.516)	(0.712)	(0.575)
Observations	56,384	56,205	56,093	56,129
R-squared	0.932	0.942	0.939	0.938
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes

Table 3.A4: Summary Statistics for the Quasi-experimental Regression Sample

The quasi-experimental sample consists 110,399 patents, among which 12,929 are single patents granted to a firm on a day. Panel A to C report the patent-level, firm-year, and industry-level characteristics of the all patents and single patents samples. The table reports means, standard deviations, and medians of the main variables employed in the empirical analysis. Variables in levels are measured in 1996 dollars. All variables are defined in Table 3.A1.

	All Patents Sample				Single Patents Sample			
Variable	N	Mean	SD	Median	N	Mean	SD	Median
Panel A: Patent-level								
Patent Value ($mln)	110,399	33.11	79.70	10.17	12,929	42.34	99.25	10.58
Backward Citations	110,399	24.51	132.96	7.00	12,929	23.51	54.87	9.00
Forward Citations	110,399	22.86	74.70	4.00	12,929	42.60	114.11	11.00
Examination Time (Year)	110,399	3.13	1.67	2.79	12,929	2.96	1.57	2.64
Panel B: Firm-level								
Assets ($bln)	4,296	4.62	15.43	0.39	3,790	2.81	10.51	0.31
Sales ($bln)	4,296	2.90	10.74	0.23	3,790	1.78	7.86	0.19
ROA	4,296	−0.09	0.37	0.02	3,790	−0.10	0.38	0.01
Market Cap. ($bln)	4,296	9.22	30.24	0.79	3,790	5.29	16.96	0.63
Market-to-Book	4,296	2.89	3.29	2.00	3,790	2.93	3.38	2.01
R&D Intensity	4,296	0.17	0.19	0.12	3,790	0.18	0.20	0.13
Book Leverage	4,296	0.13	0.16	0.05	3,790	0.12	0.16	0.04
Markup	4,296	0.10	0.12	0.06	3,790	0.09	0.12	0.05
Multiple Issuance	4,296	2.26	5.11	1.00	3,790	1.00	0.00	1.00
Panel C: Industry-level								
HHI	278	0.19	0.12	0.15	264	0.18	0.12	0.14
Competition	278	0.85	0.07	0.85	264	0.85	0.07	0.85
#Firms	278	96.20	83.49	72.00	264	98.19	82.61	77.50
Technology Gap	278	0.41	0.11	0.41	264	0.40	0.11	0.41

Table 3.A5: Additional Robustness: Quasi-experimental Regression

This table presents robustness tests of the quasi-experimental regression in Table 5 Column (3) based on horizontal merger deals that (1) do not have overlapping event windows with another deal, (2) include withdrawn horizontal mergers, and (3) are defined using TNIC peers. Overlapping deals are those with a time gap shorter than 70 days, which would result in duplicates in the identification patent sample. A merger is classified as a TNIC horizontal merger if the bidder and the target have an above-median similarity score (0.05), and rival firms are those who have an above-median similarity score with both the bidder and the target. All explanatory variables are log-transformed. Industry and month fixed effects are included in all columns. Standard errors are clustered at the industry level. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	Overlapping Deals		Include Withdrawals		TNIC Horizontal Merger	
	(1) First	(2) Last	(3) Single	(4) All	(5) Single	(6) All
Post-merger	0.044**	0.039**	0.042***	0.031***	0.019	0.071***
	(0.016)	(0.017)	(0.011)	(0.007)	(0.025)	(0.024)
Observations	6,982	7,244	14,512	124,510	2,910	11,058
Adjusted R-squared	0.922	0.921	0.916	0.892	0.915	0.910
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 3.A6: Firms' Propensity to Merge after Rival Innovation

This table presents the industry-month regression results of regressing a horizontal merger announcement dummy on patenting status of non-merging peers in the previous month. Peer patenting status is measured as the sum of patent values (Column (1) to (2)) or the adjusted number of forward citations (Column (3) to (4)) of all patents issued to the non-merging firms in the event industry in the previous month. All columns control for industry aggregate sales, competition, and HHI. All explanatory variables are log-transformed. Industry and month fixed effects are included in all columns. Standard errors are clustered at the industry level. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	Dep. Variable: Horizontal Merger Announcement Dummy			
	(1) Logit	(2) OLS	(3) Logit	(4) OLS
Patent Value $_{t-1}$	-0.144	-0.002		
	(-0.915)	(-0.501)		
Patent Citation $_{t-1}$			0.089	-0.001
			(0.780)	(-0.211)
Horizontal Merger $_{t-1}$	0.559**	0.054*	0.571**	0.054*
	(2.165)	(1.706)	(2.100)	(1.711)
Sales $_{t-1}$	0.152	-0.001	-0.020	-0.002
	(0.949)	(-0.251)	(-0.167)	(-0.799)
Competition $_{t-1}$	-2.654	-0.134	-2.560	-0.130
	(-0.513)	(-0.894)	(-0.506)	(-0.895)
HHI $_{t-1}$	4.523	0.137	4.009	0.139
	(1.234)	(1.631)	(1.102)	(1.412)
#Firms $_{t-1}$	2.134***	0.052**	2.120***	0.052**
	(4.662)	(2.661)	(4.840)	(2.646)
Observations	5,800	5,800	5,800	5,800
R-squared		0.145		0.145
Industry FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes

Appendix B

Estimation of Patent Economic Value

This section discusses the estimation procedure and assumptions of patent value; a full description can be found in section II in Kogan et al. (2017). In the absence of any other news, the difference in market value around a narrow patent issuance window should capture a portion of patent value that has not already been incorporated in the market value due to ex-ante anticipation of a successful patent. Assume the market believes the probability of successful patent application is π ex-ante, the change in market value only reflects a portion of patent value underestimated by the market and $\Delta V = (1 - \pi)\xi_j$.

Kogan et al. (2017) show that the share turnover for a patent issuing firm increases around the patent issuance day. Thus, the holding period period in the [t,t+2] window is used to estimate patent value. Market return is subtracted from firms' stock returns to capture the idiosyncratic stock returns associated with patent j: $R_j = R_i - R_m$. Nevertheless, the idiosyncratic return R_j might be contaminated by other news unrelated to the patent issuance even in a tiny window. Therefore, the return R_j around the event window contains the return v_j from the patent and a "noise" part ϵ_j that is unrelated to the patenting event: $R_j = v_j + \epsilon_j$. To filter out the noise component in market return, the authors made some distributional assumptions about v_j and ϵ_j, a:

1. Both distributions vary across firms f and time t.
2. Patent value v follows a normal distribution trancated at 0: $v_j \sim \mathcal{N}^+(0, \sigma^2_{vft})$.
3. The noise term follows a normal distribution: $\epsilon_j \sim \mathcal{N}(0, \sigma^2_{\epsilon ft})$.

With these assumptions, the filtered return from patents can be expressed as a function of the observed excess return:

$$E[v_j|R_j] = \delta_{ft} R_j + \sqrt{\delta_{ft}}\sigma_{\epsilon ft} \frac{\phi(-\sqrt{\delta_{ft}}\frac{R_j}{\sigma_{\epsilon ft}})}{1 - \Phi(-\sqrt{\delta_{ft}}\frac{R_j}{\sigma_{\epsilon ft}})}$$

where ϕ and Φ are the standard normal PDF and CDF, σ_ϵ is the standard deviation of the noise component in idiosyncratic stock return, and δ is the estimated signal-to-noise ratio, and its estimated value is 0.0145 over their sample period 1926-2010[30].

Therefore, the dollar value of the economic value of a patent j, ξ_j, can be calculated using the following equation:

$$\xi_j = (1-\bar{\pi})^{-1}\frac{1}{N_j}E[v_j|R_j]M_j$$

where $\bar{\pi}$ is the estimated probability of successful patent applications [31], N_j is the number of patents issued to the same firm on the same day, and M_j is the market capitalization of the issuing firm. $E[v_j|R_j]$ is the expected return from a successful patent conditional observed stock market reaction around the $[t, t+2]$ patent issuance window explained earlier.

Therefore, the economic value of patents is a dollar amount value based on stock market abnormal return around a 3-day patent issuance window and further adjusts for (1) the noise in return, (2) the under-expectation arise from ex-ante probability assessment of a successful patent application, and (3) firm market capitalization.

I make two modifications to KPSS's estimation of patent economic value for robustness check. First, The first modification involves calculating the idiosyncratic return as excess returns based on the capital asset pricing model (CAPM).

Kogan et al. (2017) calculate the idiosyncratic return as stock returns minus the market return to avoid measurement error of estimating stocks' market beta. This modification addresses the simplifying assumption in KPSS that all stocks

[30]I estimate the signal-to-noise ratio in the period 1976-2020, and it is 0.01419. I follow Kogan et al. (2017) for the empirical analysis.

[31]Carley et al. (2015) estimates that the unconditional probability of successful patent application π is 56% between 1991 and 2001. Kogan et al. (2017) uses this estimate for their entire sample period 1926-2010. I follow Kogan et al. (2017) in the primary analysis and address the potential measurement error in the robustness check section.

have a beta of one. Therefore, excess returns based on CAPM are used in robustness tests and the signal-to-noise ratio is reestimated using CAPM returns.

Second, the same ex-ante probability of successful patent application is applied for all patents. This assumption is less of a concern in Kogan et al. (2017) as they intend to capture the aggregate economic value of patents at firms and the economy level. Since this paper examines patent value at a more granular dimension, the assumption of $\pi_j = \bar{\pi}$ leads to measurement error. Thus, I relax this assumption by using the change in market value around patent filing date as a proxy for the ex-ante assessment of the patent value[32]. The economic value of patents under this method will be:

$$\xi_j = \frac{1}{N_f} E[v_f | R_f] M_f + \frac{1}{N_j} E[v_j | R_j] M_j$$

where the first part of this equation represents the initial market valuation of the patent on the patent filing date with the belief that the probability of successful patent application ($\pi\xi_j$) and the second part reflects a portion of patent value underestimated by the market on patent filing date ($(1 - \pi)\xi_j$).

However, one limitation of this method is that the patent filing date is only available when the patent is granted. If one or more unsuccessful patent applications are filed on the same day as a successful patent, this method assumes that the initial market valuation of the unsuccessful patent(s) is zero. In addition, patent application stock returns might not be sensitive to contemporaneous product market competition, which might introduce noise to empirical tests.

These different measures are then used to replicate Table 3.2 and Table 3.5, which are reported in Table 3.B1. Despite the difference in estimated values, all results are consistent with previous findings. The results in columns (1) and (2) using patent value based on CAPM in close to the main results. The coefficients in columns (3) and (4) are smaller in terms of economic magnitude. This can be explained by the fact that a large portion of the patent value is

[32] I use patent filing date instead of patent publication date because the publication date stock returns do not appear to be

not affected by contemporaneous product market competition, especially for the quasi-experimental regression.

Table 3.B1: Robustness to Patent Economic Value Estimation

This table presents robustness tests of the baseline and the quasi-experiment regression using alternative measures for patent economic value. All variables are defined in Table 3.A1. All columns include firm, granting year, and patent class fixed effects. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

Panel A: CAPM Excess Return	Baseline		Quasi-experiment	
	(1) Single	(2) All	(3) Single	(4) All
Log(Competition)	-0.890***	-1.367***		
	(0.153)	(0.427)		
Post-merger			0.040***	0.025***
			(0.008)	(0.007)
Observations	224,980	1,501,996	12,845	110,289
Adjusted R-squared	0.918	0.805	0.912	0.900
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes
Panel B: Filling and Issuance Date Return	**Baseline**		**Quasi-experiment**	
	(1) Single	(2) All	(3) Single	(4) All
Log(Competition)	-0.848***	-0.815**		
	(0.153)	(0.365)		
Post-merger			0.030***	0.016***
			(0.005)	(0.005)
Observations	199,283	1,405,398	10,974	105,337
Adjusted R-squared	0.924	0.783	0.919	0.818
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes

Hypothesis development for cross-sectional variation

Aghion et al. (2005) assume competition has heterogeneous effects on innovation rent for firms in different industries. Firms within an industry could either have similar or different market power. If firms within the same industry have different pricing power, they would have different pre-innovation rent. Thus, the innovation rent of an identical invention might be different for firms in the same industry. Hence, it is essential to distinguish the technologically dominant firms V.S. non-dominant firms within an industry and distinguish industries with equal and unequal pricing power. I follow Aghion et al. (2005) when developing a hypothesis for different firms in different industries.

In Aghion et al. (2005), each industry is a duopoly with only two firms and can either be a leader-laggard (L.L.) industry or neck-to-neck (N.N.) industry. In the L.L. industry, two firms have different levels of technology. The technology leader is always one step ahead of the laggard firm and has the pricing power to earn a positive profit, π_{leader}, whereas the laggard firm earns a zero profit, $\pi_{laggard} = 0$. In the N.N. industry, two firms are at the same technology level, and their profit, $\pi_{N.N.\ firm}$, is negatively associated with the degree of competition. The assumption is that $\pi_{leader} > \pi_{NN\ firm} = (1-\epsilon)\pi_{leader} > \pi_{laggard} = 0$, where ϵ is the degree of market competition. By successful innovation, a laggard firm in the L.L. industry can catch up with the leader and the duopoly becomes an N.N. industry.

For the N.N. industry, the firm that first introduces new technology becomes the leader while its peer automatically becomes the laggard firm; the duopoly becomes an L.L. industry. Thus, the incremental profit of successful innovation is positively associated with the pressure of competition for firms in an N.N. industry ($\Delta\pi_{N.N.\to leader} = \epsilon\pi_{leader}$). Thus, High competition leads to higher economic value of patents for firms in an NN industry.

For the L.L. industry, if the laggard firm innovates, it can catch up with the leader firms' technology, and the duopoly becomes the N.N. form. The incremental profit of the laggard firm's successful innovation is negatively associated with the pressure of competition ($\Delta\pi_{laggard \rightarrow N.N.} = (1-\epsilon)\pi_{leader}$).

It is noteworthy that the assumption in Aghion et al. (2013) that the technological leader in an L.L. sector never innovates is unrealistic. For example, Apple keeps introducing new iPhone models to restricting competition despite its dominant position in the cellphone industry. Therefore, dominant firms also have the incentive to innovate as they want to maintain their leadership position and avoid future competition. Assume the leader innovates to avoid being caught up by the laggard firm and forced to earn a lower profit $\pi_{N.N.}$, which is negatively associated with the degree of competition. The value of innovation for a leader firm comes from keeping its dominant position and avoiding the drop in profit. Since the potential decline in profit is positively correlated with the market competition ($\Delta\pi_{leader \rightarrow N.N.} = -\epsilon\pi_{leader}$), I conjecture that the more intense the competition, the more valuable dominant firms' innovation would be. As a result, patent value is expected to be positively associated with the market competition for the leading firm in an L.L. industry. In another words, high competition leads to higher economic value of patents for leader firms but lower economic value of patents for laggard firms in an LL industry.

Moreover, industries are not simply N.N. duopolies or L.L. duopolies with one step technology gap in the real world. To test the heterogeneous effect of competition on the value of innovation empirically, I follow Aghion et al. (2005) who use the average distance between peer firms TFP and highest TFP within the industry to gauge the industry's technological gap.

Appendix C

Table 3.C1: Sample Distribution by Fama-French 48 Industry Classification (FFIC48)

This table presents the sample distribution by the Fama-French 48 Industry Classification. The two samples are the single patent grant subsamples in the baseline and quasi-experiment regressions.

		Baseline		Identification	
FFIC48	Industry Description	#	%	#	%
1	Agric	421	0.2	0	0
2	Food	3762	1.7	0	0
3	Soda	380	0.2	0	0
4	Beer	1027	0.5	0	0
5	Smoke	919	0.4	0	0
6	Toys	2424	1.1	0	0
7	Fun	708	0.3	0	0
8	Books	451	0.2	0	0
9	Hshld	8018	3.5	7	0.07
10	Clths	955	0.4	0	0
11	Hlth	593	0.3	0	0
12	MedEq	16120	7.1	499	4.8
13	Drugs	24600	11	2031	20
14	Chems	13106	5.8	29	0.3
15	Rubbr	1068	0.5	0	0
16	Txtls	1062	0.5	0	0
17	BldMt	6452	2.9	0	0
18	Cnstr	307	0.1	0	0
19	Steel	4214	1.9	8	0.08
20	FabPr	693	0.3	0	0
21	Mach	16157	7.2	5	0.05
22	ElcEq	7326	3.2	0	0
23	Autos	10145	4.5	56	0.5
24	Aero	4188	1.9	10	0.10
25	Ships	1133	0.5	0	0
26	Guns	1859	0.8	0	0
27	Gold	44	0.02	0	0
28	Mines	534	0.2	0	0
29	Coal	37	0.02	0	0
30	Oil	6243	2.8	144	1.4
32	Telcm	2793	1.2	131	1.3
33	PerSv	252	0.1	0	0
34	BusSv	14944	6.6	2748	26
35	Comps	16485	7.3	186	1.8
36	Chips	30130	13	4420	43
37	LabEq	11530	5.1	88	0.8
38	Paper	7054	3.1	3	0.03
39	Boxes	2303	1.0	0	0
40	Trans	831	0.4	5	0.05
41	Whlsl	1815	0.8	2	0.02
42	Rtail	1479	0.7	4	0.04
43	Meals	116	0.05	2	0.02
48	Other	1239	0.5	0	0
Total		225917	100	10378	100

Table 3.C2: Sample Distribution by Cooperative Patent Classification (CPC) Section

This table presents the sample distribution by Cooperative Patent Classification (CPC) Section. The two samples are the single patent grant subsamples in the baseline and quasi-experiment regressions.

		Baseline		Identification	
CPC Section	Section Description	#	%	#	%
A	Human necessities	34536	15	1209	12
B	Performing operations; transporting	34581	15	240	2.3
C	Chemistry; metallurgy	38430	17	1392	13
D	Textiles; paper	3008	1.3	8	0.08
E	Fixed constructions	5696	2.5	17	0.2
F	Mechanical engineering	15988	7.1	62	0.6
G	Physics	47925	21	3536	34
H	Electricity	45744	20	3914	38
Y	General	9	0.004	0	0
Total		225917	100	10378	100

Table 3.C3: Sample Distribution by Cooperative Patent Issuing Year

	Baseline		Identification	
Patent Issuing Year	#	%	#	%
1976	4611	2.0	0	0
1977	4457	2.0	0	0
1978	4314	1.9	0	0
1979	3663	1.6	0	0
1980	3985	1.8	0	0
1981	4071	1.8	22	0.2
1982	3783	1.7	0	0
1983	3722	1.6	11	0.1
1984	3897	1.7	7	0.07
1985	3958	1.8	3	0.03
1986	3715	1.6	31	0.3
1987	3882	1.7	20	0.2
1988	4144	1.8	4	0.04
1989	4191	1.9	36	0.3
1990	3978	1.8	3	0.03
1991	4131	1.8	34	0.3
1992	4325	1.9	14	0.1
1993	4551	2.0	14	0.1
1994	4959	2.2	114	1.1
1995	4961	2.2	111	1.1
1996	5482	2.4	136	1.3
1997	5873	2.6	246	2.4
1998	6993	3.1	597	5.8
1999	6528	2.9	659	6.3
2000	6256	2.8	596	5.7
2001	6464	2.9	536	5.2
2002	6444	2.9	457	4.4
2003	6498	2.9	636	6.1
2004	6310	2.8	393	3.8
2005	6099	2.7	604	5.8
2006	6156	2.7	576	5.6
2007	5624	2.5	182	1.8
2008	5407	2.4	586	5.6
2009	5113	2.3	475	4.6
2010	5636	2.5	575	5.5
2011	5431	2.4	182	1.8
2012	5375	2.4	222	2.1
2013	5779	2.6	386	3.7
2014	6034	2.7	432	4.2
2015	5784	2.6	673	6.5
2016	5347	2.4	350	3.4
2017	4915	2.2	145	1.4
2018	4359	1.9	115	1.1
2019	4485	2.0	195	1.9
2020	4227	1.9	0	0
Total	225917	100	10378	100

Table 3.C4: M&A Distribution by Fama-French 48 Industry Classification (FFIC48)

This table presents the sample distribution by the Fama-French 48 Industry Classification. The two samples are the single patent grant subsamples in the baseline and quasi-experiment regressions.

