Title
The Evolution of Technological Knowledge Across Space and Time

Author
Esposito, Christopher Ross

Publication Date
2021

UNIVERSITY OF CALIFORNIA

Los Angeles

The Evolution of Technological Knowledge Across Space and Time

A dissertation submitted in partial satisfaction of the requirements for the degree Doctor of Philosophy in Geography

by

Christopher Ross Esposito

2021

ABSTRACT OF THE DISSERTATION

The Evolution of Technological Knowledge Across Space and Time

by

Christopher Ross Esposito

Doctor of Philosophy in Geography

University of California, Los Angeles, 2021

Professor David L. Rigby, Chair

Abstract

Technological change is a powerful force in economic and social life. Technological change is both an endogenous and a disruptive process because inventors create new technologies by recombining existing technological ideas and because new technologies often drive older technologies and their associated capital and skills into obsolescence. Technological disruption resonates in the economies of cities, producing both local economic growth and decline, because city-regions are a scale at which many of the factors of production are coordinated.

In the existing literature, there is broad agreement that knowledge builds on itself endogenously, and there is some recognition that innovation is disruptive with consequences for city-regions. Despite these acknowledgments, the sources of knowledge that inventors used to create historical inventions have not been systematically documented, the question of how new city-regions enter the process of endogenous knowledge production has not been resolved, the geographical distribution of breakthrough innovation has not been described nor

explained, and the mechanisms through which inventors amass the technological knowledge needed to innovate in rapidly-evolving knowledge environments have not been adequately studied.

In light of the above research gaps, this dissertation makes four contributions. First, it develops a method called *knowledge phylogenetics* and uses that method to create a long-run genealogy of technological knowledge containing over 8 million patented inventions created between 1836 and 2014. Second, it uncovers a general process that city-regions go through as they begin to become centers for innovation, involving the importation of non-local and disruptive ideas that are used to initiate local knowledge production. Third, it documents the extent to which breakthrough innovation is concentrated in large and knowledge-diverse cities, how that concentration changed over the 20th century, and how those changes resulted from asymmetric improvements in different types of communication technologies. Fourth, it calculates the productivity benefits that inventors receive from working in teams and from the experiences that they accumulate over time. In this regard, the dissertation shows that inventors do not benefit from the experience that they accumulate over time because inventors struggle to learn quickly enough to keep pace with the advances made in their knowledge fields.

The dissertation of Christopher Ross Esposito is approved.

Jacob Gates Foster

Allen J. Scott

Michael C. Storper

David L. Rigby, Committee Chair

University of California, Los Angeles

2021

Table of Contents

List of Figures

List of Tables

Acknowledgments

This research was supported by National Science Foundation Graduate Research Fellowship Grant No. DGE-1144087, a National Academy of Science Fellowship for Study at the International Institute for Advanced Systems Analysis, and a UCLA Graduate Research Mentorship Fellowship.

Chapter 5 of this dissertation is a version of Esposito, C., and van der Wouden, F., Learning, Fast and Slow: The Returns to Team Size and Maturity for High-Impact Innovation. Manuscript in Preparation. Frank van der Wouden contributed disambiguated records of U.S. patent inventors to the chapter and assisted in the conceptual and empirical development of the chapter. He is a cherished colleague and collaborator.

I thank my Ph.D. adviser and committee chair David Rigby. David's openness to unconventional ideas made crafting a dissertation under his guidance a liberating experience. I thank my committee members Jacob Foster, Allen Scott, and Michael Storper. Jacob introduced me to methods and theories that I will grow into over the course of my career. Allen brought to my research his expansive knowledge of the many and varied agglomerations that have formed and functioned across time. Michael's ability to break down arguments and expose their strong and weak parts has, I hope, made my research robust. I additionally thank my collaborators Jerry Nickelsburg and Edward Leamer. Although our study of Los Angeles' minimum wage increases is not included in my dissertation, the lessons that they shared with me in econometrics and in patience (our study has spanned several years) were central to completing my dissertation.

Scholarship is advanced by individuals, collaborators, and communities, and I am indebted to the rich community of economic geographers. I particularly would like to thank Pierre-Alex Balland, Ron Boschma, Koen Frenken, and Dieter Kogler, as well as my excellent colleagues and close friends whom I have worked alongside at UCLA, including

Frank van der Wouden, Carsten Rietmann, Christoph Roesler, Melissa Haller, Anthony Frigon, and Yiou Zhang.

Lastly, I thank my family, who nurtured my intellectual curiosity, and my partner Ashley, who is also my most challenging debate partner. Her influence has found its way onto the pages of my dissertation: into the formulation of its theories, and into the specification of its regression models.

Vita

Christopher R. Esposito
Ph.D. Candidate
Department of Geography
University of California, Los Angeles

EDUCATION

UCLA, Los Angeles, California
Master of Arts in Geography, awarded 2016
Thesis: *An Evolutionary Theory of Regional Economic Growth and Change*
Committee: Michael Storper, Lynne Zucker, David Rigby (chair)
C.Phil in Geography, awarded 2018

Colgate University, Hamilton, NY
Bachelor of Arts, 2014
magna cum laude
High honors, Economics
High honors, Geography

PEER-REVIEWED PUBLICATIONS

Esposito, C., Rigby, D. 2019. Buzz and Pipelines: The Costs and Benefits of Local and Non-Local Interaction. *Journal of Economic Geography* 19 (3). *Awarded the 2019 UCLA Geography Graduate Student Publication Award.*

Meyer, W., Esposito, C. 2014. Burgess and Hoyt in L.A.: Testing the Chicago Models in the Automobile-Age American City. *Urban Geography* 36 (2).

Meyer, W., Esposito, C. 2014. Residential Patterns in the Pre-Automobile American City. *Geographical Review* 104 (3).

PAPERS UNDER PEER REVIEW

Esposito, C. The Emergence of Knowledge Production in New Places. Revise and Resubmit at *Journal of Economic Geography*.

WORKING PAPERS

Esposito, C., Leamer, E., and Nickelsburg, J. Who Paid Los Angeles' Minimum Wage? Side-by-Side Minimum Wage Experiments in Los Angeles.

Esposito, C. Constructing and Visualizing the Tree of Technology.

Esposito, C. Creative Destruction in Space: Lessons from the Information Technology Industry and Silicon Valley.

Esposito, C. The Geography of Breakthrough Inventions.

Esposto, C., and Van der Wouden, F. Learning, Fast and Slow: The Returns to Maturity and Collaboration for High-Impact Innovation in the United States between 1836 and 1975.

AWARDS AND HONORS

National Science Foundation Graduate Research Fellow
May 2015 – 2020
$102,000 in funding plus tuition and fees

2019 UCLA Geography Graduate Student Publication Award
Awarded for "The Costs and Benefits of Local and Non-Local Interaction," published in *Journal of Economic Geography* (2019)

International Institute of Applied Systems Analysis (IIASA) Summer Fellowship
July 2018 – Aug. 2018
Funded by the U.S. National Academy of Sciences
$5,500 in funding

UCLA Graduate Research Fellow
Sept. 2015 – June 2016
$26,000 in funding plus tuition and fees

Colgate University Shannon McCune Award for Excellence in Geography
May 2014

TEACHING EXPERIENCE

M236: Regional Economics and Development. Graduate level course
Winter 2019, winter 2020, and winter 2021

GEOG 7: Introduction to GIS
Fall 2019 and fall 2020. Online course (both terms)

GEOG 5: People and Earth's Ecosystems
Spring 2019 and spring 2021

ACADEMIC AND PROFESSIONAL SERVICE

Global E-Seminar Series Co-Founder and Director
Seminars in Economic Geography (SEG)
http://www.seminarsineconomicgeography.com
April 2020-present

Pan-Social Sciences Methods Workshop Organizer
Supercomputing for Social Scientists: Getting Going on UCLA's Hoffman Supercomputer
Jan. 2020

JOURNAL REFEREE SERVICE

Research Policy (2x), *Regional Studies* (3x)

PRESENTATIONS AT SELECT CONFERENCES

Association of American Geographers | Urban Economics Association | North American Regional Science Association | Geography of Innovation | Global Conference in Economic Geography | Seminars in Economic Geography |

Chapter 1: Introduction

Since Solow (1956), the accumulation of technological knowledge has been accepted as the leading cause of long-run economic growth and change. Inventors and scientists create technological knowledge by recombining existing ideas (Nelson and Winter, 1982; Romer, 1990). Because individuals specialize in narrow areas of expertise, the recombination of ideas often involves interaction between agents (Lundvall, 1994; Powell et al., 1996).

Spatial relationships are fundamental to the accumulation of technological knowledge because spatial proximity and separation between actors influences their patterns of interaction (Balland, 2012; van der Wouden, 2020). By influencing these patterns, spatial relationships tilt the direction of technological change (Catalini, 2017). In addition, technological change is fundamental to the configuration of spatial relationships, such as inter-regional income inequality. The accumulation of technological knowledge is an unstable process—new ideas build on older ones, but new ideas may also drive older ideas into obsolescence. This instability produces shifting spatial concentrations of technological power and economic opportunity (Storper and Walker, 1991).

An evolutionary framework is powerful for studying both the influence of geography on innovation and the influence of innovation on geography. An evolutionary framework begins with the same observation as Romer (1990) and Nelson and Winter (1982), that knowledge is endogenously created by using existing ideas to create new ones. An evolutionary framework also acknowledges that ideas are often exchanged locally, so endogenous knowledge production is a place-dependent process (Storper and Venables, 2004; Martin and Sunley, 2006; Essleztichler and Rigby, 2010). However, a full evolutionary framework adds that the accumulation of technological knowledge is not merely cumulative, meaning that the state of the art is not the sum of all prior inventions. Instead, technologies interact with one-another through complex relationships (Foster et al., 2013). In Chapter 3 of this dissertation, I share the examples of the junction transistor and the vacuum tube amplifier to illustrate one of the complex ways in which technologies may interact. The junction semiconductor drove the vacuum tube amplifier into obsolescence after it was introduced, and helped to enable the relocation of the information technology industry's core innovative activities from of the Northeast of the United States to the country's south and west. Inter-technology relationships such as this example of disruption are difficult to rationalize using traditional models of endogenous knowledge production, but they are powerfully captured by ecological models of knowledge evolution (Foster et al., 2013).

The existing research on innovation and its geography does not develop a complete evolutionary framework that encompasses both endogenous knowledge production and inter-technology relationships. This critique extends to the sub-discipline of evolutionary economic geography, where the common practice of aggregating knowledge to large technological fields obscures any true evolutionary process in knowledge creation, namely the recombination of existing ideas to create new ones (Hidalgo et al., 2007; Neffke et al., 2013). In addition, the paradigm of technological relatedness, which is commonplace in evolutionary economic geography, compresses qualitatively distinct relationships between technologies into ordinal similarity metrics (Hidalgo et al., 2018). This reduction prohibits the technological relatedness literature from studying the many and varied ways in which technologies interact. As a consequence, the technological relatedness literature is unable to analyze uneven technological development, including how stocks of technological ideas influence the usefulness of other stocks of ideas, and how the stocks of ideas accumulated in one geographical region may influence the usefulness of ideas accumulated in another.

In light of the current state of the literature, the purpose of this dissertation is to develop and deploy a new evolutionary study of the creation of technological knowledge across space and time that is grounded in a micro-level theory of invention and embraces the complex interactions between technologies. The dissertation has two main parts. The first part develops a novel dataset that records the line of descent between individual inventions. These data are needed in order to study technological evolution over long periods of time. The resulting dataset links over 8 million patents granted by the U.S. Patent and Trademark Office between 1836 and 2014 to their knowledge-based antecedents. Chapter 2 of this dissertation elaborates on the motivations for this data construction exercise, describes the methods used to construct that new data, and validates the data's accuracy.

The second part of the dissertation uses the newly constructed data to address three questions that pertain to innovation and its geography. The first question, which I examine in Chapter 3, is how new geographical concentrations of innovative activity emerge. The economic geography and agglomeration theory literatures observe that innovative activities are concentrated in space, but they rarely analyze the formative process of those concentrations. A traditional application of the spatialized endogenous knowledge growth model, wherein inventors create new ideas locally by recombining existing local ideas, cannot explain the emergence of new innovative centers because new innovative centers lack sufficient local ideas to be recombined. In the process I examine in Chapter 3, inventors transfer knowledge to new places by sourcing ideas across space and by subsequently using those imported ideas to create many more ideas locally. Thus, new innovative places are born from ideas sourced from older innovative places.

The second question I examine is why the invention of breakthroughs tends to occur in spatially concentrated locations. I examine this issue in Chapter 4 where I develop a definition of technological breakthroughs as the subset of inventions that punctuate the long-run knowledge evolution by recombining existing ideas in radically imaginative and useful ways. Using data from patents awarded to U.S.-based inventors between 1900 and 1999, I show that the intensity of the spatial concentration of the production of breakthroughs has changed over time as the relative strength of long-distance communication technologies, the knowledge intensity of breakthrough inventions, and the disruptiveness of the regime of technological change have evolved.

The third question I examine is how the collection of knowledge helps inventors to create high-impact inventions. I examine this issue in Chapter 5. Inventors collect knowledge in two primary ways: they collect knowledge into organizations by building collaborative teams, and they collect knowledge over time through accumulated experiences. My analysis shows that inventors benefit from knowledge pooled in teams, but that they do not benefit from knowledge accumulated over time. This latter result is surprising, because mature inventors with more experience should have knowledge of a greater number of ideas that they can deploy to create of high-impact inventions. In this chapter, I reconcile this result by showing that the rate at which inventors learn new ideas tends to be lower than the rate at which other inventors introduce new ideas to the knowledge fields they are working in. In particular, young and relatively inexperienced inventors enter patenting professions with current skills that they use to make high-impact inventions. Older inventors may encounter difficulties learning how to use these inventions to make new inventions, and so they become less capable over time.

1.1) Bibliography

Balland, P.A. (2012). Proximity and the Evolution of Collaboration Networks: Evidence from Research and Development Projects Within the Global Navigation Satellite System Industry. *Regional Studies* 46(6). 741-756.

Catalini, C. (2017). Microgeography and the Direction of Inventive Activity. *Management Science* 64(9). 4348-4364.

Essletzbichler, J., and Rigby, D. (2010). Generalized Darwinism and Evolutionary Economic Geography. In Boschma, R., and Martin, R. (Eds.), *The Handbook of Evolutionary Economic Geography*. Edward Elgar, Cheltenham, UK.

Foster, D., Rorick, M., Gesell, T., Feeney, L., and Foster, G. Dynamic Landscapes: A Model of Context and Contingency in Evolution. *Journal of Theoretical Biology* (334). 162-172.

Hidalgo, C., Klinger, B., Barabasi, A.L., and Hausmann, R. (2007). The Product Space Conditions the Development of Nations. *Science* 317(5837).

Hidalgo, C., Balland, P.A., Boschma, R., Delgado, M., Feldman, M., Frenken, K., Glaeser, E., He, C., Kogler, D., Morrison, A., Neffke, F., Rigby, D., Stern, S., Zheng, S., and Zhu, S. (2018). The Principle of Relatedness. *International Conference on Complex Systems*.

Lundvall, B.A., and Johnson, J. (1994). The Learning Economy. *Journal of Industry Studies* 1(2). 23-42.

Martin, R., and Sunley, P. (2006). Path Dependence and Regional Economic Evolution. *Journal of Economic Geography* 6(4). 395-437.

Neffke, F., Henning, M., and Boschma, R. (2013). How Do Regions Diversify Over Time? Industry Relatedness and the Development of New Growth Paths in Regions. *Economic Geography* 87(3). 237-265.

Nelson, R., and Winter, S. (1982). *An Evolutionary Theory of Economic Change*. Harvard University Press.

Romer, P. (1990). Endogenous Technological Change. *Journal of Political Economy* 98(5). S71-S102.

Powell, W., Koput, K., and Smith-Doerr, L. (1996). Interorganizational Collaboration and the Locus of Innovation: Networks of Learning in Biotechnology. *Administrative Science Quarterly* 41(1). 116-145.

Solow, R. (1956). A Contribution to the Theory of Economic Growth. *The Quarterly Journal of Economics* 70(1). 65-94.

Storper, M., and Venables, A. (2004). Buzz: Face-To-Face Contact and the Urban Economy. *Journal of Economic Geography* 4(1). 351-370.

Storper, M., and Walker, R. (1991). *The Capitalist Imperative: Territory, Technology, and Industrial Growth*. Basil Blackwell, New York.

van der Wouden, F. (2020). A History of Collaboration in U.S. Invention: Changing Patterns of Co-Invention, Complexity, and Geography. *Industrial and Corporate Change* 29(3). 599-619.

Chapter 2: Development of Methods and Data

2.1) Introduction

The evolutionary study of technological knowledge creation has been impeded by a lack of harmonized long-run records of the sources of knowledge that inventors use to generate new ideas. Patent citation records are broadly used to study how inventors source technological knowledge, but patent citations carry two limitations. First, because many citations are added by patent examiners and the attorneys of patent applicants, the extent to which citations represent knowledge spillovers is debated (Arora et al., 2019). Second, the United States Patent and Trademark Office (USPTO) did not require patents to cite prior art before 1947, so patent citation records are unreliable before this date (Akcigit et al., 2017; Berkes, 2020).

In this chapter, I develop a new method to trace the flow of knowledge between individual inventions. My method is called *knowledge phylogenetics*. I develop this method by integrating a model of recombinatory invention as developed by management theorists into a phylogenetics algorithm (Levinthal, 1996; Fleming and Sorenson, 2001). Phylogenetics is a method commonly practiced by biologists to trace evolutionary descent between biological species and by linguists to trace evolutionary descent between languages (Fox, Fisher, and Layton, 1999; Mace and Holden, 2005; Wiley and Lieberman, 2011). The insertion of a recombinatory model of invention into the phylogenetics algorithm allows me to trace the evolutionary descent between inventions based on the heuristics that inventors use to develop new technologies.

Theorists of technological change generally agree that inventors create new technologies by combining existing ideas in new ways (Mokyr, 1990; Arthur, 2009). The recombinatory and incremental nature of invention is a natural outcome of inventors' bounded rationality (Simon, 1955). However, assembling complex technologies overwhelms the mental capacity of inventors because new inventions contain many parts that are arranged in irregular ways (Fleming and Sorenson, 2001; Broekel, 2019). Inventors thus rarely create new technologies from scratch and instead build on existing ideas. Car makers do not need to reinvent the wheel, nor do computer hardware engineers need to develop their own processor chips and video cards. Motherboards, for instance, contain all the computer chips needed to perform most of the core processing functions of computers in a pre-packaged module. Because inventors can use pre-developed modules, they can focus their attention on making connections between existing modules rather than developing modules anew (Foster and Evans, 2019).

My application of technological phylogenetics uses the subclassification codes that the USPTO assigns to patents to identify the existing ideas combined in each invention. The USPTO classifies all utility patents using its highly detailed classification scheme. At the highest level of granularity, the USPC classification scheme contains over 160,000 unique subclass codes that describe the individual components contained in each patented invention (Fleming and Sorenson, 2001). Because technological knowledge is defined by inventors' awareness of how technological elements, or components can be assembled to create functioning tools, the detailed subclassification codes listed on each patent indicates the recombinant know-how embedded in each patented technology (Fleming, 2001; Arthur, 2009). This argument is illustrated by the example of the patent granted to Thomas Edison for the incandescent light bulb (USPTO patent number 223898). Edison's bright idea was that a vacuum chamber slows the combustion of a carbon filament. The physical components Edison used to create a vacuum chamber for a filament – vacuum-tight joints to seal the bulb and a carbon filament – appear on his patent with the classification codes H01J5/24 and H01K1/14. The USPTO defines these codes as "vacuum-tight joints between insulating parts of vessel" and "incandescent bodies characterized by shape."

Because the subclassification codes listed on a patent indicate the recombinant knowledge embedded in a technology, technological phylogenetics predicts that when two timestamped patents share many subclassification codes, they are linked by flows of knowledge (Foster and Evans, 2019).

2.2) Method for Knowledge Tree Construction

In this section I describe how I use knowledge phylogenetics to construct a record of knowledge flow between inventions. I begin with the raw public files of granted patents and USPC subclass assignments available on PatentsView for all utility patents granted between 1836 and 2014.[1] This dataset contains the entire collection of the 8.7 million utility patents granted for inventions that have received intellectual property protection in the United States since the U.S. patent office was rebuilt in 1836 after the building was burned to the ground in a fire. Because many international firms and inventors seek intellectual property protection in the U.S., the dataset has relatively strong global coverage.

The USPTO assigns each patent to one or more USPC subclasses. I extract these codes when downloading the raw data from PatentsView. Most patents are assigned between 2 and 6 subclassification codes; however, a very small number of patents are assigned more than 100 codes. To make the dataset less cumbersome, I discard excess subclassification codes on patents by selecting only the first 8 codes from each patent. By selecting the first 8 codes on each patent, I retain each patent's primary subclass.

The phylogenetics algorithm begins by selecting the most recently granted patent and recording its components based on its USPC subclassification codes. I define recombinant knowledge as knowledge of components and the interactions of those components; so for each patent I generate all combinations of degree *n* of its components, where *n* is the number

[1] Download link: https://www.patentsview.org/download/

of components in a patent.[2] For example, if a focal patent (FP) contains the USPC subclassification codes A, B, C, I generate the knowledge vector expressed by Equation 2.1:

$$(2.1)\ Knowledge_{FP} = [\ A \mid B \mid C \mid AB \mid BC \mid AC \mid ABC\ ]$$

Each element in $Knowledge_{FP}$ denotes a single unit of knowledge and the length of $Knowledge_{FP}$ indicates the total quantity of knowledge units embedded in the focal patent. The knowledge units in $Knowledge_{FP}$ are used to link the FP to its knowledge-based predecessors, or "parent patents," based on the number of knowledge units that are found in both the focal patent and a possible parent patent. To identify the possible parents of a focal patent, I search for overlapping knowledge units in all patents that were granted before the focal patent. I identify the patents granted before each focal patent by sorting patents by their ID number. I do not constrain the time window during which a parent patent can serve as a source of knowledge for a child patent because inventors often build on both old and new ideas (Mukherjee et al., 2017).[3] For each possible parent that fits the simple temporal criterion, I generate a shared knowledge vector (SKnowledge) to record the knowledge units that appear in both the focal patent and in the parent. For example, if a possible parents' knowledge vector, $Knowledge_{PP}$, is given by Equation 2.2:

$$(2.2)\ Knowledge_{PP} = [\ B \mid C \mid D \mid BC \mid CD \mid BD \mid BCD\ ]$$

The shared knowledge vector of the FP and the PP, $SKnowledge_{FP,PP}$, is taken as the intersection of the $Knowledge_{FP}$ vector and the $Knowledge_{PP}$ vector:

$$(2.3)\ SKnowledge_{FP,PP} = [\ B \mid C \mid BC\ ]$$

[2] The knowledge in a technology is embedded in the individual components in that technology *and* the way those components are interconnected. For example, Edison's light bulb was created through Edison's knowledge of the existence of the bamboo filament and the vacuum-tight joints as independent components, and through his understanding that these components work synergistically when assembled together.

[3] While I do not constrain the time window, over 90% of child patents draw knowledge from parents that are less than 20 years old.

The length of the above $SKnowledge_{FP,PP}$ vector indicates that in this example the focal patent FP sourced 3 units of knowledge from the potential parent.

When an FP has multiple potential parents for an individual unit of knowledge, I assign a fractional weight to the edge based on the number of possible parents for that knowledge unit. For example, if two possible parents contain the component [B], I assume that the FP sources 0.5 units of knowledge from the [B] in the first possible parent and 0.5 units from the second. In practice, knowledge units of length 1 such as [B] tend to be found on many potential parent patents, while longer knowledge units such as [BC] tend to be found on fewer.[4] Therefore, the fractional assignment of weights emphasizes non-ubiquitous and lengthy combinations of knowledge when predicting the parents of individual patents.

The algorithm repeats the process described above for all elements in the $Knowledge_{FP}$ vector. After iterating through each of the elements in the subclass vector, the algorithm moves on to the next most recently granted patent and repeats the process. The phylogenetics algorithm outputs a directed acyclic weighted graph with over 8 million patent nodes and hundreds of millions of edges representing the knowledge flows between them.

The resulting graph is very large (50 GB) because it contains many edges between patents with very low edge weights. Therefore, I trim the size of this graph by omitting all edges between patents with weights less than one.

2.3) Methods to Detect Patent Impact

I use the resulting technology tree to identify the impact of individual patents and the knowledge sourcing strategies of inventors. To compute the impact of individual patents, I calculate their out-degree in the tree of knowledge flow. The out-degree is a continuous

[4] At the most disaggregated level of subclasses, there are about 160,000 unique codes. While it improbable that any two randomly chosen patents will share a single subclass code, the probability that two randomly chosen patents will share two or more subclasses is exponentially smaller. Moreover, the granularity of the classification scheme and its combinations allows for very specific matches between child and parent patents.

variable which records the quantity of knowledge introduced by an earlier invention that is used by subsequent inventions.

In Chapters 3 and 5, I also generate binary measures of patent impact. While the binary measures contain less variation than the continuous measure, the continuous measure has a skewed distribution until the later years of the study period because many older patents have an out-degree of 0. The binary measures do not have this skewed distribution. To compute the binary measures, I use a threshold out-degree value. Patents with an out-degree that exceeds the threshold are defined as "high-impact" patents, while patents with an out-degree that does not exceed the threshold are defined as "low-impact." In Chapter 3, the threshold is the top 10% of the impact distribution of all patents granted in the same 5-year period. In Chapter 5, I use a somewhat higher threshold – the top 5% of the impact distribution of all patents granted in the same year – because the study period of Chapter 5 extends back to 1836 and therefore includes a time period during which very many patents have an out-degree of 0.

In Figure 2.1, I plot the mean impact of patents and the value at the top 5% of the impact distribution by year between 1836 and 1975. The mean out-degree of patents hovers near 1 until the 1930s. This low mean out-degree reflects the fact that most pre-1930 patents have an out-degree of 0, indicating that they have no discernible impact on subsequent innovation. The mean out-degree rises steadily in the 1930s and exceeds 3 by 1975. The 90th percentile of the out degree distribution is slightly above 2.5 in the years leading up to 1930, after which it climbs sharply and reaches 7.5 by 1975. Right-truncation is limited in these out-degree measures because the knowledge flow graph contains all USPTO patents granted up to 2014.

Figure 2.1: Average Out-Degree of Patents by Year

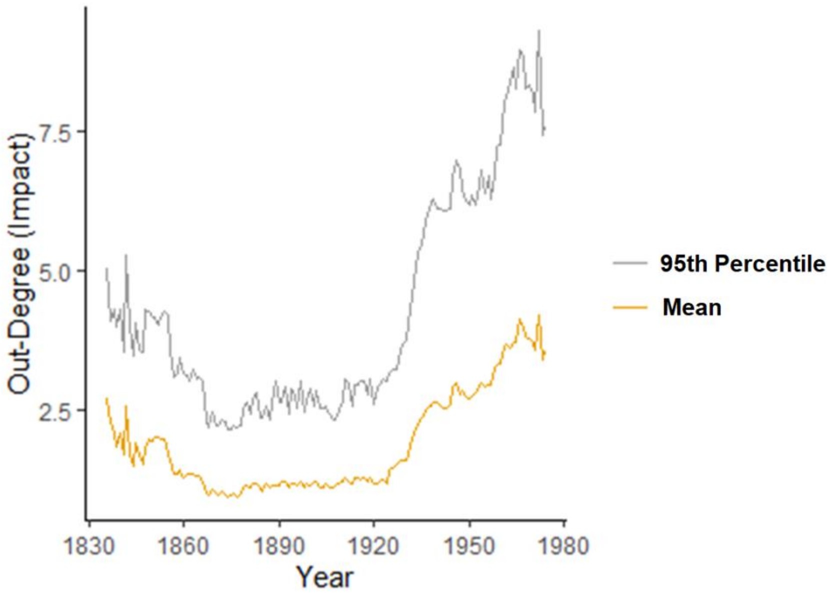

2.4) Validation of Patent Impact

I perform two validation checks to test whether the impact measures calculated using the out-degree of the technology tree corresponds to external sources. In the first validation check, I test whether patents identified by technological historians as uniquely consequential inventions have a higher mean out-degree than a comparison set of patents.

I use two sets of historian-identified patents for this purpose. The first is provided by Rogers (2011), which lists over 100 important inventions made in the U.S. between 1840 and 1920.[5] The second is provided by the Computer History Museum in San Jose, California, which lists the patents issued for inventions that were milestones in the development of the silicon engine of modern computers.[6] For each set of historian-identified patents, I create a control group of patents that are technologically similar but were not included in the historians' lists. The control group consists of all patents that were granted in same year and assigned to the same primary subclass (at which level there are over 160,000 unique

[5] While Rogers (2011) lists impactful inventions back to 1750, his pre-1840 inventions cannot be linked to patents.

[6] https://www.computerhistory.org/siliconengine/timeline/

subclasses) as the historians' patents. The mean out-degree of the historian-identified patents and their reference group are given in Table 2.1. The table shows that the mean out-degree of the two sets of historian-identified patents is greater than that of the comparison groups.

Table 2.1: Mean Out-Degree of Historian-Identified Great Patents

Historian List of Great Patents	**Mean Out-Degree**	
	Patent in Historian List	**Control Group**
Great American Patents 1840-1920 Rogers (2011)	2.6	0.77
Milestones in the Silicon Engine 1904-1983 Computer History Museum	12.8	8.9

Control group of patents are all USPTO patents granted in the same year and same primary subclass (160K unique classes) as the historian-identified patents.

In a second validation exercise, I compare the mean out-degree of patents from the technology tree to the number of forward citations they receive from subsequent inventions. Because patent forward citations are not broadly available before 1975, I perform this validation exercise using only patents granted starting in that year. In addition, both patents' forward citation count and out-degree from the technology tree suffer right-truncation in recent years. Therefore, I do not include patents granted after 2000 in this exercise.

To test the association between patents' forward citations and out-degree, I run a regression model of the forward citations received by patent *p* as a function of the out-degree of patent *p*. To compare similar types of technologies, I include a primary subclass fixed effect in the model. In addition, I include a year-specific fixed effect. The model is given by Equation 2.4 and its results are presented in Table 2.2.

$$(2.4)\ FwdCites_p = B_1 OutDegree_p + FE_{Year} + FEPrimarySubclass_p$$

Table 2.2: Regression of Forward Citations and Out-Degree of Patents (1975-2000)

B_1 Coefficient	1.16***
B_1 Standard Error	0.006
Fixed Effects?	Yes
R^2	0.39

**** Denotes statistical significance at the 99% confidence interval*

Table 2.2 shows that for patents granted between 1975 and 2000, out-degree is positively associated with forward citation counts. Therefore, the out-degree of patents calculated using the technology tree is positively associated with three external records of patent impact: forward citations, inclusion in Rogers' (2011) list of great American patents, and inclusion in the Computer History Museum's list of milestone patents in the development of the silicon engine.

2.5) Bibliography

Akcigit, U., Grigby, J., and Nicholas, T. (2017). The Rise of American Ingenuity: Innovation and Inventors in the Golden Age. Harvard Business School Working Paper 17-*063.*

Arora, A., Belenzon, S., and Lee, H. (2018). Reversed Citations and the Localization of Knowledge Spillovers. *Journal of Economic Geography* 18(3). 495-521.

Arthur, B. (2009). *The Nature of Technology: What It Is and How It Evolves*. Free Press, New York.

Berkes, E. (2020). Comprehensive Universe of U.S. Patents (CUSP): Data and Facts. Working paper.

Broekel, T. (2019). Using Structural Diversity to Measure the Complexity of Technologies. *PLoS ONE* 14(5).

Fleming, L. Recombinant Uncertainty in Technological Search. *Management Science* 47(1). 117-132.

Fleming, L., and Sorenson, O. (2001). Technology as a Complex Adaptive System: Evidence from Patent Data. *Research Policy* 30(7).

Foster, J., and Evans, J. (2019). Promiscuous Inventions. In Love, A., and Wimsatt, W. (Eds.), *Beyond the Meme: Development and Structure in Cultural Evolution*. University of Minnesota Press.

Fox, D., Fisher, D., and Leighton, L. (1999). Reconstructing Phylogeny with and Without Temporal Data. *Science* 284. 1816-1819.

Levinthal, D. (1997). Adaptation on Rugged Landscapes. *Management Science* 43(7). 934-950.

Mace, R., and Holden, C. (2005). A Phylogenetic Approach to Cultural Evolution. *TRENDS in Ecology and Evolution* 20(3). 116-121.

Mokyr, J. (1992). *The Lever of Riches: Technological Creativity and Economic Progress.* Oxford University Press.