		All Deals		Horizontal Mergers		Horizontal Mergers with Patents	
FFIC48	Industry Description	#	%	#	%	#	%
2	Food	10	0.9	1	0.2	0	0
3	Soda	1	0.09	0	0	0	0
4	Beer	3	0.3	3	0.5	0	0
6	Toys	6	0.6	4	0.7	0	0
7	Fun	7	0.7	3	0.5	0	0
8	Books	12	1.1	6	1.1	0	0
9	Hshld	19	1.8	3	0.5	2	0.5
11	Hlth	34	3.2	16	2.9	0	0
12	MedEq	56	5.2	30	5.4	29	7.6
13	Drugs	115	11	44	7.9	44	12
14	Chems	12	1.1	6	1.1	6	1.6
16	Txtls	4	0.4	1	0.2	0	0
17	BldMt	7	0.7	3	0.5	0	0
18	Cnstr	7	0.7	5	0.9	0	0
19	Steel	16	1.5	8	1.4	3	0.8
21	Mach	18	1.7	6	1.1	1	0.3
22	ElcEq	4	0.4	0	0	0	0
23	Autos	22	2.1	4	0.7	3	0.8
24	Aero	11	1.0	2	0.4	2	0.5
25	Ships	2	0.2	1	0.2	0	0
26	Guns	1	0.09	0	0	0	0
27	Gold	4	0.4	1	0.2	0	0
29	Coal	1	0.09	0	0	0	0
30	Oil	76	7.1	62	11	53	14
32	Telcm	62	5.8	34	6.1	20	5.2
33	PerSv	4	0.4	2	0.4	0	0
34	BusSv	173	16	111	20	102	27
35	Comps	78	7.3	25	4.5	24	6.3
36	Chips	120	11	77	14	75	20
37	LabEq	34	3.2	9	1.6	9	2.4
38	Paper	6	0.6	1	0.2	1	0.3
39	Boxes	1	0.09	1	0.2	0	0
40	Trans	32	3.0	25	4.5	4	1.0
41	Whlsl	28	2.6	17	3.0	2	0.5
42	Rtail	65	6.1	32	5.7	1	0.3
43	Meals	17	1.6	16	2.9	1	0.3
Total		1068	100	559	100	382	100

Table 3.C5: Baseline Regression by Fama-French 48 Industry Classification (FFIC48)

This table presents the results of the baseline regression in subsamples based on the Fama-French 48 Industry Classification. All variables are defined in Table 3.A1. Estimation is based on single patent grants on a firm-day. All columns include firm, granting year, and patent class fixed effects. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	Fama-French 48 Industry								
	(1) Autos	(2) BusSv	(3) Chem	(4) Chips	(5) Comps	(6) Drugs	(7) LabEq	(8) Mach	(9) MedEq
Log(Competition)	-0.729	0.289	-0.841*	-0.834*	-0.217	-0.087	0.645	-2.473***	-0.802*
	(0.673)	(0.561)	(0.453)	(0.392)	(0.393)	(0.198)	(0.524)	(0.507)	(0.405)
Log(Total Assets)	0.714***	0.712***	0.778***	0.582***	0.654***	0.599***	0.678***	0.718***	0.610***
	(0.039)	(0.023)	(0.039)	(0.014)	(0.022)	(0.013)	(0.019)	(0.076)	(0.036)
Net Income (Loss)	1.406***	0.115***	0.374	0.142***	0.431***	0.099**	0.073	1.109***	0.208**
	(0.286)	(0.034)	(0.245)	(0.038)	(0.095)	(0.035)	(0.154)	(0.165)	(0.070)
Market-to-Book	0.293***	0.111***	0.278***	0.116***	0.151***	0.140***	0.195***	0.110**	0.135***
	(0.037)	(0.010)	(0.041)	(0.014)	(0.013)	(0.018)	(0.025)	(0.048)	(0.016)
Book leverage	-0.934***	-0.333***	-0.934***	-0.723***	-0.928***	-0.281**	-0.589**	-0.898***	-0.605***
	(0.175)	(0.118)	(0.127)	(0.118)	(0.173)	(0.091)	(0.200)	(0.134)	(0.114)
R&D intensity	2.233**	0.782***	0.597	0.281	0.861**	0.204***	0.414	2.309*	0.336
	(0.723)	(0.199)	(0.565)	(0.160)	(0.312)	(0.034)	(0.538)	(1.147)	(0.258)
Adj. Forward Citations	0.003*	0.001	-0.001	0.001	0.004***	0.003	0.005	0.002	0.001
	(0.001)	(0.001)	(0.003)	(0.001)	(0.001)	(0.003)	(0.004)	(0.002)	(0.002)
Adj. Backward Citations	-0.004	0.002	0.004	0.002	-0.003	0.004	-0.005**	-0.005**	0.002
	(0.002)	(0.002)	(0.003)	(0.002)	(0.004)	(0.003)	(0.002)	(0.002)	(0.003)
Examination Time (Year)	0.003	0.003	-0.002	-0.001	-0.001	-0.002	0.001	0.011***	-0.001
	(0.004)	(0.002)	(0.003)	(0.002)	(0.004)	(0.003)	(0.004)	(0.004)	(0.004)
Observations	10,133	14,895	13,094	30,099	16,453	24,577	11,505	16,146	16,098
Adjusted R-squared	0.915	0.935	0.922	0.892	0.909	0.928	0.938	0.944	0.940
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 3.C6: Baseline Regression by Cooperative Patent Classification (CPC) Section

This table presents the results of the baseline regression in subsamples based on the Cooperative Patent Classification Section (1-digit). All variables are defined in Table 3.A1. Estimation is based on single patent grants on a firm-day. All columns include firm, granting year, and patent class fixed effects. Standard errors are clustered at firm and year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	Cooperative Patent Classification							
	(1) Human necessities	(2) Performing operations; transporting	(3) Chemistry; metallurgy	(4) Textiles; paper	(5) Fixed construction	(6) Mechanical engineering	(7) Physics	(8) Electricity
Log(Competition)	-0.656*	-1.229***	-0.722***	-1.180*	-0.756*	-1.446***	-0.915***	-0.878***
	(0.329)	(0.255)	(0.249)	(0.598)	(0.384)	(0.332)	(0.230)	(0.260)
Log(Total Assets)	0.596***	0.673***	0.633***	0.780***	0.612***	0.692***	0.641***	0.609***
	(0.020)	(0.025)	(0.017)	(0.044)	(0.042)	(0.034)	(0.016)	(0.016)
Net Income (Loss)	0.172***	0.551***	0.176***	0.324	0.641***	0.171*	0.177**	0.184***
	(0.056)	(0.125)	(0.049)	(0.299)	(0.117)	(0.090)	(0.071)	(0.046)
Market-to-Book	0.141***	0.165***	0.183***	0.542***	0.472***	0.205***	0.151***	0.127***
	(0.013)	(0.053)	(0.014)	(0.071)	(0.042)	(0.049)	(0.016)	(0.021)
Book leverage	-0.591***	-1.012***	-0.448***	-0.719***	-0.794***	-1.230***	-0.722***	-0.790***
	(0.080)	(0.090)	(0.117)	(0.152)	(0.131)	(0.100)	(0.095)	(0.110)
R&D intensity	0.178*	0.742**	0.355***	0.950	0.842	0.516	0.537***	0.295*
	(0.093)	(0.296)	(0.027)	(1.590)	(0.727)	(0.648)	(0.171)	(0.148)
Adj. Forward Citations	-0.002	0.004*	0.002	-0.005	0.005	0.000	0.002**	0.002**
	(0.001)	(0.002)	(0.003)	(0.004)	(0.005)	(0.003)	(0.001)	(0.001)
Adj. Backward Citations	0.004	-0.000	0.004***	-0.004	0.000	-0.001	-0.000	0.001
	(0.003)	(0.003)	(0.001)	(0.007)	(0.006)	(0.005)	(0.001)	(0.002)
Examination Time (Year)	-0.003	-0.004	-0.002	0.009	0.014**	0.005	-0.002	0.000
	(0.003)	(0.002)	(0.001)	(0.008)	(0.006)	(0.004)	(0.002)	(0.002)
Observations	34,144	34,090	38,052	2,891	5,443	15,596	47,398	45,285
Adjusted R-squared	0.942	0.937	0.927	0.956	0.962	0.928	0.925	0.911
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Patent Class FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Chapter 4

Cyanide on the Rand: Competing Methods of Technology Transfer

4.1 Introduction

A large literature documents the channels through which technology is transferred. We investigate knowledge spillovers that may arise due to physical proximity, organizational hierarchies, key personnel, or common national origins of shareholders in the South African gold mining industry. Gold was first mined near Johannesburg in 1886, and South Africa was the world's largest gold producer for most of the 20th century. The first commercial technique to extract gold from its ore with cyanide began on these goldfields in 1890. We examine South African gold production from 1887 until 1907.

There are several benefits of studying technology diffusion in our setting. First, we eliminate product market rivalry incentives due to the operation of the Gold Standard. The price of the firms' output, gold, was fixed during this period. Firms could deliver as much gold as they produced to the world's central banks and receive a fixed, nominal price per ounce. Second, we avoid the issue of equating patenting with innovation. We measure changes in mine productivity in a precise fashion, how much gold could be extracted from a tonne of ore. Third, we avoid measurement issues, such as multi-product (and multi-industry) firms which are present in studies such as Jaffe (1986). The firms in our sample mined ore and extracted gold from the ore at a single physical location and had virtually no other operations. These single-plant, single-product firms allow us to avoid the problem of the 'scarcity of directly observable measures of technological innovation' (see Kogan et al., 2017). A final advantage is the relatively high-frequency data that we possess. We observe firm-level productivity at monthly frequency, rather than the annual frequency used by studies such as Bloom et al. (2013b).

We find evidence that physical proximity and operational personnel were effective in diffusing knowledge of the recently invented cyanide process for extracting gold from its ore. In contrast, mines with common ownership (and associated organizational structures) and white-collar employees, such as directors and company secretaries, played a less important role. There is also little evidence that knowledge flowed more easily when there was a common national origin of firms' shareholders.

Many studies have found that physical proximity aids knowledge transfer (e.g., Audretsch and Feldman, 1996; Glaeser et al., 1992; Jaffe and Adams, 1996; Jaffe et al., 1993; Marshall, 1890; Porter, 1990). However, the literature is divided on how distance affects knowledge spillovers. Some find that spillovers occur at the country level (e.g., Fischer et al., 2006; Greenstone et al., 2010; Singh, 2007). Others document spillovers at the city or regional level (e.g., Glaeser et al., 1992; Porter, 1990). In contrast, we find little effect of spillovers at distances over two and a half kilometres. Our results differ from those of Lychagin (2016) and Orlando (2004) who find little attenuation of knowledge diffusion by distance. Even firms that produced an identical product, and used virtually the the same technology, show drastic decreases in knowledge spillovers even at very short distances.

Knowledge spillovers may occur at the individual, rather than the corporate, level. Pogue (2006) argues that South African mining engineers and metallurgists in this era formed a 'localized knowledge network' that aided the diffusion of the cyanide technique. Lychagin (2016) show that inventor locations, not corporate headquarters, matter for knowledge transfers. Conley and Udry (2010) document information flows, such as how much fertilizer to apply, occurs between individual Ghanaian pineapple farmers. They find that farmers who performed well above, or below, expectations provided the most useful information to their peers. In contrast, Borjas and Doran (2012) and Waldinger (2012) find little evidence of knowledge spillovers at the personal level.

Technology diffusion within an organizations often occurs (see e.g., González-Uribe, 2020; Liu, 2008)). Yet the organizational environment is important, and some organizations may lack the comfort or power to successfully transfer knowledge, even if the ability exists (see Cohen and Levinthal, 1989; Levin, 1999)). Most mining companies were members of a mining 'house' (e.g., Consolidated Goldfields of South Africa, A. Goerz and Company, Rand Mines). Mining houses had been established by early South African entrepreneurs such as Hermann Eckstein, Joseph Robinson, and the Albu brothers. The mining houses were large shareholders in the operating mines and provided technical and administrative support (see Martin, 1929). If common ownership is important to facilitate technology dissemination, mines within the same mining house should experience faster transfers than unaffiliated mines. We document that mining houses had some ability to transfer knowledge between their constituent firms. However, mining companies within a mining house rarely merged or transferred assets between each other. Therefore, the channels illustrated by González-Uribe (2020): carve-outs, spawning, and recycling, do not appear to be necessary conditions to enable information flows within an organization. Instead, mining houses appointed managers and consulting engineers to multiple mines within the 'house' to spread knowledge.

We investigate how technology was transferred in an environment that was extremely conducive to cooperation between firms, the South African goldfields at the turn of the 20th century. Spillovers are by no means guaranteed. For example, Bloom et al. (2013a) find that only 2% of neighboring firms learned any details of improved management practices in an experiment on Indian textile firms. They speculate that textile firm managers were reluctant to pass on information to rival firms who might be tempted to poach key personnel. The lack of vigorous inter-firm competition in South African gold mining make our results (of sharply declining spillovers with distance) even more striking. Mining firms did not have to worry about market competition effects (see e.g., Bloom et al., 2013b) as firms acted cooperatively for political and economic reasons. The Chamber of Mines opposed increases in taxation and advocated better conditions for European capital from the Transvaal Republic. Most mines jointly hired unskilled labour through the Witwatersrand Native Labour Association and technical knowledge was shared via the Chemical and Metallurgical Society.[1] The absence of 'market stealing' effects meant that firms in the same mining house, with their different shareholders, may have been unlikely to object to operational knowledge being transferred to other firms. Our results, of co-operative behavior between mines, stands in contrast to the noncooperative behavior found in US oil and gas drilling, see Hendricks and Porter (1996) and Lin (2013).

Our setting is useful as it allows us to distinguish between different methods of technology transfer – physical proximity, organizational affiliation, and key personnel. We observe monthly production data for 116 mines from 1887 to 1907. For each mine we compile data on the tonnes of ore mined, the ounces of gold extracted by traditional and cyanide processes, and key personnel, both operational and administrative. We observe the distance, as the crow flies, between any pair of mines. If physical proximity is critical, then closer mines should experience faster technology transfers than more distant mines. Our study of physical proximity in mining allows us to avoid the problem of agglomeration effects. Gold mines did not locate next to each other to enjoy access to a "cheaper and faster supply of intermediate goods and services, proximity to workers or consumers, or a better quality of worker-firm matches" (see Greenstone et al., 2010, p. 537). Instead gold mines located where profitable ore was discovered. As each tonne of ore contained only a few ounces of gold, treatment of the ore

[1] Commonly known as the 'Cyanide Club', see Pogue (2006) p. 82

occurred at the surface above the mine. Therefore, firm distance is effectively exogenous with respect to the level of technology employed by the mine.[2]

We observe key personnel: mine managers, engineers, directors, and company secretaries. Part of the managers' and engineers' roles were to improve the technical efficiency of extracting gold from its ore and to reduce the costs of doing so. We find strong evidence that 'hands-on' personnel successfully transferred knowledge of the cyanide process between mines. In contrast, we find no evidence that white-collar personnel were effective, by themselves, in transferring technical knowledge in contrast to the results of Baum et al. (2022). However, entrepreneurs and directors were critical in organizing mines within a mining house. Therefore organizations were indirectly responsible, via appointing common engineers and managers to multiple mines at the same time, for the spread of innovation.

Our results have implications for today's world. Bloom et al. (2013b) highlight the importance and difficulty of distinguishing technology spillovers from the countervailing market rivalry effect in economic studies. We illustrate how technology transfers operated when there were few reasons for a firm to keep knowledge secret (i.e., there were no market-rivalry concerns) without the need for assumptions or proxies. In addition, our results shed light on contemporary knowledge spillovers in areas with public good characteristics, such as open-source code, data, and research articles.

The four mechanisms we examine remain the dominant channels for knowledge spillovers. Agglomerations, conferences, and research joint ventures exist to reduce barriers posed by institutional boundaries, scattered human capital, geographic distance, and national borders. The organizational features that characterized the South African mining industry, for example, venture capital, extensive cross-ownership, and interlocking directorates, are still prevalent in the modern economy. How knowledge and innovation spillovers operate within networks is a key focus of recent finance and economic studies (see e.g., He

[2]Firms' locations were effectively fixed, however, variation in the distance to competitors was determined at the extensive margin. Unproductive mines could shut down and productive mines could open.

and Huang, 2017; and González-Uribe, 2020). Finally, innovation today is also complex and requires specific knowledge that generalist company directors cannot fully understand. Therefore, one of the channels we study, blue-collar vs white-collar, is a perennial issue.

4.2 Background

4.2.1 The Witwatersrand Goldfields

In 1886 gold was discovered in the Witwatersrand hills that ran east to west in the south of the Transvaal Republic. As the extent of the gold deposits became clear there was a rapid influx of people to the area. Small-scale diggings of individual miners soon gave way to organized companies. Mining capital mostly came from London with smaller amounts contributed by French and German investors.

The Transvaal Chamber of Mines was formed in October 1887 in the newly settled town of Johannesburg. The Chamber published monthly reports on the operations of its members. It was reformed under various names but continued in its basic mission of advocating the interests of the mine owners and sharing information between mines. A local stock exchange and telegraph also began operations in 1887. Three professional bodies were set up in Johannesburg in the early 1890s. The Association of Mine Managers, the South African Institution of Engineers, and the Chemical, Metallurgical and Mining Society of South Africa all aided the spread of technical and chemical knowledge.[3]

Mining companies were separate legal entities with their own shareholders and boards of directors. Most of the larger mines were floated by one of the nine mining houses (see Table A1 in the Online Appendix). The mining houses held large, not necessarily majority, stakes in the mining companies. The mining houses provided directors for the operating companies, organized loans, and

[3] See Pogue (2006).

supplied secretarial and engineering expertise. Smaller mines tended to be independently owned and operated.

The earlier mines began to run out of ore and shut down as their reserves were exhausted, starting in the early 1890s. As time passed, the bulk of gold production gradually shifted to the 'deep level' mines. By the end of our period, mines reached almost one kilometre below the surface. When the Boer War began in 1899 the mining industry in Johannesburg shut down. Operations were only restarted from mid-1901 onwards. By 1908 mines in the Transvaal, most of which were in the Witwatersrand area, were producing over 7 million ounces of gold per year which amounted around 35% of the world's production (see 1909 volume of *The Mining Manual*).

4.2.2 Metallurgy of Gold and the Cyanide Process

4.2.2.1 The Mercury Amalgamation Process

In 1886 the conventional method to extract gold from its ore was the mercury amalgamation process. Gold-bearing ores were first crushed and pulverised by massive coal-powered stamp batteries. The crushed ores were then reacted with mercury to form a gold-mercury amalgam. Together with unamalgamated materials the amalgam was sluiced over corrugated copper plates that were covered with mercury.[4] The amalgam adhered to the copper plates. Finally, the amalgam on the copper plates was heated to a high temperature which made the mercury vaporise, leaving behind the gold.

Particles that failed to form an amalgam, called tailings, also contained a considerable amount of gold, estimated to be around 45% of the initial amount (see Gray and McLachlan, 1933). The tailings were dumped in giant heaps or dams on the Rand because knowledge of an appropriate treatment was lacking in the 1880s (see Eissler, 1888).

[4] Some mines used mercury in the stamp batteries whereas others added mercury to the copper plates.

Mercury amalgamation was less effective for ore buried at deep levels since deep-level ore was unoxidised and contained a large amount of pyrites (iron sulphides). Gold output from the use of mercury decreased as mines delved deeper and deeper. Various attempts were made to overcome this issue. For example, unoxidised ore was ground more finely, or the ground ore was passed through a chlorine solution (Lougheed, 1989). Although chlorination was chemically effective it was quite expensive and therefore not used by mines with lower quality ore.

4.2.2.2 The Cyanide Process

The solubility of gold in a cyanide solution has been known since the Middle Ages. However, cyanide was not a profitable extraction method when it was first discovered (see Eissler (1895)). In the cyanide process, gold (Au) and the cyanide anion (CN-) in potassium cyanide (KCN) form a soluble compound: potassium dicyanoaurate (KAu(CN)2). The potassium dicyanoaurate solution is easily separated from other substances through filter vats and is then placed in a precipitation box containing zinc. Gold detaches from potassium dicyanoaurate in the precipitation box and falls to the bottom.

The cyanide process was developed in Scotland by John MacArthur and the Forrest brothers. The process was patented in the U.K. in 1887 and in the Transvaal Republic in September 1888. The cyanide process treated tailings after mercury amalgamation. Ores were treated by mercury amalgamation before they were sent for cyanide treatment. Some mines concentrated tailings before cyanide treatment. Concentrates, which had a very high percentage of gold, were treated separately with cyanide. The remaining material was divided into sands or slimes based on the size of the particles. Sands were coarser particles and slimes were finer particles. The cyanide process was first used to treat sands. It was only in 1897 that a treatment for slimes was pioneered (see Eissler, 1895).

The first commercial application of the cyanide process took place on the Witwatersrand goldfields. The first cyanide plant was built at the Salisbury mine for demonstration purposes in 1890. In 1891 the Robinson mine first used cyanide for operational purposes.[5]

Royalties had to be paid to use the MacArthur-Forrest process until November 1896, when the patent was found to be invalid by the High Court of the Transvaal Republic (Gray and McLachlan, 1933). Mining companies had sued the patent holder, the African Gold Recovery Company, alleging that knowledge of treating gold ore with cyanide had been known well before the MacArthur-Forrest patent was granted.

Opinions on the novelty of the cyanide technique varied. The *South African Mining Journal* opined on December 23, 1893 that,"there can be no question that the MacArthur-Forrest process has been of the utmost service to the Witwatersrand mining industry". The *Engineering and Mining Journal* wrote in September 1894 that the MacArthur-Forrest process was successful due to the innovative use of zinc shavings to precipitate the gold following treatment with cyanide. Mining companies were forced to pay an average of 7.5 per cent of the gold extracted to the African Gold Recovery Company (see Gray and McLachlan, 1933).

The adoption of cyanide on the Rand increased gold production by roughly 50 per cent in the first month of its use. However, it would be wrong to view cyanide as a single innovation. After the first use of cyanide there was continued development of 'microinventions' (see Mokyr, 1990). Mann (1909) notes improvements such as adding halogen salts (1892), agitation and filter pressing (1893), electric precipitation (1894), roasting of sulpha-telluride ores (1895), the addition of soluble lead (1896), tube milling (1898), vacuum filters (1903), and automatic suction filters (1905). In additional, the benefit of using oxygen in the cyanide process was not discovered by practitioners for years after the MacArthur-Forrest patent was granted. The strength of the potassium cyanide solution also had to be optimized. The initial strength ranged from 0.25 to 1.0 per cent, sometimes as high as 2.0 percent. After experimentation it was found

[5]A number of firms were formed to process mines' ore with cyanide at the mine site. The African Gold Recovery Company, the Transvaal Chemical Company, Rand Central Ore Reduction Company, and the Robinson Custom Works were the most prominent.

that a more dilute cyanide solution, of around 0.2 per cent, extracted the most gold (Gray and McLachlan, 1933).

The constantly improving cyanide technology facilitated the expansion of the gold mining industry in South Africa. The cyanide process remains a mainstay of the global gold mining sector (see Verbrugge et al., 2021). The slow but continual productivity improvements we demonstrate in mining have also been documented in the mechanization of the French textile industry during the 19th century see Juhász et al. (2023).

4.3 Theory, Data, and Testable Implication

4.3.1 Theoretical Model of Information Diffusion

We use the theoretical framework of technology diffusion by McCardle (1985) and Lippman and McCardle (1987) to guide our examination of the data. A new technology, with value p, is exogenously introduced. We interpret a new "technology" to be either the initial introduction of cyanide treatment, a 'macro invention', or the various 'micro inventions' (see Mokyr, 1990) that improved the initial performance. Firms do not directly observe the true value of the technology, p^*, and possess a prior belief, p, which is distributed $N(\mu, 1/s)$. In each period n, a firm incurs a positive cost c to gather information about p^* and receives a sequence of signals from independent and identically distributed random variables, $X_n \sim N(p^*, 1/t)$. Firms update their beliefs about p given the signals in a Bayesian fashion. Parameters s and t reflect the precision of the prior belief and the new information, respectively. The payoff of adopting the new technology is linear in p:

$$\pi(p) = -K + Ap$$

where $K > 0$ is a fixed set-up cost and $A > 0$ is the return rate. In each period, firms choose whether to adopt the new technology given their beliefs. The adoption decision is a dynamic programming problem with the following value function:

$$V_n(p) = Max\{0,\ \pi(p_n) - c,\ \beta V_{n+1}(p_{n+1})\}$$

The three terms represent the payoffs from rejecting, adopting, and waiting for another round.[6] p_n is the expected posterior belief at time n given $X_1, X_2, \ldots, X_n$, and β is the one-period discount factor. The optimal solution is that the firm will adopt the technology if the expected value of p is above a threshold.

Lippman and McCardle (1987) show that the expected time until the adoption decision is decreasing in the precision of new information, all else equal. Therefore, we expect the adoption decision of cyanide to be made earlier if a firm receives more precise information about its value. In each time period, the receiving (focal) mine aggregates data from information senders (other mines). Each receiving mine has different connections to sending mines.

Initial Adoption Decision – Mines have more precise information about the true value of p after they adopt cyanide.[7] A mine can only collect information from another mine to which it is connected in some way. The more peers that a focal mine is connected to increases the precision of the mine's information, all else equal.[8] Thus,

Hypothesis 1: A mine that is connected to more mines that have already adopted cyanide is more likely to adopt the cyanide process itself.

Alternative to Hypothesis 1: Mines that are connected to more mines that have already adopted cyanide are no more likely to adopt it.

[6]This optimal stopping model is can also be framed as a real option where a agent decides whether or not to invest in each period (see e.g., Majd and Pindyck, 1987).

[7]We assume that the true value of the cyanide method to the adopting firm is slowly revealed following Jensen (1988).

[8]The aggregation of information can also be viewed as a Bayesian updating process. If the precision of a single signal is t_0, the aggregated ("posterior") signal from n different sources is nt_0.

Technological Improvements – Following the initial adoption, mines can conduct trials to improve the new technology. A mine that conducts successful trials and implements micro-inventions experiences a positive productivity shock and has more precise information about the value of the micro-invention. Therefore, mines that receive more precise signals about micro-inventions are more likely to adopt them (see Kaustia and Knüpfer, 2012). We infer whether a mine has implemented a micro-invention from its productivity.

Hypothesis 2: A mine's productivity increases if connected mines experience positive productivity shocks.

Alternative to Hypothesis 2: A mine's productivity is unaffected if connected mines experience positive productivity shocks.

Catching Up by Laggards – The catch-up hypothesis of Abramovitz (1986) contends that less technologically advanced economies can experience faster growth than more advanced ones if their social capabilities allow them to effectively utilize existing technologies pioneered by technological leaders.

We expect mines that are 'lagging' in cyanide technology experience higher productivity growth after receiving technology signals than technologically 'leading' mines. Other temporally or spatially correlated productivity shocks are unlikely to have such differential impacts on leaders versus laggards.

Hypothesis 3: A mine with below-average cyanide productivity experiences larger productivity gains than a mine with above-average cyanide productivity if connected mines experience positive productivity shocks.

Alternative to Hypothesis 3: Below-average and above-average cyanide productivity mines experience similar productivity gains due to connected mines' productivity shocks.

4.3.2 Data

We hand collect mine-level production data for all mines in the Transvaal over the period 1887-1907 from various sources.[9] The Chamber of Mines published annual reports, which include monthly production data for each active mine in the Transvaal. We augment our data with *The Mining Manual* and *The Economist*. We observe the monthly tons of ore processed by mercury amalgamation (the primary, or non-cyanide, process) and the ounces of gold extracted. We also observe the monthly tons of sands, slimes, and concentrates processed (with the secondary, or cyanide, process) and the ounces of gold extracted by each process.

The locations of mines come from maps of the goldfields that depict each mine's surface area. We identify the approximate centroid of each mine and use the coordinates of the centroid as the mine's location. Most mines are located on the Witwatersrand, which runs for roughly 60-kilometres east to west near Johannesburg. We collect company data from *The Mining Manual*. This periodical reports the details of all companies in our sample. We collect the names of directors, company secretaries, managers, and mining engineers whenever available. We supplement *The Mining Manual* with company reports published in *The Financial Times*. We collect a mine's affiliation with a mining house from public reports, such as by the Witwatersrand Native Labour Association, companies' annual reports, and newspaper commentary.

We drop mine-month observations with missing production data. For each mine, the first and last observations in each consecutive production period are also dropped because production in these months often includes irregular activities (such as trial milling, setting up plates, and final clean-ups) that lead to extreme values. Production data are unavailable from November 1899 to April 1901 as mining activity was interrupted by the Second Boer War. Mines started resuming work in May 1901, yet production activity did not regain the pre-war level until the end of 1901. For each extraction process in a mine-month, we drop the

[9]The region we study was politically part of the Transvaal Republic in 1887. Following the Second Boer War it became the Transvaal colony in 1902 before joining the Union of South Africa in 1910.

observation unless both the ounces and tonnage of that process were reported. Last, we require mines to have reported at least 12 months of production activity after the above restrictions. The final sample contains 9659 mine-month observations, of which 7589 observations have both primary and secondary process data.

4.3.3 Cyanide Productivity

Our outcome variable, cyanide productivity ($Productivity^{Cyanide}$), measures the percentage of gold in the ore that can be extracted through cyanidation. In our baseline analysis, we focus on the largest cyanide process, the treatment of sands. We examine the productivity of slimes and concentrates in the robustness section. Since the amount of gold in its ore is not observable, we make three simplifying assumptions: (1) the percentage of gold in its ore that is recovered by mercury amalgamation is constant, (2) the proportion of ore lost after the amalgamation process, but before the cyanide process, is constant, and (3) a constant proportion of gold is sorted into sands.

If these three assumptions hold, true (but unobservable) cyanide productivity is proportional to the ounces of gold extracted via cyanidation divided by the ounces of gold extracted by mercury amalgamation (see Online Appendix for details.):

$$Productivity_{i,t}^{Cyanide} \propto \frac{Ozs_{i,t}^{Cyanide}}{Ozs_{i,t}^{Primary}}. \tag{4.1}$$

Intuitively, our measure of cyanide productivity controls for the (unobserved) amount of gold present in the ore. We use the amount of gold extracted via mercury amalgamation as a proxy for the amount of gold in the ore. Although we also use mine fixed effects, these will not be fully adequate, since the amount of gold per unit of ore tended to decline as mines became deeper. Thus, we also control for mine depth in the regression analysis.

We believe our assumptions are plausible. Mercury amalgamation was an established technology that has been around since the Middle Ages. Advances in mercury amalgamation technology are unlikely to have taken place, or are at least of second order relative to improvements in cyanide technology. However, one might worry that the chemical composition of ores might affect the recovery rate from mercury amalgamation. Since 97% of our observations are located in a 60-kilometre range, the chemical property of ores should be similar. To further alleviate this concern, we estimate spillover effects conditioning on mine fixed effects and the cumulative extraction tonnage and its squared term. Cumulative extraction acts as a proxy for mine depth.

As for the second assumption, tailings were usually sent directly to the cyanide plant after amalgamation. The amount of ore lost between the two processes was limited. The amount of ore washed or blown away was probably small, and as long as the amount lost was random, it will not cause bias in the estimation.

We examine the importance of our third assumption empirically. Our results are consistent if we relax the last assumption by aggregating the sands, slimes, and concentrates processes to measure cyanide productivity. This alternative (total) cyanide productivity is: $Productivity^{Cyanide} = \frac{Ozs^{Sands+Slimes+Conc}}{Ozs^{in\ tailings*}} \propto \frac{Ozs^{Sands+Slimes+Conc}}{Ozs^{Primary}}$.[10] The alternative measure does not require assumption (3). We use this alternative measure in the robustness section and the results are consistent with the baseline analysis.

A final concern is that mines would sometimes store tailings in a heap or a dam for later cyanide treatment rather than processing them immediately. Such behavior would lead to a timing issue in the cyanide productivity measure: tailings treated at time t are not mined at time t but could also be from earlier periods. Therefore, time t's mercury extraction productivity may be an inappropriate control for time t's cyanide productivity. To alleviate this concern, we review all remarks in the Chamber of Mine's annual reports and remove the tonnage and ounces reported as coming from accumulated tailings from the original data whenever possible. Figure B1 in the Online Appendix indicates that the tonnage

[10] $tailings*$ denotes the true, but unobserved, ounces of gold in tailings.

treated by cyanide was approximately equal to the tonnage treated by mercury amalgamation. Losses due to ore being washed or blown away appear to be small.

4.3.4 Cyanide Knowledge Pool

Our key explanatory variable, $KnowledgePool(m)_{i,t}$, is an approximation of the cyanide knowledge that can be learned by focal mine i at time t through mechanism m. Individual mines' cyanide knowledge is not directly observable. A mine's cyanide productivity per se is not a suitable proxy for cyanide knowledge because productivity may be sensitive to mine location, ore quality, or other factors. Instead, we consider the deviation of actual productivity from a regression based measure of expected cyanide productivity. This difference indicates a mine's above- or below-average ability to use the cyanide method. We estimate cyanide productivity with the following regression:

$$Productivity^{Cyanide}_{j,t} = \alpha_j + \gamma \mathcal{X}_{j,t} + \delta_t + \epsilon_{j,t}$$

where α_j is a mine fixed effect that controls for ore quality due to time-invariant mine characteristics like location.[11] $\mathcal{X}_{j,t}$ is a vector of control variables, including the cumulative extraction and its squared term, indicators for irregular mining activities like the concentrates process, the slimes process, mine stoppages, cleanups, by-products, business connections, and use of the tailings dump. δ_t is a time fixed effect that captures productivity shocks common to all mines at

[11] We include all available observations to estimate the cyanide method production function in the baseline analysis. This approach is similar to the regression-based estimation of total factor productivity (see, for example Imrohoroğlu and Tüzel (2014); Olley and Pakes (1996)). Despite the generality of this approach, using full-sample information is likely to incur issues of reverse causality. Therefore, we also predict the out-of-sample cyanide productivity at time t based on a model only using information available at $t-1$. In the out-of-sample prediction, we assume the time fixed effects are the same in t and t-1. The results based on the rolling estimation are quantitatively similar to the baseline result.

time t. The unexplained residual, $\epsilon_{j,t}$, is the mine-specific above/below average cyanide productivity which we use as our measure of cyanide knowledge.[12] We assume that only productivity above the expected level contains information that can induce learning, following Conley and Udry (2010). Therefore an individual mine's cyanide knowledge, $K_{j,t}$, is the positive residual productivity: $max\{0, \hat{\epsilon_{j,t}}\}$.[13]

From individual mines' cyanide knowledge, we use a connection intensity-weighted average across all connected mines as the Knowledge Pool measure available to mine i at time t:

$$Knowledge\ Pool^m_{i,t} = \sum_{c^m_{ijt}=1} g^m_{ijt} \times K_{j,t}$$

where c^m_{ijt} is a binary variable equal to 1 if mine i and j are connected through mechanism m at time t and g^m_{ijt} is a variable that ranges from 0 to 1 that indicates the connection intensity.[14]

We construct five knowledge pools corresponding to the five potential spillover mechanisms under examination: organizational, personnel, geographical, French, and German connections.

[12]Our approach is similar to Jenter and Kanaan (2015) who adopt an industry index model to decompose a firm's stock return into a systematic component and a firm-specific component, with the firm-specific component capturing, among other things, CEO ability. In our setting, we consider the residual productivity captures a mine's ability to apply the cyanide method. Clearly, the residual may also represent model misspecification and/or white noise. Estimation results are reported in Table A2 in the online Appendix.

[13]By setting the negative residuals to zero, we assume that mines do not learn from other mines that experienced worse than expected performance. Effectively we assume that mine's share their successful techniques with others but not their unsuccessful techniques. We relax this assumption by adopting two alternative transformations of the residual productivity: (1) exponential transformation: $K_{j,t} = exp(\hat{\epsilon_{j,t}})$, and (2) parallel shift: $K_{j,t} = \hat{\epsilon_{j,t}} + min(\hat{\epsilon_{j,t}})$. The results based on these alternative transformations are quantitatively similar to the baseline results.

[14]An unweighted knowledge pool is used in the robustness section, where $g^m_{ijt} = 1$ for all mechanisms.

4.3.4.1 Organizational Knowledge Pool

We use mining house affiliation to study the role of organizations in the diffusion of cyanide knowledge. The connection dummy, c_{ij}, equals one if two mines were members of the same mining house.[15] We assume that each mine is equally likely to learn from all other affiliated mines; thus, the connection intensity is the same for all connected mine-pairs: $g_{ij}^{Org} = 1$. The organizational knowledge pool is a summation of the cyanide knowledge of other mines in the same mining house:

$$Knowledge\ Pool_{i,t}^{Org} = \sum_{i \neq j,\, c_{ij}^{Org}=1} K_{j,t}$$

For an independent mine i that was not affiliated to a mining house, the connection dummy c_{ij}^{Org} equals zero for all j and therefore $Knowledge\ Pool_{i,t}^{Org}$ also equals zero.

4.3.4.2 Personnel Knowledge Pool

We study mine personnel interlocks to examine the role of individual networks in the diffusion of the cyanide technology. We examine the role of operational personnel (engineers and mine managers) in the baseline analysis and explore the role of white-collar personnel (company directors and secretaries) in the robustness section. We follow Zacchia (2020) to construct the knowledge pool measure through contemporary connected personnel.[16] Mine i and j are connected through a manager interlock at time t if they had any common managers at time t. The connection intensity is an increasing function of the proportion of overlapping managers in two mines: $g_{ijt}^{Manager} = f(\frac{2 \times \#CommonManager}{\#Manager_{it} + \#Manager_{jt}})$. The knowledge pool through manager interlocks is a weighted sum of cyanide knowledge of interlocked mines:

[15] There are no instances of mines changing their mining house.

[16] In Zacchia (2020) a connection exists between two inventors if they have collaborated on past patents. We use personnel who were common to multiple mines.

$$Knowledge\ Pool_{i,t}^{Manager} = \sum_{i \neq j} f(\frac{2 \times \#CommonManager}{\#Manager_{it} + \#Manager_{jt}}) \times K_{j,t}$$

where $f : [0, 1] \rightarrow [0, 1]$, $f(0) = 0$ and $f(1) = 1$ following Zacchia (2020). We apply the square-root function to the proportion of common manager in the baseline result.[17] In addition to common managers in the contemporary year, we also consider connections through a manager's past work experience. A historical connection between mines i and j exists if i's manager worked at mine j in the previous three years but not in the current year. Similarly, in the robustness section we construct two additional measures, $Knowledge\ Pool_{i,t}^{Director}$ and $Knowledge\ Pool_{i,t}^{Secretary}$, to investigate whether white-collar personnel facilitate knowledge transfer.

4.3.4.3 Geographical Knowledge Pool

The geographical knowledge pools depend on the physical distance between mines. The connection dummy c_{ijt}^{Geo} equals 1 if the distance between two mines d_{ijt} was no greater than D kilometres. We assume that a mine is more likely to learn from closer mines than more distant mines. We define the connection intensity as a monotonic function of the inverse distance between two mines: $g_{ijt}^{Geo} = f(1/(1 + d_{ijt}))$. Therefore, the geographical knowledge pool equals:

$$Knowledge\ Pool_{i,t}^{Geo} = \sum_{i \neq j,\ d_{ijt} \leq Dkm} f(\frac{1}{1 + d_{ijt}}) \times K_{j,t}$$

Similarly to the personnel knowledge pool we apply the square-root function to the connection intensity in the baseline analysis and show that the results are robust to alternative functional forms.

[17] We show in the robustness section that the empirical results are quantitatively similar base on alternative choices of function f, including (1) linear transformation: $f(x) = x$, (2) polynomial transformation: $f(x) = -x^2 + 2x$, and (3) square transformation: $f(x) = x^2$.

4.3.4.4 French and German Knowledge Pool

We define a mine as having a French and/or German connection if there was an office for that mine anywhere in France and/or Germany. Two mines are connected if they both had French(German) offices at time t.[18] The connection intensity equals one for all connected mines. Therefore, the knowledge pool through French(German) connections is an unweighted sum of cyanide knowledge of all other mines that also have French(German) offices:

$$Knowledge\ Pool_{i,t}^{French(German)} = \sum_{i \neq j} French(German)\ Office \times K_{j,t}.$$

For mines with no French and / or German connection, $Knowledge\ Pool_{i,t}^{French}$ and/or $Knowledge\ Pool_{i,t}^{German}$ equals zero.

4.4 Results

4.4.1 Summary Statistics

In Table 4.1, we present summary statistics for the 116 mines in our sample. An average mine processed 9.6 thousand tons of ore using the primary (mercury) amalgamation method and 8.1 thousand tons of sands using the cyanide process per month. The gold recovery rate, defined as the ounces of gold per ton of ore treated, averaged 0.41 for the primary process and 0.15 for the sands process. 77% of our mine-months involved the use of cyanide to treat sands, 33% used cyanide to treat slimes, and only 20% used cyanide to treat concentrates.

The average mine had a cumulative extraction of 0.48 million tonnes of ore. 9% of mines had had a business connection with an external cyanide processor

[18] South African shareholders were primarily British, but with large numbers of French and German investors. Most companies had two corporate offices, one in Johannesburg and the second in Europe. Companies with mostly British investors had a London office, mines with many French investors had an office in Paris, and a few mines with German shareholders had Berlin offices. Knowledge transfer between different mines may have been facilitated by a common ethnic origin of the shareholders (we use the presence of a corporate office in London, Paris or Berlin as our fourth knowledge pool).

in the previous year.[19] Most mines were members of one of the nine mining houses, the remaining 32% were independent mines. Figure 4.1 depicts mine locations by mining house on a modern-day map. 96% of all mines were in the Witwatersrand region of the Transvaal. The map shows that despite some location clustering, geographical variation exists for most mining houses.

The mean mine had almost 8 directors, 1.68 managers or engineers, 1.81 secretaries, and 2.16 management boards. Most mines had two official boards of directors, one in London and the other in South Africa. Some also had a French or German 'committee' to represent their shareholders' interests. 41% of all mine-months had a French connection (i.e., an office in Paris). This connection could be an individual at an office address, a bank, or sometimes a representative company, with locations scattered around Paris. It was rare for there to be a formal French committee. No French or German offices were reported before 1895. The proportion of mines with a French office grew rapidly in later years. 9% of mine-months had a German connection. In contrast to French connections, all German-connected firms used the same agent, "Deutsche Treuhand Gesellschaft" (German Trust Company), and had a formal committee of three members for all mines that they represented.