Mukherjee, S., Romero, D., Jones, B., and Uzzi, B. (2017). The Nearly Universal Link Between the Age of Past Knowledge and Tomorrow's Breakthroughs in Science and Technology: The Hotspot. *Science Advances* 3(4).

Rogers, D. (2011). *Inventions and Their Inventors, 1750-1920*. M-Y Books, London.

Simon, H. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics* 69(1). 99-118.

Wiley, E.O., and Lieberman, B. (2011). *Phylogenetics: Theory and Practice of Phylogenetic Systematics*, 2nd ed. Wiley-Blackwell, New York.

Chapter 3: The Emergence of Innovation in New Places

3.1) Abstract

This chapter investigates how new locations emerge as advantageous places for the creation of ideas. Analysis of a novel patent-based dataset that traces the flow of knowledge between inventions and across time reveals that inventors initiate technological knowledge production in new places through a three-stage process. In the first stage, about 50 years before knowledge production in a region reaches an appreciable volume, local inventors begin to experiment with a small number of promising ideas developed in other places. In the second stage, inventors use the promising ideas developed elsewhere to create a larger number of highly impactful inventions locally. In the third stage, inventors source high-impact ideas from their local environs and produce even more local inventions, albeit of lower quality. Overall knowledge production in regions peaks in this third stage, but novelty and the potential for future knowledge production growth decline.

3.2) Introduction

At the start of the 21st century the San Jose–Sunnyvale–Santa Clara Metropolitan Area, the economic core of California's Silicon Valley, ranked first among the United States' 983 metropolitan and micropolitan areas in terms of the number of patents awarded to its inventors and second in terms of per-capita income. San Jose's economic stature is particularly remarkable because it is a "young" city, even by American standards. The counties that now comprise the San Jose Metropolitan Area housed just 0.2% of the U.S. population in 1950 but expanded to 0.6% of the U.S. total by 2000. Patent production in San Jose expanded even faster over this period: San Jose inventors produced less than 1% of the country's patents in 1950 but over 8% in 2000.

While San Jose's rise is striking, nearly every innovative city in the United States began in a similar position as a location where few patentable ideas were invented. Table 3.1 shows the year that the top-15 patent-producing Core-Based Statistical Areas (CBSAs) in the United States emerged as centers for knowledge production, defined as the first five-year window that their local inventors produced 1% of the U.S.'s count of utility patents from that same five-year window. Of the 15 top-ranked CBSAs, 12 crossed the 1% threshold after 1835 when the data series begin. The San Jose, San Diego, and Austin, TX metropolitan areas, for example, all emerged as centers for innovation after the 1950s. And although "old" cities like Chicago and Detroit developed innovative economies much earlier, knowledge production in those cities too had a beginning.

Table 3.1: The Rise of New Cities as Centers of Innovation

Patenting Rank 2001-2005	Core-Based Statistical Area	Year CBSA First Produced 1% of U.S. Patents
1	San Jose	1965
2	New York	Before 1835
3	San Francisco	1865
4	Boston	Before 1835
5	Los Angeles	1905
6	Seattle	1915
7	Chicago	1855
8	Minneapolis	1890
9	San Diego	1980
10	Austin, TX	1990
11	Detroit	1865
12	Philadelphia	Before 1835
13	Houston	1955
14	Dallas	1970
15	Portland, OR	1995

How do inventors commence technological knowledge production in new places? This question is difficult to resolve using the traditional explanations of knowledge-based agglomeration from the geography of innovation literature. According to that literature, innovative activities concentrate in space (Audrestch and Feldman, 1996; Balland et al., 2020) because inventors use existing ideas to create new ideas (Nelson and Winter, 1982; Romer, 1990), and because ideas are more readily transmitted between actors located in close physical proximity or in distant but well-connected regions with established inventive milieus (Jaffe et al., 1993; Bathelt et al., 2004; Breschi and Lissoni, 2009; Kwon et al., 2020). These arguments explain how incumbent regions retain their innovativeness over time, but they leave us to puzzle over how inventors begin to produce knowledge in places that lack existing knowledge stocks or inter-regional networks to begin with.

One way inventors may initiate knowledge production in new places is by developing new domains of technological knowledge (Kuhn, 1962; Dosi, 1982). The knowledge spillovers that circulate within and between incumbent regions are domain-specific in that they contain information about how technologies with specific material or relational properties function (Arthur 2009). Because the spillovers of older ideas are less relevant to innovation in new knowledge domains, the introduction of new domains may open windows of opportunity for inventors in new locations to commence local knowledge production (Storper and Walker, 1991; Boschma and Lambooy 1999; Boschma and Frenken, 2006). Data-intensive case studies of the automotive cluster in Detroit, the steel industry in Cleveland, and textile industry in Manchester, England, have shown that the formation of major U.S. and British metropolitan economies coincided with the advent of new industries and process technologies (Brezis and Krugman, 1997; Klepper, 2007; Lamoreaux et al., 2004). However, this empirical record has not been generalized to a broader set of metropolitan regions. In addition, the existing literature tracks employment counts and the opening of new firms, and does not observe changes in the knowledge used to innovate in their industries.

In this chapter, I present evidence from over 8 million geo-located patent records granted to U.S. inventors between 1850 and 2010 to outline a common pathway through which inventors initiate technological knowledge production in new places. Namely, that pathway involves the creation and development of new knowledge domains by inventors. I identify domain-forming inventions at the patent level by calculating the extent to which a patented technology introduces new and novel ideas that are used in a large number of subsequent inventions. My results show that inventors rely heavily on knowledge embedded in novel and impactful inventions early in their home regions' innovative growth. In addition, I find that in the earliest stages of their home regions' innovative growth, inventors disproportionately rely on such knowledge from *non-local* origins, and switch to such knowledge from *local* origins as their home regions' knowledge production expands.

My analysis uses the novel dataset that traces the flow of technological knowledge between patents and across time described in Chapter 2 of this dissertation. The new dataset contains records of knowledge flow tracing back to 1850, allowing for a historical study that cannot be performed using patent citation records that have unreliable historical coverage (Berkes 2020). I combine the new records of knowledge flow with historical information on inventors' place-of-residence (Petralia et al., 2016) to study how inventors initiated knowledge production in all but the oldest U.S. cities. In addition, I test whether novel and impactful ideas are more useful for initiating local knowledge production by comparing the composition of knowledge used by inventors that reside in CBSAs that eventually become innovative centers with the composition used by inventors in CBSAs that never emerge as innovative centers. This comparative analysis reveals that inventors in ultimately successful CBSAs source a significantly larger share of knowledge from high-impact local and non-local inventions than do inventors in ultimately unsuccessful CBSAs, even 50-100 years before their home regions emerge as innovative centers.

The results of the study contribute to four literatures: the spatial patterns of knowledge transmission, agglomeration theory, evolutionary economic geography, and the regional economic lifecycle. With respect to the first literature, the decomposition of knowledge sources conducted in this study shows how local and non-local knowledge flows materialize in knowledge production growth in proximate and distant locations (Jaffe et al., 1993; Breschi and Lissoni, 2009; Kwon et al., 2020). With respect to agglomeration theory and evolutionary economic geography, the analysis reveals how spatial concentrations of knowledge production form in their earliest years, before they enter the purview of agglomeration theory and evolutionary economic geography as geographical units of observation. Finally, the findings expand the concept of the urban lifecycle (Audretsch et al., 2008) by revealing how new centers for innovation are conceived using ideas developed in other places, and by showing how even after innovation in a place declines, its ideas can flock to and flourish in new regions.

In the text that follows, I discuss how inventors create and transmit technological knowledge and the geographical implications thereof, I introduce the methods used to infer historical flows of knowledge between patents and to identify high-impact inventions that support the creation of new knowledge domains, and I present the results of the analysis, beginning with a birds-eye-view of knowledge production growth and decline in regions and continuing with a decomposition of the sources used by inventors as their regions initiate, expand, and decline in patent production. In the final section I discuss the relationship between incremental and disruptive inventions and their relationships with the emergence, evolution, and resilience of knowledge production in cities.

3.3) The Production and Transmission of Technological Knowledge

The process of invention involves the recombination of existing technological components (Romer, 1990; Weitzman, 1998). These components may be material or immaterial, ranging from screws and bolts to functions and commands in computer programs (Arthur, 2009). Recombinant technological knowledge, defined as the ability to assemble components into larger functioning systems, is exceedingly difficult to generate. Each component in a technology operates by interacting with other elements in the same system. Because of the high interdependence, inventors struggle to anticipate how their technologies will function before they assemble them (Fleming and Sorenson, 2001). Models and prototypes help inventors to simulate the interactions between components, but models and prototypes are also costly to create and time-consuming to administer (Usher, 1929; Arrow, 1962; Adler and Clark, 1991; Von Hippel and Tyre, 1995). These costs multiply when inventors design complex technologies containing many elements arranged in irregular ways (Broekel, 2019).

To ease the process of designing complex technologies, inventors rely on prior knowledge (Fleming, 2001). Inventors who already know how some assemblies of components function can focus on integrating their assemblies rather than developing them anew (Foster and Evans, 2019). The ability for inventors to build on prior knowledge, however, is limited by the breadth of their individual accumulated knowledge assets. Because technological knowledge is detailed, inventors have highly specialized areas of expertise and need to collaborate or source ideas from other inventors and scientists (Wuchty et al., 2007).

It is challenging for inventors to source recombinant technological knowledge from other inventors. In its native format, recombinant knowledge is a list of the experiences an inventor accumulates while assembling a technology (Arrow, 1962). For all but the simplest devices, that list is too detailed for an inventor to recollect let alone communicate (Polanyi, 1966). Inventors respond by compressing knowledge into diagrams and metaphors (Nonaka and Takeuchi, 1995). However, these project-oriented coding schemas can only be transmitted using supportive communication technologies. For most of the United States' industrial history, face-to-face communication held an absolute advantage in communicating messages encoded in dynamic, non-linear schemas. Face-to-face communication allows for the use of visual clues such as body language and hand gestures to convey complex points, as well as the manipulation of vocal tone to stress key aspects of a message (Storper and Venables, 2004). The interactive nature of face-to-face communication allows speakers to notice misunderstandings and to correct their presentations to improve comprehensibility (Nohria and Eccles, 1992), and to create norms, routines, and rhetorical devices that are specifically designed for the technical issues at hand (Kogut and Zander, 1992; Powell et al., 1996; Gertler, 2003).

Because close spatial proximity is a necessary condition for face-to-face communication, the ability for inventors to source recombinant knowledge is influenced by their socio-spatial environments. While inventors that reside in regions with many other inventors are able to source a wide range of ideas face-to-face, inventors in regions with sparse networks of inventors are at a competitive disadvantage. Empirical research shows that the frequency of knowledge transmission (Jaffe et al. 1993; Kwon et al., 2020) and the frequency of collaboration (Balland, 2012; van der Wouden, 2020) between inventors decline as spatial distance increases. Isolation can be momentarily remedied through travel, but the logistical and economic costs of travel also pose constraints (Torre, 2008). Inventors are unlikely to travel for work unless they or their organizations have strong incentives to undertake travel (Morrison et al., 2013). The incentive to travel is a function of the quantity and quality of the knowledge inventors expect to gather through travel or the expected value of a resulting product or invention (Cowan and Jonard, 2004). The incentive to travel to places with few knowledgeable inventors is therefore small, so most non-local flows of knowledge span between regions that already have dynamic inventive milieus (Bathelt et al., 2004; Wolfe and Gertler, 2004). Temporary face-to-face meetings such as tradeshows and conferences are another means through which inventors may source non-local knowledge. However, geographically isolated inventors tend to lack the local absorptive capacity needed to make effective use of temporary face-to-face meetings (Bathelt and Henn, 2014; Esposito and Rigby, 2019. Altogether, the odds are stacked against inventors that reside and work in peripheral locations.

Yet as Table 3.1 illustrates, early in the history of every innovative region there is a moment when local inventors overcome their constraints to sourcing knowledge and begin to produce patentable ideas. Given that the creation of recombinant knowledge is difficult and competitive, how do inventors transform peripheral regions into centers of innovation? As described by the concept of the Window of Locational Opportunity, two factors may interact to allow inventors in new locations to commence knowledge production (Storper and Walker, 1991; Boschma and Lambooy, 1999). The first factor is that historical accidents and idiosyncratic events occasionally transport ideas to regions that are not major centers for knowledge production. The second factor is that there is immense heterogeneity in usefulness of existing ideas for the creation of new ones (Dosi, 1982; Martinelli and Nomaler, 2014). Some ideas are radical, meaning that they begin at the root of all knowledge. While in practice all new ideas to a certain extent build on existing knowledge, radical inventions nonetheless explore new material or immaterial relationships and thus fall outside the dominant trajectories of knowledge development (Arthur, 2009). Because such inventions do not rely on extensive stocks of prior knowledge, they level the geographical playing field for subsequent innovation. Inventors in regions with underdeveloped knowledge bases that gain access to even a small number of impactful and novel ideas may be able to use those ideas to initiate local knowledge production.

The example of the semiconductor-based transistor demonstrates how the transportation of a single revolutionary idea to a new location can uproot the geography of knowledge production. The first two transistors, the point-touch transistor and the junction transistor, were invented at Bell Labs in the suburbs of New York City in 1947 and 1948, but the transistor's development into useful tools took place primarily in Silicon Valley, the region of California centered on the San Jose Metropolitan Area. While the design of transistors is now a complex industry around which large teams, firms, and agglomerations have organized (Balland et al., 2019), shortly after the first transistors were invented, transistor design was less knowledge-intensive. Notably, when William Shockley, Bell Labs' star engineer and inventor of the junction transistor, moved from the New York City area to Silicon Valley in 1956, he brought just one colleague from Bell Labs with him (Gertner, 2012, 181). This relocation indicates that the knowledge domain needed to design new transistors at that time was evidently simple enough for two engineers to collectively master (Wuchty et al., 2007). Because the transistor's knowledge domain was compact, idiosyncratic factors, such as Shockley's preference to be near his ailing mother, were able to influence where the new knowledge domain settled and produced innovative growth.

While Shockley's arrival in Silicon Valley brought detailed knowledge of the transistor to the region, it was just one early step in a larger process that ultimately made Silicon Valley a leading innovative center. After Shockley arrived, local inventors applied his transistor to make a large number of subsequent impactful inventions. For example, the first microprocessor (the Intel 4004) was introduced in Silicon Valley in 1971, making personal computers and hand-held computer devices possible. Additional groundbreaking inventions, such as the first dynamic RAM chip (the Intel 1003, invented in 1970) were made in Silicon Valley, as were hardware technologies such as the computer mouse in 1964 and the first commercialized computer monitor in 1973. These inventions extended the potential of the semiconductor transistor by making the transistor useful in new applications. In so doing, these inventions developed the nascent knowledge field grounded in the properties of semiconductor amplification into a full-fledged knowledge domain. The creation of these subsequent inventions in Silicon Valley caused the center of gravity of the knowledge domain to settle in the region through a virtuous, self-reinforcing cycle (Saxenian, 1994; O'Mara, 2019).

The geographical history of the semiconductor transistor thus suggests two sets of processes that occur as new places emerge as centers for innovation. The first is the creation of novel and promising ideas. These ideas do not need to be invented in the same locations where they ultimately produce innovative growth so long as they are simple enough to be transported across space through the movement of people or messages (cf. Kerr, 2010). The second process is the creation of high impact inventions locally. These follow-up inventions build a local critical mass of knowledge in the new knowledge domain, causing future knowledge production to consolidate in the emerging innovative region. The analysis section of this paper investigates these two processes and their relationship with the growth of local innovation.

3.4) Methods Overview

In this study, I use the new records of knowledge flows between inventions described in Chapter 2 to identify the earlier inventions that inventors build on when commencing knowledge production in new places, and to identify the impact of their resulting technologies. The knowledge sources of a focal patent are given by its parent patents. Knowledge sources may be local or non-local, and high or low impact. A local knowledge source is a parent patent that was invented in the same metropolitan region as its child patent. A high-impact knowledge source is a parent patent that is observed to be highly impactful, using the binary definition of impact discussed in Chapter 2.

To match patents to geographical regions, I aggregate patents to CBSAs based on the home address of the inventors of that patent. I use 2015 definitions of CBSAs, which include all metropolitan and metropolitan areas in the United States.[7] When patents have inventors living in two or more CBSAs, I fractionally assign those patents to each CBSA. I regard knowledge sourced from patents created by inventors that reside outside the U.S. as non-local knowledge.

[7] A list of 2015 CBSAs with their constitutive counties is available at the USPTO website: https://www.uspto.gov/web/offices/ac/ido/oeip/taf/cls_cbsa/cbsa_countyassoc.htm

3.5) Results: Overview

The growth of knowledge production in U.S. CBSAs tends to follow a general pattern in which regions begin knowledge production by producing a small number of ideas, expand their production of ideas over time, reach a peak in knowledge production, and thereafter tend to enter a period of decline. Using black dots in Figure 3.1, I plot the production of patents by half-decade for four representative U.S. CBSAs centered in Detroit, Cleveland, San Jose, and Austin. I selected these cities because they are or have been major centers for innovation and because they initiated patent production growth during the time period for which I have reliable data, starting in 1850. To improve the comparison of patent production across years, I express the patents produced in a CBSA in a given half-decade as a percentage of the U.S. total for that same half-decade. I also plot the number of high-impact patents produced in that CBSA using plus-signs and low-impact patents using dots, as described in Section 3.4.

Figure 3.1: High-Impact and Total Patents Produced in Four Representative CBSAs

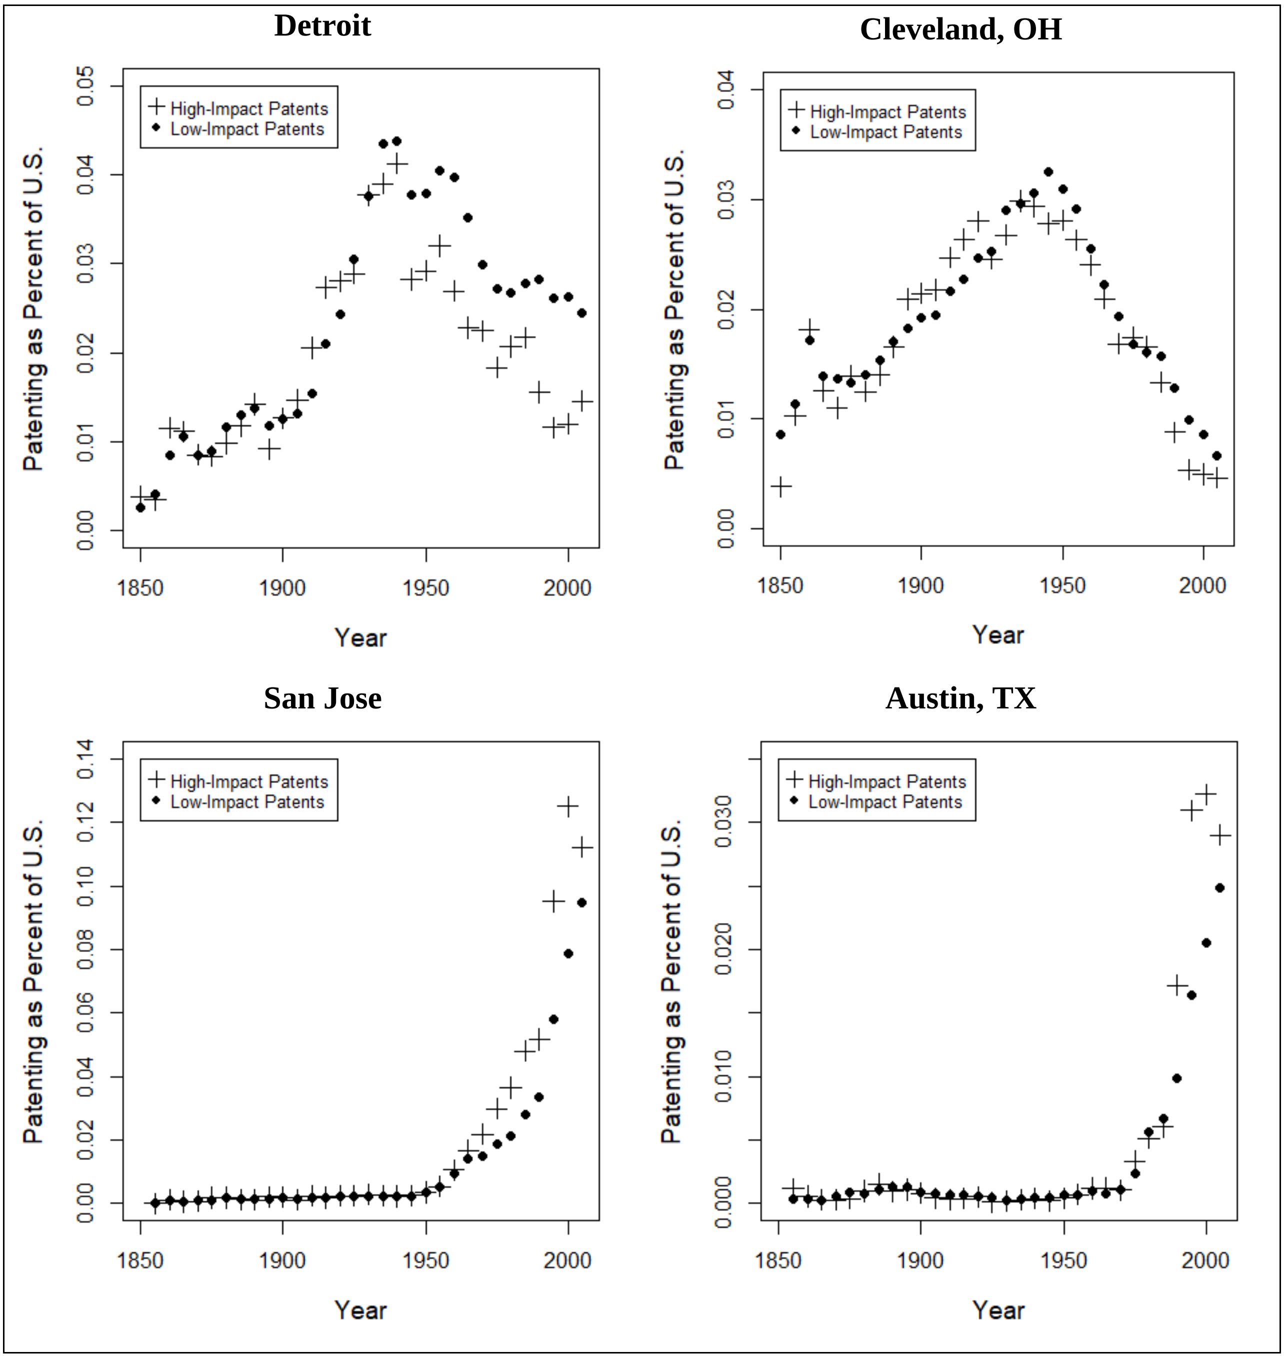

Figure 3.1 shows that the years of 1865 in Detroit, 1855 in Cleveland, 1965 in San Jose, and 1990 in Austin all bear resemblance in terms of patent production growth: during these years, knowledge production in each CBSA started to climb. Additionally, as those cities began to increase their overall production of patents, they also increased their production of high-impact patents. Generally, their production of high-impact patents grew

faster than their production of low-impact patents: when the plus-signs rise in Figure 3.1, the black dots rise even faster.

The examples of Detroit, Cleveland, San Jose, and Austin in Figure 3.1 thus suggest two general patterns. First, overall patent production in cities experiences both a rise and a decline; second, the rise and decline of patenting in cities is preceded by the rise and decline of the production of high-impact patents. To test if these patterns are found generally across U.S. CBSAs, I compute the average percentage of U.S. total and high-impact patents that each CBSA produces at each stage in its patenting growth. To compare cities that underwent knowledge growth during different periods of time (such as San Jose and Detroit), I align the time dimension of their patenting based on their CBSA "age," defined as the first five-year period a CBSA produces 1% or more of the U.S. total stock of patents. For example, I assume that San Jose in 1965 (when it first produced 1% of U.S. patents) was at the same stage of its growth trajectory as Detroit was in 1865. Extending that reasoning, 35 years later, San Jose was at the same stage of its growth process as Detroit was in 1900. Formally, I calculate the $Age_{c,t}$ of a CBSA c in half-decade t by subtracting the observation year from the year it first crosses the 1% patenting threshold, as captured by Equation 3.1:

$$(3.1)\ Age_{c,t} = Year_{c,t} - ThresholdYear_c$$

After aligning the curves of each CBSA based on $Age_{c,t}$, I compute aggregate patent production curves by averaging the percent of U.S. overall and high-impact patents in a given half-decade that are produced in cities with a given $Age_{c,t}$ value, as in Figure 3.2. In the chapter appendix, I show that the general findings of Figure 3.2 are robust when 0.5% and 5% patenting thresholds used to identify CBSA age.

Figure 3.2: Average Patent Production in U.S. CBSAs by CBSA Age

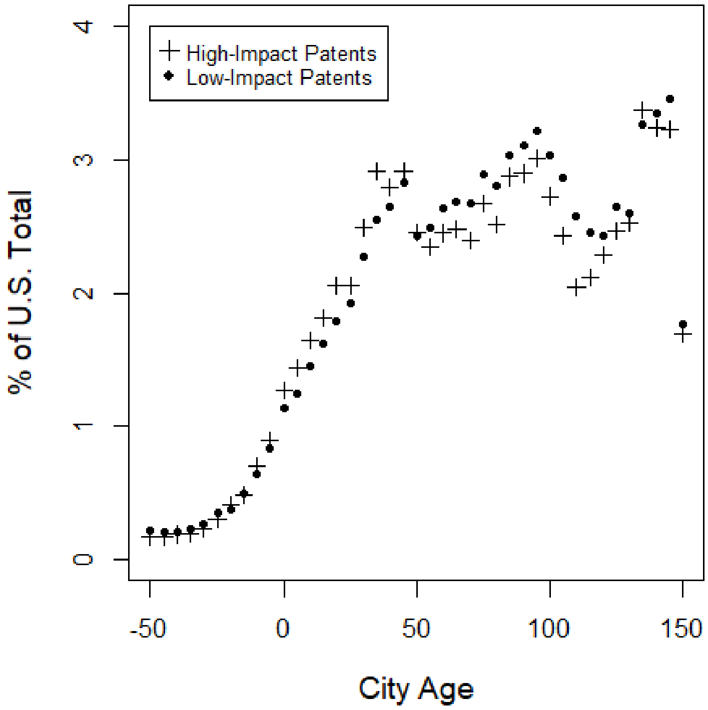

Note: Only CBSAs that produce at least 1% of U.S. patents during one or more 5-year periods of their history are included in the analysis. Thus, 19 CBSAs are included in Figure 3.2.

Figure 3.2 generates three observations. First, patenting growth in cities is related to CBSA age as defined in Equation 4.1. Second, the production of high-impact patents in cities generally increases before the production of low-impact patents; similarly, the production of high-impact patents starts to decline before low-impact patenting goes down. Third, the temporal order of the growth of high-impact and low-impact patent production documented in Figure 3.2 suggests that a causal relationship runs from the local production of high-impact patents to increased low-impact patenting later on.

Because high-impact patenting rises before low-impact patenting does, Figure 3.2 suggests that the creation of local high-impact inventions creates a local knowledge base that inventors subsequently use to create further inventions. However, Figure 3.2 does not explicitly identify flows of knowledge running from these high-impact inventions to the subsequent low-impact inventions; this identification is made strictly based on temporal ordering. In addition, Figure 3.2 leaves unresolved the issue of how inventors source the knowledge used to make the very first local inventions. In the following section, I address these two by explicitly studying the composition of knowledge that inventors use as overall patenting in their home region begins and expands.

3.6) Results: The Sources of Knowledge Production Growth in Cities

In this section, I examine how inventors source impactful knowledge locally and non-locally as knowledge production commences, expands, and declines in their home regions. To undertake this analysis, I develop a knowledge-source accounting framework that reveals the relative importance of various types of knowledge flows for instigating knowledge production. Most research on the geographical patterns of knowledge flows, such as Jaffe et al. (1993), Breschi and Lissoni (2007) and Arora et al. (2018), infer the geographical consequences of local and non-local knowledge flows based on the friction posed by distance. For reasons discussed in Section 3.3, even if geographical distance exerts strong frictions on the spreading of ideas across distances, non-local ideas can nonetheless be used to initiate knowledge production in new places if those ideas are transformative. The source-accounting framework developed here overcomes this analytical challenge by directly calculating the relative importance of local, non-local, high-impact, and low-impact knowledge sources for inventors to initiate local knowledge production.

In Table 3.2, I show how the types of knowledge that inventors' sources evolve as their home regions become innovative. To do so, I compute the percentage of each type of knowledge source that the inventors of patents use at three stages of knowledge production growth in their home regions. Those stages are defined using the CBSA age measures defined in the previous section. CBSAs in the first stage (CBSA age -100 to 0) are nascent innovative centers, CBSAs in the second stage (CBSA age 0 to 50) are maturing innovative centers, and CBSAs in the third stage (CBSA age 50 to 150) are mature or declining innovative centers.

As discussed in the methods section, I distinguish between four types of knowledge sources that inventors used by inventors: non-local high-impact patents (NL.High), non-local low-impact patents (NL.Low), local high-impact patents (L.High), and local low-impact patents (L.Low). High-impact inventions introduce novel and useful new knowledge and thus represent stages in the formation of new knowledge domains. The values in Table 3.2 use the 1% patenting threshold to define CBSA age; the appendix presents similar results using 0.5% and 5% patenting thresholds.

Table 3.2: Composition of Knowledge Sources by Source Type and by CBSA Age

	Source Type			
Age Range	**NL.High**	**NL.Low**	**L.High**	**L.Low**
-100 to 0	67%	28%	3%	2%
0 to 50	59%	27%	9%	5%
50 to 150	57%	26%	11%	6%

Note: Table 2 uses the 1% patenting threshold to define CBSA age. Only the 19 CBSAs that break the 1% threshold are included in the table.

Table 3.2 shows that inventors disproportionately use knowledge from non-local high-impact sources (NL.High) when knowledge production in their home regions first begins. Non-local high-impact sources account for 67% of all knowledge sources when CBSA age is between -100 and 0, while it accounts for 57% of all knowledge sources when CBSA age is between 50 and 150. While inventors' reliance on non-local high-impact sources declines as their home regions grow, their reliance on local high-impact sources (L.High) grows over time: local high-impact sources account for just 3% of all sources used by inventors when their CBSA age is between -100 and 0, but it accounts for 11% of all sources when CBSA age is between 50 and 150.

Table 3.2 also shows important changes in inventors' use of low-impact knowledge sources as their home regions emerge as innovative centers. The percent of knowledge sourced from non-local low-impact (NL.Low) sources declines slightly over time, from 28% in the -100 to 0 CBSA age range to 26% in the 50-150 age range. More importantly, inventors' reliance on local low-impact sources (L.Low) increases from 2% in the -100 to 0 age range to 6% in the 50 to 150 age range. This increased reliance on local low-impact sources coincides with the period when patent production in the CBSA peaks (Figure 3.2).

The overall patterns identified in Table 3.2 hold up when CBSA age is not aggregated to multi-decade periods of time. Figure 3.3 shows the composition of knowledge sources used by inventors at each unique value of CBSA age. The figure uses the 1% patenting threshold to compute CBSAage, but in the appendix I show that the results do not change significantly when 0.5% or 5% thresholds are used to compute CBSA age. In Figure 3.3, I also overlay the percentage of all U.S. patents produced in CBSAs at a given age value using a solid line to show how changes in the composition of knowledge sources used by inventors correspond with the rise, levelling off, and decline of local patent production. Figure 3.3 indicates that inventors most heavily rely on knowledge from non-local high-impact inventions as local patent production commences and increase their use of knowledge embedded in local high-impact inventions as local knowledge production reaches its peak.

Figure 3.3: Composition of Knowledge Sources and Patent Production by CBSA Age

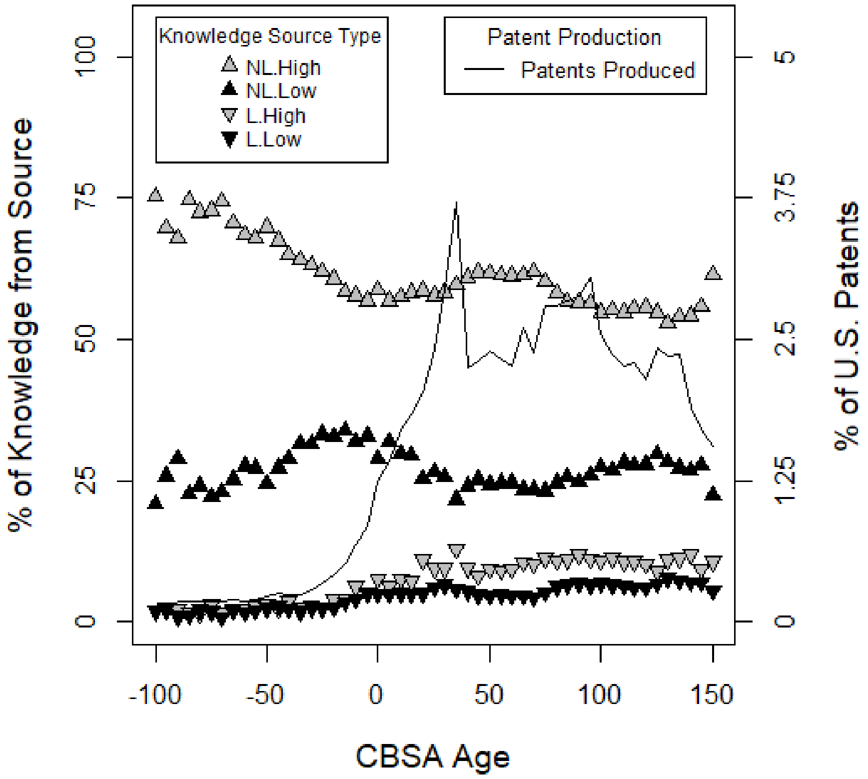

While Table 3.2 and Figure 3.3 show that the knowledge types of sources used by inventors evolve as their home regions emerge as innovative centers, it remains undetermined whether the types of sources used by inventors are important for inventors to initiate knowledge production in their home regions. An ideal causational analysis would study two plots of land, both located far away from any existing innovative centers. At the start of the study, both plots are completely empty of people and knowledge. Later in time, the first plot receives an infusion of inventors with knowledge of the inner-workings of novel and impactful technologies. The second plot also receives an infusion of inventors, but their knowledge is less novel and less useful for subsequent invention. Several decades later, we would check back to see which plot of inventors produced more inventions.

Of course, such an experiment is impossible to administer. Nonetheless, a revealing identification can be generated by comparing the knowledge sources used by inventors in

CBSAs that break the 1% patenting threshold with the knowledge sources used by inventors in CBSAs that never break the 1% threshold.

To compare the types of knowledge sources in "successful cities" (CBSAs that break the threshold) and "unsuccessful cities" (CBSAs that never do), I compute the frequency by which their inventors use each type of source at each CBSA age value. Unsuccessful cities do not have CBSA age values because they never break the patenting threshold. However, a robust comparison can be generated by comparing the composition of knowledge sources used in successful cities with the composition used in USCs in the same 5-year period of observation. In a simple one-city example, assume that I seek to compare San Jose when its CBSA age was -15 with all unsuccessful cities at that same moment in history. San Jose was age -15 in 1950, so I compare the knowledge sources used by San Jose's inventors in 1950 with the sources used by inventors in unsuccessful cities in 1950.

I generalize the method described above to all 983 CBSAs in my dataset. To do so, I begin by generating the vector $Years_{Age}$, which records the 5-year periods in which successful cities are observed at a given age value. For example, if a total of three successful cities in reach age 10, the first in 1900 and the second and third in 1995, then the Years vector is given by:

$$(3.2)\ Years_{Age=10} = [1900, 1995, 1995]$$

I calculate the composition of knowledge types used by inventors in unsuccessful cities at a given age value by averaging the composition used in unsuccessful cities over the $Years_{Age}$ vector. Let the count of patent parents used by all inventors in unsuccessful cities (abbreviated as *USC)* in a given 5-year period and of a given source type be defined as $PatParents_{Years,USC,Type}$. The average propensity for inventors in unsuccessful cities of a given age to source knowledge of a given source type is:

$$(3.3)\ \%Parents_{USC,Type,Age} = \frac{\sum_{Years}\sum_{USC} PatParents_{Years,USC,Type}}{\sum_{Years}\sum_{USC} PatParents_{Years,USC}}$$

In Equation 3.3, the *Years* subscript is an index of the $Years_{Age}$ vector.[8] To provide an example of how Equation 3.3 is computed, let us make three assumptions:

(1) $Years_{Age=10} = [1900, 1995, 1995]$, as in Equation 3.2

(2) $PatParents_{Years,USC} = [50,100,100]$, denoting that patents in unsuccessful cities sourced knowledge from 50 parent patents in 1900 and 100 parents in 1995

(3) $PatParents_{Years,USC,Type=NL.HIGH} = [5,8,8]$, denoting that patents in unsuccessful cities sourced knowledge from 5 high-impact parents in 1900 and 8 in 1995

In this example, percentage of parents unsuccessful cities sourced from NL.HIGH parent patents at Age=10 is computed as:

$$(3.4)\ \%Parents_{USC,Type=NL.HIGH,Age=10} = \frac{5+8+8}{50+100+100} = 8.4\%$$

Finally, to compare the composition of knowledge sources used by inventors in successful cities with the composition used in unsuccessful cities, I compute the successful cities composition premium as the difference between the compositions used in successful cities (*SC*) and unsuccessful cities (*USC*):

$$(3.5)\ SCPremium_{Type,Age} = \%Parents_{SC,Type,Age} - \%Parents_{USC,Type,Age}$$

In Figure 3.4, I create scatterplots of $SCPremium_{Type,Age}$ by plotting it against CBSA age. I overlay Loess regression fit lines with a search distance of 100% to identify general trends in the data across age values and to generate 95% confidence intervals. The chapter appendix provides similar charts using 0.5% and 5% threshold values.

[8] The 1995 value is double-counted because it appears twice in the $Years_{Age=10}$ vector, which amounts to taking weighted means.

Figure 3.4: Knowledge Sources used by Cities that Break the Patenting Threshold in Excess of Knowledge Sources used by Cities that Never Break the Patenting Threshold

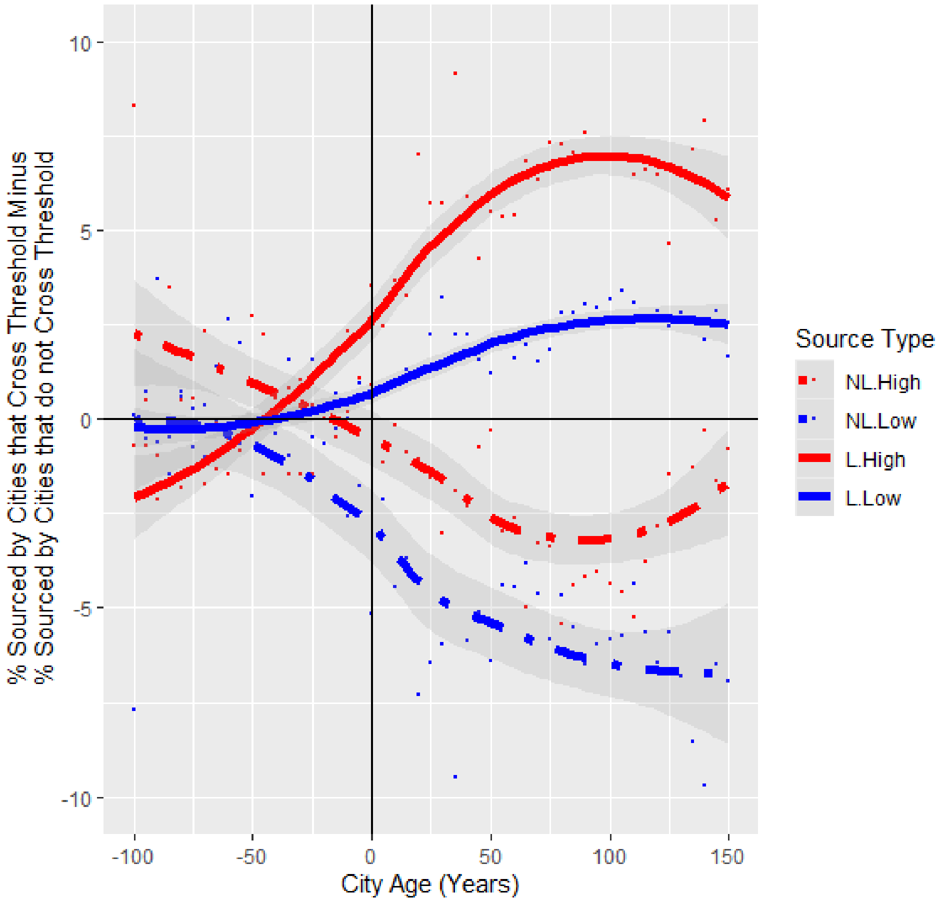

Figure 3.4 shows that inventors in successful CBSAs use a different composition of knowledge sources than do inventors in unsuccessful CBSAs. These differences occur during and after successful CBSAs cross the patenting threshold. Early in their CBSA's innovative growth, when CBSA age is between -100 and -50, inventors in successful cities source about 2.5% more of their knowledge from NL.HIGH patents than do inventors in unsuccessful cities, as indicated by the dotted red line that is located above the x-axis. The shaded confidence interval indicates that this difference is statistically significant at the 95% level. Between age -50 and 0, inventors in successful cities increase a growing share of their knowledge from local high-impact inventions, as indicated by solid red line that rises well-above the x-axis. Eventually, when city age passes 50 years, inventors in successful cities source a growing share of their knowledge from local low-impact inventions. However, the

solid red line remains far above the solid blue line for the duration of the chart, indicating that inventors in successful cities rely much more heavily on local high-impact knowledge than do inventors in unsuccessful cities for the full duration of their region's innovative growth and decline.

3.7) Discussion

This paper has used data on knowledge production and knowledge sourcing to study how inventors initiate and expand knowledge production in new places. The creation of new technological knowledge is difficult but inventors are aided by the existing ideas that they can build on. However, inventors in regions without histories of knowledge production can access few ideas through face-to-face communication. In turn, inventors tend to create new technologies in established milieus.

Occasionally new and impactful ideas are generated that are less reliant on existing stocks of knowledge. Certainly, many of the subsequent inventions enabled by these breakthroughs are realized within immediate environs where the breakthroughs are initially made. However, it is possible for inventors to commence knowledge production in new places by importing promising new ideas from other regions and by developing those ideas to create larger technological systems. As those technological systems grow, the associated knowledge is distributed across expanding networks of actors that co-locate in space to interact through face-to-face communication. New innovative clusters and new knowledge domains are thus endogenously developed.

While a wealth of research examines how geographical proximity enhances the generation and transmission of technological knowledge, those literatures have paid less attention to how these social and physical are created. As a consequence, two related questions about how regions emerge as innovative centers remain unresolved. The first question is why inventors in some locations manage to source high-impact non-local ideas but fail to sustain long-run local innovation (Scott and Storper, 1987). The results of this study indicate that in addition to importing promising ideas from afar, inventors need to introduce high-impact ideas locally. Yet, this finding begs a related question: why do inventors in some regions introduce more high-impact inventions than in others? For Saxenian (1996) and Storper et al. (2015), competitive advantage in the creation of impactful new technologies is derived from dexterous local institutions that allow regional actors to develop new methods for organization and coordination. To test this hypothesis, researchers should study the association between the cross-regional variation in the fluidity of local inventor networks and the local creation and capture of high-impact inventions.

Finally, while this study demonstrates how high-impact ideas provide opportunities for inventors to commence knowledge production in new places, too little is known about why inventors integrate some new and impactful ideas into existing technological systems while in other cases they develop technologies that drive existing technological systems into obsolescence. These technological relationships have regional consequences as certain inventions promote regional diversification and resilience (Neffke et al., 2013; Rigby, 2015; Boschma, 2015) while other inventions render regional knowledge bases obsolete and establish new geographical outposts for innovation (Scott and Storper, 1987; Storper and Walker, 1991; Boschma and Lambooy, 1999). Ecological models of symbiotic and adversarial relationships between technologies (Foster et al., 2014) and network-based models of disruption (Funk and Owen-Smith, 2016) are encouraging starting points to unpack these conflicting sources of evolution and revolution in the geography of innovation.

3.8) Appendix

The production of high and low-impact patents in CBSAs by CBSA Age.

Figure 3.5: Age Defined Using 0.5% Threshold

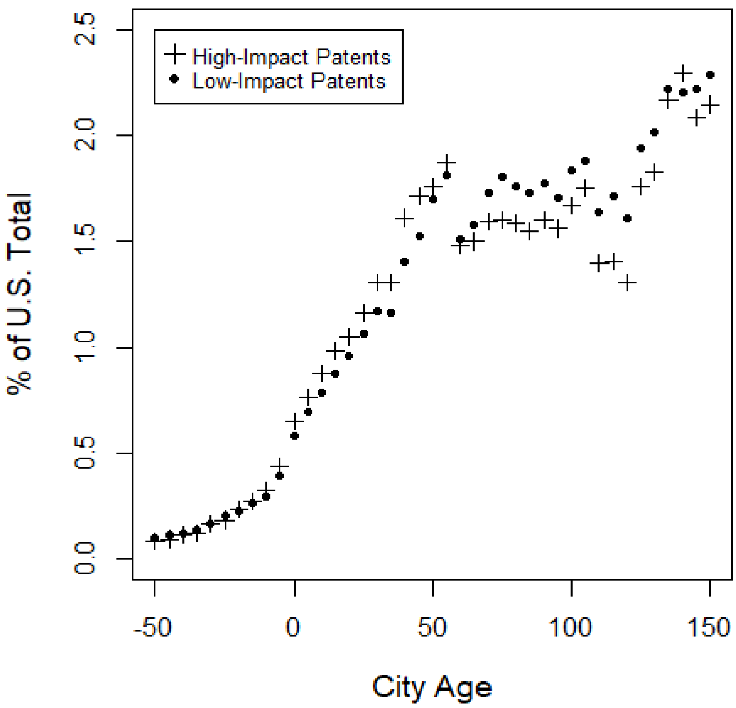

Figure 3.6: Age Defined Using 5% Threshold

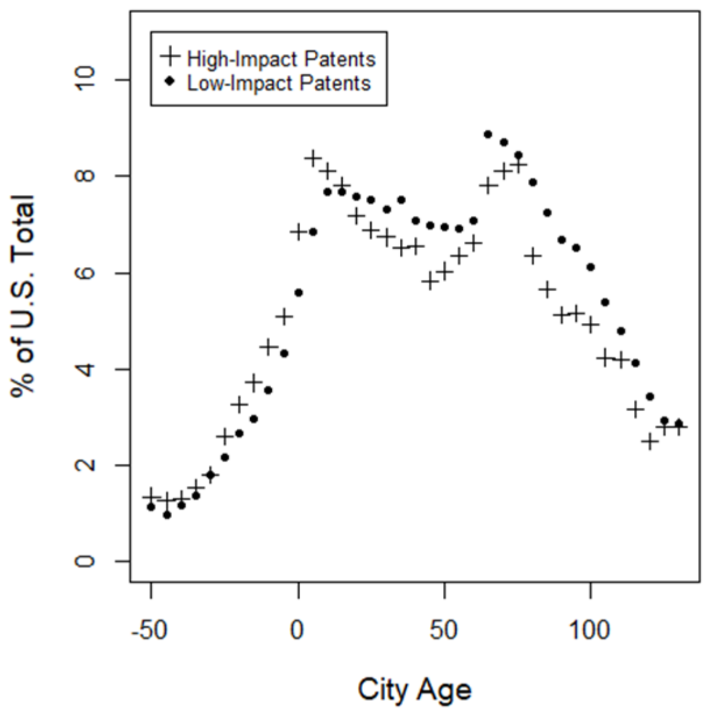

Composition of knowledge sources by ranges of CBSA age.

Table 3.3: Age Defined Using 0.5% Threshold

	Source Type			
Age Range	**NL.High**	**NL.Low**	**L.High**	**L.Low**
-100 to 0	68%	26%	4%	2%
0 to 50	69%	27%	9%	5%
50 to 150	57%	27%	10%	6%

Table 3.4: Age Defined Using 5% Threshold

	Source Type			
Age Range	**NL.High**	**NL.Low**	**L.High**	**L.Low**
-100 to 0	70%	23%	5%	3%
0 to 50	55%	28%	10%	6%
50 to 150	55%	23%	15%	7%

Composition of knowledge sources and patent production by CBSA Age using 0.5% and 5% thresholds in calculating CBSA Age.

Figure 3.7: Age Defined Using 0.5% Threshold

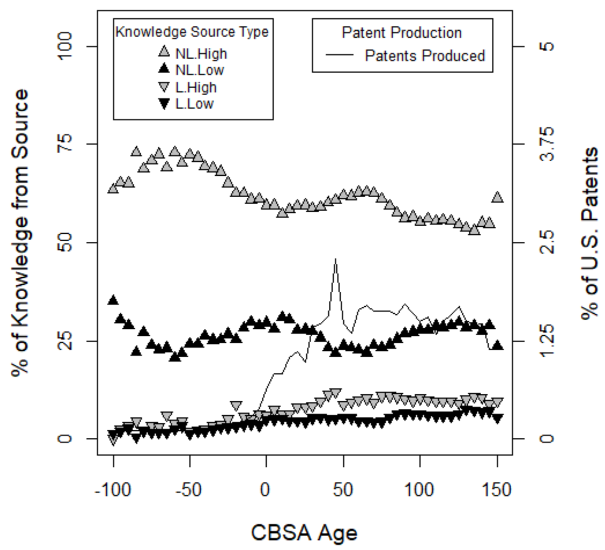

Figure 3.8: Age Defined Using 5% Threshold

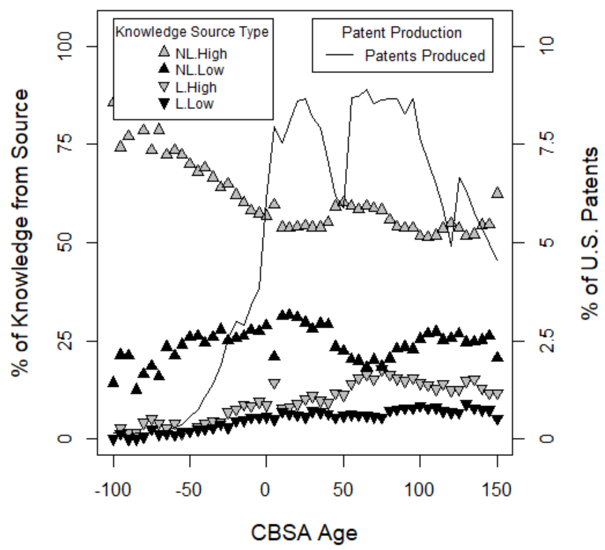

Composition of knowledge sources used by cities that break the patenting threshold in excess of knowledge sources used by cities that never break the patenting threshold.

Figure 3.9: Age Defined Using 0.5% Threshold

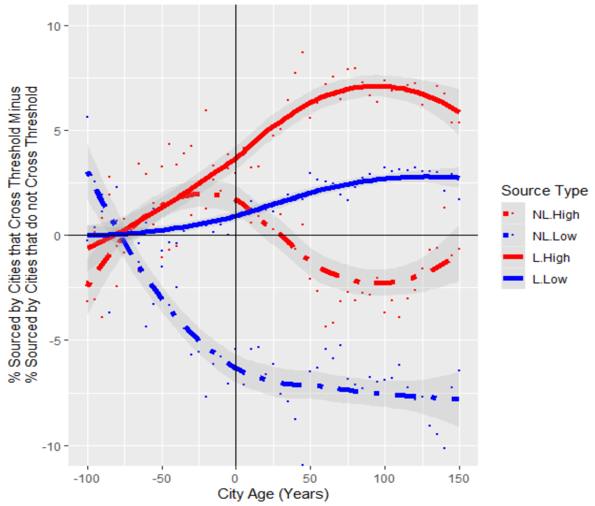

Figure 3.10: Age Defined using 5% Threshold

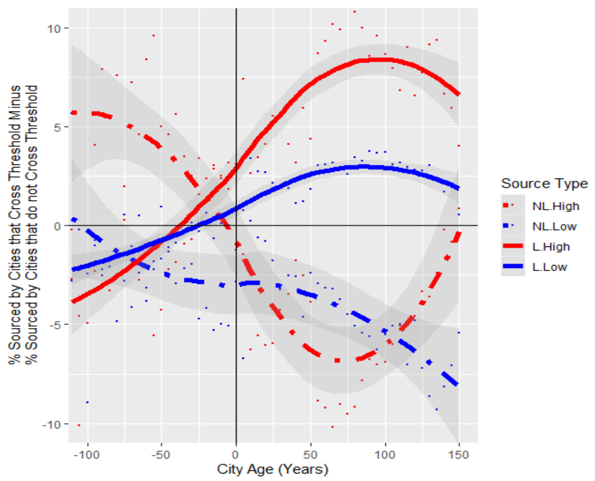

3.10) Bibliography

Adler, P., and Clark, K. (1991). Behind the Learning Curve: A Sketch of the Learning Process. *Management Science* 37(3). 267-281.

Arora, A., Belenzon, S., and Lee, H. (2018). Reversed Citations and the Localization of Knowledge Spillovers. *Journal of Economic Geography* 18(3). 495-521.

Arrow, K. (1962). The Economic Implications of Learning by Doing. *Review of Economic Studies* 29(3). 155-173.

Arthur, B. (2009). *The Nature of Technology: What It Is and How It Evolves*. Free Press, New York.

Audrestch, D., and Feldman, M. (1996). R&D Spillovers and the Geography of Innovation and Production. *American Economic Review* 86(3). 630-640.

Audretsch, D., Falk, O., Feldman, M., and Heblich, S. (2008). The Lifecycle of Regions. *CEPR Discussion Papers* 6757.

Balland, P.A. (2012). Proximity and the Evolution of Collaboration Networks: Evidence from Research and Development Projects Within the Global Navigation Satellite System Industry. *Regional Studies* 46(6). 741-756.

Balland, P.A., and Rigby, D. (2018). The Geography of Complex Knowledge. *Economic Geography* 93(1). 1-23.

Balland, P.A., Jara-Figueroa, C., Petralia, S., Steijn, M., Rigby, D., and Hidalgo, C. (2020). Complex Economic Activities Concentrate in Large Cities. *Nature Human Behavior* (4). 248-254.

Bathelt, H., and Henn, S. (2014). The Geographies of Knowledge Transfer Over Distance: Toward a Typology. *Environment and Planning A: Economy and Space* 46(6). 1403-1424.

Bathelt, H., Malmberg, A., and Maskell, P. (2004). Clusters and Knowledge: Local Buzz, Global Pipelines, and the Process of Knowledge Creation. *Progress in Human Geography* 28. 31-56.

Berkes, E. (2020). Comprehensive Universe of U.S. Patents (CUSP): Data and Facts. Working Paper.

Boschma, R. (2015). Towards an Evolutionary Perspective on Regional Resilience. *Regional Studies* 49(5). 733-751.

Boschma, R., and Frenken, K. (2006). Why Is Economic Geography Not an Evolutionary Science? Towards an Evolutionary Economic Geography. *Journal of Economic Geography* 6(3). 273-302.

Boschma, R., and Lambooy, J. (1999). Evolutionary Economics and Economic Geography. *Journal of Evolutionary Economics* 9. 411-429.

Breschi, S., and Lissoni, F. (2009). Mobility of Skilled Workers and Co-Invention Networks: An Anatomy of Localized Knowledge Flows. *Journal of Economic Geography* 4(9). 439-468.

Brezis, E., and Krugman, P. (1997). Technology and the Life Cycle of Cities. *Journal of Economic Growth* 2(4). 369-383.

Broekel, T. (2019). Using Structural Diversity to Measure the Complexity of Technologies. *PLoS One* 14(5).

Cowan, R., and Jonard, N. (2003). The Dynamics of Collective Invention. *Journal of Economic Behavior and Organization* 52. 512-532.

Dosi, G. (1982). Technological Paradigms and Technological Trajectories. *Research Policy*. 11(3). 147-162.

Esposito, C., and Rigby, D. (2019). Buzz and Pipelines: The Costs and Benefits of Local and Nonlocal Interaction. *Journal of Economic Geography* 19(3): 753-773.

Fleming, L. (2001). Recombinant Uncertainty in Technological Search. *Management Science* 47(1). 117-132.

Fleming, L., and Frenken, K. (2007). The Evolution of Inventor Networks in the Silicon Valley and Boston Regions. *Advances in Complex Systems* 10(1). 53-71.

Fleming, L., and Sorenson, O. (2001). Technology as a Complex Adaptive System: Evidence from Patent Data. *Research Policy* 30(7). 1019-1039.

Foster, D., Rorick, M., Gesell, T., Feeney, L., and Foster, G. (2013). Dynamic Landscapes: A Model of Context and Contingency in Evolution. *Journal of Theoretical Biology* 334. 162-172.

Foster, J., and Evans, J. (2019). Promiscuous Inventions. In Love, A., and Wimsatt, W. (Eds.), *Beyond the Meme: Development and Structure in Cultural Evolution.* University of Minnesota Press, Minneapolis.

Funk, R., and Owen-Smith, J. (2016). A Dynamic Network Measure of Technological Change. *Management Science* 63(3). 587-600.

Gertler, M. (2003). Tacit Knowledge and the Economic Geography of Context, or The Undefinable Tacitness of Being (There). *Journal of Economic Geography* 3(1). 75-99.

Gertner, J. (2012). *The Idea Factory: Bell Labs and the Great Age of American Innovation*. Penguin Random House, New York.

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.

Kemeny, T., and Storper, M. (2020). Superstar Cities and Left-Behind Places: Disruptive Innovation, Labor Demand, and Interregional Inequality. *International Inequalities Institute, London School of Economics*, Working Paper 41.

Kerr, W. (2010). Breakthrough Inventions and Migrating Clusters of Innovation. *Journal of Urban Economics* 67(1). 46-60.

Klepper, S. (2007). Disagreements, Spinoffs, and the Evolution of Detroit as the Capital of the U.S. Automobile Industry. *Management Science* 53(4). 616-631.

Kogut, B., and Zander, U. (1992). Knowledge of the Firm, Combinative Capabilities, and the Replication of Technology. *Organization Science* 3(3). 383-397.

Kuhn, T. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.

Kwon, H., Lee, J., Lee, S., and Oh, R. (2020). Knowledge Spillovers and Patent Citations: Trends in Geographic Localization 1976-2015. *Economics of Innovation and New Technology*. In press.

Lamoreaux, N., Levenstein, M., and Sokoloff, K. (2004). Financing Invention During the Second Industrial Revolution: Cleveland, Ohio, 1870-1920. *NBER Working Paper #10923*.

Leamer, E., and Storper, M. (2001). The Economic Geography of the Internet Age. *Journal of International Business Studies 32*(4). 641-665.

Lundvall, B. (1988). Innovation as an Interactive Process: From User-Producer Interaction to the National Innovation Systems. In Dosi, G., Freeman, G., Nelson, R., Silverberg, G., and Soete, L. (Eds.), *Technology and Economic Theory*. Printer Publishers, London.

Martin, R., and Sunley, P. (1996). Paul Krugman's Geographical Economics and its Implications for Regional Development Theory: A Critical Assessment. *Economic Geography* (3). 259-292.

Martinelli, A., and Nomaler, O. (2014). Measuring Knowledge Persistence: A Genetic Approach to Patent Citation Networks. *Journal of Evolutionary Economics* 24(3). 623-652.

Morrison, A., Rabellotti, R., and Zirulia, L. (2013). When Do Global Pipelines Enhance the Diffusion of Knowledge in Clusters? *Economic Geography* 89. 77-96.

Neffke, F., Henning, M., and Boschma, R. (2013). How Do Regions Diversify Over Time? Industry Relatedness and the Development of New Growth Paths in Regions. *Economic Geography* 87(3). 237-265.

Nelson, R., and Winter, S. (1982). *An Evolutionary Theory of Economic Change*. Harvard University Press.

Nohria, N., and Eccles, R. (1992). *Networks and Organizations: Structure, Form, and Action.* Harvard Business School Press.

Nonaka, I., and Takeuchi, H. (1995). *The Knowledge Creating Company: How Japanese Companies Create the Dynamics of Innovation*. Oxford University Press, New York.

O'Mara, M. (2019). *The Code: Silicon Valley and the Remaking of America*. Penguin Press, New York.

Petralia, S., Balland, P., and Rigby, D. (2016). Unveiling the Geography of Historical Patents in the United States from 1836 to 1975. *Scientific Data* (3).

Polanyi, M. (1966). *The Tacit Dimension*. Doubleday, New York.

Powell, W., Koput, K., and Smith-Doerr, L. (1996). Interorganizational Collaboration and the Locus of Innovation: Networks of Learning in Biotechnology. *Administrative Science Quarterly* 41(1). 116-145.

Rigby, D. (2015). Technological Relatedness and Knowledge Space: Entry and Exit of US Cities from Patent Classes. *Regional Studies* 49(11). 1922-1937.

Romer, P. (1990). Endogenous Technological Change. *Journal of Political Economy* 98(5). S71-S102.

Saxenian, A. (1996). *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*. Harvard University Press.

Scott, A., and Storper, M. (1987). High Technology Industry and Regional Development: A Theoretical Critique and Reconstruction. *International Social Science Journal* 1(12). 215–232.

Sorenson, O., Rivkin, J., and Fleming, L. (2006). Complexity, Networks, and Knowledge Flow. *Research Policy* 35(7). 994-1017.

Storper, M., Kemeny, T., Makarem, N., and Osman, T. (2015). *The Rise and Fall of Urban Economies: Lessons from San Francisco and Los Angeles*. Stanford University Press.

Storper, M., and Venables, A. (2004). Buzz: Face-To-Face Contact and the Urban Economy. *Journal of Economic Geography* 4(1). 351-370.

Storper, M., and Walker, R. (1991). *The Capitalist Imperative: Territory, Technology, and Industrial Growth*. Basil Blackwell, New York.

Torre, A. (2008). On the Role Played by Temporary Geographical Proximity in Knowledge Transmission. *Regional Studies* 42(6). 869-889.

Usher, A. (1929). *A History of Mechanical Inventions*. Dover Publications, New York.

van der Wouden, F. (2020). A History of Collaboration in U.S. Invention: Changing Patterns of Co-Invention, Complexity, and Geography. *Industrial and Corporate Change* 29(3). 599-619.

Von Hippel, E., and Tyre, M. (1995). How Learning by Doing Is Done: Problem Identification in Novel Process Equipment. *Research Policy* 24. 1-12.

Weitzman, M. (1998). Recombinant Growth. *Quarterly Journal of Economics* 113(2). 331-360.

Wolfe, D., and Gertler, M. (2004). Clusters from the Inside and Out: Local Dynamics and Global Linkages. *Urban Studies* 41(5). 1071-1093.

Wuchty, S., Jones, B., and Uzzi, B. (2007). The Increasing Dominance of Teams in Production of Knowledge. *Science* 316.

Chapter 4: The Geography of Breakthrough Innovation in the United States

4.1) Abstract

Over the 20th century, the geography of breakthrough innovation in the United States – defined as the spatial distribution of the production of patents that are both novel and impactful – underwent three broad transformations. At the start of the 20th century, breakthrough innovation was concentrated in populous metropolitan areas where diverse technological knowledge circulated. However, by the 1930s breakthroughs were created less frequently across the entire country. Consequently, their invention had a less distinct geography. The large-scale creation of breakthroughs resumed in the 1960s and was once again concentrated in large and knowledge-diverse metropolitan areas. However, the invention of breakthroughs during the latter part of the century also frequently involved long-distance collaborations between inventors. In this chapter, I document these historical changes to the geography of breakthrough innovation and advance a model to explain why these changes occurred. The model proposes that the geography of breakthroughs is established by four factors: (1) the prevailing knowledge intensity of breakthrough inventions, (2) the distance-based frictions incurred by the technologies used for collaboration, (3) the distance-based frictions incurred by the technologies used for knowledge-sourcing, and (4) the disruptiveness of the regime of technological change. With respect to these parameters, I provide evidence to suggest that (a) the knowledge-intensity of breakthrough inventions increased over the 20th century, (b) the technologies used to collaborate across distance improved markedly over the 20th century, (c) the technologies used to source knowledge across distance did not improve over the 20th century, and (d) the regime of technological change was disruptive at the beginning and at the end of the 20th

century but incremental in the middle of the 20th century. I additionally show how the combination of these four parameters accurately predicts the empirical geography of breakthrough innovation during the beginning, middle, and end of the 20th century. I conclude the paper by discussing lessons that the 20th century's geography of breakthrough innovation provide for anticipating possible futures for the geography of innovation in the 21st century, including in the years beyond COVID-19.