Table A3 in the Online Appendix reports the proportion of connected mine-pairs over time: about 7% of mine-pairs were connected through a common mining house, 10% were connected through geographical proximity if we set $D = 2.5km$. These two networks were relatively stable over time. Mines with manager interlocks were less than 1% of the sample in the early 1890s, but this number gradually grew to 7% by 1907.

Figure 4.2 is a visual representation of active mines over time. Mines might temporarily shut down for several months or completely shut down for good. The number of active mines and cyanide-using mines each year is plotted in Figure 4.3. Approximately 60 mines were active each year, excluding the Boer War period. The number of cyanide-using mines increased rapidly after the

[19] Such as the African Gold Recovery Company, the Transvaal Chemical Company etc.

introduction of the cyanide process in 1890. After 1895, there were very few mines not using cyanide.

In Figure 4.4 we show how the cyanide process contributed to total gold output over time with output data aggregated at the annual level. Mines using cyanide increased their gold output by approximately 50%, primarily due to gold recovered from sands. The proportion of gold produced by the cyanide process increased gradually during our sample period. We also observe an increase in the proportion of gold recovered from slimes, after the slimes technique debuted in 1898.

4.4.2 The Adoption of the Cyanide Process

We first examine how technology diffuses at the extensive margin. We model mines' decisions to adopt the sands process in the early 1890s and the slimes process in the late 1890s. Tailings were sorted into sands and slimes, and the first use of cyanide was to treat sands. A micro-invention was developed around 1897, making the cyanide treatment of slimes feasible. We model the adoption of the sands and slimes processes, given a mine's knowledge of cyanide from its neighbors, using both ordinary least squares (OLS) and a hazard model. The knowledge pools depend upon neighboring mines' use, or not, of the sands and slimes cyanide processes.

In Table 4.2 we show the technology spillover effects of the adoption of the sands and slimes processes. The adoption decisions were affected by the number of other mines in the same mining house that had adopted that process. We find that one extra cyanide-using mine in the same mining house increased the probability of sands cyanide adoption by 2 percentage points per month (column 1), and this positive association is confirmed by the hazard model (column 2). Mining houses were the instigators of cyanide adoption, physical proximity was irrelevant. Managerial/engineering connections were not important for the adoption of cyanide, but in the pre-1895 period less than one percent of mines had common operational personnel (see Table A3 in the Online Appendix).[20]

For the slimes process adoption we find that the adoption decisions were positively affected by the slimes cyanide usage of mines with the same manager. We find that one extra slimes cyanide-using mine with a common manager increased the probability of slimes adoption by 5.4 percentage points per month (column 3), and the hazard model finds qualitatively similar results. Although mining house affiliation alone does not matter, mine managers and engineers could be assigned to multiple mines within the mining house to facilitate technology transfer. In the early 1890s, manager interlocks were less common than in the late 1890s. We also find evidence that knowledge of the slimes process was propagated through German connections but not French connections. One extra German-connected mine that used slimes cyanide increased the probability that a German mine would adopt the slimes process by 3.8 percentage points per month (column 3), and the hazard model also suggests significant technology spillovers among German-connected mines (column 4).

4.4.3 Cyanide Process and Productivity Gain

We now examine the impact of the cyanide process on the gold recovery rate. Figure 4.5 shows that the total recovery rate (the mercury amalgamation recovery rate plus the cyanide recovery rate) increased substantially immediately after using cyanide. In contrast, the primary recovery rate remained relatively stable after the use of cyanide. This suggests the gold output from cyanide is an additional output, rather than gold output being substituted from mercury to the cyanide method.

Table 4.3 regresses the gold recovery rate on an indicator variable that indicates whether the mine has adopted cyanide. Columns (1) to (3) show that an additional 0.26-0.28 ounces of gold were recovered from each ton of ore treated by

[20] We ignore French and German connections for the sands process adoption as very few mines had those connections before 1895.

cyanide. This increase is more than 50% of the median total recovery rate one year prior to the adoption (median = 0.52 and mean = 0.61 in the [-12, -1] adoption window). However, selection issues may be important. The increase in the total recovery rate could be driven by a mine's decision to adopt cyanide when it reached high quality ore veins. Therefore, we test whether the primary recovery rate was affected by the adoption of cyanide. Columns (4) to (6) show no significant difference in the primary recovery rate after the adoption of the cyanide process. Therefore, the productivity gain was mainly obtained from cyanidation and selection effects appear slight.

4.4.4 The Diffusion of Cyanide knowledge

Table 4.4 presents our main results on cyanide knowledge spillovers. We test whether an increase in other mines' cyanide productivity affected a mine's productivity through the regression:

$$Productivity_{i,t}^{Cyanide} = \sum_{m \in \mathcal{M}} \beta_m Knowledge\ Pool_{i,t-1}^{m} + \gamma \mathcal{X}_{i,t} + \alpha_i + \delta_t + v_{it} \quad (4.2)$$

where $\mathcal{M}$ is the set of the five mechanisms of interest and $\mathcal{X}_{i,t}$ is a vector of control variables, including the (quadratic form of the) cumulative extraction of ore that captures the variation in ore quality due to mine depth. We also condition on lagged own cyanide knowledge, the use of by-products, slimes process, and concentrates processes, business connections, mine stoppages, mine clean-ups, the use of cyanide on accumulated tailings, and French/German connections. We use mine fixed effects α_i to capture differences in (unobservable) ore type driven by mine location.[21] Month-year fixed effect δ_t are used to control for common productivity shocks in each month, such as floods, droughts, and labor shortages.

[21] Differences could be the amount of gold in the ore or the chemical composition of the ore. Some ore types were more amenable to cyanide usage than others.

Columns (1) to (5) test the five mechanisms separately. We find strong knowledge spillovers through organizations (mining houses), manager interlocks, geographical proximity (when the distance parameter D is set to 2.5 kilometres), and a shared German connection. We discover little evidence that French connections facilitated knowledge transfer. The different results for French and German connections are likely due to French connections being weaker than German ones. All German connections indicate a shared office in Berlin with a dedicated committee to represent German shareholders. In contrast, French connections were more tenuous, there were many agents and offices for South African mines in Paris. These French agents may not have been in frequent contact with each other.

In the single knowledge pool regressions the coefficients for a particular knowledge pool may be confounded by other knowledge pools. For example, some mines with a common manager are located near each other. Therefore, the manager knowledge pool may also partly capture geographical spillovers if they exist. In column (6) we include all five spillover measures and test their joint significance. After controlling for other confounding knowledge spillover mechanisms, the coefficient of the organizational knowledge pool becomes indistinguishable from zero. However, all other mechanisms remain statistically significant and of a similar magnitude. This finding does not suggest that mining houses played no role in facilitating knowledge spillovers. Instead, mining houses enabled the diffusion of new practices and procedures among affiliated mines by appointing a common manager, or consulting engineer, to oversee multiple mines. The joint significance test of the coefficients on these five knowledge pool variables has an F statistic of 8.79; hence, at least one of these knowledge pools is essential for knowledge transfer.

We now quantify the economic magnitude of these estimated coefficients. In column (6), all else equal, a one standard deviation increase in the external managerial/geographical/German knowledge pool is associated with an 0.011/0.013/0.008 increase in cyanide productivity. If we consider the estimate of 0.013 productivity spillover for a median mine in the cyanide-using sample, it is equivalent to a

3% increase in the cyanide productivity (median = 0.43, mean = 0.48) and a 0.9% increase in the total gold recovery rate. These figures relate to approximately an extra 36 ounces of gold that could be recovered per month. The spillover value of this incremental monthly output is nearly $70,000.[22]

4.4.4.1 The Role of Personnel and Common ownership

We examine how spillovers depend upon a mine's personnel in Table 4.5. We separate personnel into white-collar (directors and secretaries) and blue-collar (mine managers and engineers). Technical knowledge might have been shared between mines by white-collar workers at board meetings or via written reports. In Table 4.5 we find that neither current director interlocks nor a common secretary were effective in facilitating technology transfer. Mine managers and engineers were involved in the daily mining operations, and this hands-on experience appears to have been necessary to transfer cyanide expertise.

Our work also touches on the role of common ownership and innovation (see Vives 2020; Anton et al 2021; Schmalz 2018). Mines within the same mining house had high levels of common ownership, with the mining house itself providing much of the equity capital. Once we condition on geography, operational personnel, and French/German connections, we find no effect of common ownership (i.e., Knowledge Pool (Org)) on knowledge spillovers (see Table 5, column 4). Company directors are also ineffective in disseminating knowledge between mines.

4.4.5 Robustness to Alternative Explanations

4.4.5.1 Market Rivalry Effect

One issue in identifying knowledge spillovers is the potential confounding of technology spillovers with market rivalry effects (Arora et al., 2021; Bloom et al.,

[22] This amount is calculated using the gold price on 17/03/2022 at 2022 dollar value.

2013b). This issue is less of a concern in our setting because most of the world, and Great Britain in particular, was on the gold standard (Eichengreen and Flandreau, 1985). The Bank of England had been required by law, since 1821, to exchange a fixed quantity of gold for bank notes. Hence, any increase in total gold output, via knowledge spillovers for example, could not adversely affect the market price or the quantity of gold demanded.

4.4.5.2 Endogenous Manager Assignment

A concern about our interpretation of managerial spillovers is that a common manager might be appointed to two or more mines when those mines have unusually high, or low, cyanide productivity. Such a selection effect might wrongly appear to be a technology spillover. To partially address this concern we examine historical managerial interlocks based on our personnel data. A historical interlock exists if a mine manager/engineer worked at another mine in the previous three years but not in the current year. If selection effects are present then historical manager interlocks may still be associated with our measure of spillovers. However, if personnel are useful for transferring state-of-the-art knowledge of cyanide usage, then historical connections are unlikely to be useful. Table 4.5 shows that historical managerial interlocks have statistically insignificant and near zero effects on spillovers. This test shows no evidence of selection effects.

4.4.5.3 Spatially Correlated Productivity Shocks

We face limited concerns about reverse causality since we test how a reference group's past productivity affects the current productivity of a mine. Moreover, we control for lagged own above expectations productivity. Our specification partially addresses concerns about unobservable factors that might drive simultaneous productivity shocks across mines, such as a mining area hitting richer or poorer veins of ore. However, one might still worry that the observed positive correlation between productivity and neighbouring mines' lagged productivity surprises could be driven by both spatially and serially correlated unobservable productivity shocks, such as water shortages or floods. We cannot rule out this possibility. Nonetheless, we turn to a battery of robustness checks to demonstrate that spatially correlated unobservable factors do not solely drive the positive correlation of productivity shocks between neighbouring mines.

In Table 4.6, we compare geographical knowledge spillovers between cyanide technology leaders and laggards. We expect that mines already at the technology frontier have less to learn from their neighbours whereas laggard mines have more room to improve. If the positive coefficient on geographical knowledge pool only captures spatially correlated shocks, the coefficient should be similar for technology leaders and laggards. Thus, we run a split sample test based on cyanide productivity and cyanide experience. Column (1) presents the results for above median cyanide productivity mines. The coefficient estimate for the geographical knowledge pool is 0.001 and statistically insignificant. In column (2), the estimate for the same coefficient for the bottom 50% cyanide productivity mines is 0.018, which is both economically and statistically significant. Columns (3) and (4) suggest that a one standard deviation increase in the geographical knowledge pool is associated with an increase in cyanide productivity of 0.007 for technology early adopters and 0.014 for later adopters. In unreported results, we include interaction terms between the bottom/novice mine dummy with all control variables to test whether the difference in knowledge spillover effects between top and bottom (veteran and novice) mines is statistically significant. We find bottom cyanide productivity mines benefit significantly more from their geographical neighbours' cyanide knowledge than top cyanide productivity mines. Overall, we conclude that less experienced mines enjoyed faster productivity growth through learning than technological leaders. This result suggests that spatially correlated shocks are unlikely to be the main explanation for our results.

4.4.5.4 Placebo Tests

We run several placebo tests to further rule out the possibility that our findings are caused by unobservables. First, we run a placebo test using the primary (mercury amalgamation) process. The use of mercury amalgamation to recover gold can be traced back to the 4th century BC, and it was the standard gold extraction process at the start of mining operations on the Witwatersrand. Thus, we consider mercury amalgamation to be an established technology for gold mines. The literature suggests that knowledge spillovers only exist when there is something new to learn (Conley and Udry, 2010; Foster and Rosenzweig, 1995, 2010; Juhász et al., 2023). Therefore a mine's peers should not generate knowledge spillovers for the mercury amalgamation process, since there was little new to learn.

Table 4.7 shows that only the geographical knowledge pool is statistically significant in explaining spillovers of mercury knowledge when the five mechanisms are tested separately. In column (6) all coefficient estimates of knowledge pools are separately insignificant at conventional levels. Furthermore, a joint significance test fails to reject the null hypothesis that all coefficients on the five knowledge pools are equal to zero (F = 1.02). The limited evidence of mercury amalgamation knowledge transfers suggests the positive and significant cyanide spillovers observed in our baseline results do not just reflect a correlation of productivity shocks among connected mines.

We now conduct 1000 Monte Carlo simulations of mine connections. In each simulation, a connection dummy is randomly assigned to all possible mine pairs so that 10% of mine pairs are connected. A random connection intensity is then assigned to the connected mine pairs. The connection and connection intensity are bilateral and constant over time. We construct a pseudo-knowledge pool based on this randomly generated network analogous to $knowledge\ pool(m)$ and estimate its coefficient using the baseline model. Figure 4.6 presents the fitted kernel density of the coefficient estimates of the randomized knowledge pool which is centered on zero. Two vertical lines, indicating the estimated manager spillovers and geographical spillovers, are to the far right of the fitted density curve. No simulation has an estimated coefficient higher than 0.01, which suggests the probability of a false-positive is near zero.

4.4.6 Robustness Results

4.4.6.1 Model Specification

In Table 4.8 we examine the assumptions we have made about: the timing of the knowledge transfer; the definition of the cyanide process; the measure of cyanide productivity; the construction of the cyanide knowledge pool; and the model specification. We vary one assumption at a time and test whether the results are consistent with our baseline results.

In panel A we relax the assumption regarding the timing of spillovers. Our baseline results use a lagged one-month positive productivity surprise. We now test whether mines could learn from cyanide innovations that took place several months ago. The cyanide knowledge pools in columns (1) to (3) are based on the average 3, 6, and 12 month lagged above expectation productivity. The *Lagged own knowledge* control variable is defined using the average own above expectation productivity over the same period. Geographical knowledge spillovers diminish in magnitude when the cyanide innovation is from a more distant period, but the magnitudes of manager spillovers are similar in all three columns. One explanation is that close geographical proximity speeds up knowledge transfer. Employees of geographically close mines are likely to interact, so news about cyanide technology might be easily transmitted within a short time. Workers or managers in commonly managed mines, on the other hand, may visit the mining house headquarters (in downtown Johannesburg) less frequently and information exchange may be delayed.

In panel B, we examine knowledge spillovers of all the cyanide processes, not just the sands process. Our results suggest that including the slimes, concentrates, or both knowledge pools do not change our conclusion. The results in column (3) also confirm that our assumption that a constant proportion of gold in tailings is sorted into sands does not affect the results, since all tailings are included in column (3).

In panel C, we test the robustness to three alternative measures for cyanide productivity: (1) gold ounces from sands per ton of sands tailings treated, (2) gold ounces from sands per ton of crushed ore treated, and (3) gold ounces from sands per ton of sands treated scaled by the gold yield from amalgamation per ton of crushed ore treated. Compared to the baseline measure, the measures in columns (1) and (2) do not adjust for the underlying ore quality. However, all results include mine fixed effects which mitigates the lack of a direct control for ore quality. The measure in column (3) addresses the concern that not all ore crushed in month t may have been treated as tailings in that month. This metric assumes that the proportion of gold sorted into sands from tailings is proportional to the ratio between the tonnage of sands and crushed ores (i.e. $\frac{Ozs^{in\ sands}}{Ozs^{in\ tailings}} \propto \frac{Tons^{Sands}}{Tons^{Primary}}$). However, partially treated sands tailings could contain more gold than untreated sands. Moreover, the ratio could be sensitive to the moisture level of the sands and crushed ores. The results based on these alternative measures are qualitatively similar to the baseline results.

In panel D column (1), we test a sub-sample that consists of mines registered before 1890 to address the concern that mine locations might be endogenously determined by cyanide knowledge spillover potential. In this sub-sample, mine locations were determined and registered before the cyanide method was used on the Witwatersrand in 1890. Thus, these mines' locations are unlikely to have been affected by characteristics that affect cyanide productivity, which alleviates the endogenous location choice concern. The coefficient estimate of the geographical spillovers is 0.013 with a standard error of 0.04, consistent with the baseline result. The manager spillover and German connection spillover are still positive but less statistically significant in this sub-sample. As about 40% of our sample mines were registered after 1890, the sample restriction may lead to a loss of information that would be useful to identify the manager and German spillovers. In column (2), we predict the expected cyanide productivity using a rolling estimation window with all observations up until month $t-1$, rather than using full sample data as we do in our baseline results. Our results are robust to the knowledge pool constructed based on the dynamic estimation of the expected cyanide productivity. Column (3) modifies the knowledge pool to address the concern that an increase in the knowledge pool could be due to an increase in the number of connections instead of an increase in 'true' cyanide knowledge. Hence, we construct a 'max' knowledge pool with only the highest productivity surprise of a connected mine. This knowledge pool measure is less sensitive to market entry/exit and the size of the connections network. We find quantitatively similar results relative to the baseline specification. Last, in column (4), we assume mines have a Cobb-Douglas production function with the external knowledge pool as an additional factor of production. The Cobb-Douglas functional form implies the relationship between productivity and the knowledge pool should be estimated in a log-log regression form. The log-log model finds coefficient estimates on the manager knowledge pool and geographical knowledge pool to be 0.007 and 0.008, respectively, suggesting a 1% increase in a neighbouring mine's lagged productivity is associated with a 0.007% increase in productivity.

Last, in Panel E, we apply alternative weights for connected mines' cyanide knowledge. Column (1) examines the assumption regarding manager interlock intensity and geographical connection intensity by forcing all connection intensities to equal one. Columns (2) to (4) present the results based on alternative transformations of the connection intensity measure following Zacchia (2020), including linear, square, and polynomial transformations. All three transformations provide results consistent with the baseline analysis, suggesting that our results are not sensitive to the choice of the transformation function of the connection intensity.

4.4.6.2 Geographical Spillovers

Geographical knowledge spillovers are based on the idea that people are more likely to exchange information when they are located nearby and the knowledge spillover effect should diminish with the geographic distance.[23] We now explore the sensitivity of spillovers to distance. Table 4.9 presents the spillover effects from neighbouring mines: (1) within a 2.5 kilometres radius from each other, (2) mines outside a 2.5 km radius but within a 5 km radius, and (3) mines outside a 5 km radius but within a 10 km radius. Geographical spillovers are extremely localized. The is no statistical evidence for knowledge spillovers from mines located more than 2.5 km away, and the point estimates for more distant mines are virtually zero.

4.4.6.3 Multiple Spillover Mechanisms

In Table 4.10 we investigate if knowledge spillovers were intensified when two mines were connected through multiple mechanisms at the same time. We construct knowledge pools where mines in the pool must simultaneously meet two criteria. In column (1) we cannot identify any spillover enhancements between mines in the same mining house that also share a common manager at the usual levels of statistical significance.[24] This result is somewhat expected given the high correlation between *Manager Knowledge Pool* and $Org * Manager\ Knowledge\ Pool$: all mines connected through a manager interlock were also connected through the same mining house, except for a few independent mines that shared a common manager. Columns (2) and (3) show that close geographical proximity enhances the spillovers coming from managers and organizations. Knowledge spillovers were increased by roughly one quarter to one third if both criteria for connections were met.

[23]For example, Jaffe et al. (1993) found stronger knowledge spillovers at the local level than the state level.

[24]Although the organizational, managerial, and interacted org*manager pools are not significant at conventional levels by themselves, the joint significance test yields an F statistic of 4.68, which implies that at least one of these three mechanisms is important in diffusing cyanide knowledge.

4.5 Conclusion

We investigate a unique setting and use a novel archival dataset of the South African gold mining industry at the turn of the 20th century. We find strong evidence that technology diffusion took place at both the extensive and intensive margins. We observe substantial productivity gains immediately after adoption of the cyanide ore treatment process. The new method allowed marginally profitable mines to survive and new mines to open. We also find that continuous learning permitted further productivity gains in the longer term. Common ownership and an organizational hierarchy made the transition to the cyanide technology more likely. Mining houses implemented a rollout of the new technology throughout their operating mines. Mining houses also facilitated further development of the technology by appointing engineers and mine managers to oversee multiple affiliated mines. In contrast, we find that administrative personnel such as directors and secretaries were ineffective in transmitting technological knowledge. Mines also learned from their geographical neighbours, but the spillovers were very localized.