4.2) Introduction

Innovation is a critical determinant of the competitiveness of firms and the aggregate economic prosperity of the residents of cities (Nelson and Winter, 1982; Moretti, 2012; Chetty et al. 2014). For these reasons, a widespread effort in urban economics, economic geography, and innovation science seeks to uncover the types of spatial environments that enhance creativity and promote innovation. Such analyses often focus on the spatial concentration of actors in regions with high population densities and ready access to flows of diverse ideas which circulate in those regions (Duranton and Puga, 2001; Mewes, 2019; Berkes and Gaetani, 2020). This line of research, however, is complicated by the fact that innovation has thrived in regions with very different local agglomeration densities. No two places have been more influential for the development of the agglomeration-based theory of innovation than Jane Jacobs' (1962; 1969) neighborhood of Greenwich Village in New York City and AnnaLee Saxenian's (1994) Silicon Valley, but the territorial form of these two agglomerations are vastly different: while ideas spilled across Greenwich Village's narrow streets and alleyways, Silicon Valley is currently a suburban landscape, and during Silicon Valley's initial phase of innovative growth, the region was borderline rural (O'Mara, 2018). Moreover, despite the current tendency for innovative activities to concentrate in large and dense metropolitan areas (Balland et al., 2020), important historical inventions such as the airplane and the cotton gin were made outside dense urban environments (Mokyr, 1990). The rise of non-local collaboration further complicates the relationship between agglomeration and innovation: the average distance between co-inventors of patents tripled in the United States between 1900 and 2015 (van der Wouden, 2020; Clancy 2020). The prevalence of innovation in urban, suburban, and rural environments, as well as the rise of inter-regional collaborations between inventors, demonstrates that there is no singular territorial form of economic activity which absolutely optimizes creativity and innovation.

Nonetheless, certain types of environments have proven to be advantageous for creative invention during specific periods of U.S. history. Anecdotal records and patent data indicate that rural innovation was prominent during the 18th and 19th centuries (Mokyr, 1990; Gordon, 2016; Balland et al., 2020; Mewes, 2019). Both patent data and employment records suggest that a big-city advantage for complex, high-impact, and well-compensated innovative activities emerged at the start of the 20th century (Desmet and Rossi-Hansberg, 2009; Bettencourt et al., 2007; Kemeny and Storper, 2020; Balland et al., 2018; Mewes, 2019; van der Wouden, 2019; Berkes and Gaetani, 2020). The agreement between patent data and employment records on the spatial concentration of innovative activities breaks down in the middle of the 20th century, when patent records indicate that innovation remained concentrated in large cities (Balland et al., 2018; Mewes, 2019) but employment records indicate that innovative occupations spread out across space (Desmet and Rossi-Hansberg, 2009; Kemeny and Storper, 2020). Finally, there is consensus that a strong big-city advantage for innovative activities emerged at the end of the 20th century. During this latter period, the productivity spillovers radiating from firms fully decayed across a distance as short as 250 meters (Arzaghi and Henderson, 2008; Baum-Snow et al., 2020). However, during the late 20th century non-local collaboration between the inventors of patents also became increasingly frequent (Bathelt et al., 2004; van der Wouden, 2019). With these caveats in mind, in Table 4.1 I synthesize these sources by presenting a stylized timeline of the broad shifts in the spatial distribution of innovative activities across regions of the United States.

Table 4.1: Geographical Distribution of Innovative Activities by Historical Period Suggested by Existing Literature

Before 1900	Early 20th Century	Mid-20th Century	Late 20th Century
Dispersed across space	Strongly concentrated in major metropolitan areas	Spatially dispersed high-wage employment; spatially concentrated patenting	Spatially concentrated high-wage employment; spatially concentrated patenting in major metropolitan areas but involving inter-regional collaborations

Source: Author's elaboration of sources cited in the paragraph above.

While suggestive, the timeline in Table 4.1 is only partly helpful for understanding how the relationship between agglomeration and creative innovation evolved over time. The outcomes summarized in Table 4.1 – overall patenting and high-wage employment – vary in terms of the extent to which they demand creative insight, as patents are often awarded to incremental inventions and high wages are often paid to workers for factors unrelated to creativity. This limitation to the existing literature motivates the first objective of this chapter: to describe how the relationship between agglomeration density and breakthrough innovation evolved in the United States over the 20th century. In this study, I emphasize breakthroughs because breakthroughs' combination of novelty and outward impact implies that they are outcomes of a highly creative process.

To describe the evolution of the relationship between agglomeration and breakthrough innovation, I analyze data from about 4 million patents granted between 1900 and 1999 to study how the propensity for inventors residing in knowledge-diverse metropolitan areas in the U.S. to create breakthrough inventions changed over time. I define knowledge-diverse regions as the Core-Based Statistical Areas (CBSAs) where local inventors patent in a wide array of patent classes.[9] Inventors that reside in knowledge-diverse metropolitan areas are able to source a broader set of ideas through face-to-face communication than are all other inventors and thus may invent breakthroughs more frequently. Because the frequency of non-local collaboration between inventors increased over the 20th century (van der Wouden, 2020; Clancy, 2020), I also examine changes in the propensity for inventors engaged in non-local collaborations to invent breakthroughs. My analysis generates three findings. First, in the early 20th century, inventors residing in knowledge-diverse regions were disproportionately more likely than inventors residing in non-knowledge-diverse cities (hereafter, *knowledge homogeneous regions*) to develop breakthroughs. Second, in the mid-20th century, inventors located in knowledge-diverse regions were no more likely than inventors in knowledge-homogeneous regions to develop breakthroughs. Third, at the end of the 20th century, inventors that both resided in knowledge-diverse regions *and* engaged in non-local collaborations were more likely than all other types of inventors, including those that resided in knowledge-diverse regions and collaborated locally, to develop breakthroughs.

The chapter's second objective is to explain why these changes to the geographical distribution of breakthrough innovation occurred. For this purpose, I develop a spatial model of endogenous knowledge production which emphasizes four factors that collectively determine the geographical distribution of breakthrough innovation. These factors are the knowledge intensity of breakthrough inventions, the distance-based frictions incurred by collaborative technologies, the distance-based frictions incurred by knowledge-sourcing technologies, and the disruptiveness of the prevailing regime of technological change. I elaborate on these four model parameters in Section 4.5. In Section 4.6, I show that the theoretical distribution of breakthrough innovation predicted by the model closely aligns with the observed distribution.

[9] Empirically, I define knowledge-diverse metropolitan areas using a year-specific variable, so a metropolitan area that is not knowledge-diverse in one year may be knowledge-diverse in later years.

The balance of the chapter consists of a review of the literatures on innovation and its geography (Section 4.3), a description of the methods used in the study (Sections 4.4 and 4.5), and an empirical analysis of the evolving geography of breakthrough innovation (Section 4.6). The introduction of the theoretical model, its calibration to empirical data, and the discussion section occupy Sections 4.7 and 4.8. In Section 4.9, I discuss how the model developed in this article revises a common perspective on why economic activities dispersed across space during the mid-20th century, and I share lessons that this historical revision imply for the future of the agglomeration of breakthrough innovation, including in the years after COVID-19.

4.3) Invention, Breakthroughs, and Location

Economic geography theory argues that strong distance-based frictions in endogenous knowledge production cause innovative activities to concentrate in space (Jaffe et al., 1993; Berkes and Gaetani, 2020). Those distance-based frictions arise because the quantity of information embedded in dense technological knowledge exceeds the bandwidth of long-distance communication technologies. In this regard, long-distance communication technologies compete poorly with face-to-face communication. The combination of vocal, visual, and physical cues that are possible with face-to-face communication allows partners to transmit a large quantity of information in a short amount of time and is a primary reason why inventors located in knowledge-rich locations hold an enduring advantage in sourcing ideas and developing of new technologies (Gertler, 2003; Storper and Venables, 2004).

The advantages of face-to-face communication are likely to be particularly pertinent during the creation of breakthrough inventions. This is because the creation of breakthroughs involves extensive technological search, which is performed more efficiently when the communication technology used to conduct that search has sufficient bandwidth.

Breakthroughs involve extensive search because they are novel and highly impactful. Breakthroughs are novel in that they combine existing ideas in dramatically imaginative ways, often departing from the well-worn search paths (Uzzi et al., 2013). Breakthroughs are highly-impactful because they combine existing ideas with a high level of complementarity and so enable a large quantity of subsequent innovation (Fleming et al., 2001). Together, the two criteria of novelty and impact imply that inventors need to identify well-functioning and radically new combinations of ideas in order to develop breakthroughs. Most of the combinations of ideas that inventors can create do not fit both criteria of novelty and complementarity, so inventors have to search widely amongst the set of combinatorial possibilities to identify the few that do (Youn et al., 2015).

Duranton and Puga (2001) and Berkes and Gaetani (2020) develop formal models where distance-based search costs cause innovating actors to co-locate in space. In Duranton and Puga's (2001) model, innovating firms agglomerate to minimize distance-based transaction costs while they search for inputs that are complementary to their production process. While Duranton and Puga's (2001) model makes the identification of complementary pairings of inputs and outputs endogenous to the process of agglomeration, they exogenously introduce novelty in their model by assuming that all firms enter the market with a new production process. Berkes and Gaetani's (2020) model explains the creation of novelty in densely populated cities through the increased exposure of inventors in dense cities to intra-industry spillovers. While their model makes the creation of novelty endogenous to the provision of local intra-industry spillovers, the model does not describe how complementary ideas are generated locally because all ideas in the model are made available to every firm after they are invented, regardless of a firm's location. Moreover, in Berkes and Gaetani (2020), novel combinations of ideas are generated locally, but complementary combinations are generated globally. Thus, while neither model generates an explicit prediction for how distance-based frictions affect the search for combinations of ideas that are both novel and complementary, they collectively propose that the creation of novel and complementarity combinations are both positive functions of the heterogeneity of ideas that circulate in local environs.

Despite widespread interest in breakthrough innovation, few empirical studies have investigated the spatial concentration of breakthrough innovation. One exception is Grashof et al. (2019), who find that breakthrough inventions made in Germany are disproportionately created by firms that are located geographically inside innovative clusters but whose inventors are in the periphery of their clusters' collaborative networks.[10] From these results, the authors infer that both local and non-local interactions between inventors are important for the creation of breakthroughs. Additional studies have separately examined the geographical distribution of the creation of novel inventions and impactful inventions, but they have not studied the geographical distribution of novel *and* impactful inventions in conjunction. Balland et al. (2020) show that overall patenting in the United States is concentrated in populous metropolitan areas and that this association is stronger for novel patents.[11] Mewes (2019) also studies the spatial concentration of overall patenting and novel patenting in the U.S. and finds both types of innovation to be concentrated in metropolitan areas with diverse local knowledge stocks. However, neither Balland et al. (2020) nor Mewes (2019) analyze the impact of those patents on subsequent invention. Berkes and Gaetani (2020) perform a similar analysis using U.S. counties as their unit of observation. In addition, Berkes and Gaetani (2020) test the aggregate relationship between the novelty of patents and the impact of patents, measured using forward citations. They find that novel inventions in

[10] Grafhof et al. (2019) refer to breakthrough inventions as "radical inventions". They define "radical inventions" as patents that are both novel *and* impactful, which is the definition of breakthroughs adopted by this paper.

[11] Balland et al. (2020) define novel patents as "complex" patents. Their measurement of "complexity", which measures the newness of the subclassification codes on patents, closely resembles this paper's definition of novelty.

the U.S. are disproportionately introduced in counties with high population densities and that novel patents are on average more impactful than non-novel inventions on subsequent innovation. However, Berkes and Gaetani (2020) do not analyze whether patents which are both novel and impactful are more often made in high-density counties. Finally, Castaldi et al. (2015) examine the knowledge-based characteristics of U.S. states that are more likely to produce high-impact patents (measured using forward citation counts). They find that inventors in states with diverse stocks of circulating unrelated ideas tend to produce high-impact inventions more frequently. However, Castaldi et al., (2015) do not analyze the novelty content of these inventions. In addition, Castaldi et al.'s (2015) study is at the state level, within which population and knowledge density can vary strongly. Thus, while each of these four studies of U.S. invention suggest that agglomeration economies are important for overall patenting, novel patenting, and high-impact patenting, they do not analyze the relationship between agglomeration and the production of patents that are both novel *and* impactful. As a result, the geography of breakthrough innovations in the U.S. has yet to be systematically described.

In addition to these issues related to the identification of breakthrough inventions, the geography of breakthrough innovation may contain important variations across time. As discussed earlier, the advantages that knowledge-diverse regions provide for the creation of breakthroughs is a function of the knowledge intensity of breakthrough innovation and the distance-based frictions incurred by the technologies used to collaborate and source knowledge. The states of these parameters are likely to change over time as the nature of the process of innovation and the state of communication technologies evolves (Lamoreaux and Sokoloff, 1996; Wuchty et al., 2007; Storper and Leamer 2001). In addition, the disruptiveness of the dominant regime of technological change may increase or decrease over time. Schumpeter (1934; 1942) proposes that there are periods of more incremental technological change and periods of more disruptive technological change. In incremental regimes, technological knowledge is primarily advanced through the introduction of incremental or “normal” inventions while in disruptive regimes the state of technological knowledge primarily advances through the creation of novelty. In an incremental regime, the novel inventions introduced to the economy may not be impactful enough to warrant the distinction of breakthrough inventions. Therefore, no geography of breakthrough innovation will establish itself during incremental regimes of technological change.

Two historical studies analyze the geographical concentration of innovation in the U.S. over an extended time period (Mewes, 2019; Balland et al., 2020). Both studies use USPTO patent records to measure innovative output and find that the spatial concentration of overall patenting increased between 1850 and 2000. While Balland et al. (2020) finds that the increased concentration is even stronger for novel patents (measured by the age of the subclassification codes assigned to patents), Mewes (2019) does not identify a significant difference between increased agglomeration of overall patenting and novel patenting using a slightly different measure of novelty. Again, neither study examines changes in the geographical concentration of breakthrough inventions.

Finally, there is growing recognition that the geography of innovation is more complex than a binary typology of spatial concentration or dispersion or an ordinal gradient spanning the two. In particular, non-local collaboration allows inventors to bridge separate inventive milieus, experiment with underexplored combinatorial possibilities, and possibly introduce high-impact inventions (Bathelt et al., 2004; Esposito and Rigby, 2018). While past studies have documented the increase in the prevalence of non-local collaborations (van der Wouden, 2020; Clancy, 2020), the relationship between non-local collaboration and the invention of breakthroughs has not been systematically studied.

4.4) Methods: Identifying Breakthrough Inventions

Breakthroughs inventions are the subset of inventions that are both novel and highly impactful. To empirically identify breakthroughs, one must assess individual inventions along both of these dimensions. Past research has defined novel inventions as those which generate entirely new ideas or recombine existing ideas in new ways. To this end, Uzzi et al. (2013) compute the atypicality of the knowledge combinations in individual inventions using z-scores, which calculate the extent to which each combination of knowledge units in a given invention deviates from the combinations inventors have made in the past. Kim et al. (2016) and Mewes (2019) apply this method to U.S. patent records using subclassification codes as indicators of the knowledge components in each invention, while Berkes and Gaetani (2020) apply z-scores to patent citations in a similar manner. Atypicality measured at the pairwise level between all ideas combined in an invention can be aggregated to the invention level to compute an invention's novelty.

In this study, I identify novel inventions by assessing the atypicality within them. I calculate the atypicality of all combinations of knowledge units in each invention by calculating z-scores for all USPTO patents issued between 1900 and 1999.[12] I use the coarse-grained subclasses assigned to each patent for this purpose, at which scale there are about 16,000 unique USPC subclasses (Kim et al., 2016). Because z-scores require a sufficient pre-history of patenting to accurately measure the mean frequency of the combination of any two subclasses, I compute Z-scores for the combinations of subclasses on patents granted starting in 1900 (Mewes, 2019). The z-score of the combination of subclass i with subclass j on a patent is given by Equation 4.1:

$$(4.1)\ Z_{i,j} = \frac{o_{i,j} - u_{i,j}}{\sigma_{i,j}}$$

[12] I source raw patent data and their USPC subclasses from the publicly-accessible Patents View website: https://www.patentsview.org/

In Equation 4.1, $o_{i,j}$is the number of past co-occurrences of subclasses i and j on all previously-granted patents. The term $u_{i,j}$ gives the expected number of past co-occurrences of subclasses i and j if inventors were to combine subclasses randomly. Its value is computed using Equation 4.2:

$$(4.2)\ u_{i,j} = \frac{n_i * n_j}{N}$$

In Equation 4.2, n_i and n_i are the respective cumulative number of patents that contain subclasses i and j on all prior patents, and N is the cumulative count of all prior patents. Finally, the variance of the subclass pairing, $\sigma^2_{i,j}$, is given by Equation 4.3:

$$(4.3)\ \sigma^2_{i,j} = u_{i,j}\left(1 - \frac{n_i}{N}\right)\left(\frac{N - n_j}{N - 1}\right)$$

$Z_{i,j}$ is positive when two subclasses are combined more frequently than expected given a random process, and negative when two subclasses are combined less frequently than expected given a random process. To generate a simple interpretation of the extent to which a combination is atypical, I follow Mewes (2019) and define atypical combinations as those with negative Z-scores. In addition, I define novel patents as those which introduce 1 or more atypical combinations of subclasses. I define all patents which do not introduce an atypical combination of subclasses as a "normal" patent.

The second criterion of breakthroughs is that they have outsized impact on subsequent innovation. To identify high-impact inventions, researchers often count the number of forward citations received by patents (Cremers et al., 1999; Hall, Jaffe, and Trajtenberg, 2001). In Chapter 2 of this dissertation, I developed related approach by tracing the flow of knowledge between individual patents based on the co-occurrence of fine-grained subclassification codes found on patents invented at different moments of time. There are two advantages to the latter method. First, citation records are unreliable for patents granted before 1947 (Berkes 2020) but the subclass codes used by my method are available for all

USPTO utility patents starting in 1836. Second, the same subclassification codes used to compute patent impact can also be used to assess the novelty profile of patents using Z-scores (Kim et al., 2016; Mewes, 2019). Thus, subclassification codes allow the novelty and impact of individual patents to be assessed using a common data input. Therefore, I follow the methods described in Chapter 2 to compute the impact of each patent on subsequent invention.

In Table 4.2, I present a typology of patents that vary in terms of impact and novelty. While in the subsequent analyses I treat the number of knowledge-based descendents of a patent as a continuous variable, for simplicity in Table 4.2 I convert patent impact into a binary measure by defining high-impact inventions as those that are in the top decile of the impact distribution for the same cohort year. The first quadrant of the 2x2 matrix describes the patents that do not introduce novelty and have low impact. Patents of this type are failed conservative experiments and they account for 72.4% of all USPTO utility patents granted between 1900 and 1999. The second quadrant of the matrix describes the inventions that do not introduce novelty but are nonetheless highly impactful. These incremental improvements account for 5.8% of USPTO patents granted 1900-1999. The third quadrant describes novel inventions that have low impact. These failed radical experiments account for 19.1% of USPTO patents 1900-1999. Finally, quadrant 4 describes the small percentage of inventions that are both novel and highly-impactful. These breakthrough inventions are rare, comprising just 2.7% of all USPTO patents 1900-1999. They are created when inventors deviate from the status quo in useful ways.

Table 4.2: Typology of Inventions by Novelty and Impact

	Low-Impact	High-Impact
Normal	(1) Failed conservative experiments 72.4% of Patents	(2) Incremental improvements 5.8% of Patents
Novel	(3) Failed radical experiments 19.1% of Patents	(4) Breakthroughs 2.7% of Patents

Patents granted 1900-1999. High-Impact patents are those in the top decile of their cohort year in terms of impact on subsequent invention. Because of integer cutoffs, the High-Impact column does not sum to 10%.

4.5) Methods: The Geography of Breakthroughs

After classifying each patent based on the typology in Table 4.2, I link patents to the metropolitan areas where they are invented. To do so, I use place-of-residence data provided van der Wouden (2020) for all U.S. inventors between 1836 and 1975, and I use place-of-residence data publicly available on the PatentsView website for all U.S. inventors between 1976 and 1999. I use constant-boundary 2015 definitions of metropolitan areas for this purpose. Because the innovative potential of inventors is expected to be greater for inventors residing in knowledge-diverse metropolitan areas (Duranton and Puga, 2001; Berkes and Gaetani, 2020), I measure the local knowledge diversity of the regions in which each patent is produced. I measure local knowledge diversity by counting the number of unique USPC coarse-grained subclassification codes assigned to the patents produced by inventors that reside in each core-based statistical area (CBSA) in a given year. Next, I transform the raw counts of local knowledge diversity into a binary variable by defining knowledge-diverse CBSAs as those where inventors produced patents in 10% or more of the USPC course-grained subclassification codes assigned to all U.S. patents in a given year. All CBSAs that do not meet the diversity criterion are labeled "knowledge-homogeneous cities". For example, in 1950 the USPTO assigned patents using 7,454 unique course-grain subclass codes, so in 1950 diverse CBSAs were those that produced patents with at least 745 unique subclasses. In 1950, 13 CBSAs met the diversity criterion.[13]

A core argument of this paper is that the generation of novelty is not important *per se*, because many novel inventions have minor downstream impact. Indeed, in Figure 4.1 I aggregate the total number of patents produced, the number of novel patents produced, and the number of high-impact patents produced to the CBSA level and show that there is no relationship between the production of novelty and the knowledge diversity of the regions where those patents are invented. In particular, Figure 4.1 shows that the concentration of total patenting in knowledge-diverse cities (solid line) is identical to the concentration of novel patenting in knowledge-diverse cities (dashed line). Thus, my data affirm the conclusion reached by Mewes (2019) that novel patenting is no more concentrated in knowledge-diverse cities than overall patenting is. On the other hand, Figure 4.1 shows that high-impact patenting (dotted line, defined as in Table 4.2) is more concentrated in knowledge-diverse cities than is overall patenting. Because these results show that the creation of novelty does not benefit from location in knowledge-diverse regions but the creation of high-impact inventions does, the results raise the possibility that the key advantage of local knowledge diversity is the enhanced ability to identify complementary combinations of ideas. Therefore, in the proceeding empirical analysis, I take the creation of novelty as a given, and examine how the impact of novel patents varies with the local knowledge diversity of the regions in which those novel patents are invented.

[13] In 1950, the knowledge-diverse CBSAs were (in descending order), New York, Chicago, Los Angeles, Philadelphia, Cleveland, Boston, Detroit, Pittsburgh, Cincinnati, San Francisco, Washington DC, Milwaukee, and Bridgeport CT.

Figure 4.1: Percentage of Patents Produced in Knowledge-Diverse Cities by Patent Type

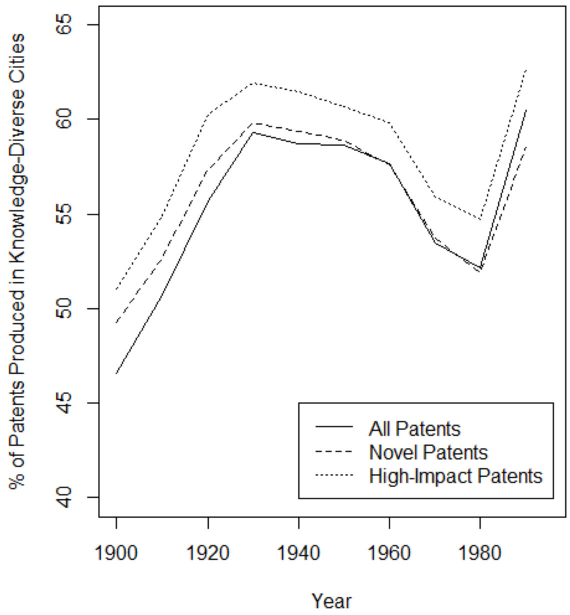

4.6) Results: The Geography of Breakthrough Innovation

For the reasons discussed above, I test whether the average impact of novel patents varies based on the diversity of the local knowledge stock in which they are invented. I perform this test in Figure 4.2 by creating time charts of the average impact of four types of patents: novel patents invented in knowledge-diverse cities (Nov | Div), novel patents invented in knowledge-homogeneous cities (Nov | Homog), normal patents invented in knowledge-diverse cities (Norm | Div), and normal patents invented in knowledge-homogeneous cities (Norm | Homog). Because there are about 4 million observations in the dataset, a scatterplot is infeasible so I instead show fit lines with 95% confidence intervals. The large number of observations also renders LOESS regression infeasible, so I produce the fit lines using a Generalized Additive Model (GAM) with a cubic spline smoothing parameter (Wood et al. 2017). I use this same plotting method for all subsequent figures.

Figure 4.2: Average Patent Impact by Novelty and Local Knowledge Diversity

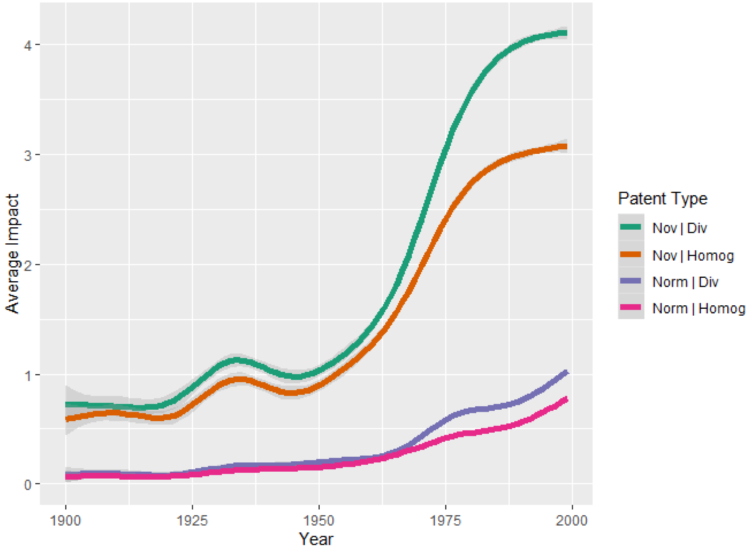

Figure 4.2 generates three inferences. First, across all years, novel patents invented in knowledge-diverse cities (Nov|Div) were on average the most impactful type of patents, followed by novel patents invented in knowledge-homogeneous cities (Nov|Homog) inventions. Second, the average impact of all types of inventions increased over time. Third, the increases in average impact were larger in knowledge-diverse cities: the impact of Nov|Div patents increased relative to Nov|Homog, and the impact of Norm|Div increased relative to Norm|Homog.

The increase in the average impact of Nov|Div patents relative to Nov|Homog patents suggests that the invention of breakthrough patents increasingly concentrated in knowledge-diverse cities over time. However, there are two reasons to exercise caution when interpreting this raw data. First, the large increases in average impact for all types of patents over time make it difficult to identify differential trends. Second, patents vary in terms of the number of subclasses assigned to them. Patents with more subclasses have higher impact values by virtue of their larger subclass count. The latter consideration arises because the method used to identify knowledge-based descendants searches for overlapping subclasses and combinations of subclasses on patents. Patents assigned many subclass codes therefore have more opportunities for knowledge-based descendants.

To take these two considerations into account, I compute the predicted impact of patents by adjusting for the year a patent is granted and the number of subclasses assigned to it. To compute predicted impact, I regress raw patent impact against a year*subclass count factor variable. I collect the residuals from the regression and plot them against the each patents' grant year, broken out by patent type to derive predicted impact values by patent type. The regression model used to predict these impact values is given by Equation 4.4:

$$(4.4)\ Impact_p = Year_p * NrSubclasses_p + E_p$$

In the dataset, there are 99 years and the number of subclasses assigned to patents ranges from 2 to 7, creating 594 unique values of the interaction factor variable. The predicted impact values, broken out by patent type and CBSA type, are presented in Figure 4.3.

Figure 4.3: Predicted Patent Impact by Novelty and Local Knowledge Diversity of CBSA of Invention

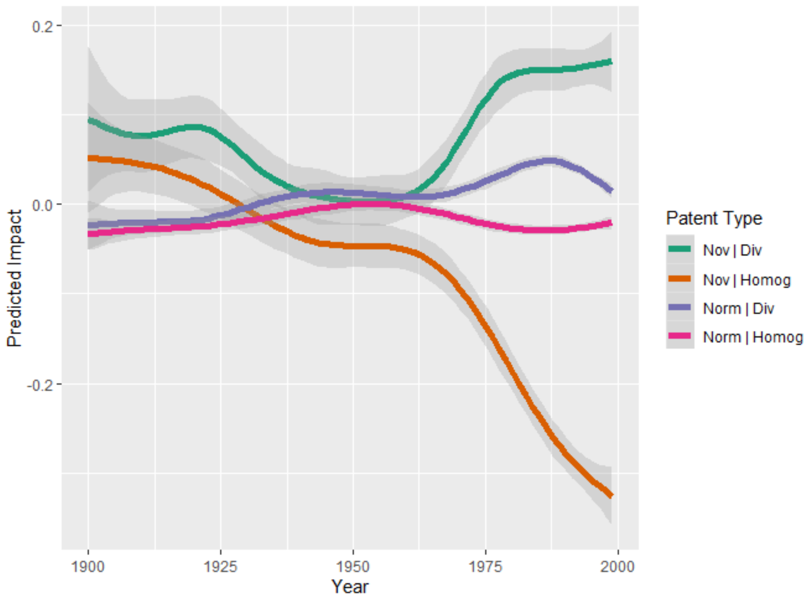

Note: The regression used to estimate predicated impact is given in Equation 4.4.

Figure 4.3 shows that the concentration of patenting by patent and city type went through three distinct periods during the 20th century. The first phase spanned from 1900 to 1930. During that period, novel patents were more impactful than normal patents. In addition, starting in 1910 novel patents invented in knowledge-diverse cities were significantly more impactful than novel patents invented in knowledge-homogeneous cities. The second period began around 1930 and lasted until approximately 1965. During this period, the predicted impact of novel inventions declined. By 1950, novel patents invented in knowledge-diverse cities were no more impactful than normal patents, and novel patents invented in knowledge-homogeneous cities were much less impactful than normal patents. The third period began around 1965 when the predicted impact of novel inventions made in knowledge-diverse cities increased above that of normal patents. During this period, the predicted impact of novel patents invented in knowledge-homogeneous cities declined. This latter result shows that by the end of the 20th century, breakthrough innovation was concentrated in these knowledge-diverse cities.

While Figure 4.3 shows that breakthrough innovation concentrated in knowledge-diverse cities at the end of the 20th century, the propensity for teams of inventors to collaborate non-locally also increased during that time period (Van der Wouden, 2019; Clancy, 2020). The increase in non-local collaboration suggests that the classical model of local innovation resulting from high distance-based communication costs became more complex (c.f. Duranton and Puga, 2001; Storper and Venables, 2004; Berkes and Gaetani, 2020), so in Figure 4.4 I examine the relationship between the engagement of inventors in non-local collaborations and the creation of breakthroughs. To study this relationship, I compare the average impact of patents invented by single-location and multi-locational teams of inventors. Multi-locational patents are defined as patents invented by teams of inventors that resided in different metropolitan areas when their patent was granted. In addition, I decompose teams based on the knowledge diversity of their home cities by differentiating between multi-locational teams that reside in knowledge-diverse and knowledge-homogeneous cities. To ease interpretation, I momentarily omit all teams with inventors in both knowledge-diverse and knowledge-homogeneous cities (I analyze these mixed teams in the appendix). Finally, I omit all patents invented by lone inventors.

Figure 4.4: Average Patent Impact of Collaborative Patents by Type of Collaboration

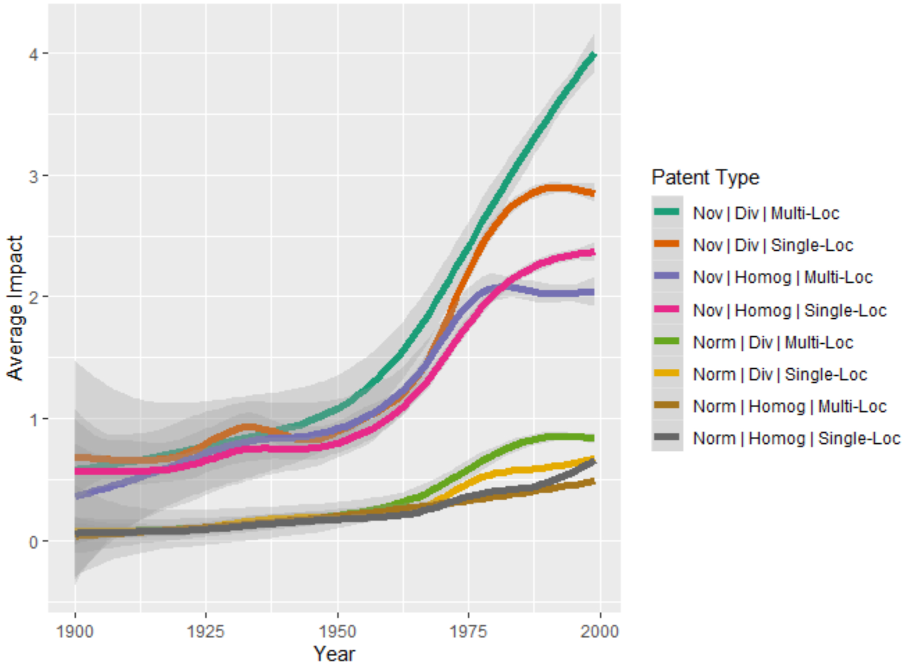

Figure 4.4 shows that the average impact of novel patents produced by teams in knowledge-diverse cities and in knowledge-homogeneous cities were statistically identical until 1960. In addition, before 1980 there was no significant difference in the average impact of novel patents produced by single-location or multi-locational teams. However, after 1980 novel patents produced by teams in knowledge-diverse cities became significantly more impactful than novel patents produced by single-location teams in knowledge-diverse cities or by teams of any type located in knowledge-homogeneous cities. The theoretical model developed in Section 4.7 will show how this spatial pattern can emerge when the state of collaborative and knowledge-sourcing technologies fit certain conditions.

Finally, in Figure 4.5 I display the predicted impact of patents based on their novelty, the knowledge diversity of their inventors' CBSAs, and whether their collaborative teams are multi-locational. As before, I compute the predicted impact of patents by regressing patent impact against the Year*NrSubclasses factor variable as in Equation 4.4 and aggregate the residuals by year and patent type.

Figure 4.5: Predicted Patent Impact of Collaborative Patents by Type of Collaboration

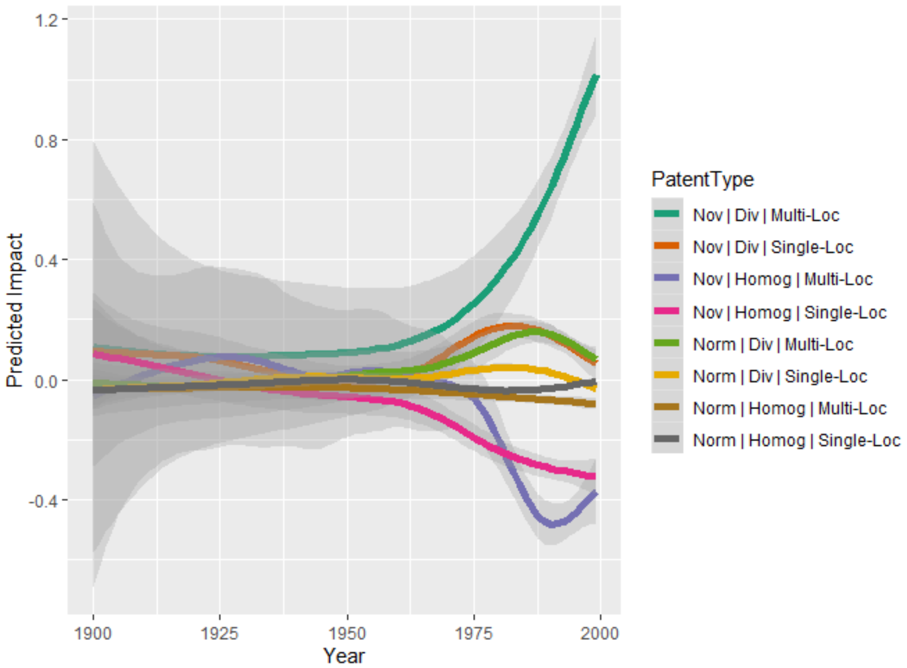

Note: Regression to estimate predicted impact is given in Equation 4.4.

Figure 4.5 shows that the predicted impact of all types of collaborative patents was identical until 1975. After 1975, the predicted impact of novel patents created by multi-locational teams residing in knowledge-diverse cities increased far above the predicted impact of any of the other types of patent. Thus, Figure 4.5 shows that the increasing concentration of breakthrough innovation in knowledge-diverse cities documented in Figure 4.3 was driven by inventors that collaborated with non-local teammates. Moreover, breakthrough innovation at the end of the 20th century was maximized through a combination of large clusters integrated into long-distance networks.

4.7) Interpretation of the Causes of Changes in the Geography of Breakthroughs

What explains the changes in the geography of breakthrough inventions documented in the text above? In this section, I develop a model in which the geography of breakthroughs at any moment in time is established by the articulation of four factors: the disruptiveness of the regime of technological change, the knowledge-intensity of breakthrough invention, the state of long-distance collaboration technology, and the state of long-distance knowledge-sourcing technology. Equation 4.5 writes this model as a series of interactions to express the interdependency between each of the factors:

$$(4.5)\ GeographyOfBreakthroughs_t = Disruptive_t * KnowedgeIntensity_t * LongDistaceCollabTech_t * LongDistaceSourchingTech_t$$

In Equation 4.5, $GeographyOfBreakthroughs_t$ can take one of four discrete values: the production of breakthrough inventions can be perfectly concentrated in a single metropolis, uniformly distributed across space, concentrated in a small number of major cities connected by non-local collaborations, or undefined. The geography of breakthroughs is undefined during time periods when no breakthroughs are introduced to the economy.

The model described by Equation 4.5 has four factors on the right-hand side of the equation, the first of which is the disruptiveness of the regime of technological change (Schumpeter, 1934; Schumpeter, 1942). This factor captures the extent to which technological knowledge advances through the creation of novel inventions versus normal or incremental ones. During periods of time when the disruptiveness of the regime of technological change is high, a large number of breakthroughs are created to the economy and so their invention can take on a distinct geography. When the disruptiveness of the regime of technological change is low, few breakthroughs are developed so their spatial distribution is indistinct – or more technically, undefined.

The second factor in the model is the knowledge intensity of breakthrough inventions (Wuchty et al., 2007; Balland et al., 2018; Bloom et al., 2020). This factor is defined as the prevailing returns that sourcing a larger number of knowledge-based inputs has on the creation of high-impact novelty. When the knowledge intensity of breakthroughs is high, the impact of novel inventions responds positively to the use of a large number of ideas when creating them.

The third factor in the model is long-distance collaboration technology. Long-distance collaboration technologies are the devices used by inventors to collaborate with co-inventors that reside in other regions, such as letters, email, videoconferencing, and long-distance travel. The robustness of long-distance collaboration technology is defined as its information loss relative to face-to-face collaboration. By assumption, face-to-face collaboration suffers no information loss but is only possible between inventors that reside in the same region (Storper and Venables, 2004). Therefore, when completely robust, long-distance collaboration technologies are perfect substitutes for face-to-face communication and inventors are able to create novelty through collaborations across distance with no loss in impact.

The fourth factor is long-distance knowledge-sourcing technology. Long-distance knowledge-sourcing technologies are the items inventors use to source ideas from non-local regions in which they do not have active collaborators. They include scientific articles, patent documents, and physical technological devices that can be reverse-engineered. The robustness of long-distance knowledge-sourcing technologies is defined by their information loss relative to the information loss encountered by inventors who source local knowledge. By assumption, sourcing local knowledge, which benefits from physical presence and embeddedness in an environment of shared norms, is assumed to entail no information loss (Gertler, 2003). Therefore, when long-distance knowledge-sourcing technologies are completely robust, they are perfect substitutes for local knowledge sourcing and inventors are able to source ideas from all locations in the world in the creation of novelty without a reduction in impact.

For simplicity, I assume that all independent variables are binary. Therefore, when $Disruptive_t = 1$, technological change is advanced through novel inventions as opposed to normal inventions; when $KnowledgeIntensity_t = 1$, the knowledge-intensity of breakthrough inventions is high, meaning that inventors need to source a large quantity of knowledge in order to create high-impact novelty; when $LDCollabTech_t = 1$, inventors are able to collaborate non-locally with no loss of information; when $LDSourcingTech_t = 1$, inventors are able to source non-local ideas with no loss of information. Because equation 6 has four binary parameters, a 2x2x2x2 hypercube is required to demonstrate each of the equation’s possible outcome states. To ease the exploration of the equilibria, I present the possible states of the model in Table 4.3 using a 2x8 matrix. In the matrix, changes in the disruptiveness of the regime of technological change are shown along the matrix’s columns, and changes in knowledge intensity, long-distance collaboration technology, and long-distance learning technology are shown along its rows. The cells of the matrix with white backgrounds show the predicted geography of breakthrough innovation given a set of conditions of disruptiveness, knowledge intensity, long-distance collaboration technology, and knowledge-sourcing technology.

Table 4.3: Geography of Breakthroughs under Conditions of Technological Disruptiveness, Knowledge Intensity, Long-Distance Collaboration Technology, and Long-Distance Knowledge-Sourcing Technology

Knowledge Intensity	L.D. Collab Tech	L.D. K-Sourcing Tech	Disruptiveness of Tech Regime: Weak	Disruptiveness of Tech Regime: Strong
Weak	Weak	Weak	Undefined	Dispersed
		Strong	Undefined	Dispersed
	Strong	Weak	Undefined	Dispersed
		Strong	Undefined	Dispersed
Strong	Weak	Weak	Undefined	Perfectly Concentrated
		Strong	Undefined	Dispersed
	Strong	Weak	Undefined	Multi-Nodal
		Strong	Undefined	Dispersed

Table 4.3 shows Equation 4.5's predictions for the geography of breakthrough invention given any combination of disruption, knowledge intensity, long-distance collaboration technology, and long-distance knowledge sourcing technology. For example, the top-left white-background cell in Table 4.3 contains the value "Undefined", indicating that the geography of breakthroughs is undefined when the disruptiveness of the regime of technological change, the knowledge intensity of breakthroughs, and long-distance collaboration and knowledge-sourcing technologies are all weak. The geography is undefined under these conditions because the non-disruptive nature of the regime of technological change implies that few to no breakthroughs are introduced into the economy. More generally, the geography of breakthroughs is undefined whenever the technological regime is not disruptive.

More interesting geographies of breakthroughs emerge when the disruptiveness of the regime of technological change is strong. For example, under the conditions of a strongly disruptive regime of technological change but weak long-distance collaboration and knowledge-sourcing technologies and weak knowledge intensity, the geography of breakthroughs is dispersed across space. The reason the geography is dispersed under these conditions is because the knowledge intensity of breakthrough innovation is low. Low knowledge intensity means that inventors do not need to source many ideas to develop high-impact novelty, so the weak states of collaboration and knowledge-sourcing technologies are not major impediments to the creation of breakthroughs by spatially-isolated inventors. These conditions may describe U.S. invention during the 19th century, when anecdotal evidence indicates that many breakthroughs were invented in the countryside (Mokyr, 1990).

When the knowledge-intensity of breakthrough innovation is strong and the regime of technological change is disruptive, the resulting geography of breakthrough invention depends on the strength of long-distance collaboration and knowledge-sourcing technologies. When long-distance collaboration and knowledge-sourcing technologies are both weak, the geographical distribution is perfectly concentrated in a single region because inventors need a large quantity of knowledge to develop breakthroughs, but they can only gain clear informational signals through co-presence with other inventors or objects. If, *ceteris paribus,* long-distance knowledge-sourcing technology becomes stronger, the spatial distribution disperses because inventors can source ideas from any region in the world without loss in impact. The most interesting state of the model, however, is when the knowledge intensity of breakthroughs, the disruptiveness of the regime of technological change, and long-distance collaboration technology are all strong but long-distance learning technology is weak. This state results in a multi-nodal geography of breakthroughs because inventors (a) need large quantities of knowledge to invent breakthroughs, (b) collaborate with minimal information loss with inventors in other locations, and (c) experience a large amount of information loss when they source knowledge from locations where they do not have active collaborators on-the-ground. The multi-nodal structure emerges as inventors establish multi-locational teams to source diverse ideas held by actors in various locations.

4.8) Measurement of the Parameters of the Theoretical Model to Data

How did the state of the disruptiveness of the regime of technological change, the knowledge intensity of breakthroughs, and the distance-based frictions incurred by collaborative and learning technologies evolve over the 20th century? In this section, I review evidence from patent records to document changes in the model's parameters over the study period.

I begin by describing the evolution of the knowledge intensity of breakthrough innovation, which is measured as the added benefit of additional knowledge sources for helping inventors to create high-impact novelty. I measure the number of prior knowledge sources that each patent draws ideas from as the in-degree of the knowledge flow network described in Chapter 2. Next, I transform the number of knowledge sources used by the inventors of each patent into a binary variable by defining patents with "many knowledge-based parents" as the patents in the top decile of a given year in the knowledge source count distribution. I define patents as having "few knowledge-based parents" if they fall in the bottom 90% of the knowledge source count distribution from their same year. The empirical goal is to examine if the average impact of patents with many knowledge-based parents increased more than the average impact of patents with few knowledge-based parents over the 20th century.

As in the previous analyses, I adjust for the relationships between the number of subclasses on a patent and the year it is granted by regressing patent impact against a Decade*SubclassCount factor variable, as in Equation 4.4. I collect the residuals from the model, aggregate them to groups based on the novelty and knowledge-intensity of patents, estimate GAM-function fit lines to these predicted impact values, and plot those fit lines with 95% confidence intervals by year. The resulting fit lines are shown in Figure 4.6. A similar figure using raw impact values is presented in the chapter's appendix.

Figure 4.6: Predicted Patent Impact by Novelty and Number of Patent Parents

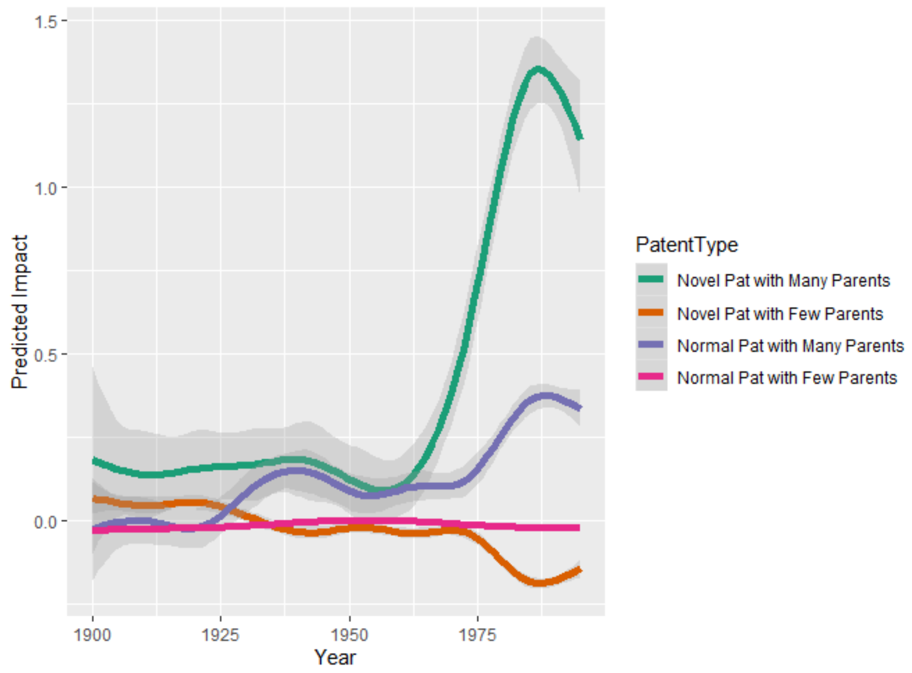

Note: Regression to estimate predicted impact is given in Equation 4.4.

In Figure 4.6, the predicted impact of novel patents with many patents (green line) is slightly but statistically-significantly higher than the predicted impact of novel patents with few patents (orange line) until about 1965. Thereafter, the predicted impact of novel patents that source knowledge from many parent patents increases sharply while the predicted impact of novel patents with few parents declines. Therefore, the knowledge intensity of breakthroughs was moderate until 1965 but very high after 1965. The relationships identified in Figure 4.6 are also present within technological fields, as shown in Figure 4.13 in the appendix, which illustrates the predicted impact of the four types of patents with the inclusion of an aggregate technology class fixed effect.

Next, I investigate changes in the strength of long-distance communication technologies. There are two types of long-distance communication technologies: long-distance collaboration technology, and long-distance knowledge-sourcing technology. I measure the strength of each type of long-distance communication technology based on the revealed ability for inventors to create high-impact novelty while collaborating with distant teammates or while sourcing knowledge from distant environs. Figure 4.5 presented evidence that long-distance collaborative technology was weak before 1960 but grew stronger thereafter. In particular, the average impact of novel patents invented by multi-locational teams in knowledge diverse cities climbed well above that of novel patents invented by single-location teams starting in the 1960s.

To assess the strength of long-distance knowledge-sourcing technologies, I test whether novel patents created by inventors who source knowledge locally are more impactful than novel patents created by inventors who source knowledge non-locally. In administering this test, I develop a definition of local knowledge sourcing that defines knowledge sourced from locations in which multi-locational teams have on-the-ground collaborators as local. Only one teammate needs to reside in a given CBSA for the knowledge sourced from that CBSA to be considered local. In particular, I label patents as "patents which source knowledge with proximity" as the patents for which an above-average number of their knowledge sources are from CBSAs in which the patent's inventors reside. All other patents are defined as patents that "source knowledge without proximity". I recompute the average number of local knowledge sources within each year, so in any given year half of all granted patents are defined as patents that source knowledge with proximity. As in the previous analyses, I account for changes in the average impact of patents across time and across patents assigned a different number of subclasses by regressing the impact of patents against a Year*NrSubclasses factor variable as in equation 4 to compute predicted impact. I plot the predicted values, aggregated by patent type, in Figure 4.7.

Figure 4.7: Predicted Patent Impact by Novelty and Extent to which a Patent Sources Knowledge with Proximity

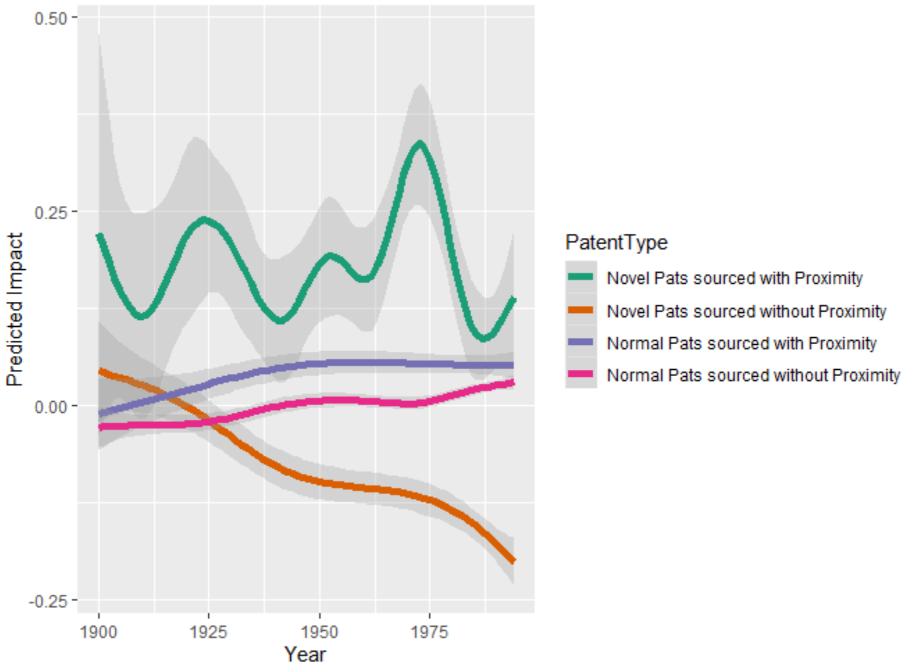

Note: Regression to estimate predicated impact is given in Equation 4.4.

Figure 4.7 indicates that novel patents using knowledge sourced with proximity were more impactful than novel patents using knowledge sourced without proximity during the full study period. Moreover, the green line is always significantly above the orange line. The persistent advantage of sourcing knowledge with proximity for creating high-impact novel patents is also robust to the inclusion of fixed effects for the aggregate classification code of patents (Figure 4.14). These results suggest that minimal progress was made over the 20th century to improve the ability for inventors to source knowledge from locations where they do not have active collaborators. When viewed alongside Figure 4.5's finding that breakthroughs were disproportionately produced by multi-locational teams toward the end of the 20th century, Figure 4.7 suggests that multi-locational teams have emerged in response to the inability for inventors to source knowledge from regions where they do not have collaborators.

Finally, I document changes in the disruptiveness of the regime of technological change over the 20th century. I measure the disruptiveness of the technological change regime by comparing the average impact of novel patents relative to that of normal patents. Again, I control for changes in the impact of patents across decades and across patents with different numbers of subclass codes by plotting predicted impact values using Equation 4.4. The predicted impact values, presented in Figure 4.8, show that novel patents were more impactful than normal patents during the early 20th century. Thereafter, the average impact of novel patents decline and eventually fall below that of normal patents.

Figure 4.8: Predicted Patent Impact by Novelty

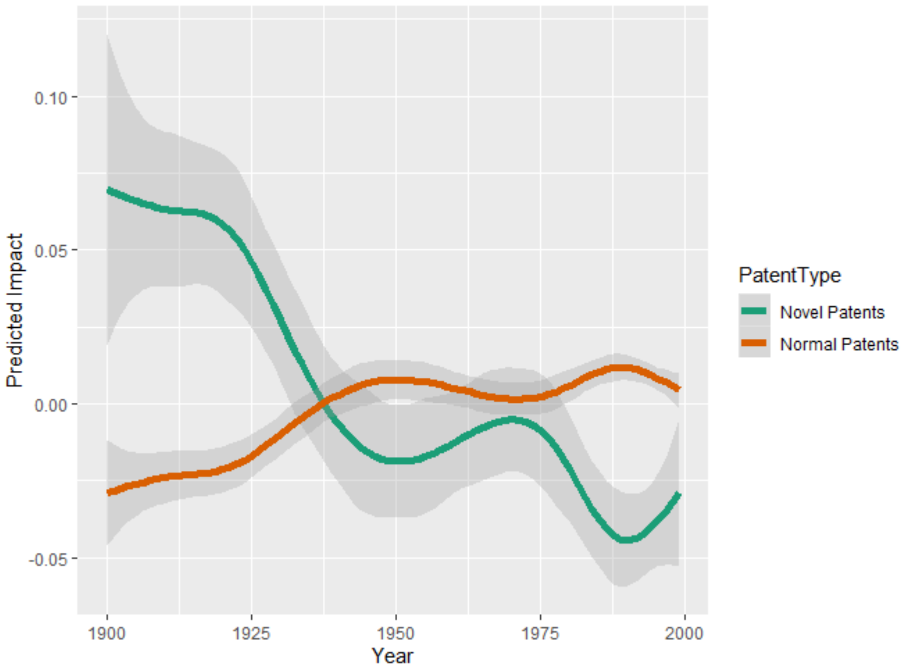

The decline in the impact of novel patents relative to normal patents between 1900 and 1950 indicates that technological change was less disruptive during the middle and end of the 20th century. The low level of disruptiveness at the end of the century seems at odds with the earlier finding that the average impact of novel patents invented in knowledge-diverse cities rebounded during the 1960s (Figure 4.3). One possible explanation to reconcile these two findings is that novelty was produced in both knowledge-diverse and knowledge-homogeneous cities at the end of the 20th century and that the novelty produced in knowledge-homogeneous cities was particularly low-impact.

To test whether inventors in knowledge-homogeneous cities developed fundamentally different types of novelty than did inventors in knowledge-diverse cities, I analyze how the average impact of novel inventions evolved relative to normal inventions within broad technological fields. If inventors in knowledge-homogeneous cities develop a large quantity of low-impact novelty in technological fields that are generally low impact, then controlling for the mean impact of each technological field will project these values out of the data. Therefore, in Figure 4.9 I perform the same analysis as in Figure 8 before but add a fixed effect for the primary USPC class of each patent to the regression model. The regression model is given by Equation 4.6, where FE_{C438} designates fixed effects at the primary class level, at which scale there are 438 unique classes:

$$(4.6)\ Impact_p = Year_p * NrSubclasses_p + FE_{C438} + E_p$$

Figure 4.9: Predicted Patent Impact by Novelty with Aggregate Technology Class Fixed Effects

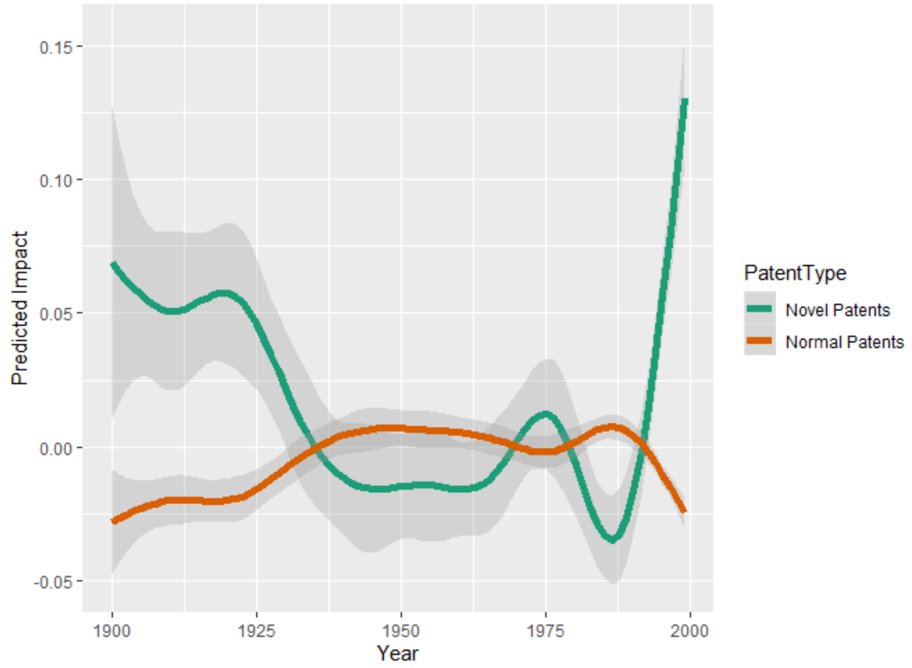

Note: The regression used to estimate predicated impact is given in Equation 4.6.

In Figure 4.9, the average impact of novel patents declines during the first several decades of the century, bottoms out in 1955, and jumps after 1985. The relationship presented in Figure 4.9 is robust to the use of a more detailed course-grained subclass fixed effect (Figure 4.15) and indicates that the regime of technological change became very disruptive within technological classes at the end of the 20th century, but that disruption did not extend beyond technology classes. Moreover, while many novel and impactful technologies were introduced between 1985 and 1999, they were not sufficiently impactful to shift the entire economy into a disruptive regime of technological change. This finding is similar to Gordon's (2016) inference that the information technology revolution failed to revolutionize a broad an expanse of the economy in contrast with the breakthrough inventions of the early 20th century. Thus, we may conclude from Figure 4.9 that a narrow sector of the economy became disruptive between 1985 and 1999.

To conclude the analysis, in Table 4.4 I assemble together the observed state of the knowledge-intensity of breakthrough innovation, the state of collaborative technologies and learning technologies, and the disruptiveness of the regime of technological change to generate the empirically-predicted state of the geography of breakthrough innovation for the early, mid, and late 20th century.

Table 4.4: Observed States of Model Parameters and Model-Predicted Geography of Breakthrough Inventions by Time Period

	Approximate Time Period		
Parameter	**1900-1930**	**1930-1970**	**1970-1999**
Knowledge Intensity	Moderate	Moderate	High
Disruptiveness	High	Low	High *within* sectors
LD Collaboration Tech	Weak	Weak	Strong
LD Knowledge-Sourcing Tech	Weak	Weak	Weak
Predicted Geography of Breakthroughs	Weakly concentrated	Undefined	Multi-Nuclei

Note: the empirical observation of the four model parameters are described in the text. The predicted geography of breakthrough inventions is presented in Table 4.3.

To test the model, the predicted geographies of breakthrough innovation from Table 4.4 can be compared to the observed geographies documented in Figure 4.3 and Figure 4.5. Notably, the states of breakthrough innovation predicted in Table 4.4 closely correspond to the empirical distributions found in Figures 4.3 and 4.5. During the first part of the 20th century, the weakness of long-distance collaboration and knowledge-sourcing technologies, high disruptiveness, and moderate knowledge intensity of breakthroughs implies a weakly concentrated geography of breakthrough innovation. Figure 4.3 bears out this prediction by showing that the predicted impact of novel patents was slightly higher for patents invented in knowledge-diverse cities than for patents invented in knowledge-homogeneous cities. During the mid-20th century (approximately 1930-1970), long-distance collaboration and knowledge-sourcing technologies were still comparatively poor and the knowledge intensity of breakthroughs was moderate. While these factors *ceteris paribus* would predict a spatially-concentrated geography of breakthrough innovation, the disruptiveness of the regime of technological change was low. Because few breakthroughs were invented during this time period, the geography of breakthrough innovation was undefined. This proposition is empirically confirmed in Figure 4.3 where the average impact of novel patents is shown to be no higher than the average impact of normal patents, regardless of the local knowledge diversity in which the novel patents are invented. Finally, at the end of the 20th century the combination of a high knowledge intensity of breakthroughs, strong long-distance collaboration technology, weak long-distance knowledge-sourcing technology, and a high disruptiveness of technological change within sectors (as measured using technology classification codes) predicts a multi-nuclei geography of breakthrough innovation. That geography is emerges because inventors concentrate in knowledge-diverse metropolitan areas to source knowledge locally (because long-distance knowledge-sourcing technologies are poor) but nonetheless collaborate with non-local teammates by taking advantage of the stronger long-distance collaboration technology. The multi-location team is thus rationalized as a response to the asymmetric improvements to long-distance knowledge sourcing technology and long-distance collaboration technology over the 20th century. The geography predicted by these parameters corresponds to the observed distribution described in Figure 4.5, where high-impact novelty was shown to be produced by multi-location teams with co-inventors residing in multiple knowledge-diverse cities.

An important caveat regarding the geography of breakthroughs at the end of the 20th century is that the breakthroughs produced during this period were not very impactful outside the sectors in which they were invented. This finding, evident in Figures 4.8 and 4.9, indicates that the knowledge-diverse cities where most of the breakthroughs of the late 20th century were made were engaged in relatively esoteric technological problems. Over time, these inventions may have diffused throughout the economy and instigate an economy-wide period of disruptive technological change. However, such a transformation had not taken hold by the end of the 20th century.

4.9) Discussion

The spatial concentration of innovation is not an inherent quality of density, agglomeration, or urbanization (c.f. Duranton and Puga, 2001; Mewes, 2019; Balland et al., 2020; Berkes and Gaetani, 2020). Instead, innovation spatially concentrates, disperses, or adopts a non-ordinal distribution as more fundamental changes institutions and communication technologies take hold. These institutional and technological factors determine inventors' general need to interact in order to create inventions, and the frictions involved in sustaining those interactions across distance.

The focus of this paper was to document changes in the spatial distribution of breakthrough innovation in the United States evolved over the 20th century and to advance an explanation for why those changes occurred. To this end, I began the paper by describing how the advantages afforded by locating in knowledge diverse cities and participating in multi-locational collaborations for creating high-impact novelty changed over time. Thereafter, I proposed a model in which breakthrough inventions are generated through interactions sustained by collaboration technologies and knowledge-sourcing technologies that incur different levels of distance-based frictions and within regimes of technological change that vary in terms of their disruptiveness and knowledge-intensity. Finally, I showed that the model predicts geographical distributions of breakthrough innovation which closely align with the observed distributions in the United States over the 20th century.

Explicit recognition of how institutional and technological contingencies shape spatial distributions of innovation can help to revise existing understandings of why certain geographies have emerged historically. One example is the mid-20th century, which is broadly understood to be an era during which economic activities in the U.S. spread out across space (Rosen, 1979; Roback, 1982; Glaeser and Tobio, 2007; Glaeser, 2008). According to these sources, the spreading out of economic activities in the middle of the century was caused by improvements in communication technologies, decreases in transportation costs, and the high cost of housing and labor in densely-populated locations. However, the results from this study suggest that a reduction in the disruptiveness of the regime of technological change also may have supported the dispersal of economic activities during the mid-20th. As documented in Figure 4.3, breakthrough innovation did not disperse across space during the mid-20th century; instead, few to any breakthroughs were invented across the entire country during that time period. In the knowledge economy, the advantages of agglomeration are larger for firms that compete in environments riddled by uncertainty and rapid change (Duranton and Puga, 2001; Delgado et al., 2015; Lin, 2012; Berger and Frey, 2016). If the regime of technological change was less disruptive during the mid-20th century, firms may have faced minimal pressure to locate in dense agglomerations (Kemeny and Storper, 2020).

This historical insight may prove helpful for predicting future changes to the geography of breakthrough innovation. The COVID-19 pandemic has shifted many strongly agglomerated high-skilled service jobs to remote work (Dingel and Neiman, 2020). Recent advancements in communication technologies are generally thought to have reduced the costs associated with sharing knowledge across space. Some authors argue that this widespread temporary adoption of remote work will shift the economy to a new spatial equilibrium of geographical dispersion (Catalini et al., 2018; Dong et al., 2018; Agrawal et al., 2017; Clancy, 2020). While the future may break from the past and a geographically-dispersed distribution of breakthrough innovation may indeed prevail, this study emphasizes that there is no historical precedent from the 20th century in the United States for such a dispersal of breakthrough innovation in the absence of a reduction in the disruptiveness of the regime of technological change.