Understanding which channel is most efficient in diffusing knowledge in a clean setting can provide insight for management's strategic planning and for policymakers to determine how best to promote knowledge exchange activities. We find strong effects of common personnel, but little effect for geographic proximity. Therefore, a firm should not anticipate that locating in a research hot spot will allow it to "soak up" the knowledge present in that location; instead, it will need to engage in specific personnel interactions to benefit from any co-location decision. One caveat of our finding is that our study centers on intra-industry knowledge spillovers and the impact on efficiency, rather than profitability.

4.6 Figures

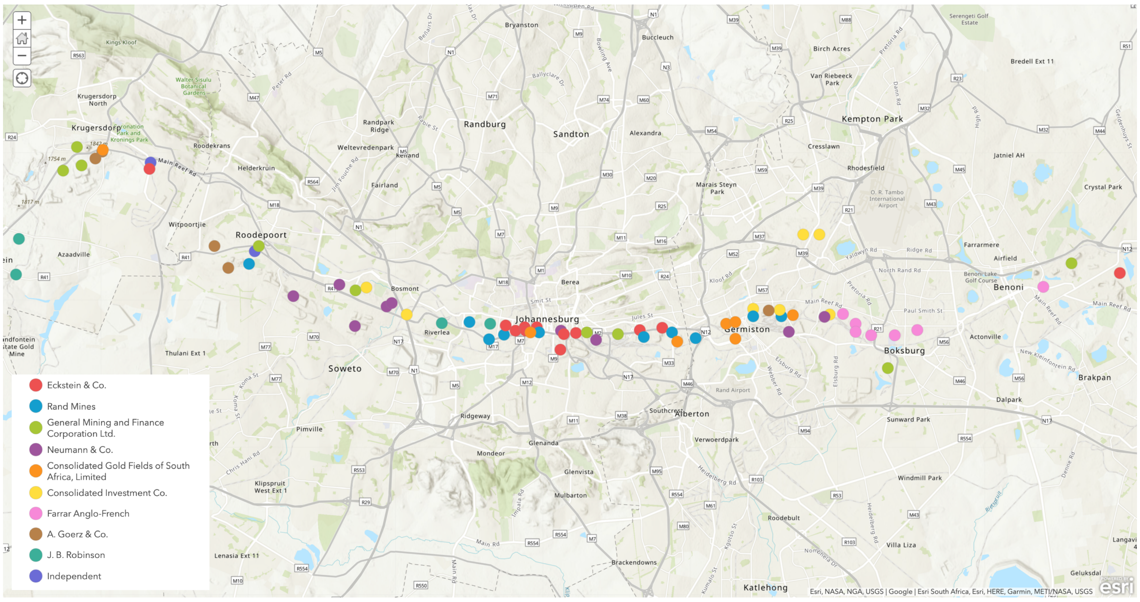

Figure 4.1: Mine Locations (Witwatersrand District), by Mining House

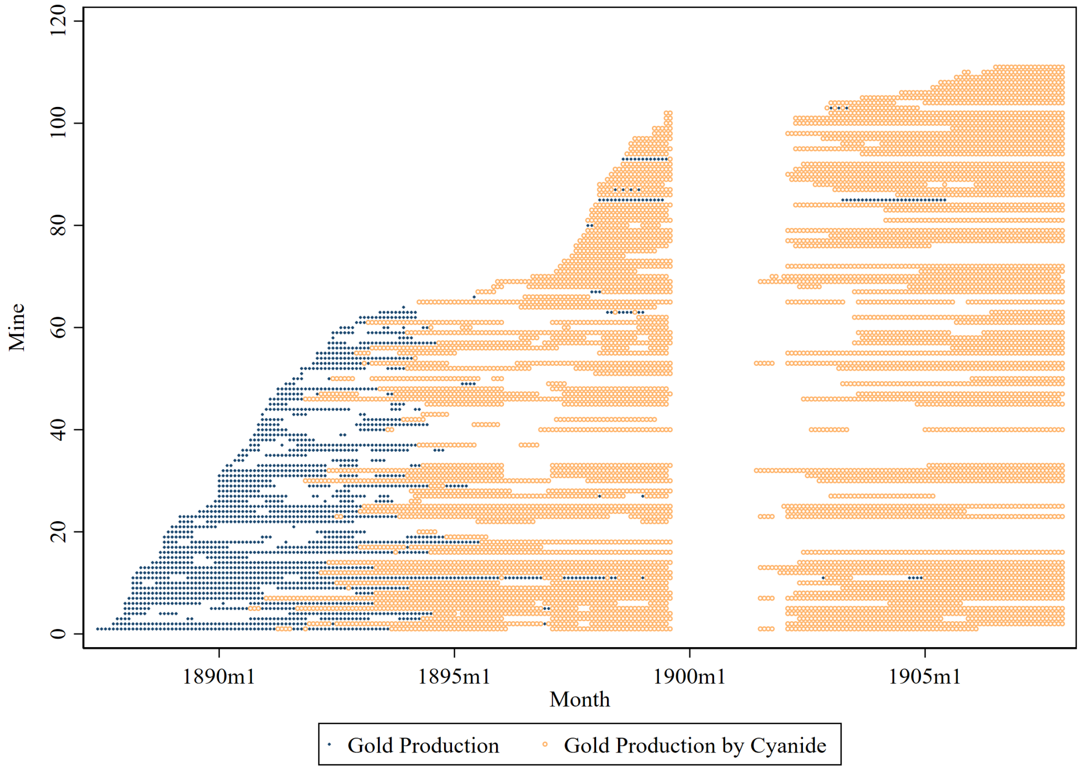

Figure 4.2: Months of Gold Production, by Mine

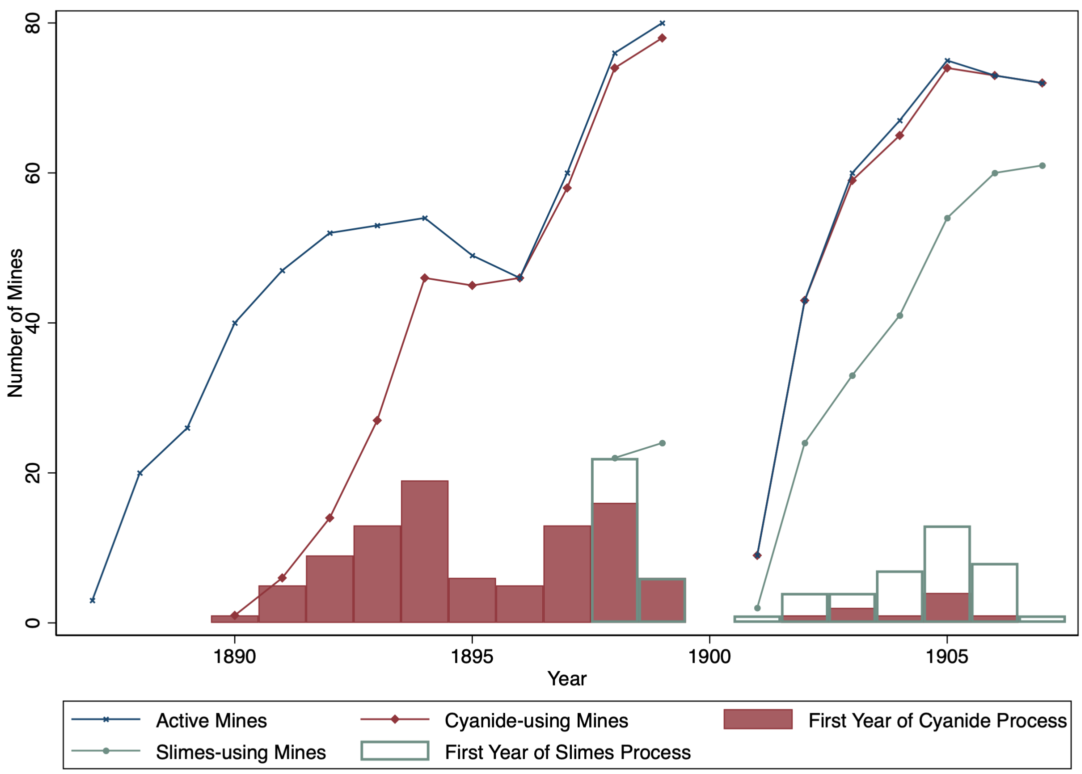

Figure 4.3: Cyanide Process Adoption

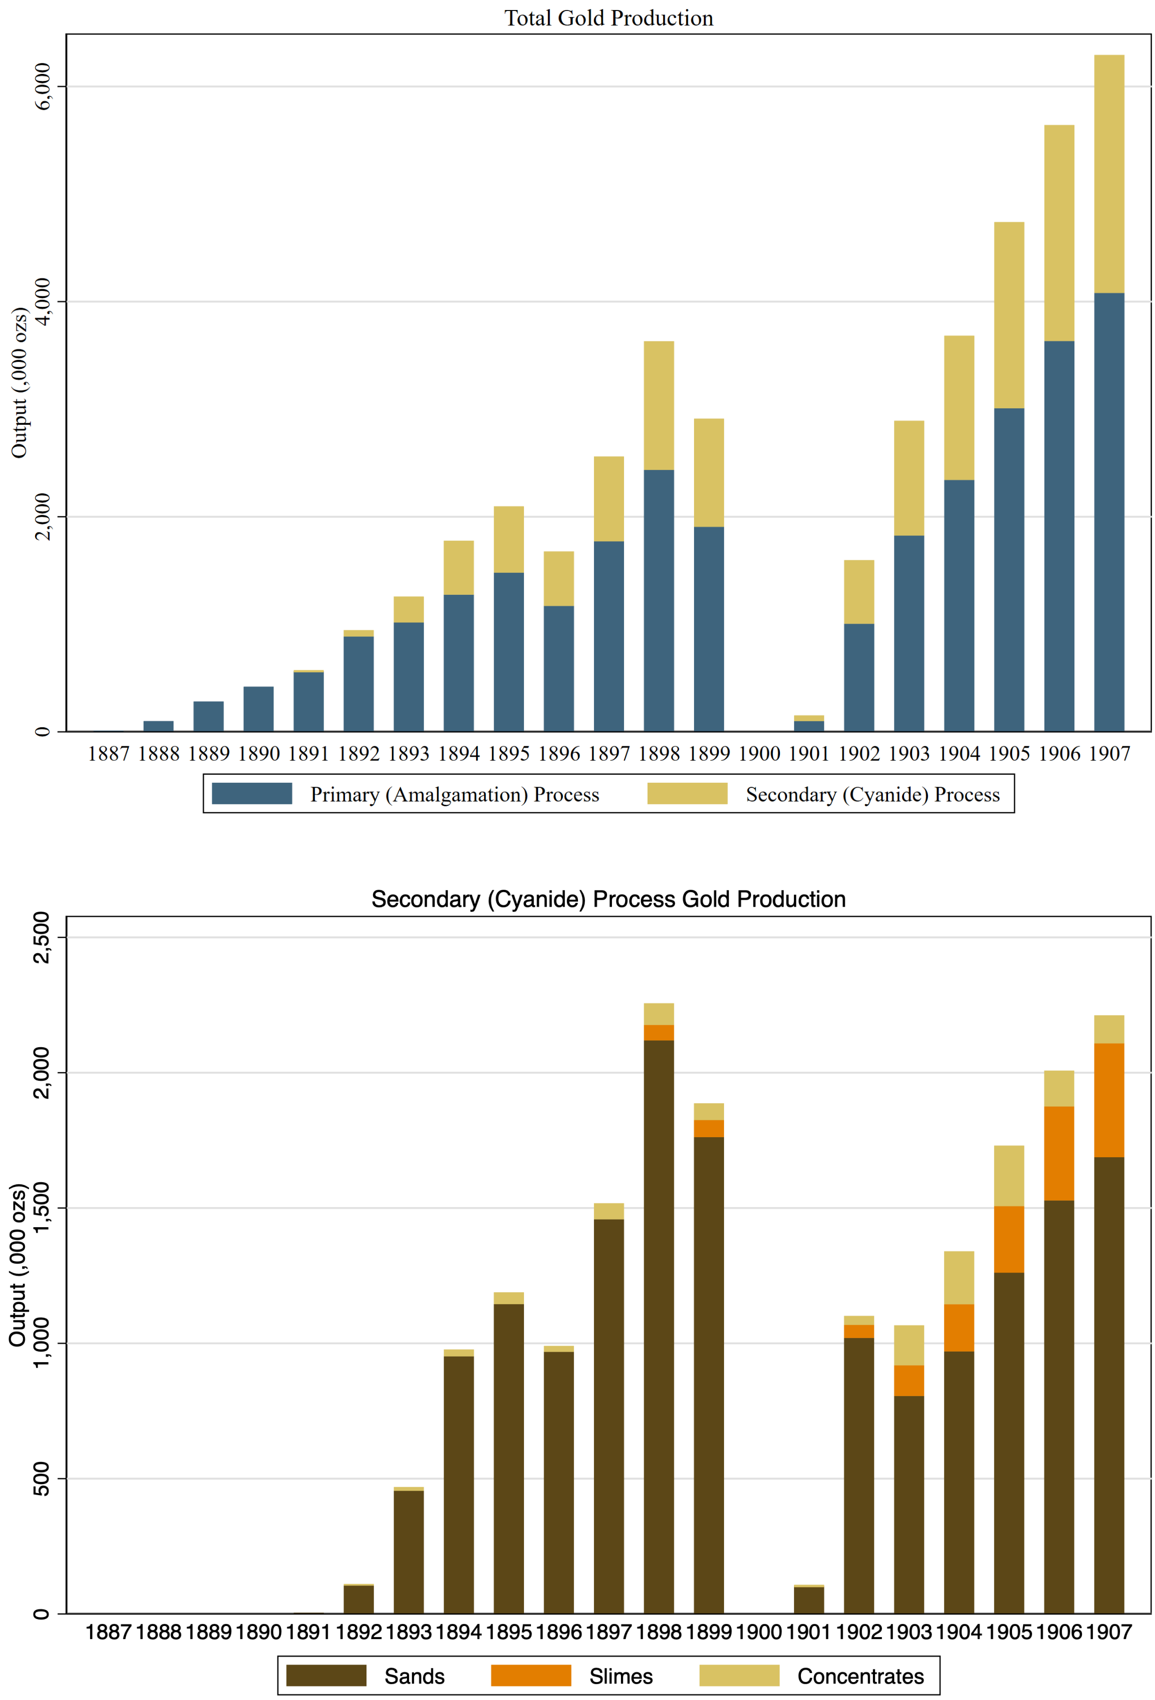

Figure 4.4: Gold Output, by Process

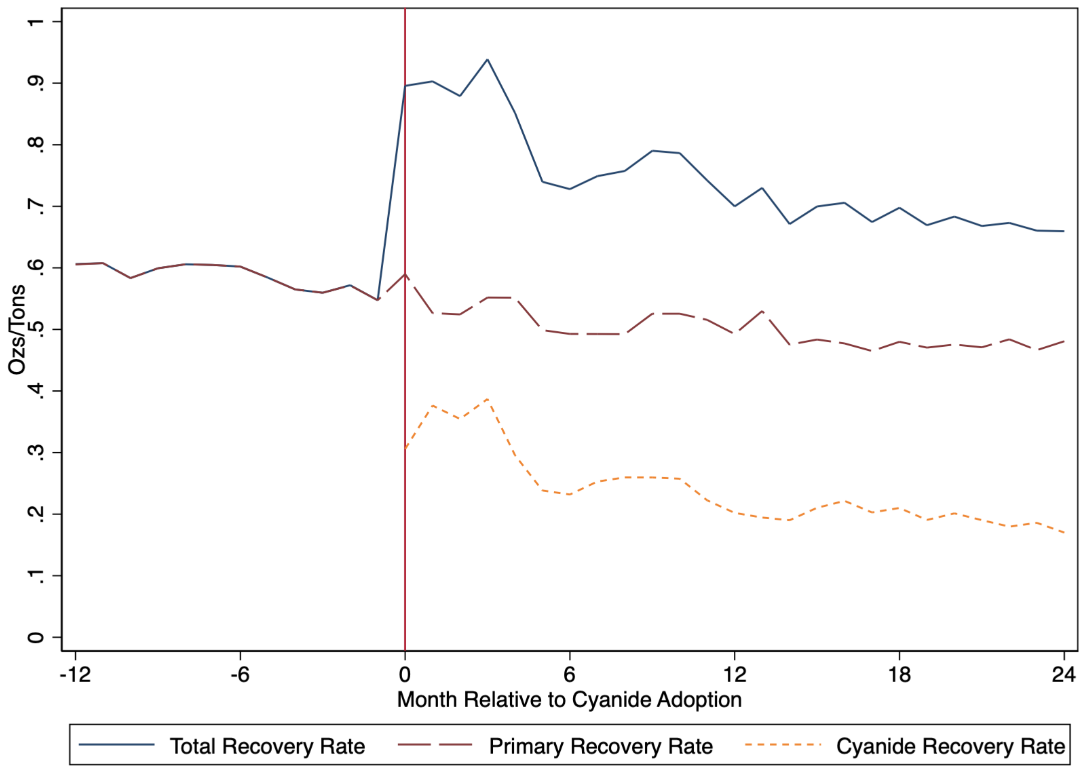

Figure 4.5: Gold Recovery Rates around Cyanide Adoption (Simple Average)

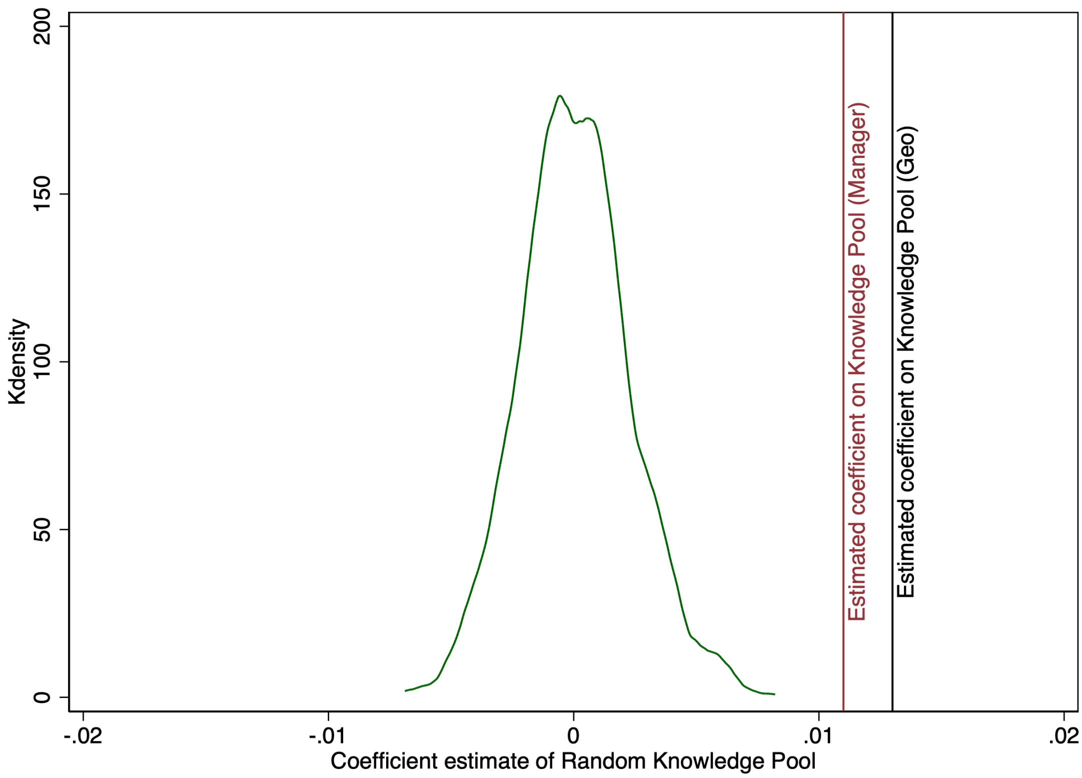

Figure 4.6: Randomized Inference (N = 1000)

Table 4.1: Summary Statistics

Each observation is a mine-month. *Primary Recovery Rate* is the ratio of $Ozs^{Primary}$ to $Tons^{Primary}$. *Cyanide Recovery Rate* is the ratio of $Ozs^{Cyanide}$ to $Tons^{Primary}$. *Total Recovery Rate* is the sum of *Primary Recovery Rate* and *Cyanide Recovery Rate*. $Productivity^{Cyanide}$ is as derived in subsection 4.3.3, and is proportional to $Ozs_{i,t}^{Cyanide}/Ozs_{i,t}^{Primary}$. *Cyanide Process* equals one if any of the Sands, Slimes or Concentrates processes were used. *Sands/Slimes/Concentrates Process* equals one if a mine used the sands/slimes/concentrates process in that month and zero otherwise. *CumulativeExtraction* is the cumulative tons of ores used in the primary process in millions of tons. *Business Connection* equals one if a mine had a business relationship with an external cyanide processor in the past 12 months and zero otherwise. *Stoppage* equals one if a mine experienced a temporary stoppage that month and zero otherwise. *Clean-up* equals one if a mine's output was affected by irregular events, such as from the clean-up of the gold extraction equipment and zero otherwise. *By-products* equals one if the gold output contained irregular outputs, such as from by-products, and zero otherwise. *Tailings Used* equals one if previously treated ore was reused in the cyanide process that month and zero otherwise. *Independent mine* equals one if the mine did not belong to a mining house and zero otherwise. *Witwatersrand Mine* equals one if the mine is located in the Witwatersrand area and zero otherwise. *French/German Connection* equals one if a mine had an office in that location and zero otherwise.