In turn, any forecast of the post-COVID-19 geography of breakthrough innovation needs to pay careful attention to a possible decline in technological disruption. Notably, market concentration in firms in the United States has reached its highest value since the 1970s (Autor et al., 2017; Grullon et al., 2019). The ongoing increase in market concentration may either cause, or be a result of, a slowdown in technological change as the competencies of incumbent firms are less frequently disrupted by new product or process technologies. If technological change is increasingly advanced through incremental inventions, as was the case during the mid-20th century, then companies and industries may de-agglomerate following COVID-19 not just because of the widespread adoption of Skype and Zoom, but also because of the advantages of co-location will be less important in a period of greater technological stability.

The current literature on the effect of market concentration and the geographical distribution of economic activity has not yet investigated this relationship between oligopolistic market structure and the demand for co-location (Manduca, 2019; Feldman et al. 2020). Instead, that literature focuses on how the rents accrued by oligopolistic firms concentrate wealth in those firms' immediate spatial environs. The policy response advocated by the existing literature is to increase antitrust enforcement in order to reduce inter-regional income inequality. Increasing antitrust enforcement may reduce inter-regional income inequality by shrinking the monopoly rents bestowed on "superstar metros". However, increasing competition in innovative industries through antitrust enforcement may also *increase* inter-regional income inequality by stimulating faster and more disruptive technological change. Management theory, network theory, and product cycle theory all emphasize that small firms better adapt to disruptive technological change than large ones

(Acs and Audretsch, 1988; Feldman and Audretsch 1999). In addition, economic geography has strongly argued that cities and regions informally coordinate production amongst small firms when market conditions are fast-moving and riddled with uncertainty (Scott, 1988; Saxenian, 1994; Levinthal, 1997; Storper et al., 2016). If the organizational ties of firms are broken through antitrust enforcement, an alternative organization of inter-inventor coordination is likely to emerge. Historically, in the absence of organizational ties, that coordination has been achieved through co-location.

In conclusion, the analysis in this paper generates three core insights for interpreting and forecasting the geography of breakthrough innovation. First, the geography of breakthrough innovation changes over time as social, economic, and technological conditions evolve. Second, by identifying changes to the broader social, economic and technological conditions and by modeling their interrelationships, research can inform and improve predictions for past and future distributions of the geography of breakthrough innovation. Third, breakthrough innovation in the post-COVID-19 era is likely to involve high knowledge intensity, powerful collaborative technologies, high market concentration, and a possible reduction in the disruptiveness of the regime of technological change. Careful measurement and modeling of these four parameters is needed for researchers and policy makers to understand and rectify the new geographical and technological challenges that are bound to emerge.

4.10) Appendix

Multi-Locational Collaboration Type

The following figures examine the average impact of novel and normal patents that are created through non-local collaborations based on the knowledge diversity of respective cities. For simplicity, I restrict the data to collaborative teams located in two metropolitan areas. This generates 3 types of collaborative possibilities: collaborations between inventors located in two knowledge-diverse cities (Div-Div), collaborations between inventors located in one diverse and one homogeneous city (Mixed), and collaborations between inventors located in two homogenous cities (Homog-Homog).

Figure 4.10: Average Impact of Multi-Locational Patents by Collaboration Type

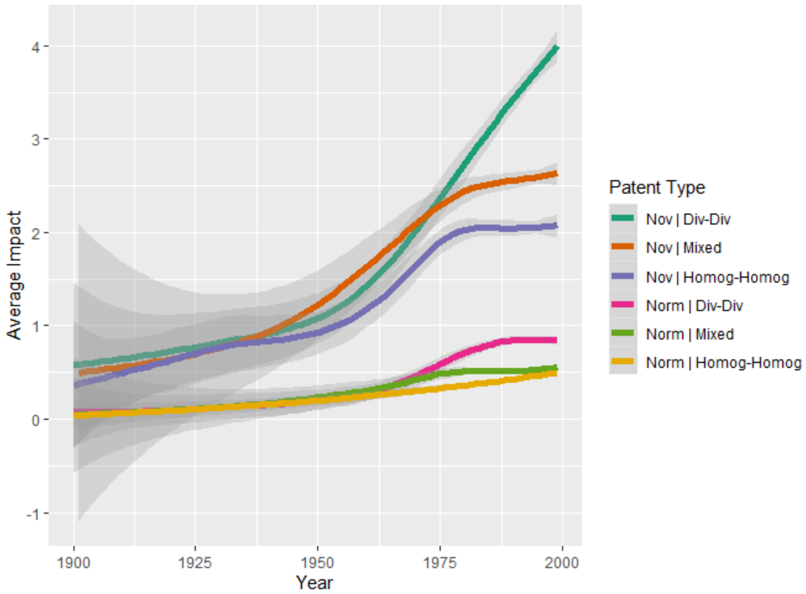

To compute the residual impact of inventions, I collect residuals from the following model and display them in Figure 4.11:

$$Impact_p = Year_p * NrSubClasses_p + E_p$$

Figure 4.11: Predicted Impact of Multi-Locational Patents by Collaboration Type

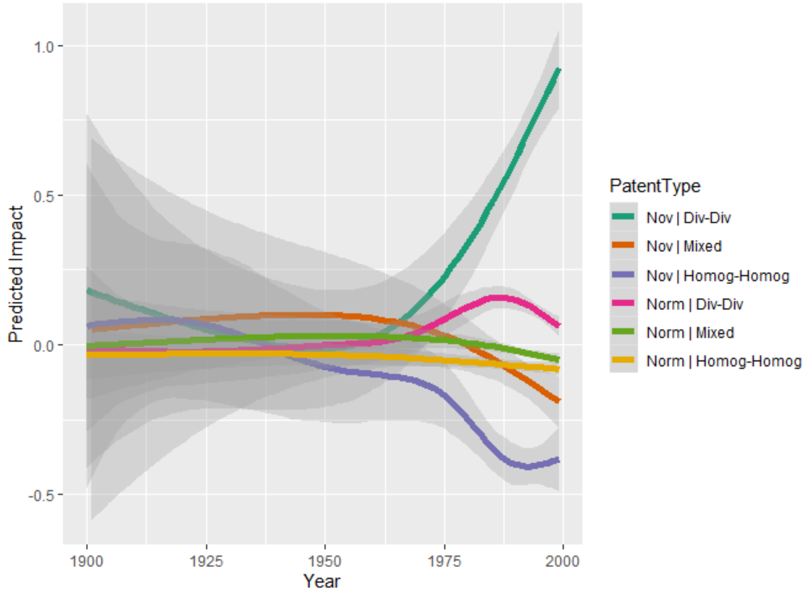

Knowledge Intensity and Impact of Patents

Figure 4.12 plots the raw patent impact of novel and normal patents with many and few patents by year.

Figure 4.12: Predicted Impact by Novelty and Knowledge Intensity

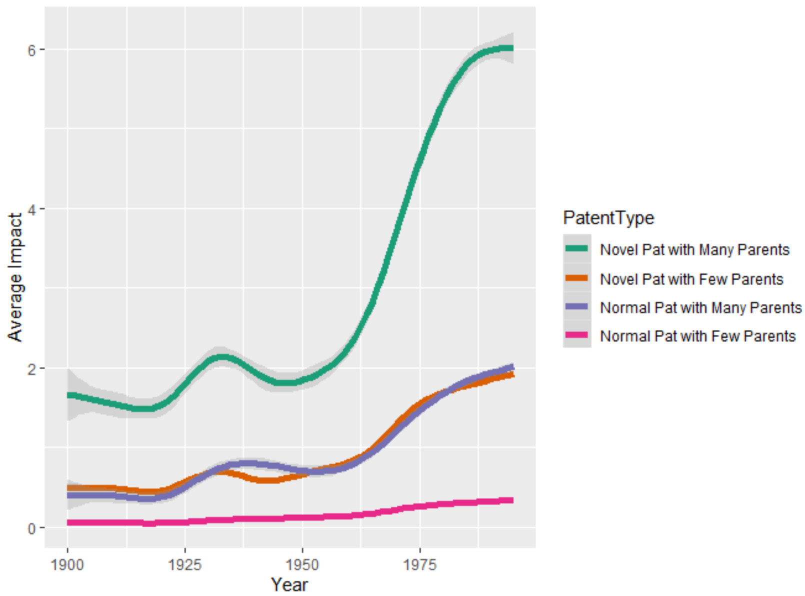

Figure 4.13 plots the residuals from the following model:

$$Impact_p = Decade_p * NrSubClasses_p + PrimaryClassFE_p + E_p$$

Figure 4.13: Predicted Impact by Novelty and Knowledge Intensity

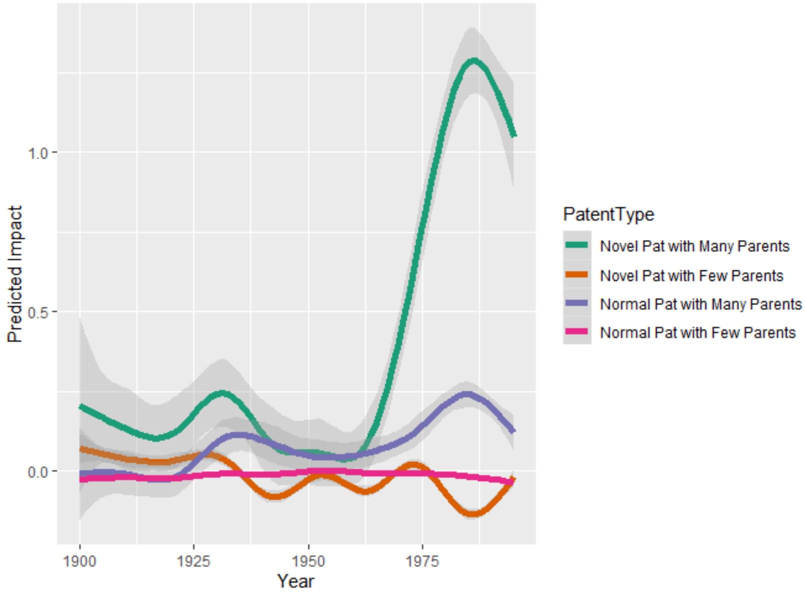

Residual Impact of Locally-Sourced and Non-Locally Sourced Patents,

Figure 4.14 plots the residuals from the following model:

$$Impact_p = Decade_p * NrSubClasses_p + PrimaryClassFE_p + E_p$$

In the model, C438 is a factor variable designating the primary class that a patent is assigned to. In the USPC classification scheme, there are 438 unique classes.

Figure 4.14: Residual Impact of Locally-Sourced Knowledge

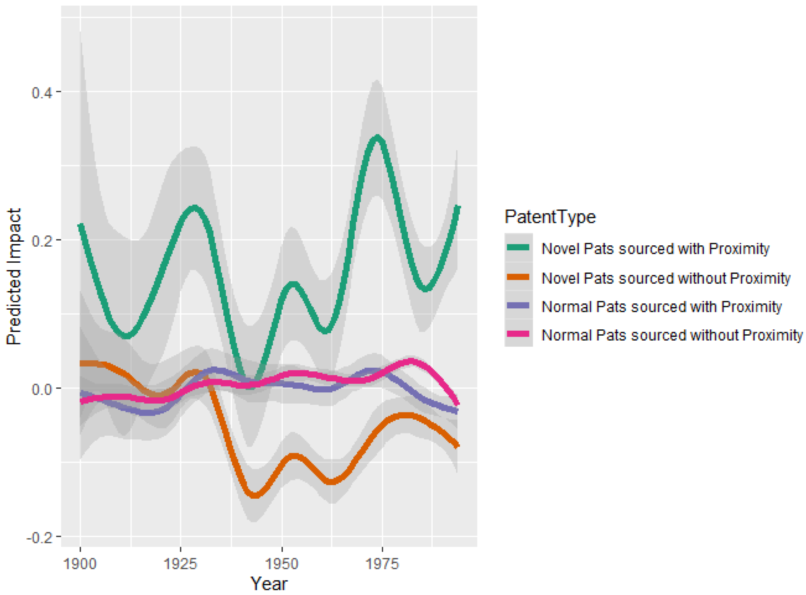

Disruptiveness of Regime of Technological Change

Figure 4.15 plots the residuals from the following model:

$$Impact_p = Decade_p * NrSubClasses_p + SublcassFE_p + E_p$$

The subclass fixed effect is a factor variable designating the primary subclass that a patent is assigned to. In the USPC classification scheme, there are about 16,000 unique classes.

Figure 4.15: Residual Impact of Novel and Normal Patents

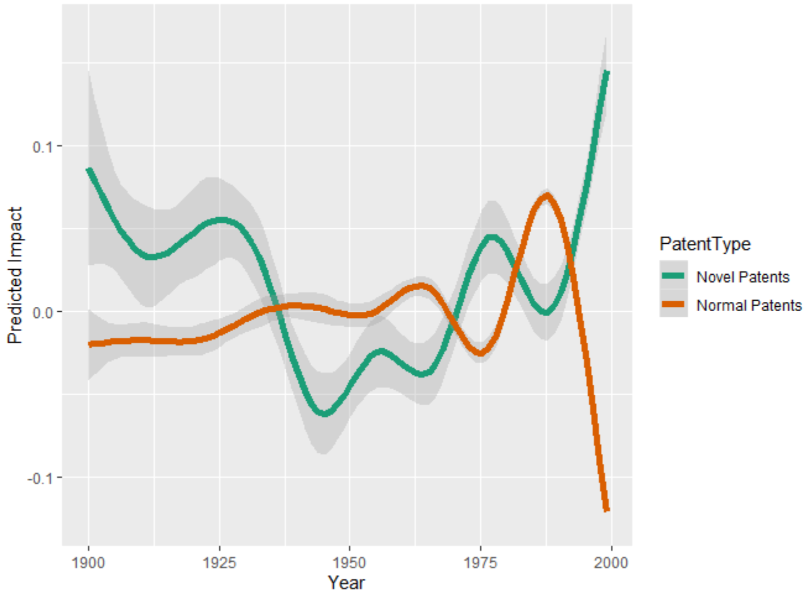

4.11) Bibliography

Acs, Z., and Audretsch, D. (1988). Innovation in Large and Small Firms: An Empirical Analysis. *American Economic Review* 78(4). 678-690.

Agrawal, A., Galasso, A., and Oettl, A. (2017). Roads and Innovation. *Review of Economics and Statistics* 99(3). 417-434.

Arthur, B. (2009). *The Nature of Technology: What It Is and How It Evolves*. Free Press, New York.

Arzaghi, M., and Henderson, V. (2008). Networking off Madison Avenue. *The Review of Economic Studies* 74(4). 1011-1038.

Autor, D., Dorn, D., Katz, L., Patterson, C., and Van Reenen, J. Concentrating on the Fall of the Labor Share. *American Economic Review: Papers and Proceedings* 107(5). 180-185.

Balland, P.A., and Rigby, D. (2018). The Geography of Complex Knowledge. *Economic Geography* 93(1).

Balland, P.A., Jara-Figueroa, C. Petralia, S., Steijn, M., Rigby, D., and Hidalgo, C. (2020). Complex Economic Activities Concentrate in Large Cities. *Nature Human Behavior* (4).

Bathelt, H., Malmberg, A., and Maskell, P. (2004). Clusters and Knowledge: Local Buzz, Global Pipelines, and the Process of Knowledge Creation. *Progress in Human Geography* 28.

Baum-Snow, N., Gendrom-Carrier, N., and Pavan, R. (2020). Local Productivity Spillovers. Working Paper.

Berkes, E., and Gaetani, R. (2020). The Geography of Unconvential Innovation. *The Economic Journal.* Forthcoming.

Bettencourt, L., Lobo, J., Helbing, D., Kühnert, C., and West, G. (2007). Growth, Innovation, Scaling, and the Pace of Life in Cities. *Proceedings of the National Association of Scientists* 104(17). 7301-7306.

Castaldi, C., Frenken, K., and Los, B. (2015). Related Variety, Unrelated Variety, and Technological Breakthroughs: An Analysis of U.S. State-Level Patenting. *Regional Studies* 49(5). 767-781.

Catalini, C., Fons-Rosen, C., and Gaule, P. (2018). How do Travel Costs Shape Collaboration? *NBER Working Paper Series #24780*.

Clancy, M. (2020). The Case for Remote Work. *Iowa State University Economics Working Papers #20007.*

Cremers, K., Harhoff, D., Narin, F., Scherer, F., and Vopel, K. (1999). Citation Frequency and the Value of Patented Inventions. *Review of Economics and Statistics* 81(3). 511-515.

Desmet, K., and Rossi-Hansberg, E. (2009). Spatial Growth and Industry Age. *Journal of Economic Theory* 144(6).

Dingel, J., and Neiman, B. (2020). How Many Jobs Can Be Done at Home? *Becker Friedman Institute Working Paper Series*.

Dong, X., Sheng, S., and Kahn, M. (2018). The Role of Transportation Speed in Facilitating High Skilled Teamwork. *NBER Working Paper Series #24539.*

Dosi, G. (1982). Technological Paradigms and Technological Trajectories. *Research Policy* 11(3). 147-162.

Duranton, G., and Puga, D. (2001). Nursery Cities: Urban Diversity, Process Innovation, and the Life Cycle of Products. *American Economic Review* 91(5). 1454-1477.

Esposito, C. (2020). The Emergence of Knowledge Production in New Places. *PEEG Working Paper #2046*.

Esposito, C., and Rigby, D. (2019). Buzz and Pipelines: The Costs and Benefits of Local and Nonlocal Interaction. *Journal of Economic Geography* 19(3). 753-773.

Feldman, M., and Audetsch, D. (1999). Innovation in Cities: Science-Based Diversity, Specialization, and Localized Competition. *European Economic Review* 43(2). 409-429.

Feldman, M., Guy, F., and Iammarino, S. (2020). Regional Income Disparities, Monopoly and Finance. *Cambridge Journal of Regions, Economy and Society* 14(1). 25-49.

Fleming, L. (2001). Recombinant Uncertainty in Technological Search. *Management Science* 47(1). 117-132.

Fleming, L., and Sorenson, O. (2001). Technology as a Complex Adaptive System: Evidence from Patent Data. *Research Policy* 30(7).

Florida, R. (2002). *The Rise of the Creative Class*. Basic Books, New York.

Gertler, M. (2003). Tacit Knowledge and the Economic Geography of Context, or The Undefinable Tacitness of Being (There). *Journal of Economic Geography* 3(1). 75-99.

Glaeser, E. (2008). *Cities, Agglomeration, and Spatial Equilibrium.* Oxford University Press, New York.

Glaeser, E. (2011). *Triumph of the City: How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier, and Happier*. Penguin Books, New York.

Glaeser, E., and Tobio, K. (2008). The Rise of the Sunbelt. *Southern Economic Journal* 74(3). 609-643.

Gordon, R. (2016). *The Rise and Fall of American Growth*. Princeton University Press.

Grashof, N., Hesse, L., and Fornahl, D. (2019). Radical or Not? The Role of Clusters in the Emergence of Radical Innovations. *European Planning Studies* 27(10). 1904-1923.

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

Hall, B., Jaffe, A., and Trajtenberg, M. (2001). The NBER Patent Citations Data File: Lessons, Insights, and Methodological Tools. *NBER Working Paper #8498.*

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.

Kemeny, T., and Storper, M. (2020). Superstar Cities and Left-Behind Places: Disruptive Innovation, Labor Demand, and Interregional Inequality. *LSE International Inequalities Institute Working Paper #41*.

Kim, D., Cerigo, D., Jeong, H., and Youn, H. (2016). Technological Novelty Profile and Invention's Future Impact. *EPJ Data Science* 5(8).

Kuhn, T. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.

Lamoreaux, N., and Sokoloff, K. (1996). Long-Term Change in the Organization of Inventive Activity. *Proceedings of the National Association of Scientists* 93(23). 12686-12692.

Manduca, R. (2019). Antitrust Enforcement as a Federal Policy to Reduce Regional Economic Disparities. *Annals of the American Academy of Political and Social Science* 685(1). 156-171.

Mewes, L. (2019). Scaling of Atypical Knowledge Combinations in American Metropolitan Areas from 1836 to 2010. *Economic Geography* 95(4).

Mokyr, J. (1992). *The Lever of Riches: Technological Creativity and Economic Progress.* Oxford University Press, New York.

O'Mara, M. (2019). *The Code: Silicon Valley and the Remaking of America*. Penguin Press, New York.

Powell, W., Koput, K., and Smith-Doerr, L. (1996). Interorganizational Collaboration and the Locus of Innovation: Networks of Learning in Biotechnology. *Administrative Science Quarterly* 41(1). 116-145.

Saxenian, A. (1996). *Regional Advantage*. Harvard University Press.

Schumpeter, J. (1934). *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle*. Harvard University Press.

Schumpeter, J. (1942). *Capitalism, Socialism, and Democracy*. Harper & Brothers, New York.

Storper, M., and Leamer, E. (2001). The Economic Geography of the Internet Age. *Journal of International Business Studies* 32(4). 641-665.

Storper, M., and Venables, A. (2004). Buzz: Face-To-Face Contact and the Urban Economy. *Journal of Economic Geography* 4(1). 351-370.

Sorenson, O., Rivkin, J., and Fleming, L. (2006). Complexity, Networks, and Knowledge Flow. *Research Policy* 35(7). 994-1017.

Uzzi, B., Mukherjee, S., Stringer, M., and Jones, B. (2013). Atypical Combinations and Scientific Impact. *Science* 342(6157). 468-472.

Van der Wouden, F. (2020). A History of Collaboration in U.S. Invention: Changing Patterns of Co-Invention, Complexity, and Geography. *Industrial and Corporate Change* 29(3).

Wood, S., Pya, N., and Saefken, B. (2017). Smoothing Parameter and Model Selection for General Smooth Models. *Journal of the American Statistical Association* 111(516). 1548-1563.

Wuchty, S., Jones, B., and Uzzi, B. (2007). The Increasing Dominance of Teams in Production of Knowledge. *Science* 316.

Youn, H., Strumsky, D., Bettencourt, L., and Lobo, J. (2015). Invention as a Combinatorial Process: Evidence from U.S. Patents. *Journal of the Royal Society Interface* 12. 1-8.

Chapter 5: Learning, Fast and Slow: The Returns to Maturity and Team Size for High-Impact Innovation in the United States between 1836 and 1975

5.1) Abstract

This paper analyses changes in the returns to maturity and collaboration for U.S. patent inventors between 1836 and 1975. I combine two novel datasets which allow me to trace the careers of inventors, record instances of collaboration, and measure the impact of patents on subsequent invention across a long period of U.S. history. The study advances two stylized facts. First, I show that the returns to maturity, measured as the marginal effect that an additional year of patenting experience has on inventors for creating high-impact inventions, have been negative for the inventors of U.S. utility patents since the 1920s. Second, I show that the returns to collaboration, measured as marginal effect that collaborating with a larger number of co-inventors has on an inventor's probability of creating a high-impact patent, were negligible until the 1920s and became positive thereafter. I develop a simple model to interpret these findings. The model proposes that teams "learn" quickly by rapidly pooling together the knowledge of their co-inventors while individual inventors learn slowly through experiential search and discovery. When the knowledge frontier is rapidly expanding, individual inventors are not able to learn quickly enough to keep up with the new ideas that are being introduced, so maturity is associated with lower impact. Collaboration, which allows teams to gather knowledge more quickly, is better-suited to innovation in fast-growth knowledge environments. I find empirical support for the model through an analysis that relies on variation in the growth rate of knowledge across technological fields. In this respect, I find that (a) the average maturity level of inventors is lower in fast-growth technological fields, (b) the returns to maturity are most negative in

rapidly expanding technological fields, (c) the average size of inventor teams is larger in fast-growth technological fields, and (d) the returns to team size are greatest in fast-growth technological fields.

5.2) Introduction

The ability of inventors to create high-impact inventions is largely determined by their accumulated technological knowledge. Inventors expand the breadth and effectiveness of their knowledge through two ways. First, inventors learn over time by search and discovery, and so their accumulated knowledge assets expand as they mature (Arrow, 1956; Fleming, 2001). Second, inventors collaborate with other inventors and scientists. Collaboration allows inventors to utilize the ideas developed by their collaborating partners without needing to fully develop those ideas themselves (Katz and Martin, 1997).

Empirical research has shown that collaboration is positively associated with the creation of impactful patents and scientific publications (Jones et al., 2008; van der Wouden, 2020; Wuchty et al., 2007). However, evidence of the relationship between the maturity of inventors the creation of high-impact inventions is mixed. One way researchers measure the returns to maturity is by calculating the age at which scientists' and inventors' creative output peaks. Jones (2010) shows that the average age at which Noble Laureates and historian-identified great inventors created their top achievements increased by 6 years during the 20th century. However, Sarada et al.'s (2017) analysis of the careers of patent inventors finds no relationship between inventor age and the impact of their patents, measured by the number of forward citations their patents receive. An alternative measure of maturity is the number of years that have elapsed since an inventor or scientist produces her or his first patent or article. Using the second measure, Sinatra et al. (2016), Liu et al. (2018), and Liu et al. (2021) find no relationship between scientists' accumulated years of experience and the forward citations received by their articles. Akcigit et al. (2018) apply the career progress measure of maturity to the inventors of U.S. patents 1947-2010 and find a *negative* relationship between their level of maturity and their patents' forward citations. Observing that both collaboration and the accumulation of experiential knowledge over time are both means by which inventors expand their knowledge-based capabilities, the above literature prompts the question: why do inventors receive positive returns to collaboration but not to maturity?

In this paper, I analyze a novel dataset of the inventors of over 8 million U.S. utility patents granted between 1836 and 1975 and show that (a) positive returns to collaboration, defined as the number of co-inventors of a patent, emerged in the 1920s and expanded thereafter, and (b) no positive returns to maturity, defined as the number of years elapsed since an inventor's first patent was granted, emerged during the study period. I generate these findings by combining two datasets. The first dataset, introduced by van der Wouden (2020), identifies each patent produced by U.S.-based patent inventors between 1836 and 1975 and allows me to examine variation in the impact of patents across inventors' careers and to adjust for inventor-level variation in ability. The second dataset records the impact of individual patents granted between 1836 and 2010 on subsequent patenting by applying the method outlined in Chapter 2 of this dissertation. This second dataset provides a new source for information about the quality of historical patents. By combining the two datasets, I am able to compute inventor-adjusted returns to inventor maturity and inventor collaboration over nearly 150 years of U.S. history.

One possible reason scientists and inventors receive positive returns to collaboration but not to maturity is that collaboration allows inventors to extend the capabilities of their knowledge sets quickly, while individual inventors who accumulate knowledge through the accumulation of experience do so slowly. The slow pace of experiential knowledge accumulation causes the returns to maturity to be insignificant or negative because, as inventors slowly accumulate experience over time, other inventors introduce their own inventions and in so doing increase the quantity of knowledge that inventors must learn in order to remain competent. Generally, inventors are not able to accumulate experience as quickly as other inventors expand the knowledge frontier, so maturity and impact are inversely related. Although collaboration with other inventors entails costs associated with building trust and shared norms (Levin and Cross, 2004), it nonetheless allows inventors to collect knowledge more rapidly than they could through the accumulation of experience. To extend the capability of their knowledge assets, inventors may thus benefit more from collaboration than from experience when they seek to innovative in knowledge fields that are rapidly expanding.

In Section 5.5 of this chapter, I develop a simple model to flesh out this logic. In the model, inventors create new ideas by combining the ideas that they already know. The impact of their resulting ideas is determined by the level of complementarity between the ideas that they combine. Inventors begin their careers with an intrinsic stock of ideas accumulated during their educational training. To their intrinsic knowledge, inventors add ideas learned through experience accumulated over the course of time during their patenting careers.

The challenge for inventors is that as they accumulate ideas through experience, other inventors introduce new ideas. The introduction of new ideas by others reduces the percentage of the total number of ideas known by individual inventors. Because individual inventors' knowledge stocks decline in a relative sense over time, the likelihood that they have knowledge of the most complementary combinations of ideas decreases and the impact of their inventions declines. Collaborating inventors are able to use the knowledge of their collaborative partners without having to develop that knowledge anew, so collaboration is a relatively effective method for inventors to mobilize knowledge when many new ideas are being introduced to the world. To validate the model empirically, I exploit cross-sectional variation in the rate of knowledge growth across technology fields. Technology fields are the groups of ideas, linked by common materiality or domain, between which there are strong recombinatory complementarities (Arthur, 2007). I identify the technology fields of patents using aggregate USPC patent classification codes, and I identify fast and slow-growth technology fields based on the growth rate of patents in each field. My analysis demonstrates that the average number of years of prior patenting experience of inventors is lower in fast growth technology fields, that the average size of inventor teams is higher in fast growth knowledge fields, that the returns to maturity are lower in fast growth knowledge fields, and that the returns to collaboration are larger in fast growth fields.

In addition to these cross-sectional results, I also show that the acceleration in the rate of knowledge growth explains the increase in the returns to collaboration and the decrease in the returns to maturity across time. To this end, I demonstrate that (a) the difference in the returns to maturity in fast and slow growth fields emerged after 1900, and (b) the difference in the returns to team size in fast and slow growth fields emerged after 1920. Moreover, I argue that the rate of knowledge growth expanded after 1920, which reduced the returns to maturity and increased the returns to collaboration.

This study contributes to the current literatures on inventor collaboration, inventor maturity, and the knowledge-intensity of innovation. With regard to team formation, past studies have documented long-run increases in the size of teams of inventors and scientists (Lamoreux and Sokoloff, 1996; Wuchty et al., 2007; van der Wouden, 2020); however, the novel dataset on the impact of inventions used in this study allows me to explore historical changes in the returns to inventor collaboration by linking collaboration to impact data. In this regard, I show that positive returns to collaboration first emerged in the 1920s.

With regard to inventor maturity, three studies have found no relationship between the number of years elapsed since the publication of scientists' first article and the impact of their articles (Sinatra et al., 2016; Liu et al., 2018; Liu et al., 2021) and one study has demonstrated the average impact of patents is lower for inventors with more years of patenting experience (Akcigit et al., 2018). A limitation to the latter study is that it makes use of patent forward citations to assess the impact of inventions, which are not available for patents granted before

1947. The novel dataset I use to measure the impact of patents allows me to perform a related analysis that extends back to 1836. My analysis shows that the returns to inventor maturity were statistically insignificant before 1900 and were negative thereafter.

Collectively, the two core findings of this paper – that inventors started to receive positive returns to collaboration in the 1920s and negative returns to maturity in the first decade of the 20th century – make a third contribution to the literature on the reduction in R&D productivity (Gordon, 2014; Bloom et al., 2020). The literature on the R&D productivity decline argues that long-run increases in the complexity of new technologies has caused the process of invention to involve the coordination of more extensive and more costly organizations of inventors, firms, and research centers (Powell et al., 1996; Wuchty et al., 2007; Clancy, 2017; Balland et al., 2020; Esposito, 2021; van der Wouden, 2020). The analysis performed in the current study help to scrutinize these assertions. In particular, I show that knowledge production has become not just more knowledge-intensive, but also more reliant on recent and fast-advancing knowledge. This result implies that big and expensive-to-fund research and development teams do not form because ideas are "getting harder to find" (Bloom et al., 2020), but because there are many new ideas that are being rapidly introduced. Big teams are able to organize and collect these recent ideas more quickly than inventors are able to learn them. Moreover, instead of being a cause of the decrease in R&D productivity, the increase in team size appears to be an adaptation to the accelerating rate of innovation.

The remainder of the paper consists of a description of the construction and structure of the datasets in Section 5.3, an empirical analysis of changes in the returns to collaboration and experience in Section 5.4, the development of the formal model in Section 5.5, the empirical validation of the model in Section 5.6, and a reflection on directions for further research in Section 5.7.

5.3) Data Construction and Description

The empirical analysis in this paper uses two long-run datasets of 8 million utility patents that cover about 150 years of U.S. inventive activity. The first dataset originates from van der Wouden (2020) and links disambiguated inventors to patents. These data were constructed by searching the text of historical patents made available in Petralia et al. (2016) to identify potential first and family names of inventors. Once a potential name is found, three groups of algorithms are run. The first group searches each patent for the text found before and after the potential name to determine if the potential name is in fact the name of the inventor of the patent. When a patent has multiple co-inventors, the group of algorithms determines whether the potential names are in fact names of inventors. The second group of checks the accuracy of the identified names outputted from the first group of algorithms against an external comparison data set. This second group of algorithms record 30 features associated with the found name(s) which act as predictor variables for a subsequent supervised machine learning exercise. In this exercise, information from labeled (true known cases) inventor names on patents originating from Google Patents and EspaceNet is used to train gradient boosting models to predict whether the previously-identified names are truly inventor names. The third group of algorithms disambiguates the inventors by creating a unique ID code for each unique inventor. This step disambiguates all inventors by pair-wise comparing all inventors based on the similarity of 12 characteristics (e.g. first name, middle name, last name, location, year, technology). The disambiguated inventor data from Lai et al. (2012) is used to train a supervised machine learning model for this purpose. This model is used to predict which identified inventors are likely to be the same individuals, resulting in over 1.9 million disambiguated unique inventors of U.S. utility patents between 1836 and 1975.

The disambiguated inventor dataset is used to identify inventor career trajectories and to compute the number of collaborators on a patent. Inventor career trajectories are built by sorting the inventor-patent data by grant year and patent number, after which I measure experience as the elapsed time between the inventor's first and current patent, rounded to years. The number of collaborators is taken as the number of unique inventors listed on the patent. Figure 5.1 plots the share of patents that have two or more co-inventors by year. The share of patents that are collaborative rose steadily across the study period.

Figure 5.1: Share of Collaborative Patents by Year

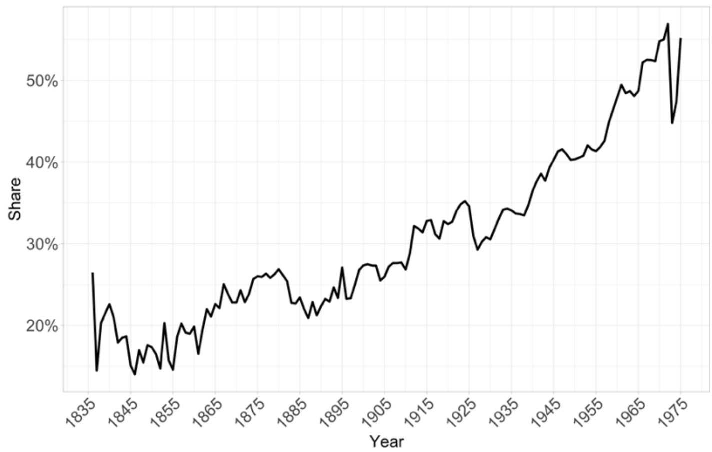

In Figure 5.2, I plot the average career length (in years) of inventors, broken out by cohort year. Career length is defined as the number of years that elapse between the granting of an inventor's first patent and the granting of an inventor's final patent, and cohort years are defined as the year that an inventor's first patent is granted. The figure shows that the average career length hovered between 5 and 12 years for inventors that started their careers between 1836 and 1885 and then gradually dropped to 3 years in 1905. This evolution of inventor career lengths broadly corresponds with earlier observations of shifts in inventor careers and the organization of inventive activity (Lamoreux and Sokoloff, 1996; Lamoreaux and Sokoloff, 2005; Lamoreaux et al., 2011; Lamoreaux et al. 2013). The dataset contains one spike in average career length in 1911, when the average career length jumps to 8 years and declines thereafter. The primary identification strategy used in this paper exploits within-year variation in maturity levels of inventors and is thus robust to random temporal shocks. Finally, because I cannot trace inventor careers beyond 1975, careers are right-truncated toward the end of the study period.

Figure 5.2: Mean Career Length by Inventor Cohort Year

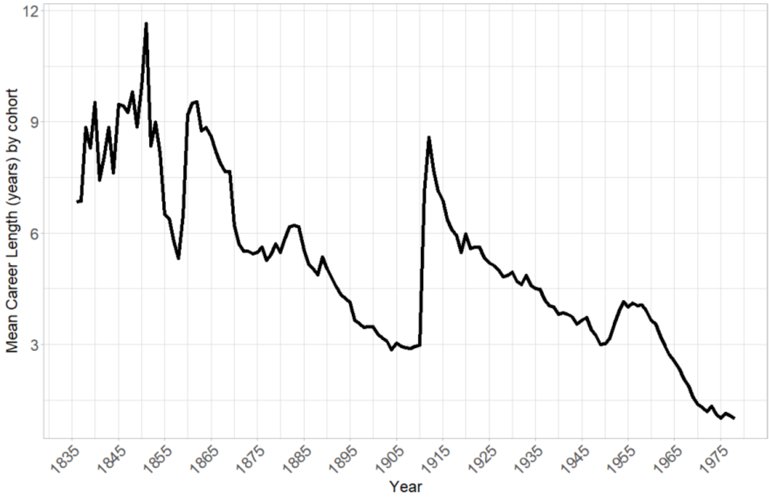

I measure inventor maturity as the number of years elapsed since the granting of an inventor's first patent. In the subsequent analyses, I interchangeably use a binary and a continuous variable to define inventor maturity. When I use the continuous definition, I take its natural logarithm. When I use a binary variable, I use a 5-year cutoff value; under this definition, all inventors with 0-4 years of patenting experience are defined as "premature" while all inventors with 5 or more years of experience are defined as "mature".

The second dataset provides information about each patent's impact on subsequent invention. This dataset is created by tracing the flow of technological knowledge between patents using the method described in Chapter 2 of this dissertation. Using the patent impact measures computed from that method, I define high-impact inventions as patents with an out-degree in the top 5% of their grant year cohort. In the appendix, I replicate the key analyses of this chapter using a top 10% threshold to define high-impact patents. In that analysis, I arrive at results that are identical to the results generated using the 5% threshold, with one notable exception. I discuss this exception in the appendix.

The final descriptive exercise, Figure 5.3, pulls together the data on team size, inventor maturity, and patent impact to show how these variables relate to one-another. Each column in the left panel plots the distribution of patent impact by inventor maturity for patents granted between 1836 and 1919 (left chart) and 1920-1975 (right chart). Each column in the right panel plots the distribution of patent impact by team size for patents granted between 1836 and 1919 (left chart) and 1920-1975 (right chart). I break the data out by these two time periods because, as I will later demonstrate, the relationships between team size, experience, and patent impact change substantively after 1920.

Figure 5.3: Distribution of Patent Impact by Inventor Maturity and Number of Inventors by 35-Year Periods

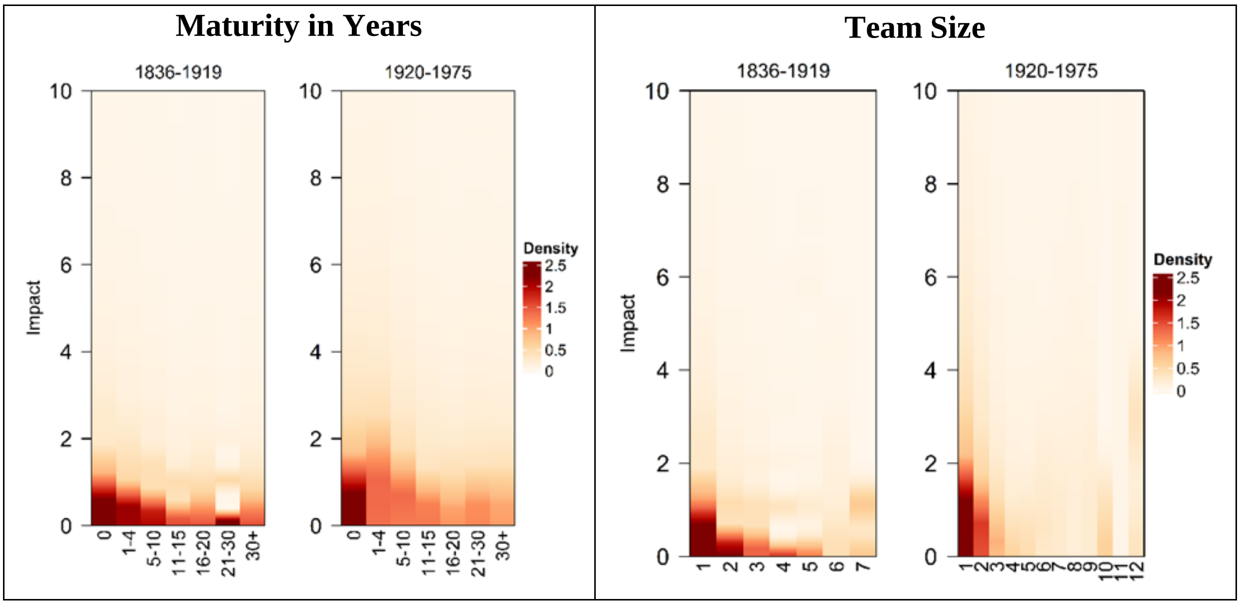

In the left panel of Figure 5.3, the columns associated with 0 years of experience indicate that most of the patents by inventors with 0 years of patenting experience have very low impact. This is signaled by the very high density for patents created by inventors with 0 years of experience at the patent impact level of 0. For columns corresponding to more years of experience (i.e. 5+ years of experience), the density of patent impact is less strongly concentrated at the low end of impact, suggesting that the impact is a little more evenly spread out across the spectrum. Comparing the 1836-1919 maturity chart with the 1920-1975 maturity chart shows that the shading remains largely the same across these two time periods. Therefore, the relationship between inventor maturity and patent impact does not change much over the timeframe of the dataset. However, comparing the 1836-1919 team size chart with the 1920-1975 team size chart (Figure 5.3, right panel) shows that the relationship between team size and patent impact changes strongly between the two time periods. In particular, the impact distribution is centered near 0 for large teams in 1836-1919, but is dispersed for large teams between 1920 and 1975. Moreover, between 1920 and 1975, the variation in the impact of patents was large for big teams

While the descriptive insights garnered from Figure 5.3 are informative, the potential selection of less-able inventors to discontinue patenting or of more-able inventors to collaborate more frequently motivates the analysis in the following section.

5.4) Returns to Collaboration and Experience for High-Impact Invention

I begin the analysis of the returns to collaboration and experience in Figure 5 by plotting the mean probability premature inventors (< 5 years of patenting experience) and mature inventors (5+ years) create a high-impact patent between 1836 and 1975. I additionally overlay 95% confidence intervals on the chart. The left panel of Figure 5.4 uses the complete dataset containing all inventors, while the right panel uses the subset of inventors that reach maturity, meaning their careers last 5 years or longer. In this figure, the "0-4 Years" category shows the probability that an inventor who ultimately reaches maturity invents a high-impact patent while still premature.

Figure 5.4: Probability of High-Impact Patent by Inventor Experience

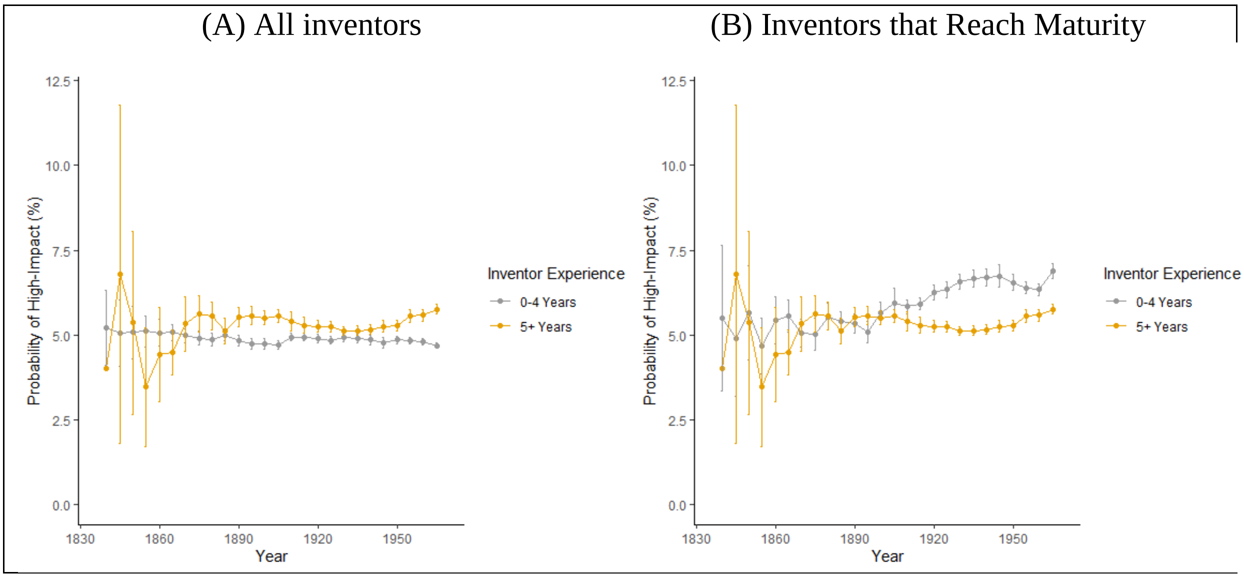

The left panel of Figure 5.4 shows that mature inventors are more likely to create high-impact patents starting in the 1880s. The probability that mature inventors will produce a high-quality patent is roughly 5.5%, while for premature inventors that probability is approximately 4.5%. This initial result, however, is driven by a selection effect in which inventors with lesser ability tend to have shorter careers. The exit of less able inventors early in their careers places a downward bias on the observed frequency of high-impact patents by premature inventors. The right panel of Figure 5.4 confirms that this selection effect is substantial. In the right panel, I restrict the dataset to inventors that patent over the course of 5 or more years. By removing the selection effect, the right panel shows that the probability of creating a high-impact invention was unrelated to inventor maturity until 1910, and that after 1910 premature inventors were more likely to create high-impact patents than mature ones. In the appendix, I show that result is unaffected when the maturity of inventors is computed using only patents created by lone inventors.

In Figure 5.5, I plot the mean probability that inventor teams and lone inventors create a high-impact patent, broken out into five-year periods. As before, I overlay 95% confidence intervals on the chart. The figure shows that the probability of creating a high-impact patent was 5% for both teams of inventors and lone inventors until around 1920.[14] After this year, the probability that a patent created by a team was high-impact rose to about 6% while the likelihood for lone inventors dropped to about 4%. Moreover, by the end of the study period in 1975, patents invented by teams were about 50% more likely to be high-impact were patents generated by lone inventors.

Figure 5.5: Probability of High-Impact Patent for Teams and Individuals

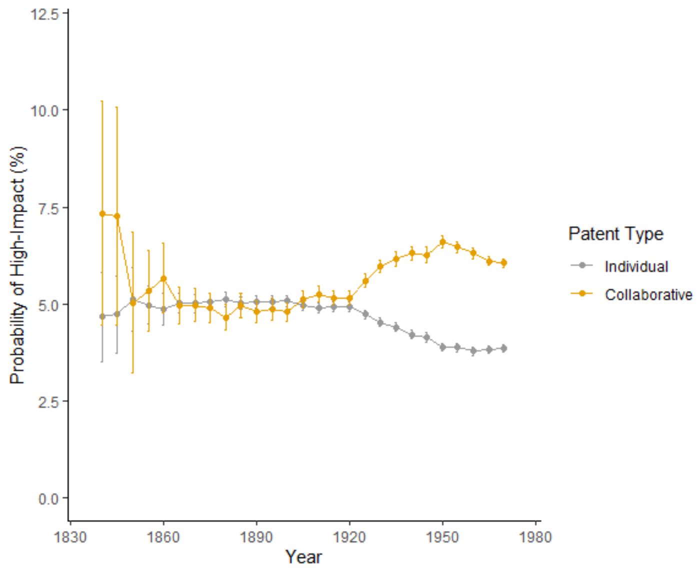

Figures 5.4 and 5.5 do not adjust for inventor-level heterogeneity in ability. To adjust for differences in inventor ability in computing the returns to inventor experience, I develop a linear probability regression model in which the probability that a patent p is high impact (HI_p) is a function of inventor i fixed effects, a factor variable recording the half-decade during which a patent is granted, and the binary variable $Mature_p$, which records a 1 if an inventor has 5 or more years of patenting experience prior to the granting of patent p and 0 otherwise. To allow the returns to collaboration to vary across time, I interact the $Mature_p$ binary variable with the half-decade factor variable. The regression model is given by Equation 5.1:

[14] The probability of creating a high-impact patent can exceed 5% across all patents because the 5% cutoff often falls on integer values.

$$(5.1)\ Prob(HI_p) = B_1 Mature_p * HalfDecade_p + B_2 Mature_p + B_3 HalfDecade_p + FE_i + E_p$$

The dataset is sparse before 1875 because relatively few patents were granted during that time period. Therefore, I aggregate all patents granted before 1875 into a single time unit (coded as "Pre-1875). I aggregate patents granted after 1875 to the half-decade. The pre-1875 patent set serves as the reference group for the regression. With these considerations in mind, the coefficients associated with the $Mature * HalfDecade_p$ term are plotted in the left panel of Figure 5.6. In the right panel, I produce a similar chart but replace the binary *Maturity* variable with a continuous measurement of inventor maturity by computing the log number of years of experience of the inventors of each patent.

Figure 5.6: Predicted Probability of Creating a High-Impact Patent by Inventor Maturity

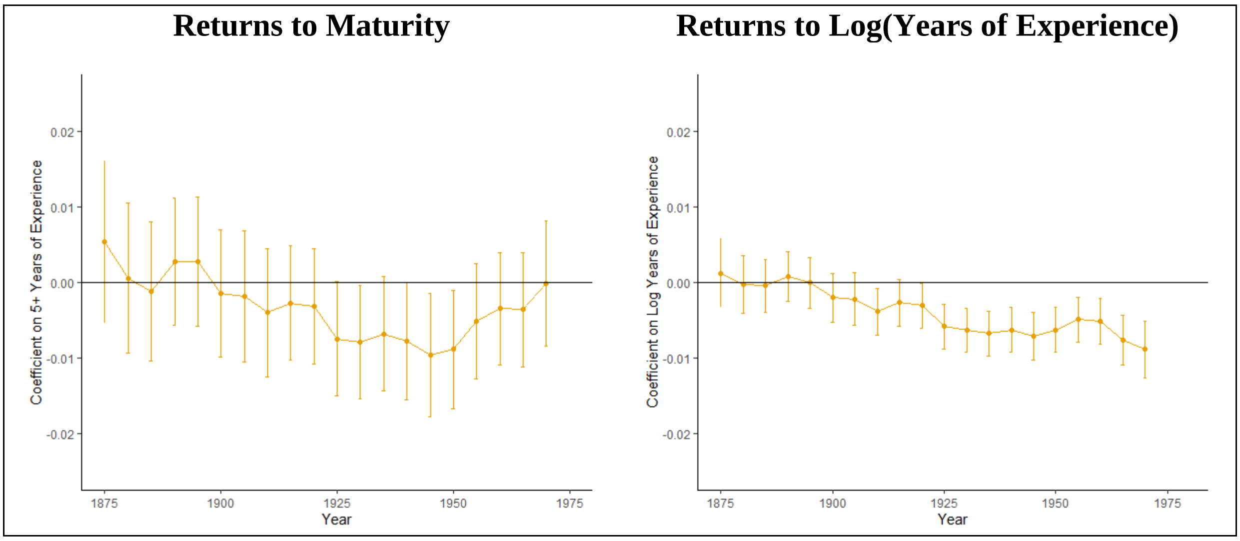

Figure 5.6 shows that the returns to maturity were indistinguishable from 0 between 1875 and the first few decades of the 20th century. The returns to maturity were negative between 1925 and 1955, after which they became statistically insignificant. The returns to log

years of experience became negative starting in 1905 and remained negative through the end of the study period. The differences in the way these two variables are measured can account for the slight differences in their returns across time. In the appendix, I find no evidence of positive returns to maturity when I restrict the dataset to solo-invented patents.

To adjust for individual-level ability when calculating the returns to collaboration, I develop a similar regression model containing inventor i fixed effects, a half-decade factor variable, and a binary variable $Collaboration_p$ which records a value of 1 if patent p is invented by a team of 2 or more co-inventors. I interact the $Collaboration_p$ binary variable with the half-decade factor variable to allow the returns to collaboration to vary across time. As before, I group all patents granted before 1875 into a single time period which serves as the reference group for the estimation. The regression model is given by Equation 5.2:

$$(5.2)\ Prob(HI_p) = B_1 Collaboration_p * HalfDecade_p + B_2 Collaboration_p + B_3 HalfDecade_p + FE_i + E_p$$

I plot the coefficients associated with the $Collaboration_p * HalfDecade_p$ term in Figure 5.7. The left panel uses the binary $Collaboration_p$ variable, while the right panel replaces this measure with a continuous measure of collaboration computed by taking the natural log of the number of co-inventors on a patent.

Figure 5.7: Predicted Probability of Creating a High-Impact Patent by Collaboration

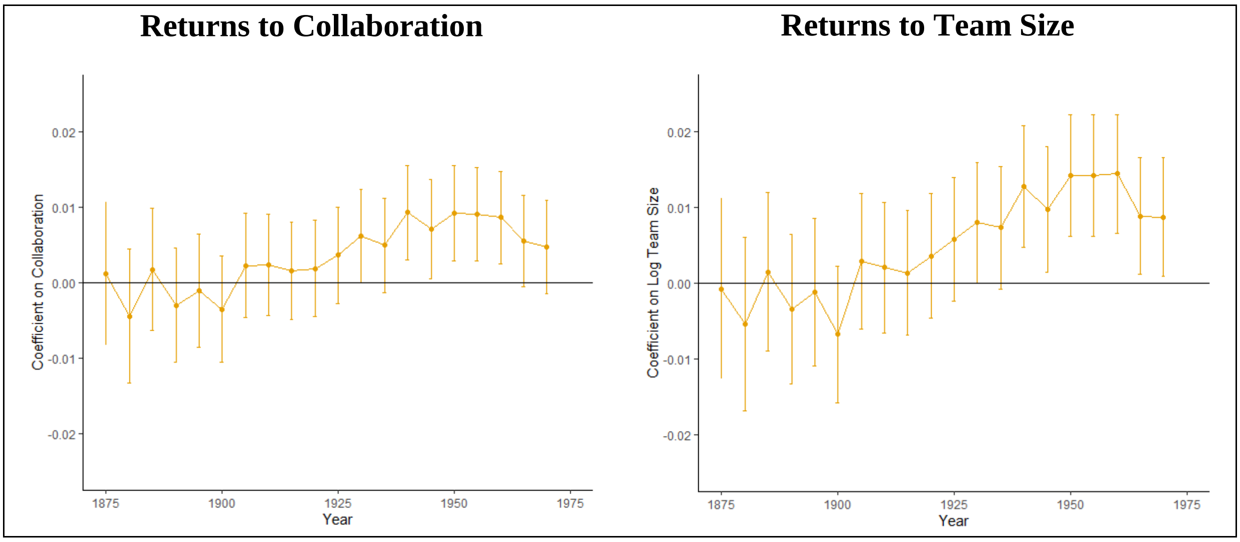

Figure 5.7 shows that the returns to collaboration and team size were statistically insignificant before 1930. After 1930, the returns to collaboration and team size became positive and significant. The returns to team size remained positive and significant through the end of the study period, while the returns to collaboration become insignificant using a 95% confidence interval in the final years of the study. In an un-reported set of results, I find that the returns to collaboration remained significant and positive through the end of the study period when using a 90% confidence interval.

To summarize the empirical investigation, I find that inventors received no returns to increases in team size and maturity before 1900. However, after 1900 inventors received negative returns to maturity and after 1920 inventors received positive returns to team size. These disparate returns are surprising because both increasing the size of a team and accumulating experience over time as inventors' careers progress are means through which inventors' knowledge-based capabilities expand. The expansion of those knowledge-based capabilities should translate into high-impact innovation. In the following section, I develop a model to assist in interpreting why only team size is positively related to creating high-impact patents.

5.5) Theory: The Returns to Experience in Fast- and Slow-Moving Knowledge Fields

Building on a model of recombinant knowledge creation (Weitzman, 1998), I assume that inventors create new ideas by combining existing ideas. I allow new ideas to vary in terms of their impact. The impact of a new idea *p* is a function of the complementarity of the existing ideas that are combined in the creation of the new idea (Fleming and Sorenson, 2001). Equation 5.3 defines the impact of patent p as a function of the complementarity between its components and a stochastic element, *Z*:

$$(5.3)\ Impact_p = f\left(Complementarity_p + Z_p\right)$$

To invent a new idea, an inventor *i* maximizes the impact of an invention subject to *i*'s knowledge constraint by searching amongst their accumulated ideas and by combining the set with the strongest complementarities. The knowledge constraint of inventor *i* defines the existing ideas that the inventor has knowledge of and is able to combine. The knowledge constraints of inventors relax when they accumulate knowledge. Inventors accumulate knowledge in two ways. First, inventors learn during their educational careers prior to becoming patenting inventors. This knowledge is intrinsic from the standpoint of the model. I denote *i*'s intrinsic knowledge by K_i, which is fixed for the duration of *i*'s patenting career. Second, inventors accumulate knowledge through interactions with texts, materials, and with other inventors during their patenting careers. The term $E_{i,t}$ denotes this time-varying experiential knowledge. The accumulated knowledge of inventor *i* is thus a function of *i*'s intrinsic and experiential knowledge:

$$(5.4)\ Knowledge_{i,t} = f\left(K_i + E_{i,t}\right)$$

Because inventors create new ideas by combining the maximally-complementary existing ideas within their knowledge constraint, the complementarity of a new idea created by inventor *i* is an increasing function of the number of ideas that *i* knows:

$$(5.5)\ Complementarity_p = f\left(K_i + E_{i,t}\right)$$

Equation 5.5 can be substituted into Equation 5.3 to relate the impact of an invention p to the knowledge accumulated by inventor i in the time leading up to time t plus the stochastic element Z_p:

$$(5.6)\ Impact_p = f\left(K_i + E_{i,t} + Z_p\right)$$

In Equation 5.6, the absolute impact of patent *p* is a function of the absolute quantity of the knowledge held by the inventor. The *relative impact* of *p*, measured as the impact of *p* relative to all other ideas that could possibly be created during *t*, is determined by the extent of inventor i's accumulated knowledge stock $(K_i + E_{i,t})$ relative to the knowledge stocks of all other inventors. Under the assumption that inventors can collaborate with other inventors without cost,[15] the relative impact of p is determined by the extent of *i*'s accumulated knowledge stock relative to the number of ideas known by all other inventors, denoted by U_t. Therefore, the relative impact of *p* is given by the percentage of all ideas known by inventors in time *t* by inventor *i* plus the stochastic element Z_p:

$$(5.7)\ RelativeImpact_p = f\left(\frac{K_i+E_{i,t}}{U_t} + Z_p\right)$$

In Equation 5.7, U_t is the sum of all intrinsic knowledge and experiential knowledge across all inventors. Because new inventors enter the patenting universe over time (i.e. new inventors start their careers), the quantities of universe-wide experiential knowledge *and* intrinsic knowledge are time-varying:

$$(5.8)\ U_t = \sum_i^n\left(K_{i,t} + E_{i,t}\right)$$

Substituting Equation 5.8 into Equation 5.7 and taking the differences of each variable with respect to time relates the changes in the impact of ideas created by individual inventors to their flow of experiences, the flow of new inventors into the universe, and the flow of experiences by incumbent inventors:

[15] We will relax the assumption of costless collaboration in the following section. For now, it is helpful to note that a non-zero cost of collaboration reduces the value of U_t by segmenting the knowledge stocks of individual inventors.

$$(5.9)\ \Delta RelativeImpact_{p,i} = f\left({}^{\Delta E_i}\!\!\Big/_{(\Delta K + \Delta E)} \right)$$

Equation 5.9 states that the returns to inventor maturity, defined as the change in the relative impact of an inventor's inventions over time ($\Delta RelativeImpact_{p,i}$), is determined by three factors: the change in an inventor's experience over time (ΔE_i), the change in the number of inventors in the universe (ΔK), and the change in the number of experiences in the universe (ΔE). For simplicity, in the subsequent discussion I will compact the denominator to a single term, ΔU, denoting the change in the number of new ideas in the universe without distinguishing whether those ideas are added by intrinsic or experiential knowledge.

The discussion so far has assumed that the complementarity between ideas is generated by a random process. I now assume that ideas are organized in fields of knowledge, indexed by c. Ideas from the same field have a fixed probability ρ of being complements, while ideas from different fields have zero probability of being complements. This assumption can be relaxed such that ideas from different fields would have some probability between 0 and ρ of being complements; however, the simple binary classification is sufficient to define the model and derive its propositions. The assumption that ideas are grouped into knowledge fields also generates the consequence that inventors who innovate in large fields with many recombinant possibilities are less likely to have knowledge of the most-complementary combinations in their field. This latter statement is expressed in Equation 5.9, where the impact of idea p relative to all other ideas that can theoretically be generated in field c is a function of the number of ideas known by i in field c divided by the number of ideas in field c known by all inventors:

$$(5.10)\ RelativeImpact_{p,c} = f\left(\frac{K_{i,c} + E_{i,c,t}}{U_{c,t}} + Z_p \right)$$

Taking differences of each variable in Equation 5.10 over time expresses the change in the relative impact of the ideas produced by i in field c as a function of the change in the knowledge accumulated by i in field c and the change in the number of ideas in field c that are added by other inventors:

$$(5.11)\ \Delta RelativeImpact_{p,i,c} = f\left({\Delta E_{i,c}}/{\Delta U_c} \right)$$

The denominator of Equation 5.10 indicates that the returns to inventor maturity, $\Delta RelativeImpact_{p,i,c}$, vary across knowledge fields. In particular, the returns to maturity are a decreasing function of ΔU_c, the rate at which all inventors introduce new ideas to the field. When $\Delta U_c > \Delta E_{i,c}$, new ideas are added to fields faster than individual inventors learn those ideas. In these fast-growth knowledge fields, the relative impact of the new ideas declines as their inventors mature.

Finally, I assume that inventors are awarded patents for ideas if the relative impact of their ideas is above the threshold value τ. This is the threshold above which inventions are deemed to be "useful", which is one of the criteria used by the USPTO when determining whether a patent application should be granted.[16] All ideas that fall below this threshold are not granted patents. In addition, an idea will become a high-impact patent if its relative impact is above the threshold value γ, such that $\gamma > \tau$.

Combining these assumptions with the specification of Equation 10 and defining fast-growth knowledge fields as the fields in which ΔU_c is large generates the first two propositions:

- **Proposition 1:** The probability that inventors create patents (ideas with relative impact exceeding τ) declines more rapidly over their careers in fast-growth knowledge fields.

[16] See: https://www.uspto.gov/patents/basics

- **Proposition 2:** The probability that inventors create high-impact patents (ideas with relative impact exceeding γ) declines more rapidly over their careers in fast-growth knowledge fields.

5.6) Theory: The Returns to Collaboration in Fast and Slow-Moving Fields

Collaboration allows inventors to pool their intrinsic and experiential knowledge. Collaborating with a new partner also entails a cost associated with building trust and shared norms. This cost is measured in units of time and is denoted by *r*. Inventors must choose between investing their time toward accumulating experiential knowledge or toward establishing collaborative relationships. The collaboration between inventor *i* and inventor *j* results in the shared knowledge constraint:

$$(5.12)\ SharedKnowledge_{i+j} = f\left(K_i + \left(E_{i,t} - r\right) + K_j + \left(E_{j,t} - r\right)\right)$$

A collaboration between *i* and *j* produces an idea *p* with a relative impact determined by their shared knowledge constraint, the number of ideas known by all inventors in the universe U_t, and the stochastic element Z_p:

$$(5.13)\ RelativeImpact_{p,i+j} = f\left(\frac{K_i+E_{i,t}+K_j+E_{j,t}-2r}{U_t} + Z_p\right)$$

Inventors collaborate with partners if the returns to collaboration, measured in terms of the increase in the relative impact of the resulting invention, exceed the costs of collaboration. For inventor *i*, this condition holds when the shared knowledge constraint, including the cost incurred by establishing the collaborative relationship, is less restrictive than *i*'s individual knowledge constraint:

$$(5.13)\ \frac{K_i+K_j+2E_{i\vee j,t}-r}{U_t} > \frac{K_i+E_{i,t}}{U_t}$$

This inequality is more likely to be satisfied when the rate at which all inventors introduce new ideas to the universe (ΔU_c) is large. The reason for this is that a large growth rate of new ideas implies that inventors accumulate knowledge through experiential learning more slowly than the introduction of new ideas to the universe erodes their awareness of the highest-impact combinations of ideas. When ΔU_c is very large, the growth rate of i's individual knowledge constraint is determined solely by K_i and i and j's shared knowledge constraint is determined by $K_i + K_j - r$. So long as j's intrinsic knowledge, K_j, exceeds the cost of forming a collaboration (r), i is incentivized to collaborate with j when ΔU_c is large. For reasons discussed above, ΔU_c also varies across technological fields. Therefore, I arrive at the model's third proposition:

- **Proposition 3:** Collaborative teams are more likely to create a patent (an idea with relative impact exceeding τ) in fast-growth knowledge fields.

For similar reasons to those described above, I arrive at the model's fourth proposition:

- **Proposition 4:** Collaborative teams are more likely to create a high-impact patent (an idea with relative impact exceeding γ) in fast-growth knowledge fields.