Variable	N	Mean	SD	Median	Min	Max
Tons ('000) Primary	9659	9.58	8.46	7.24	0.12	66.2
Ozs ('000) Primary	9659	3.03	2.47	2.38	0.07	18.5
Tons ('000) Cyanide	7458	8.10	5.94	6.30	0.38	47.1
Ozs ('000) Cyanide	7458	1.51	1.06	1.24	0.03	9.52
Primary Recovery Rate	9659	0.41	0.32	0.31	0.03	7.08
Cyanide Recovery Rate	7458	0.15	0.10	0.13	0.01	2.22
Total Recovery Rate	9659	0.54	0.33	0.45	0.06	7.08
Cyanide Productivity	7458	0.48	0.22	0.43	0.16	1.41
Cyanide Process	9659	0.80	0.40	1	0	1
Sands Process	9659	0.77	0.42	1	0	1
Slimes Process	9659	0.33	0.47	0	0	1
Concentrates Process	9659	0.20	0.40	0	0	1
Cumulative Extraction	9659	0.48	0.59	0.27	0	4.50
Business Connection	9659	0.09	0.28	0	0	1
Stoppage	9659	0.005	0.07	0	0	1
Clean-up	9659	0.003	0.06	0	0	1
By-products	9659	0.01	0.11	0	0	1
Tailings Used	9659	0.04	0.19	0	0	1
Independent Mine	9659	0.32	0.46	0	0	1
Witwatersrand Mine	9659	0.96	0.20	1	0	1
#Personnel	9499	11.1	3.21	11	1	22
#Directors	9423	7.98	2.54	8	1	17
#Managers	8045	1.68	0.71	2	1	4
#Secretaries	9428	1.81	0.40	2	1	2
French Connection	9306	0.41	0.49	0	0	1
German Connection	9306	0.09	0.29	0	0	1

Table 4.2: Diffusion of the Cyanide Process

We run OLS and Hazard models of a mine's adoption of the sands/slimes cyanide process. *Knowledge Pool*(m) is the connection intensity weighted sum of the number of connected sands/slimes-using mines. We drop the French and German knowledge pools in the first two columns because only one mine had a French connection and no mines had German connections in the pre-cyanide period. Observations after the first month of sands/slimes adoption are dropped. For slimes process adoption, we further restrict the sample to cyanide using mines after January 1898, when the slimes process was first reported to the Chamber of Mines, and before 1900, when the second Boer war started. We condition on cumulative extraction and cumulative extraction squared, and lagged own-mine total recovery rate. Robust standard errors are clustered at the mine and month-year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

Dependent Variable	Sands Adoption		Slimes Adoption	
Sample Period	1887m1-1895m12		1898m1-1899m12	
	(1) OLS	(2) Hazard	(3) OLS	(4) Hazard
Knowledge Pool (Org)	0.020***	0.239***	-0.004	-0.559
	(0.007)	(0.090)	(0.003)	(0.389)
Knowledge Pool (Manager)	-0.021	-0.335	0.054**	1.719***
	(0.057)	(0.821)	(0.026)	(0.562)
Knowledge Pool (Geo)	-0.008	-0.221	0.012	0.269
	(0.006)	(0.176)	(0.012)	(0.176)
Knowledge Pool (French)			-0.009**	-0.496**
			(0.003)	(0.203)
Knowledge Pool (German)			0.038**	2.282***
			(0.017)	(0.813)
Observations	1,292	1,296	869	869
Adjusted R-squared	0.063		0.059	
Mine FE	No	No	No	No
Month FE	Yes	Yes	Yes	Yes

Table 4.3: Cyanide Process Adoption and Gold Recovery Rates

We run OLS regressions of gold recovery rates on *Cyanide Process* and controls, with no allowance for knowledge spillovers. Columns 3 and 6 only use a mine's observations from 12 months before to 12 months after the adoption of cyanide. All variables are as defined in Table 4.1. Mine and month fixed effects are used in all regressions. Robust standard errors are clustered at the mine and month-year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	Total Recovery Rate			Primary Recovery Rate		
	(1) Entire Sample	(2) Entire Sample	(3) Adoption ± 1-Year	(4) Entire Sample	(5) Entire Sample	(6) Adoption ± 1-Year
Cyanide Process	0.268***	0.265***	0.280***	0.033	0.031	-0.001
	(0.053)	(0.050)	(0.041)	(0.030)	(0.028)	(0.016)
Cumulative Extraction		-0.054	-0.863		-0.047	-0.053
		(0.127)	(0.865)		(0.119)	(0.669)
Cumulative Extraction Sqr		0.021	0.270		0.016	-0.048
		(0.020)	(1.666)		(0.019)	(0.994)
Concentrates Process		0.015	0.045		-0.005	-0.035
		(0.027)	(0.041)		(0.021)	(0.025)
Slimes Process		-0.008	0.102**		-0.021	0.019
		(0.025)	(0.044)		(0.021)	(0.018)
Business Connection		-0.009	0.074***		-0.013	0.057**
		(0.019)	(0.027)		(0.017)	(0.023)
Stoppage		-0.029	-0.010		-0.025	-0.030
		(0.025)	(0.046)		(0.022)	(0.044)
Clean-up		0.093	0.254		0.036	0.117**
		(0.086)	(0.186)		(0.036)	(0.050)
By-products		0.005	0.033		0.002	0.016
		(0.015)	(0.050)		(0.011)	(0.022)
Tailings Used		-0.084**			-0.054**	
		(0.036)			(0.026)	
Observations	9,659	9,659	1,551	9,659	9,659	1,551
Adjusted R-squared	0.577	0.579	0.755	0.622	0.624	0.849

Table 4.4: Cyanide Productivity and Knowledge Spillovers

We regress cyanide productivity on a knowledge pool(s), mine controls, and mine and month fixed effects. Column (1) to (5) correspond to five knowledge pools: (1) organizational spillovers through membership of a common mining house, (2) spillovers through common mine managers/engineers, (3) spillovers from geographical neighbours, (4) spillovers from mines with French connections, and (5) spillovers from mines with German connections. We define each knowledge pool as the weighted sum of the cyanide productivity of all mines in the pool: $\sum_{c^m_{ijt}=1} g^m_{ijt} K_{jt}$. c^m_{ijt} is a binary variable equal to one if mine j is in mine i's knowledge pool; g^m_{ijt} is the connection intensity measure; and K_{jt} is the above expectation cyanide productivity of mine j in month t. A mine j is in mine i's knowledge pool if: both mines are in the same mining house (column 1), both mines share a common manager and/or engineer (column 2), the physical distance between i and j is no greater than 2.5 kilometres (column 3), and both mines have a French/German connection (columns 4/5). Knowledge pools are standardized to be mean zero with a standard deviation of one for ease of interpretation hereafter. All variables are defined in Table 4.1. Robust standard errors are clustered at the mine and month-year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Knowledge Pool (Org)	0.010***					0.001
	(0.003)					(0.003)
Knowledge Pool (Manager)		0.015***				0.011***
		(0.003)				(0.004)
Knowledge Pool (Geo)			0.013***			0.013***
			(0.003)			(0.003)
Knowledge Pool (French)				0.005		0.003
				(0.008)		(0.008)
Knowledge Pool (German)					0.010***	0.008**
					(0.004)	(0.004)
Lagged Own Knowledge	0.862***	0.854***	0.863***	0.878***	0.867***	0.835***
	(0.061)	(0.064)	(0.062)	(0.066)	(0.063)	(0.058)
Cumulative Extraction	-0.074**	-0.082***	-0.086**	-0.084**	-0.078**	-0.087***
	(0.035)	(0.030)	(0.036)	(0.038)	(0.037)	(0.031)
Cumulative Extraction Sqr	0.009	0.009*	0.013**	0.011*	0.011*	0.012**
	(0.006)	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)
Concentrates Process	-0.100***	-0.101***	-0.102***	-0.100***	-0.103***	-0.104***
	(0.010)	(0.011)	(0.010)	(0.010)	(0.010)	(0.010)
Slimes Process	-0.047***	-0.043***	-0.046***	-0.045***	-0.041***	-0.041***
	(0.013)	(0.011)	(0.012)	(0.013)	(0.013)	(0.011)
Observations	7,032	6,516	7,032	7,032	7,032	6,516
Adjusted R-squared	0.720	0.733	0.721	0.718	0.720	0.737

Table 4.5: Cyanide Knowledge Spillovers by Personnel Roles

We test for spillovers if the connected person is a (1) previous manager, (2) director, or (3) company secretary. An individual is defined as a previously connected mine manager/engineer if they also worked at another mine in the past three years but not in the current year. The knowledge pools of previous managers, directors, and secretaries are calculated as the weighted sum of the above expectation cyanide productivity of other mines which share historical managers, current directors, and current secretaries with the focal mine. We control for spillovers from all other mechanisms, and we include all control variables as in Table 4.4. All variable definitions are in Table 4.1. Mine fixed effects and month fixed effects are included in all columns. Robust standard errors are clustered at the mine and month-year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	(1) Previous Manager	(2) Director	(3) Secretary	(4) All
Knowledge Pool (Previous Manager)	-0.000			-0.000
	(0.002)			(0.002)
Knowledge Pool (Director)		0.007		0.005
		(0.006)		(0.006)
Knowledge Pool (Secretary)			0.004	0.002
			(0.004)	(0.004)
Knowledge Pool (Org)	0.001	-0.001	-0.001	-0.002
	(0.003)	(0.004)	(0.004)	(0.004)
Knowledge Pool (Manager)	0.011***	0.011***	0.011***	0.011***
	(0.004)	(0.004)	(0.003)	(0.003)
Knowledge Pool (Geo)	0.013***	0.013***	0.013***	0.013***
	(0.003)	(0.003)	(0.003)	(0.003)
Knowledge Pool (French)	0.004	0.003	0.003	0.003
	(0.008)	(0.008)	(0.008)	(0.008)
Knowledge Pool (German)	0.008**	0.007*	0.007*	0.007*
	(0.004)	(0.004)	(0.004)	(0.004)
Observations	6,516	6,516	6,516	6,516
R-squared	0.747	0.748	0.748	0.748

Table 4.6: Technology Leaders and Laggards

We test for knowledge spillovers at technology leaders/laggards. In columns (1) and (2) we divide all mines, year by year, into two groups based on their cyanide productivity: top mines are those with above-median average cyanide productivity in the previous year, and bottom mines are all others. We restrict the sample to mines with at least six months' productivity data in the previous year. In columns (3) and (4) we divide all mines into two groups based on the time of cyanide adoption: veteran mines are those that adopted the cyanide process before or in 1895 and novice mines are those that adopted it after 1895. We estimate columns (3) and (4) with observations from January 1896 onwards. All variable definitions are in Table 4.1. We include all control variables as in Table 4.4. Mine fixed effects and month fixed effects are included in all columns. Robust standard errors are clustered at the mine and month-year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1%, respectively.

	Cyanide Productivity		Cyanide Experience	
	(1)	(2)	(3)	(4)
	Top	Bottom	Veteran	Novice
Knowledge Pool (Org)	-0.003	-0.007	-0.001	-0.002
	(0.005)	(0.006)	(0.003)	(0.009)
Knowledge Pool (Manager)	0.014***	0.014**	0.004	0.015*
	(0.005)	(0.006)	(0.003)	(0.008)
Knowledge Pool (Geo)	0.001	0.018**	0.007***	0.014**
	(0.003)	(0.007)	(0.003)	(0.005)
Knowledge Pool (French)	-0.013	0.000	-0.002	0.037
	(0.009)	(0.012)	(0.007)	(0.023)
Knowledge Pool (German)	0.003	0.007*	0.006**	0.015*
	(0.003)	(0.003)	(0.003)	(0.009)
Observations	2,480	2,420	3,433	2,364
R-squared	0.819	0.728	0.773	0.778

Table 4.7: Knowledge Spillovers in the Mercury Amalgamation Process

We replicate Table 4.4 using the primary mercury amalgamation process recovery rate as the dependent variable, defined as $Ozs^{primary}/Tons^{Primary}$. Knowledge pools are constructed with above expectation primary recovery rates. All variable definitions are in Table 4.1. We include all control variables as in Table 4.4. Mine fixed effects and month fixed effects are included in all columns. Robust standard errors are clustered at the mine and month-year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Knowledge Pool (Org)	-0.002					-0.004
	(0.005)					(0.010)
Knowledge Pool (Manager)		0.002				0.002
		(0.003)				(0.004)
Knowledge Pool (Geo)			0.019***			0.013
			(0.007)			(0.011)
Knowledge Pool (French)				-0.008		-0.009
				(0.009)		(0.009)
Knowledge Pool (German)					-0.002	0.000
					(0.003)	(0.003)
Observations	8,998	7,318	8,998	8,998	8,998	7,318
R-squared	0.880	0.879	0.882	0.880	0.880	0.880

Table 4.8: Robustness

We replicate Table 4.4 using alternative definitions of cyanide productivity and cyanide knowledge pools. All variable definitions are in Table 4.1. Mine fixed effects and month fixed effects are included in all columns. Robust standard errors are clustered at the mine and month-year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

Panel A	K_{jt} = n-month Avg. above Expectation Productivity		
	(1) 3-month Avg. Knowledge	(2) 6-month Avg. Knowledge	(3) 12-month Avg. Knowledge
Knowledge Pool (Org)	-0.003	-0.006	-0.010*
	(0.004)	(0.004)	(0.006)
Knowledge Pool (Manager)	0.010**	0.012**	0.012**
	(0.004)	(0.005)	(0.005)
Knowledge Pool (Geo)	0.011***	0.006**	0.002
	(0.003)	(0.003)	(0.003)
Knowledge Pool (French)	0.003	0.001	0.005
	(0.007)	(0.007)	(0.008)
Knowledge Pool (German)	0.006*	0.006*	0.005
	(0.003)	(0.004)	(0.004)
Observations	6,240	5,821	5,033
Adjusted R-squared	0.745	0.722	0.701

Panel B	Multiple Cyanide Processes		
	(1) Sands & Slimes	(2) Sands & Concentrates	(3) Sands, Slimes, & Concentrates
Knowledge Pool (Org)	0.001	0.003	0.003
	(0.003)	(0.003)	(0.003)
Knowledge Pool (Manager)	0.012***	0.012***	0.013***
	(0.004)	(0.004)	(0.004)
Knowledge Pool (Geo)	0.012***	0.012***	0.010***
	(0.003)	(0.003)	(0.003)
Knowledge Pool (French)	0.003	-0.001	-0.000
	(0.008)	(0.008)	(0.008)
Knowledge Pool (German)	0.006	0.010**	0.008**
	(0.004)	(0.004)	(0.004)
Observations	6,516	6,516	6,516
Adjusted R-squared	0.763	0.699	0.742

Panel C	Cyanide Productivity Measure		
	(1) $Ozs^{Cyanide}/Tons^{Cyanide}$	(2) $Ozs^{Cyanide}/Tons^{Primary}$	(3) $\frac{Ozs^{Cyanide}/Tons^{Cyanide}}{Ozs^{Primary}/Tons^{Primary}}$
Knowledge Pool (Org)	0.002	0.002	0.003
	(0.002)	(0.002)	(0.004)
Knowledge Pool (Manager)	0.003***	0.004***	0.009**
	(0.001)	(0.001)	(0.004)
Knowledge Pool (Geo)	0.001	0.005***	0.007**
	(0.002)	(0.002)	(0.003)
Knowledge Pool (French)	0.004	0.001	0.010
	(0.004)	(0.003)	(0.010)
Knowledge Pool (German)	0.002	0.002	0.006
	(0.001)	(0.001)	(0.004)
Observations	6,516	6,516	6,516
Adjusted R-squared	0.784	0.714	0.763

Table 8: Robustness (Continued)

Panel D	Knowledge Pool Construction			
	(1) Incumbent Mines Registered before 1890	(2) Dynamic Estimation	(3) Max Knowledge Pool $\max(K_{jt} \mid c_{ijt} = 1) \times g_{ijt}$	(4) Log-log Model $\sum c_{ijt} g_{ijt} log(K_{jt})$
Knowledge Pool (Org)	0.001	0.001	0.002	0.000
	(0.003)	(0.003)	(0.003)	(0.002)
Knowledge Pool (Manager)	0.005	0.008**	0.012***	0.007***
	(0.004)	(0.004)	(0.003)	(0.002)
Knowledge Pool (Geo)	0.013***	0.016***	0.011***	0.008***
	(0.004)	(0.003)	(0.002)	(0.002)
Knowledge Pool (French)	-0.001	0.002	-0.004	0.004
	(0.010)	(0.008)	(0.007)	(0.005)
Knowledge Pool (German)	0.006	0.010**	0.010**	0.005*
	(0.004)	(0.004)	(0.004)	(0.003)
Observations	4,107	6,516	6,516	6,516
Adjusted R-squared	0.738	0.736	0.738	0.732

Panel E	Alternative Connection Intensity g_{ijt}			
	(1) Unweighted Knowledge Pool $g_{ijt} = 1$	(2) Linear Transformation $f(x) = x$	(3) Square Transformation $f(x) = x^2$	(4) Polynomial Transformation $f(x) = -x^2 + 2x$
Knowledge Pool (Org)	0.000	0.002	0.004	0.001
	(0.003)	(0.003)	(0.003)	(0.003)
Knowledge Pool (Manager)	0.012***	0.010***	0.007**	0.011***
	(0.004)	(0.004)	(0.003)	(0.004)
Knowledge Pool (Geo)	0.013***	0.013***	0.011***	0.013***
	(0.003)	(0.003)	(0.003)	(0.003)
Knowledge Pool (French)	0.003	0.004	0.004	0.003
	(0.008)	(0.008)	(0.008)	(0.008)
Knowledge Pool (German)	0.008**	0.008*	0.007*	0.008**
	(0.004)	(0.004)	(0.004)	(0.004)
Observations	6,516	6,516	6,516	6,516
Adjusted R-squared	0.737	0.736	0.735	0.737

Table 4.9: Cyanide Knowledge Spillovers by Geographical Proximity

We construct three geographical knowledge pools using mines that were located within a 0-2.5km, 2.5-5km, and 5-10 km radius, respectively. We control for spillovers from all other mechanisms, and all variables included in Table 4.4. All variable definitions are in Table 4.1. Mine fixed effects and month fixed effects are included in all columns. Robust standard errors are clustered at the mine and month-year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	(1) 0-2.5km	(2) 2.5-5km	(3) 5-10km	(4) All
Knowledge Pool (Geo2.5)	0.013***			0.014***
	(0.003)			(0.003)
Knowledge Pool (Geo5)		0.000		0.002
		(0.003)		(0.003)
Knowledge Pool (Geo10)			0.000	0.005
			(0.003)	(0.003)
Knowledge Pool (Org)	0.001	0.001	0.001	0.001
	(0.003)	(0.003)	(0.003)	(0.003)
Knowledge Pool (Manager)	0.011***	0.013***	0.013***	0.011***
	(0.004)	(0.004)	(0.004)	(0.004)
Knowledge Pool (French)	0.003	0.001	0.001	0.003
	(0.008)	(0.008)	(0.008)	(0.008)
Knowledge Pool (German)	0.008**	0.007*	0.007*	0.008**
	(0.004)	(0.004)	(0.004)	(0.004)
Observations	6,516	6,516	6,516	6,516
R-squared	0.747	0.745	0.745	0.748

Table 4.10: Amplification Effects in the Network

We test for differential impacts of the cyanide knowledge pool by allowing for the interaction of a mine's managerial, organizational, and geographical networks. The *KnowledgePool* variables in column (1) to (3) are constructed using mines that (1) were in the same mining house and had at least one common manager (2) were located within 2.5km and had at least one common manager (3) were in the same mining house and were located within 2.5 km relative to the focal mine. We control for spillovers from all other mechanisms, and all variables included in Table 4.4. All variable definitions are in Table 4.1. Mine fixed effects and month fixed effects are included in all columns. Robust standard errors are clustered at the mine and month-year levels. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	(1) Org*Manager	(2) Geo*Manager	(3) Org*Geo
Knowledge Pool (Org*Manager)	0.003		
	(0.009)		
Knowledge Pool (Geo*Manager)		0.006	
		(0.004)	
Knowledge Pool (Org*Geo)			0.005*
			(0.003)
Knowledge Pool (Org)	0.000	0.001	-0.001
	(0.004)	(0.003)	(0.004)
Knowledge Pool (Manager)	0.009	0.008**	0.011***
	(0.008)	(0.004)	(0.004)
Knowledge Pool (Geo)	0.013***	0.011***	0.011***
	(0.003)	(0.003)	(0.003)
Knowledge Pool (French)	0.003	0.004	0.004
	(0.008)	(0.008)	(0.008)
Knowledge Pool (German)	0.008**	0.008**	0.008**
	(0.004)	(0.004)	(0.004)
Observations	6,516	6,516	6,516
R-squared	0.747	0.748	0.748

4.1 Appendix

Table 4.A1: Mining House Affiliation

Mining House	Founded	Entrepreneurs	Operating Mines
A. Goerz and Co.	1897	Adolf Goerz	Geduld Proprietary, Lancaster, Lancaster West, May Consolidated, Princess Estate, Roodepoort Central Deep
Consolidated Goldfields of South Africa	1887		Jupiter, Knights Deep, Luipaardsvlei, Nigel Deep, Robinson Deep, Simmer and Jack, Simmer and Jack East, Simmer Deep
Johannesburg Consolidated Investment Co.	1889	Barney Barnato	Balmoral, Knight's, New Croesus, New Spes Bona, Buffelsdorn, Consolidated Langlaagte, Ginsberg, Glencairn, Langlaagte Royal, New Heidelberg Roodepoort, New Primrose, New Rietfontein, Unified Main Reef, Rietfontein A, Roodepoort Kimberley
Eckstein and Co. (Corner House Group)	1888	Hermann Eckstein, Alfred Beit	Bonanza, City and Suburban, Crown Reef, New Heriot, New Modderfontein, Ferreira, French Rand, Henry Nourse, Robinson, Robinson Central Deep, Village Deep, Village Main Reef
Farrar Anglo-French		George Farrar	Angelo, New Blue Sky, New Kleinfontein, Cason, Driefontein Consolidated, East Rand Proprietary, Langlaagte United, New Comet
General Mining and Finance Corporation	1895	George Albu, Leopold Albu	Aurora West, Cinderella Deep, New Goch, Meyer and Charlton, Rand Nigel, Roodepoort United Main Reef, Van Ryn, Van Ryn West, Violet, West Rand Consolidated, West Rand Mines, Steyn Estate
J. B. Robinson Group		Joseph Robinson	Langlaagte Block B, Porges Randfontein, Langlaagte Estate, Langlaagte Star, North Randfontein, Robinson Randfontein, South Randfontein
Neumann and Co.		Sigismund Neumann	Bantjes, Consolidated Main Reef, Knights Central, Main Reef, Main Reef West, Treasury, Vogelstruis Consolidated, Witwatersrand Deep, Wolhunter
Rand Mines		Hermann Eckstein, Alfred Beit	Crown Deep, Durban Roodepoort Deep, Ferreira Deep, Geldenhuis Deep, Glen Deep, Jumpers Deep, Langlaagte Deep, Nourse Deep, Paarl Central, Rose Deep

Table 4.A2: Cyanide Productivity Estimation

We run OLS regressions of cyanide productivity on its potential determinants. All variables are as defined in Table 4.1. Mine and month fixed effect are used in all regressions. t-statistics based on robust standard errors are clustered at the mine and month-year levels are in parentheses. Levels of significance are indicated by *, **, and *** for 10%, 5%, and 1% respectively.

	(1) Sands	(2) Sands & Slimes	(3) Sands & Conc.	(4) Sands & Slimes & Conc.
Cumulative Extraction	-0.049***	-0.079***	-0.043**	-0.073***
	(0.017)	(0.018)	(0.017)	(0.018)
Cumulative Extraction Squared	0.005	0.010**	0.002	0.007
	(0.004)	(0.004)	(0.004)	(0.004)
Concentrates Process	-0.104***	-0.105***	0.012**	0.010
	(0.006)	(0.006)	(0.006)	(0.006)
Slimes Process	-0.060***	0.045***	-0.058***	0.046***
	(0.007)	(0.007)	(0.007)	(0.008)
Business Connection	-0.001	-0.002	-0.005	-0.007
	(0.008)	(0.008)	(0.008)	(0.008)
Stoppage	0.025	0.013	0.028	0.016
	(0.027)	(0.029)	(0.027)	(0.029)
Clean-up	0.085***	0.085***	0.079**	0.077**
	(0.030)	(0.032)	(0.031)	(0.033)
By-products	0.016	0.013	0.009	0.007
	(0.015)	(0.016)	(0.015)	(0.016)
Tailings Used	-0.022**	-0.036***	-0.030***	-0.044***
	(0.010)	(0.010)	(0.010)	(0.011)
Observations	7,454	7,454	7,454	7,454
Adjusted R-squared	0.526	0.566	0.490	0.551

Table 4.A3: Percentage of Connected Mine Pairs over Time

This table presents the proportion of connected mine pairs for each of the five potential mechanisms from 1887 to 1907.

	% of Connected Active Mine Pairs				
Year	Org	Geo	Manager	French	German
1887	5.3	73.7		0	0
1888	8.4	30.1	0	0	0
1889	7.7	17.8	.5	0	0
1890	5.1	13.9	.4	0	0
1891	4.9	12.8	.2	0	0
1892	5	12.1	.3	0	0
1893	6.7	12	.3	0	0
1894	6	11.9	.5	0	0
1895	6	11.5	.8	.1	1
1896	8.8	13.9	2.3	4.4	.2
1897	7.7	11.5	2.4	16.3	1.1
1898	6.8	9	2	20.2	.7
1899	7	8.5	2.5	29.5	1.1
1901	20.6	42.5	0	97.2	7.7
1902	12.8	18.3	6.5	48.3	2.2
1903	8.3	12.2	4.7	42.2	2.3
1904	7.4	10.1	4.5	40	2.1
1905	8	10	5.9	43.6	2
1906	7.7	8.7	7	45	2
1907	7.6	8.5	6.7	42.6	1.9

Table 4.A4: Correlation Table

Knowledge Pool	Org	Mgr	Geo	FR	GE	Dir	Sec
Organization	1						
Manager	0.64	1					
Geographical	0.18	0.20	1				
French Committee	0.17	0.24	0.11	1			
German Committee	0.25	0.22	-0.00	0.32	1		
Director Interlock	0.58	0.50	0.24	0.43	0.12	1	
Secretary Interlock	0.55	0.51	0.17	0.35	0.17	0.73	1

Table 4.A5: Baseline Results with Bootstrapped Standard Errors

This table presents the baseline estimation with bootstrapped standard error. We block bootstrap the estimation sample with 1000 Monte Carlo simulations to compute the bootstrapped standard errors.

	(1)	(2)	(3)	(4)	(5)	(6)
Knowledge Pool (Org)	0.010					0.001
	(0.007)					(0.005)
Knowledge Pool (Manager)		0.015				0.011
		(0.008)				(0.007)
Knowledge Pool (Geo)			0.013			0.013
			(0.007)			(0.006)
Knowledge Pool (French)				0.005		0.003
				(0.005)		(0.009)
Knowledge Pool (German)					0.010	0.008
					(0.006)	(0.005)
Observations	7,032	6,516	7,032	7,032	7,032	6,516
Adjusted R-squared	0.720	0.733	0.721	0.718	0.720	0.737

We derive the cyanide productivity based on the following three simplifying assumptions:

1. Percentage of gold that can be recovered by mercury amalgamation is a
2. Percentage loss of gold ore from primary to cyanide process is b
3. Percentage of gold in tailings that are sorted into sands is c

Based on the first assumption, the gold in crushed ores can be written as:[25]

$$Ozs^{in\ ores*} = \frac{Ozs^{Primary}}{a}$$

Then, given assumption 2, the gold left in the tailings would be $(1-b)$ multiplied by the difference between the gold in crushed ores and the gold recovered from crushed ores:

$$Ozs^{in\ tailings*} = (1-b) \times (Ozs^{in\ ores*} - Ozs^{Primary}) = \frac{(1-a)(1-b)}{a} Ozs^{Primary}.$$

With the last assumption, the gold sorted into sands would be c multiplied by the gold in tailings:

$$Ozs^{in\ sands*} = c \times Ozs^{in\ tailings*} = \frac{c(1-a)(1-b)}{a} Ozs^{Primary}.$$

Therefore, cyanide productivity can be written as:

$$Productivity^{Cyanide} = \frac{Ozs^{Cyanide}}{Ozs^{in\ sands*}} = \frac{a}{c(1-a)(1-b)} \times \frac{Ozs^{Cyanide}}{Ozs^{Primary}} \propto \frac{Ozs^{Cyanide}}{Ozs^{Primary}}.$$

Since a, b, and c are constant, the cyanide productivity is proportional to the ratio between the gold yield in the cyanide process and the gold yield in the primary process. We use this ratio as the empirical measure of cyanide productivity.

To address the problem that, in some cases, tailings and slimes are only partially treated or not treated, further adjustments are made as a robustness check. The

[25] Mine and time subscripts are dropped for simplicity

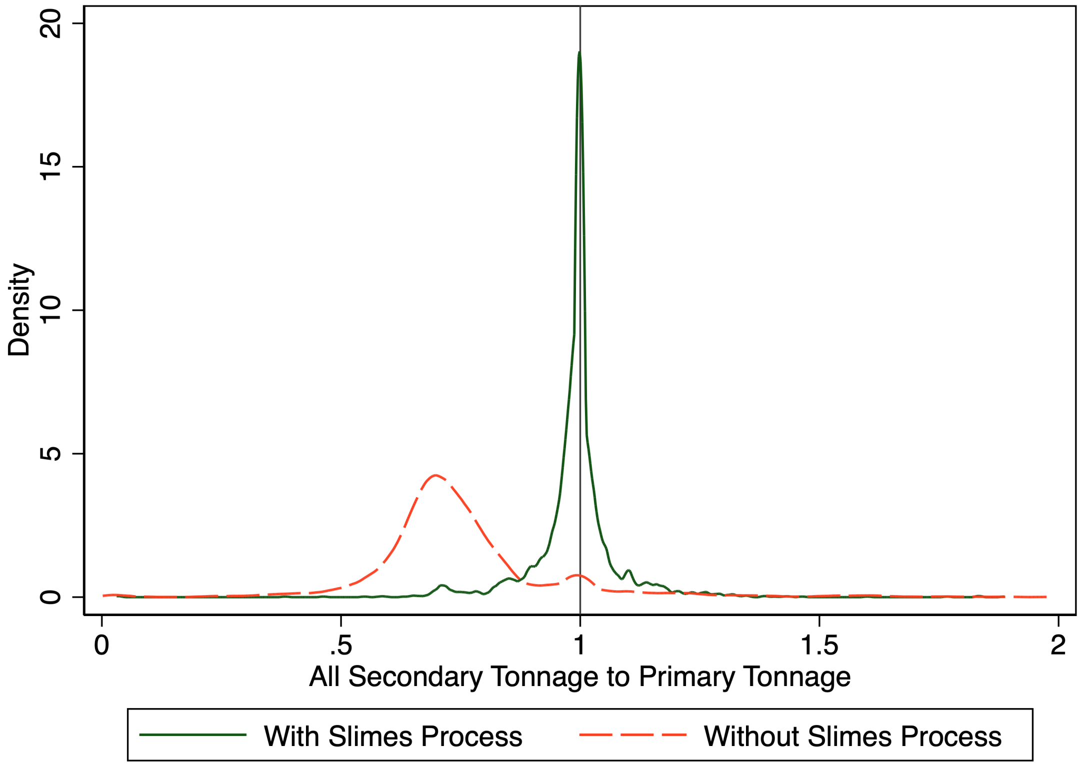

Figure 4.B1: Proportion of Tailings Went to Secondary Processes

assumption is that the recovery rate of the partial treatment is the same as direct treatment of tailings from the primary process, $\frac{Ozs^{Cyanide(observed)}}{Tons^{Cyanide}} = \frac{Ozs^{Cyanide}}{Tons^{Primary}}$[26]. After replacing the cyanide output with the gold output from direct treatment, the adjusted cyanide productivity measure becomes:

$$Productivity_{i,t}^{Cyanide} = \frac{\alpha}{1-\alpha} \times \frac{Ozs_{i,t}^{Cyanide} / Tons_{i,t}^{Cyanide}}{Ozs_{i,t}^{Primary} / Tons_{i,t}^{Primary}}$$

[26]The tonnage processed by cyanide process is the sum of sands tonnage and slimes tonnage. If a mine did not use slimes process in month t, we use the sands tonnage divided by 0.68 to get the approximate cyanide tonnage as we want to distinguish delayed treatment and lack of appropriate technology to treat. In our sample, the average and median sands tonnage to cyanide tonnage are both 68%.

Chapter 5

Summary

In conclusion, this thesis explores the impact of corporate innovation and technology spillovers on economic gains, with a specific focus on the influence of product market competition and the mechanisms of technology transfer. The literature review examined the relationship between competition and innovation, as well as the effects of knowledge spillovers on firms' productivity and performance.

In the first essay (Competition and the Value of Innovation), I find that intense product market competition leads to a lower economic value of patented innovation, contrary to the existing literature suggesting an inverted-U relationship between competition and patent scientific value. Patent value increases by an average of 4.8% after horizontal M&A announcements, considered anti-competitive events. More pronounced positive effects on patent value are documented for merger deals that are expected to have stronger anti-competitive effects. In contrast, there is no significant change in patent value after merger deals that are not expected to affect product market competition. The results highlight the potential benefits of granting firms some extent of market power.

In the second essay (Cyanide on the Rand: Competing Methods of Technology Transfer), my coauthor and I investigate the efficiency gains resulting from competing mechanisms of technology spillovers in the context of the cyanide method of gold extraction in South Africa in the 1890s. The study found knowledge spillovers at both the extensive and intensive margins, with mines managed by the same mining house more likely to adopt the cyanide process. The adoption of cyanide led to significant increases in the gold recovery rate, enabling marginally profitable mines to survive and new mines to open. The research also identified organizational and personnel networks as effective channels for transmitting technological know-how, while geographical spillovers were found to be localized.

Overall, this thesis enhances the understanding of the relationship between competition, innovation, and knowledge spillovers. The implications for policymakers and practitioners emphasize the importance of market dynamics and technology transfer in promoting economic growth and productivity. Future research could delve deeper into the interplay between competition, innovation, and technology spillovers, employ alternative innovation measures, analyze knowledge spillovers in diverse contexts, consider policy implications, and explore international and cross-industry perspectives. Such endeavors will advance knowledge in innovation, competition, and economic growth, deepening our understanding of technology spillovers' influence on economic outcomes.

Bibliography

Abebe, G., McMillan, M., and Serafinelli, M. (2022). Foreign direct investment and knowledge diffusion in poor locations. *Journal of Development Economics*, 158:102926.

Abramovitz, M. (1986). Catching Up, Forging Ahead, and Falling Behind. *The Journal of Economic History*, 46(2):385–406.

Abrams, D., Akcigit, U., and Popadak, J. A. (2020). Patent Value and Citations: Creative Destruction or Strategic Disruption? *NBER Working Paper*, #19647.

Aghion, P., Bloom, N., Blundell, R., Griffith, R., and Howitt, P. (2005). Competition and Innovation: an Inverted-U Relationship. *Quarterly Journal of Economics*, 120(2):701–728.

Aghion, P., Harris, C., Howitt, P., and Vickers, J. (2001). Competition, imitation and growth with step-by-step innovation. *Review of Economic Studies*, 68(3):467–492.

Aghion, P. and Howitt, P. (1992). A Model of Growth Through Creative Destruction. *Econometrica*, 60(2):323–351.

Aghion, P., Van Reenen, J., and Zingales, L. (2013). Innovation and Institutional Ownership. *American Economic Review*, 103(1):277–304.

Akcigit, U., Celik, M. A., and Greenwood, J. (2016). Buy, Keep, or Sell: Economic Growth and the Market for Ideas. *Econometrica*, 84(3):943–984.

Alcácer, J. and Gittelman, M. (2006). Patent citations as a measure of knowledge flows: The influence of examiner citations. *Review of Economics and Statistics*, 88(4):774–779.

Ali, A., Klasa, S., and Yeung, E. (2008). The limitations of industry concentration measures constructed with compustat data: Implications for finance research. *Review of Financial Studies*, 22(10):3839–3871.

Anand, J. and Singh, H. (1997). Asset redeployment, acquisitions and corporate strategy in declining industries. *Strategic Management Journal*, 18(S1):99–118.

Arora, A., Belenzon, S., and Lee, H. (2018). Reversed citations and the localization of knowledge spillovers. *Journal of Economic Geography*, 18(3):495–521.

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

Arrow, K. (1962). J., Economic Welfare and the Allocation of Resources to Inventions. *National Bureau Committee for Economic Research (Ed.), The Rate and Direction of Inventive Activity. Princeton, NJ*.

Atanassov, J. (2013). Do Hostile Takeovers Stifle Innovation? Evidence from Antitakeover Legislation and Corporate Patenting. *Journal of Finance*, 68(3):1097–1131.

Audretsch, D. B. and Feldman, M. P. (1996). R & D Spillovers and the Geography of Innovation and Production. *The American Economic Review*, 86(3):630–640.

Austin, D. (1993). An event-study approach to measuring innovative output: The case of biotechnology. *American economic review*, 83(2):253–258.

Autant-Bernard, C. (2001). The Geography Of Knowledge Spillovers And Technological Proximity. *Economics of Innovation and New Technology*, 10(4):237–254.

Autor, D. H., Dorn, D., Hanson, G. H., Pisano, G., and Shu, P. (2020). Foreign Competition and Domestic Innovation: Evidence from U.S. Patents. *American Economic Review: Insights*, 2(3):357–374.

Ayerst, S., Ibrahim, F., MacKenzie, G., and Rachapalli, S. (2023). Trade and Diffusion of Embodied Technology: An Empirical Analysis. *Working Paper*.

Azoulay, P., Graff Zivin, J. S., and Wang, J. (2010). Superstar extinction. *The Quarterly Journal of Economics*, 125(2):549–589.

Babina, T. and Howell, S. (2022). Entrepreneurial Spillovers from Corporate R&D. *Journal of Labor Economics*, Accepted Articles.

Bahar, D., Choudhury, P., Kim, D. Y., and Koo, W. W. (2023). Innovation on Wings: Nonstop Flights and Firm Innovation in the Global Context. *Management Science*, 69(10):5695–6415.

Banal-Estañol, A., Duso, T., Seldeslachts, J., and Szücs, F. (2022). R&D spillovers through RJV cooperation. *Research Policy*, 51(4):104465.

Barber, B. and Diestre, L. (2022). Can firms avoid tough patent examiners through examiner-shopping? Strategic timing of citations in USPTO patent applications. *Strategic Management Journal*, 43(9):1854–1871.

Baum, C. F., Lööf, H., Stephan, A., and Viklund-Ros, I. (2022). Innovation by start-up firms: The role of the board of directors for knowledge spillovers. *Research Policy*, 51(1):104375.

Bena, J. and Li, K. (2014). Corporate Innovations and Mergers and Acquisitions. *Journal of Finance*, 69(5):1923–1960.

Bernstein, J. I. and Nadiri, M. I. (1989). Research and development and intra-industry spillovers: An empirical application of dynamic duality. *Review of Economic Studies*, 56(2):249–269.

Bloom, N., Draca, M., and Van Reenen, J. (2016). Trade induced technical change? The impact of chinese imports on innovation, IT and productivity. *Review of Economic Studies*, 83(1):87–117.

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013a). Does management matter? Evidence from India. *Quarterly Journal of Economics*, 128(1):1057–1106.

Bloom, N., Schankerman, M., and van Reenen, J. (2013b). Identifying Technology Spillovers and Product Market Rivalry. *Econometrica*, 81(4):1347–1393.

Blundell, R., Griffith, R., and Van Reenen, J. (1999). Market share, market value and innovation in a panel of British manufacturing firms. *Review of Economic Studies*, 66(3):529–554.

Borjas, G. J. and Doran, K. B. (2012). The collapse of the Soviet Union and the productivity of American mathematicians. *The Quarterly Journal of Economics*, 127(3):1143–1203.

Braguinsky, S., Ohyama, A., Okazaki, T., and Syverson, C. (2021). Product Innovation, Product Diversification, and Firm Growth: Evidence from Japan's Early Industrialization. *American Economic Review*, 111(12):3795–3826.

Brav, A., Jiang, W., Ma, S., and Tian, X. (2018). How does hedge fund activism reshape corporate innovation? *Journal of Financial Economics*, 130(2):237–264.

Breschi, S., Lissoni, F., and Miguelez, E. (2017). Foreign-origin inventors in the USA: Testing for diaspora and brain gain effects. *Journal of Economic Geography*, 17(5):1009–1038.

Bwalya, S. M. (2006). Foreign direct investment and technology spillovers: Evidence from panel data analysis of manufacturing firms in Zambia. *Journal of Development Economics*, 81(2):514–526.

Cao, S. S., He, J., Lin, Z., and Ren, X. (2022). Venture Labor: A Nonfinancial Signal for Start-up Success. *SSRN Electronic Journal*.

Carley, M., Hedge, D., and Marco, A. (2015). What Is the Probability of Receiving a US Patent? *Yale JL & Tech.*, 17:203.

Chava, S. and Roberts, M. R. (2008). How does financing impact investment? The role of debt covenants. *Journal of Finance*, 63(5):2085–2121.

Chen, I. J., Hsu, P. H., Officer, M. S., and Wang, Y. (2020). The Oscar goes to...: High-tech firms' acquisitions in response to rivals' technology breakthroughs. *Research Policy*, 49(7):104078.

Chu, Y., Tian, X., and Wang, W. (2019). Corporate innovation along the supply chain. *Management Science*, 65(6):2445–2466.

Clarke, R. N. (1989). SICs as Delineators of Economic Markets. *The Journal of Business*, 62(1):17–31.

Cohen, L., Diether, K., and Malloy, C. (2013). Misvaluing Innovation. *Review of Financial Studies*, 26(3):635–666.

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

Comin, D. and Hobijn, B. (2010). An exploration of technology diffusion. *American Economic Review*, 100(5):2031–2059.

Conley, T. G. and Udry, C. R. (2010). Learning about a new technology: Pineapple in Ghana. *American Economic Review*, 100(1):35–69.

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

Dasgupta, P. and Stiglitz, J. (1980). Industrial Structure and the Nature of Innovative Activity. *The Economic Journal*, 90(358):266–293.

David, P. A. (1990). The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. *The American Economic Review*, 80(2):355–361.

De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics*, 135(2):561–644.

Duflo, E. and Saez, E. (2003). Social Impact on Retirement Plan Decisions. *The Quarterly Journal of Economics*, 118(3):815–842.

Dupas, P. (2009). What Matters (And What Does Not) in Households' Decision to Invest in Malaria Prevention? *American Economic Review*, 99(2):224–230.

Dutz, M. A. (1989). Horizontal mergers in declining industries: Theory and evidence. *International Journal of Industrial Organization*, 7(1):11–33.

Eckbo, B. E. (2014). Corporate takeovers and economic efficiency. *Annual Review of Financial Economics*, 6:51–74.

Ederer, F. and Manso, G. (2011). Incentives for innovation: Bankruptcy, corporate governance, and compensation systems. In *Handbook on Law, Innovation and Growth*, pages 90–111. Edward Elgar Publishing.

Ederer, F. and Manso, G. (2013). Is pay for performance detrimental to innovation? *Management Science*, 59(7):1496–1513.

Eichengreen, B. J. and Flandreau, M. (1985). *The Gold Standard in Theory and History*. Routledge, London.

Eissler, M. (1888). *The Metallurgy of Gold*. Crosby Lockwood and Son, London.

Eissler, M. (1895). *The Cyanide Process for the Extraction of Gold*. Crosby Lockwood and Son, London.

Eriksson, R. H. (2011). Localized Spillovers and Knowledge Flows: How Does Proximity Influence the Performance of Plants? *Economic Geography*, 87(2):127–152.

Fama, E. F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. Technical Report 2.

Fama, E. F. (1991). Efficient Capital Markets: II. *The Journal of Finance*, 46(5):1575–1617.

Fang, V. W., Tian, X., and Tice, S. (2014). Does Stock Liquidity Enhance or Impede Firm Innovation? *Journal of Finance*, 69(5):2085–2125.

Farrell, J. and Shapiro, C. (1990). Horizontal mergers: an equilibrium analysis. *The American Economic Review*, pages 107–126.

Fathollahi, M., Harford, J., and Klasa, S. (2022). Anticompetitive effects of horizontal acquisitions: The impact of within-industry product similarity. *Journal of Financial Economics*, 144(2):645–669.

Fee, C. E. and Thomas, S. (2004). Sources of gains in horizontal mergers: Evidence from customer, supplier, and rival firms. *Journal of Financial Economics*, 74(3):423–460.

Fischer, M. M., Scherngell, T., and Jansenberger, E. (2006). The geography of knowledge spillovers between high-technology firms in Europe: Evidence from a spatial interaction modeling perspective. *Geographical Analysis*, 38(3):288–309.

Foster, A. D. and Rosenzweig, M. R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, 103(6):1176–1209.

Foster, A. D. and Rosenzweig, M. R. (2010). Microeconomics of technology adoption. *Annual Review of Economics*, 2:395–424.

Futia, C. A. (1980). Schumpeterian Competition. *The Quarterly Journal of Economics*, 94(4):675–695.

Ghosh, A. and Morita, H. (2017). Knowledge transfer and partial equity ownership. *RAND Journal of Economics*, 48(4):1044–1067.

Giannetti, M., Liao, G., and Yu, X. (2015). The Brain Gain of Corporate Boards: Evidence from China. *Journal of Finance*, 70(4):1629–1682.

Glaeser, E. L., Kallal, H. D., Scheinkman, J. A., and Shleifer, A. (1992). Growth in Cities. *Journal of Political Economy*, 100(6):1126–1152.

Gomes-Casseres, B., Hagedoorn, J., and Jaffe, A. B. (2006). Do alliances promote knowledge flows? *Journal of Financial Economics*, 80(1):5–33.

González-Uribe, J. (2020). Exchanges of innovation resources inside venture capital portfolios. *Journal of Financial Economics*, 135(1):144–168.

Gort, M. (1969). An Economic Disturbance Theory of Mergers. *Quarterly Journal of Economics*, 83(4):624–642.

Gray, J. and McLachlan, J. A. (1933). A history of the introduction of the MacArthur-Forrest cyanide process to the Witwatersrand goldfields. *Journal of the Southern African Institute of Mining and Metallurgy*, 33(12):375–397.

Greenhalgh, C. and Rogers, M. (2006). The value of innovation: The interaction of competition, R&D and IP. *Research Policy*, 35(4):562–580.

Greenstone, M., Hornbeck, R., and Moretti, E. (2010). Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. *Journal of Political Economy*, 118(3):536–598.

Griliches, Z. (1979). Issues in Assessing the Contribution of R&D to Productivity Growth. *Bell Journal of Economics*, 10(1):92–116.

Griliches, Z. (1992). The Search for R&D Spillovers. *The Scandinavian Journal of Economics*, 94:29–47.

Grullon, G., Larkin, Y., and Michaely, R. (2019). Are US Industries Becoming More Concentrated? *Review of Finance*, 23(4):697–743.

Gu, L. (2016). Product market competition, R&D investment, and stock returns. *Journal of Financial Economics*, 119(2):441–455.

Hall, B. H. and Harhoff, D. (2012). Recent research on the economics of patents. *Annual Review of Economics*, 4:541–565.

Hall, B. H., Jaffe, A., and Trajtenberg, M. (2005). Market value and patent citations. *RAND Journal of Economics*, 36(1):16–38.

Hall, B. H. and Lerner, J. (2010). *The financing of R&D and innovation*, volume 1. Elsevier B.V.

Hankir, Y., Rauch, C., and Umber, M. P. (2011). Bank M&A: A market power story? *Journal of Banking and Finance*, 35(9):2341–2354.

Harhoff, D., Narin, F., Scherer, F. M., and Vopel, K. (1999). Citation frequency and the value of patented inventions. *Review of Economics and Statistics*, 81(3):511–515.

Hashmi, A. R. (2013). Competition and innovation: The inverted-u relationship revisited. *Review of Economics and Statistics*, 95(5):1653–1668.

He, J. and Huang, J. (2017). Product market competition in a world of cross-ownership: Evidence from institutional blockholdings. *Review of Financial Studies*, 30(8):2674–2718.

He, J. J. and Tian, X. (2013). The dark side of analyst coverage: The case of innovation. *Journal of Financial Economics*, 109(3):856–878.

He, J. J. and Tian, X. (2020). Institutions and Innovation. *Annual Review of Financial Economics*, 12:377–398.

Helmers, C., Patnam, M., and Rau, P. R. (2017). Do board interlocks increase innovation? Evidence from a corporate governance reform in India. *Journal of Banking and Finance*, 80:51–70.

Hirshleifer, D., Hsu, P. H., and Li, D. (2013). Innovative efficiency and stock returns. *Journal of Financial Economics*, 107(3):632–654.

Hoberg, G. and Phillips, G. (2010). Product market synergies and competition in mergers and acquisitions: A text-based analysis. *Review of Financial Studies*, 23(10):3773–3811.

Hoberg, G. and Phillips, G. (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5):1423–1465.

Hu, Z. (2023). Redeploying the Jockeys : The Intermediary Role of Venture Capital in Entrepreneurial Labor Markets. *Working Paper*.

Im, H. J., Park, Y. J., and Shon, J. (2015). Product market competition and the value of innovation: Evidence from US patent data. *Economics Letters*, 137:78–82.

Imrohoroğlu, A. and Tüzel, (2014). Firm-Level Productivity, Risk, and Return. *Management Science*, 60(8):2073–2090.

Jaffe, A. and Adams, J. D. (1996). Bounding effects of R&D. *NBER Working Paper*, No. 5544.

Jaffe, A., Trajtenberg, M., and Henderson, R. (1993). Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *Quarterly Journal of Economics*, 108(3):577–598.

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

Jaffe, A. B., Trajtenberg, M., and Fogarty, M. S. (2002). The Meaning of Patent Citations: Report on the NBER/Case Western Reserve Survey of Patentees. In Jaffe, A. and Trajtenberg, M., editors, *Patents, Citations, and Innovations: A Window on the Knowledge Economy*, page 379–401. MIT Press, Cambridge, MA.

Javorcik, B. S. and Spatareanu, M. (2011). Does it matter where you come from? Vertical spillovers from foreign direct investment and the origin of investors. *Journal of Development Economics*, 96(1):126–138.

Jensen, R. (1988). Information Capacity and Innovation Adoption. *International Journal of Industrial Organization*, 6:335–350.

Jenter, D. and Kanaan, F. (2015). CEO Turnover and Relative Performance Evaluation. *The Journal of Finance*, 70(5):2155–2184.

Juhász, R., Squicciarini, M., and Voigtländer, N. (2023). Technology Adoption and Productivity Growth: Evidence from Industrialization in France. *NBER Working Paper*, 27503.

Kamien, M. I. and Schwartz, N. L. (1975). Market Structure and Innovation : A Survey. *Journal of Economic Literature*, 13(1):1–37.

Kaustia, M. and Knüpfer, S. (2012). Peer performance and stock market entry. *Journal of Financial Economics*, 104(2):321–338.

Keil, J. (2017). The trouble with approximating industry concentration from Compustat. *Journal of Corporate Finance*, 45:467–479.

Kepler, J., Naiker, V., and Stewart, C. R. (2021). Stealth Acquisitions and Product Market Competition. *SSRN Electronic Journal*.

Kerr, W. R. (2008). Ethnic Scientific Communities and International Technology Diffusion. *The Review of Economics and Statistics*, 90(3):518–537.

Kerr, W. R. and Nanda, R. (2015). Financing Innovation. *Annual Review of Financial Economics*, 7:445–462.

Kogan, L. and Papanikolaou, D. (2019). Technological Innovation, Intangible Capital, and Asset Prices. *Annual Review of Financial Economics*, 11:221–242.

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

Kremer, M. and Miguel, E. (2007). The Illusion of Sustainability. *The Quarterly Journal of Economics*, 122(3):1007–1065.

Kuhn, J., Younge, K., and Marco, A. (2020). Patent citations reexamined. *The RAND Journal of Economics*, 51(1):109–132.

Lerner, J. and Seru, A. (2022). The Use and Misuse of Patent Data: Issues for Finance and Beyond. *The Review of Financial Studies*, 35(6):2667–2704.

Levin, D. Z. (1999). Transferring Knowledge within the Organization in the R&D Arena. *PhD Thesis, Northwestern University*.

Lin, C. Y. C. (2013). Strategic decision-making with information and extraction externalities: A structural model of the multistage investment timing game in offshore petroleum production. *Review of Economics and Statistics*, 95(5):1601–1621.

Lindsey, L. (2008). Blurring Firm Boundaries: The Role of Venture Capital in Strategic Alliances. *Journal of Finance*, 63(3):1137–1168.

Lippman, S. A. and McCardle, K. F. (1987). Does Cheaper, Faster, or Better Imply Sooner in the Timing of Innovation Decisions? *Management Science*, 33(8):1058–1064.

Liu, Z. (2008). Foreign direct investment and technology spillovers: Theory and evidence. *Journal of Development Economics*, 85(1-2):176–193.

López, L. and Vives, X. (2019). Overlapping ownership, R&D spillovers, and antitrust policy. *Journal of Political Economy*, 127(5):2394–2437.

Lougheed, A. L. (1989). The discovery, development, and diffusion of new technology: The cyanide process for the extraction of gold, 1887–1914. *Prometheus*, 7(1):61–74.

Loury, G. C. (1979). Market structure and innovation. *Quarterly Journal of Economics*, 93(3):395–410.

Lychagin, S. (2016). Spillovers, absorptive capacity and agglomeration. *Journal of Urban Economics*, 96:17–35.

Ma, S. (2021). Technological Obsolescence. *Working Paper*.

Majd, S. and Pindyck, R. S. (1987). Time to build, option value, and investment decisions. *Journal of Financial Economics*, 18(1):7–27.

Mann, H. T. (1909). The history and the development of the cyanide industry. *Masters Theses, University of Missouri*.

Mann, W. (2018). Creditor rights and innovation: Evidence from patent collateral. *Journal of Financial Economics*, 130(1):25–47.

Manski, C. F. (1993). Identification of endogenous social effects the reflection problem. *Review of Economic Studies*, 60(3):531–542.

Manso, G. (2011). Motivating Innovation. *Journal of Finance*, 66(5):1823–1860.

Marino, A., Mudambi, R., Perri, A., and Scalera, V. G. (2020). Ties that bind: Ethnic inventors in multinational enterprises' knowledge integration and exploitation. *Research Policy*, 49(9):103956.

Marshall, A. (1890). *Principles of Economics*. Palgrave Macmillan UK, London, 1st edition.

Martin, J. (1929). Group Administration in the Gold Mining Industry of the Witwatersrand. *The Economic Journal*, 39(156):536–553.

Masulis, R. W., Wang, C., and Xie, F. (2012). Globalizing the boardroom-The effects of foreign directors on corporate governance and firm performance. *Journal of Accounting and Economics*, 53(3):527–554.

Matray, A. (2021). The local innovation spillovers of listed firms. *Journal of Financial Economics*, 141(2):395–412.

McCardle, K. F. (1985). Information Acquisition and the Adoption of New Technology. *Management Science*, 31(11):1372–1389.

Mitchell, M. L. and Mulherinb, J. . H. (1996). The impact of industry shocks on takeover and restructuring activity. *Journal of Financial Economics*, 41(2):193–229.

Mokyr, J. (1990). *The lever of riches: Technological creativity and economic progress*. Oxford University Press.

Murata, Y., Nakajima, R., Okamoto, R., and Tamura, R. (2014). Localized knowledge spillovers and patent citations: A distance-based approach. *Review of Economics and Statistics*, 96(5):967–985.

Nickell, S. J. (1996). Competition and corporate performance. *Journal of Political Economy*, 104(4):724–746.

Nuvolari, A., Tortorici, G., and Vasta, M. (2023). British-French Technology Transfer from the Revolution to Louis Philippe (1791-1844): Evidence from Patent Data. *Journal of Economic History*, 83(3):833–873.

Olley, G. S. and Pakes, A. (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, 64(6):1263–1297.

Orlando, M. J. (2004). Measuring Spillovers from Industrial R&D: On the Importance of Geographic and Technological Proximity. *The RAND Journal of Economics*, 35(4):777–786.

Oster, E. and Thornton, R. (2012). Determinants of Technology Adoption: Private Value and Peer Effects in Menstrual Cup Take-Up. *Journal of the European Economic Association*, 10(6):1263–1293.

Pakes, A. (1985). On patents,R&D,and the stock market. *Journal of Political Economy*, 93:390–409.

Perry, M. K. and Porter, R. H. (1985). Oligopoly and the incentive for horizontal merger. *The American Economic Review*, 75(1):219–227.

Phillips, G. M. and Zhdanov, A. (2012). R&D and the incentives from merger and acquisition activity. *Review of Financial Studies*, 26(1):34–78.

Phillips Sawyer, L. (2019). US Antitrust Law and Policy in Historical Perspective. In Butler, J., editor, *Oxford Research Encyclopedia of American History*. Oxford University Press.

Pierce, J. R. and Schott, P. K. (2016). The surprisingly swift decline of US manufacturing employment. *American Economic Review*, 106(7):1632–1662.

Pogue, T. E. (2006). The Evolution of Research Collaboration in South African Gold Mining: 1886-1933. *PhD Thesis, Universiteit Maastricht*.

Porter, M. E. (1990). *The Competitive Advantage of Nations*. Free Press, New York.

Raith, M. (2003). Competition , Risk , and Managerial Incentives. *American Economic Review*, 93(4):1425–1436.

Rhodes-Kropf, M. and Viswanathan, S. (2004). Market valuation and merger waves.

Scherer, F. M. (1967). Research and development resource allocation under rivalry. *The Quarterly Journal of Economics*, 81(3):359–394.

Schnitzer, M. and Watzinger, M. (2022). Measuring the Spillovers of Venture Capital. *Review of Economics and Statistics*, 104(2):276–292.

Schumpeter, J. (1942). *Capitalsims, Socialsm, and Democracy*, volume 20. Harper & Row, New York.

Serafinelli, M. (2019). "Good" firms, worker flows, and local productivity. *Journal of Labor Economics*, 37(3):747–792.

Seru, A. (2014). Firm boundaries matter: Evidence from conglomerates and R&D activity. *Journal of Financial Economics*, 111(2):381–405.

Shahrur, H. (2005). Industry structure and horizontal takeovers: Analysis of wealth effects on rivals, suppliers, and corporate customers. *Journal of Financial Economics*, 76(1):61–98.

Shleifer, A. and Vishny, R. W. (2003). Stock market driven acquisitions. *Journal of Financial Economics*, 70(3):295–311.

Singh, J. (2005). Collaborative networks as determinants of knowledge diffusion patterns. *Management Science*, 51(5):756–770.

Singh, J. (2007). Asymmetry of knowledge spillovers between MNCs and host country firms. *Journal of International Business Studies*, 38(5):764–786.

Singh, J. and Marx, M. (2013). Geographic constraints on knowledge spillovers: Political borders vs. spatial proximity. *Management Science*, 59(9):2056–2078.

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

Stoyanov, A. and Zubanov, N. (2012). Productivity spillovers across firms through worker mobility. *American Economic Journal: Applied Economics*, 4(2):168–198.

Syverson, C. (2011). What determines productivity. *Journal of Economic Literature*, 49(2):326–365.

Tambe, P. and Hitt, L. M. (2014). Job hopping, information technology spillovers, and productivity growth. *Management Science*, 60(2):338–355.

Thompson, P. (2006). Patent citations and the geography of knowledge spillovers: Evidence from inventor-and examiner-added citations. *The Review of Economics and Statistics*, 88(2):383–388.

Thompson, P. and Fox-kean, M. (2005). Patent Citations and the Geography of Knowledge Spillovers : A Reassessment. *The American Economic Review*, 95(1):450–460.

Trajtenberg, M. (1990). A penny for your quotes: patent citations and the value of innovations. *RAND journal of economics*, 21(1):172–187.

Verbrugge, B., Lanzano, C., and Libassi, M. (2021). The cyanide revolution: Efficiency gains and exclusion in artisanal- and small-scale gold mining. *Geoforum*, 126:267–276.

von Graevenitz, G., Graham, S. J., and Myers, A. F. (2022). Distance (still) hampers diffusion of innovations. *Regional Studies*, 56(2):227–241.

Wagner, S., Hoisl, K., and Thoma, G. (2014). Overcoming localization of knowledge—the role of professional service firms. *Strategic Management Journal*, 35:1671–1688.

Waldinger, F. (2012). Peer effects in science: Evidence from the dismissal of scientists in Nazi Germany. *Review of Economic Studies*, 79(2):838–861.

Wang, K., Yu, X., and Zhang, B. (2022). Panda Games Corporate Disclosure in the Eclipse of Search. *Management Science*, 69(6):3263–3284.

Yuan, R. and Wen, W. (2018). Managerial foreign experience and corporate innovation. *Journal of Corporate Finance*, 48:752–770.

Zacchia, P. (2020). Knowledge Spillovers through Networks of Scientists. *Review of Economic Studies*, 87(4):1989–2018.

Zheng, Y. and Wang, Q. (2020). Shadow of the great firewall: The impact of Google blockade on innovation in China. *Strategic Management Journal*, 41(12):2234–2260.