5.7) Learning and Collaborating in Slow and Fast-Moving Knowledge Fields

The model generated the following predictions:

1. Inventors in fast growth knowledge fields are less mature than inventors in slow growth fields
2. The returns to maturity are lower in fast growth fields than in slow growth fields
3. Inventors in fast growth knowledge fields collaborate more frequently than in slow growth fields
4. The returns to collaboration are larger in fast growth fields than in slow growth fields

To empirically identify fast growth and slow growth knowledge fields, I compute the patent growth rate in each aggregate USPC technology class and in each 5-year time period. There are 438 unique technology classes at the level of the USPC schema that I utilize. The growth rate of patents in each class *c* is computed as the number of patents assigned to class *c* in the 5-year period *t* divided by the number of patents assigned to class *c* in all time periods up to *t-1*:

$$(5.14)\ GR_{c,t} = \frac{P_{c,t}}{\sum_{t=0}^{t-1} P_{c,i}}$$

Next, I transform $GR_{c,t}$ into a binary variable by defining class-time period pairs as "fast growth classes" if their growth rate is above the average across all classes in the same 5-year period. All other class-time period pairs are defined as "slow growth classes". Therefore, in each 5-year period half of all classes are fast growth and half are slow growth.

I begin the empirical test of the model by analyzing the model's proposition, that inventors in fast growth classes have fewer years of patenting experience than inventors in slow growth classes. Figure 5.8 plots the mean experience of repeat inventors of patents in fast and slow growth technological fields by 5-year periods. In the figure, I restrict the sample to repeat inventors (inventors granted at least two patents) because inventors that patent only 1 patent during their careers do not accumulate meaningful experience. As before, I overlay 95% confidence intervals in the charts.

Figure 5.8: Average Years of Experience of Repeat Inventors by Class Growth Rate

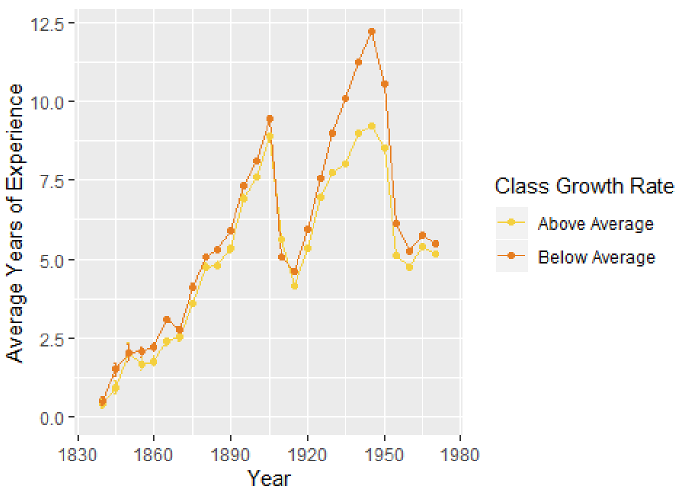

95% confidence intervals are overlaid in the figure.

Figure 5.8 shows that the average number of years of experience level of repeat inventors in fast growth technological fields was lower than that of the inventors of patents in slow growth fields for the full study period. The 95% confidence intervals, which are tightly concentrated around the mean values, indicate that these differences are statistically significant across all time periods. The one exception is the period spanning from 1910 to 1914 during which there is a discontinuity in the inventor career panels, as discussed in Section 5.3. With this caveat aside, Figure 5.8 confirms the model's first proposition that mature inventors tend to patent in slow growth fields.

The second proposition from the model is that the returns to maturity are more strongly negative in fast growth fields. To test this proposition, in Figure 5.9 I plot the probability that four types of patents – patents in fast growth fields invented by premature inventors, patents in fast growth fields invented by mature inventors, patents in slow-growing classes invented by premature inventors, and patents in slow-growing classes invented by mature inventors – are high-impact. The left panel of Figure 5.9 uses the entire dataset, while the right panel accounts for the selection of low-ability inventors to discontinue patenting by restricting the dataset to inventors who reach maturity. In the right panel, premature inventors are those who started their patenting careers recently but go on to reach maturity.

Figure 5.9: Probability of High-Impact Patent by Maturity and Class Growth Rate

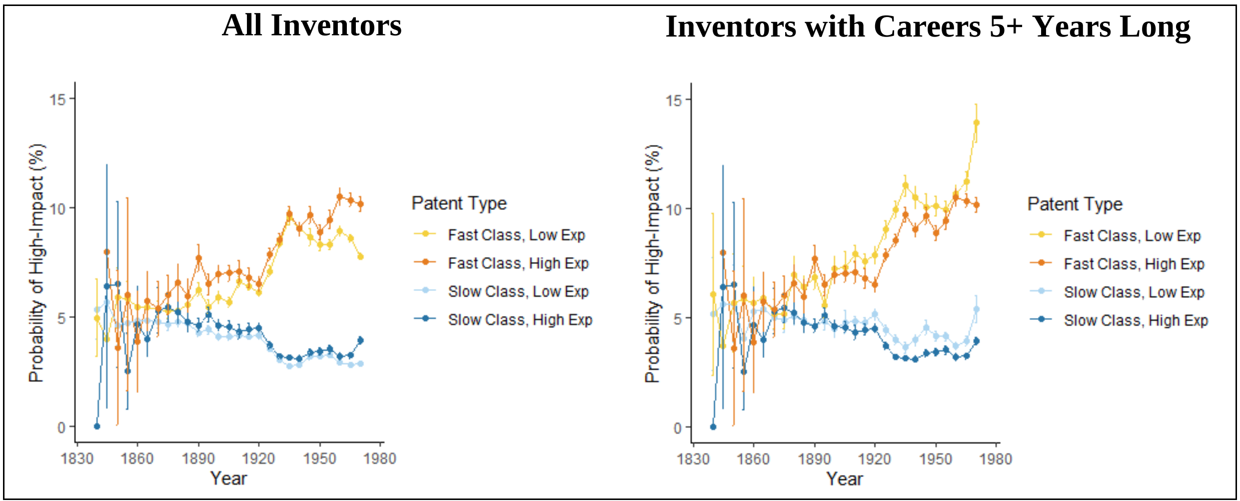

95% Confidence intervals are overlaid in the chart.

Figures 5.9 shows that the probability that a patent will be high-impact was about the same for patents produced in fast and slow growth classes, regardless of the level of maturity of their inventors, until 1880. After 1880, patents in fast growth were more likely to be highly impactful. The right panel of Figure 5.9, which accounts for inventor selection by omitting inventors that do not reach maturity, shows that the probability of inventing a high-impact patent in a fast-moving class was greater for premature inventors than for mature ones. The right panel also shows that the probability of creating a high-impact patent in a slow growth class was greater for premature inventors than for mature ones.

Although informative, Figure 5.9 does not account for inventor-level heterogeneity including the possible self-selection of inventors into particular technology classes. In addition, Figure 5.9 does not test whether the difference in the returns to mature are significantly different in fast and slow-growing technology fields. To account for these potential issues, I develop a regression model where the probability that a patent is high-impact is a function of inventor maturity measured as log years of experience, the binary variable *FastClass* which records a 1 if a class is fast and 0 if it is slow growth, the interaction of maturity and *FastClass*, and inventor and year fixed effects. I am primarily interested in B_3, the coefficient on the interaction of inventor maturity and *FastClass*. I hypothesize that B_3 will be negative and significant, indicating that the returns to experience are lower in fast growth classes. The model is given by Equation 5.15:

$$(5.15)\ Prob(HighImpact_p)$$
$$= B_1 Log(YearsOfExperience)_p + B_2 FastClass_p$$
$$+ B_3 Log(YearsOfExperience)_p * FastClass_p + FE_i + FE_t + E_p$$

Because the earlier results showed that the returns to maturity change over time (see Section 5.4), I run the model described by Equation 5.15 separately for patents granted before and after 1920. I hypothesize is that B_3 will be statistically insignificant for patents issued before 1920 but negative and significant for patents issued after 1920. The results are presented in Table 5.1. As before, I estimate the model as a linear probability model because the large number of fixed effects makes logit or probit specification infeasible.

Table 5.1: Regression Estimates of Effect of Years of Experience on Prob(High Impact)

	Patents Granted 1836-1919		Patents Granted 1920-1975	
$Log(YearsOfExperience)$	-0.00164** (0.000736)		-0.00384*** (0.000448)	
$FastClass$	0.00920*** (0.00129)	0.0102*** (0.00143)	0.0214*** (0.000937)	0.0226*** (0.000994)
$Log(YearsOfExperience) * FastClass$	0.000135 (0.000692)		-0.000931** (0.000464)	
$Mature$		0.000153 (0.00133)		-0.00323*** (0.000753)
$Mature * FastClass$		-0.000531 (0.00190)		-0.00289** (0.00116)
Inventor and year fixed effects	Yes	Yes	Yes	Yes
R^2	0.315	0.276	0.281	0.227
Inventor Subset	All Inventors	Inventors who reach maturity	All Inventors	Inventors who reach maturity
NOBS (unique inventors)	211,298	135,813	447,847	254,434

Standard errors clustered at inventor level in parenthesis.

Table 5.1 generates several insights. First, for patents granted before 1920 (left-hand column), maturity when measured using log years of experience was associated with a lower probability of high impact. In addition, *FastClass* was positively associated with high-impact patenting. The interaction of maturity and *FastClass* was insignificant, indicating that the returns to maturity were statistically the same in fast and slow growth classes before 1920. The right-hand column of Table 5.1 shows that, between 1920 and 1974, maturity was also negatively associated with high impact and that *FastClass* was also positively associated with high-impact during the later years of the study. However, after 1920 the interaction of maturity, specified as both a binary and a continuous logged variable, and *FastClass* was significantly negative. Therefore, I conclude that the model's second proposition is confirmed by the empirical analysis: the returns to maturity were significantly lower in fast growth classes than in slow-growth classes after 1920.

The third proposition of the model is that inventors in fast growth classes are more likely to collaborate than inventors in slow growth classes. I test this proposition in Figure 5.10 by plotting the average team size of patents in the two types of classes. To separate out possible changes in the propensity to collaborate as a binary outcome from the propensity to form larger collaborative teams, I also produce a version of the figure in which I omit all solo-inventor patents in the right panel of the figure.

Figure 5.10: Average Team Size by Class Growth Rate with 95% Confidence Intervals

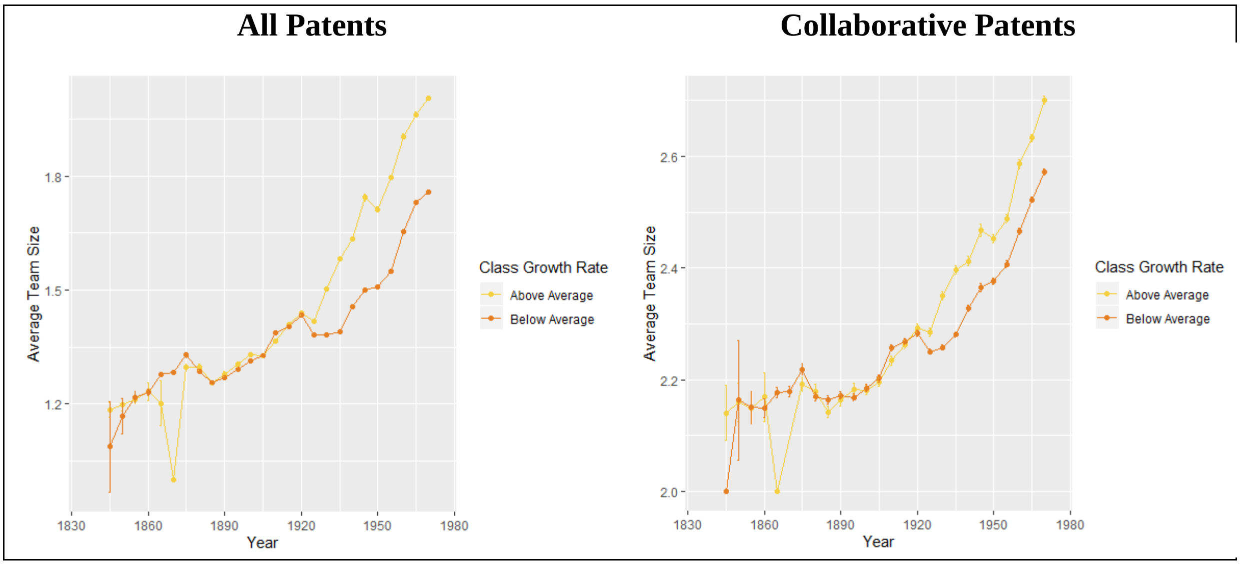

95% confidence intervals overlaid in the chart. The confidence intervals are tightly clustered around the means.

Figure 5.10 shows that the average team size was the same for patents in fast and slow growth classes until 1920. After 1920, the average team size of patents in fast-growth classes increased significantly above that of patents in slow-growth patents. The 95% confidence intervals are very small, indicating that the difference in average team size after 1920 is statistically significant. These results are found in both panels in Figure 5.10, including the right-hand panel which omits all single-inventor patents. Thus, Figure 5.10 confirms the model's third proposition.

The final proposition from the model is that the returns to collaboration are greater in fast growth classes than in slow growth classes. I start to explore this proposition in Figure 5.11 plotting by computing the probability that patents of four types -- patents in fast growth classes invented by individuals, patents in fast growing classes invented by teams, patents in slow growth classes invented by individuals, and patents in slow growth classes invented by teams – are high-impact. I overlay 95% confidence intervals in the chart.

Figure 5.11: Probability of High-Impact Patent by Collaboration and Class Growth Rate

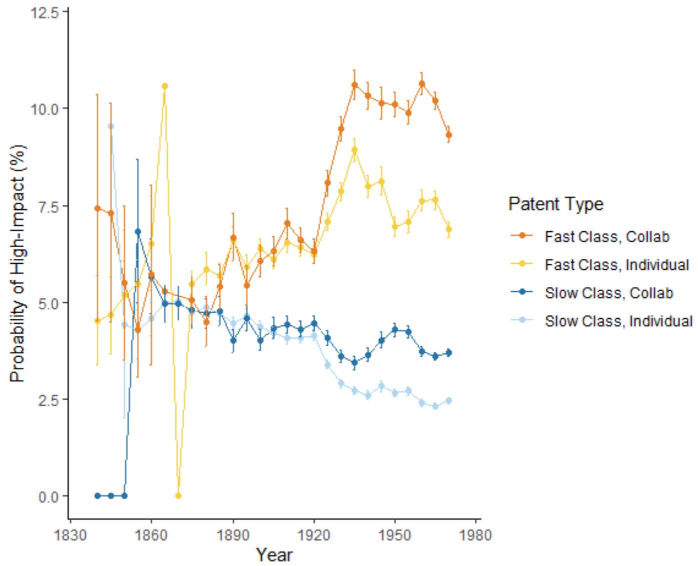

Figure 5.11 shows that the probability of creating a high-impact patent did not vary across fast and slow growth classes until 1890. After 1890, patents in fast growth classes were more likely to be high-impact than patents in slow growth classes. In addition, after 1920 patents invented by teams in fast growth classes were more likely to be high-impact than patents invented by individuals in fast growth classes, and patents invented by teams in slow-growth classes were more likely to be high impact than patents invented by individuals in slow growth classes.

Figure 5.11 may contain two issues. First, inventors with greater ability may select into fast growth classes or into collaboration more frequently. In addition, Figure 5.11 does not test whether the returns to collaboration vary by class growth rate. To address these issues, I develop a regression model where the probability that a patent is high-impact is a function of the team size that invented the patent, the *FastClass* binary variable described above, a term which interacts team size with *FastClass*, and inventor and year fixed effects:

$$(5.16)\ Prob(HighImpact_p) = B_1 Log(TeamSize)_p + B_2 FastClass_p + B_3 Log(TeamSize)_p * FastClass_p + FE_i + FE_t + E_p$$

As before, I am primarily interested in B_3. I hypothesize that B_3 will be positive and significant, indicating that the returns to collaboration are larger in fast growth technology classes. In addition, I run the model described by Equation 5.16 separately for patents granted before and after 1920 because I found in Figure 5.7 that 1920 is the critical year when positive returns to collaboration emerged. I estimate Equation 5.16 using a linear probability model and provide the regression results in Table 5.2.

Table 5.2: Regression Estimates of Effect of Team Size on Prob(High Impact)

	Patents Granted 1836-1919		Patents Granted 1920-1975	
$Log(TeamSize)$	-0.000490 (0.00127)		0.00661*** (0.000587)	
$FastClass$	0.00897*** (0.000968)	0.00907*** (0.00101)	0.0171*** (0.000696)	0.0178*** (0.000706)
$Log(TeamSize) * FastClass$	0.00213 (0.00212)		0.00703*** (0.00108)	
$Collab$		-0.000196 (0.000985)		0.00416*** (0.000498)
$Collab * FastClass$		0.00252 (0.00171)		0.00449*** (0.000952)
Inventor and year fixed effects	Yes	Yes	Yes	Yes
R^2	0.315	0.315	0.281	0.281
Inventor Subset	All Inventors	All Inventors	All Inventors	All Inventors
NOBS (unique inventors)	211,298	211,298	447,847	447,847

Standard errors clustered at inventor level in parenthesis.

Table 5.2 generates three important results. First, team size was not associated with high impact before 1920, but *FastClass* was. Second, the interaction team size and *FastClass* was insignificant for patents granted before 1920, indicating that the returns to team size were no different in fast-growth fields before 1920. Third, the interaction term was positive and significant for patents granted after 1920. This finding confirms the model's fourth proposition, that the returns to collaboration were greater in fast-moving classes after 1920.

5.8) Discussion

This paper demonstrated that inventors receive positive returns to collaboration and negative returns to maturity. Both sets of returns emerged during the first three decades of the 20th century. This timing is consequential. Historians of technology and the economy have identified the start of the 20th century as a transformative era. Technology historians describe shifts in the organization of innovation that took place in the early 1900s, including the invention of the corporate R&D lab and the emergence of firms as the principal coordinating actors of innovative activities (Lamoreaux and Sokoloff, 1996; Neffke et al., 2021). Economic historians have identified the 1920s as the final decade of rapid innovation in the United States. Following his herculean work in data collection and archival research, Gordon (2014) concludes that the period of rapid growth in the U.S. standard of living and business productivity during the middle of the 20th century resulted from the diffusion of technologies that were invented during the first three decades of the 20th century. According to his analysis, the rate of innovation in the United States slowed substantially after the 1930s.

It is plausible that the changes to the organization of innovative activities and the slowdown in the rate of invention in the early 20th century are joint outcomes of changes in the environment within which inventors operate. Before the turn of the 20th century, invention required a relatively small amount of knowledge. Empirical evidence supports this point: the opening of Carnegie library in towns and small cities in the U.S. at the dawn of the 20th century caused local patenting to increase between 8% and 13%, suggesting that the independent study of library texts was sufficient for invention during that era (Berkes and Nencka, 2021). However, as the 20th century unfolded, invention became harder to do. Bloom et al. (2020) show that the productivity of R&D declined by on average 4% per year between 1930 and 2015.

The increase in the knowledge-intensity of invention has generated several related consequences. First, the average education level of the inventors of patents has increased (Junge and Ejermo, 2014). Second, innovative activities have become increasingly concentrated in populous counties and metropolitan areas where inventors were able to source a wider range of ideas through face-to-face communication (Mewes, 2019; Balland et al., 2020). Third, inventors have reorganized by creating larger teams, by partnering and collaborating with non-family members, and by organizing into firms and corporate R&D labs (Lamoreaux and Sokoloff, 1996; van der Wouden, 2019; Neffke et al., 2021). A commonality between each of these adaptations is that they allowed inventors to access and employ more extensive stocks of knowledge.

Interestingly, the increase in the knowledge intensity of invention during the 20th century did not result in an increase in the average experience level of inventors (Liu et al., 2018; Liu et al., 2021) or in inventors' average age (Sinatra et al., 2016; c.f. Jones, 2010). In this study, I also showed that the returns to inventor maturity became sharply negative during that time period. The explanation I proposed for why maturity is inversely related to innovation is that the rate of knowledge growth has accelerated over the last century. Experience accumulated over time is an ineffective means for knowledge accumulation when the sought-after knowledge is rapidly evolving. For this reason, inventors adapt to accelerating knowledge growth by developing methods to accumulate knowledge rapidly, such as by forming teams, by organizing into firms, and by agglomerating in dense cities.

There is, to my knowledge, one alternative plausible hypothesis that could explain the negative returns inventors receive to maturity. Assume a model in which inventors begin their careers with a finite set of ideas, and that they patent their most impactful ideas early in their careers. As inventors mature, only their lowest impact ideas will remain to be patented. While plausible, this alternative model does not have empirical support. Notably, the model predicts that inventors who, for idiosyncratic reasons, take time off from patenting would not see any decrease in the impact of their inventions once they resume patenting, because they would not lose any of their ideas during those patenting breaks. To the contrary, Haller (2021) shows that, upon their return, inventors who take breaks from patenting create much lower impact inventions when compared with otherwise-similar inventors. This result is incompatible with a model in which the negative returns to maturity arise because inventors run out of ideas, but it is fully consistent with the model outlined in Section 5.5 of this chapter, in which the negative returns to maturity arise because other inventors expand the knowledge frontier and render mature inventors' knowledge stocks obsolete.

Finally, I conclude by discussing an issue associated with team formation in fast-advancing knowledge fields. The formation of teams allows inventors to expand their individual knowledge capabilities *only if* inventors already know a certain quantity of knowledge. Therefore, inventors must accumulate knowledge in some way in order to share it; otherwise the benefits to sharing knowledge would be nil. In this study, I demonstrated that the returns to inventor maturity are negative, so it is unlikely that inventors accumulate their sharable knowledge during their patenting careers. A more plausible explanation is that inventors develop highly specific and current technological skills and methods during their educational training. The effectiveness of inventors' knowledge is thus greatest when they commence their patenting careers, and as their industry evolves over time their skills become dated. Therefore, the positive returns to collaboration result from the sharing of educational backgrounds, and not professional experiences.

5.9) Appendix

Exclusion of Collaborative Patents in Calculation of Experience

Figure 5.12: Likelihood of Creating a High-Impact Patent by Inventor Maturity, Excluding Collaborative Patents

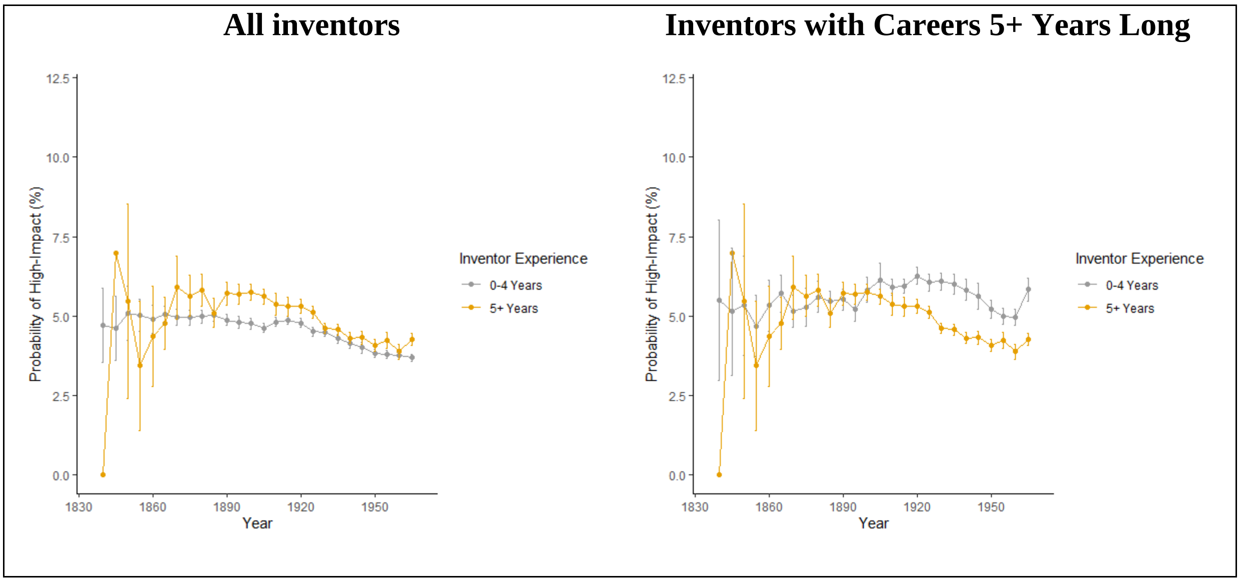

Returns to Experience Excluding Co-Invented Patents

Figure 5.13: Predicted Probability of Creating a High-Impact Patent, Excluding Co-Invented Patents

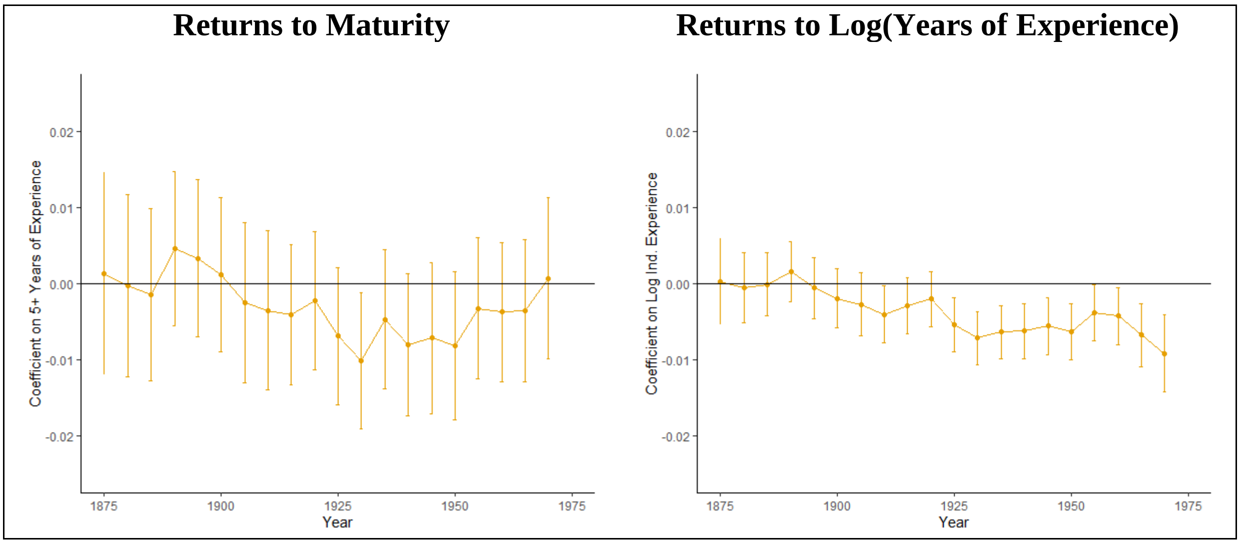

Analysis using 10% Threshold

In this sub-section, I replicate the main analyses using a top-10% threshold to define high-impact inventions. The results are similar to those using the top-5% threshold in the main text, with the exception of Table 5.3. I discuss this exception in the text below.

I begin in Figure 5.14 by plotting the probability that premature and mature inventors create a high-impact patent.

Figure 5.14: Probability of High-Impact Patent by Individuals and Teams

All Inventors | **Inventors that Reach Maturity**

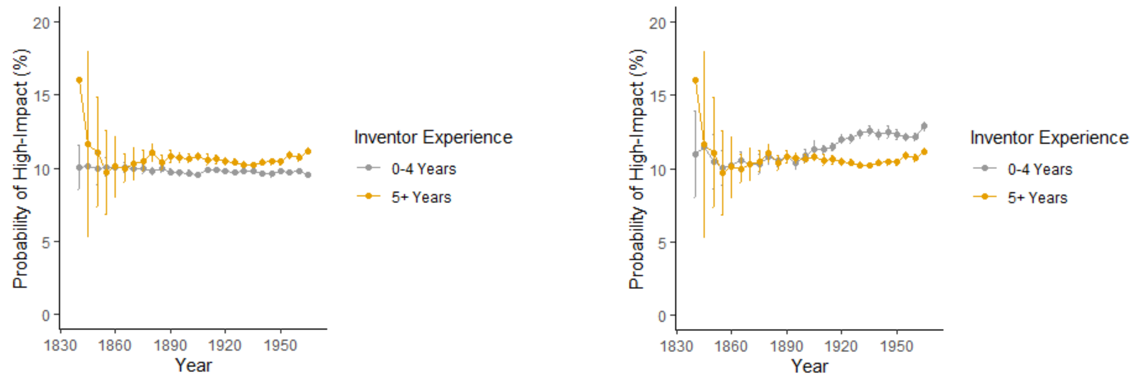

In Figure 5.15, I plot the probability that individuals and teams create a high-impact patent.

Figure 5.15: Probability of High-Impact Patent by Individuals and Teams

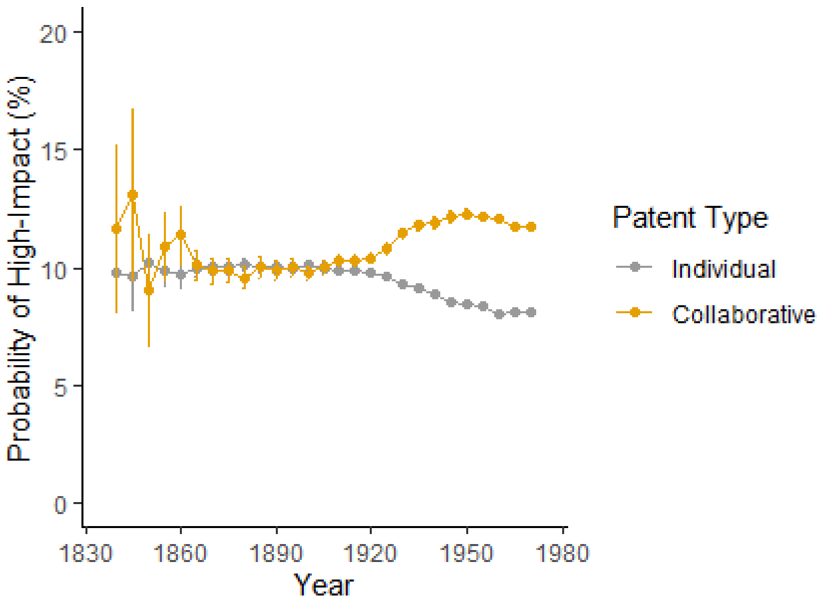

Figures 5.16 and 5.17 show the returns to experience and collaboration after adjusting for inventor-level variation in ability. The regressions used for these figures are the same as Equations 5.1 and 5.2 in the main text.

Figure 5.16: Returns to Maturity

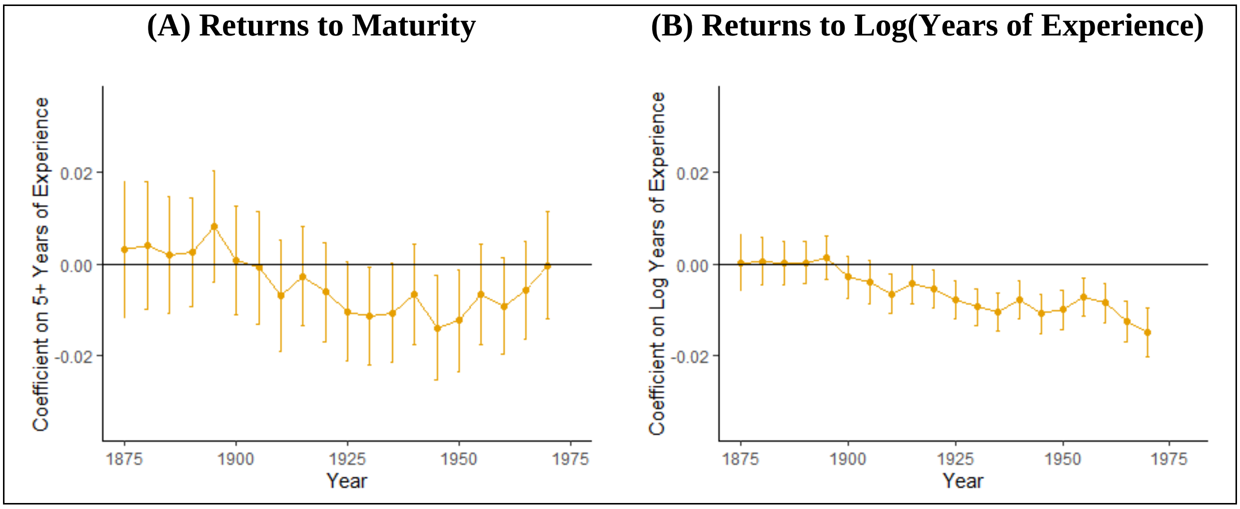

Figure 5.17: Returns to Collaboration and Team Size

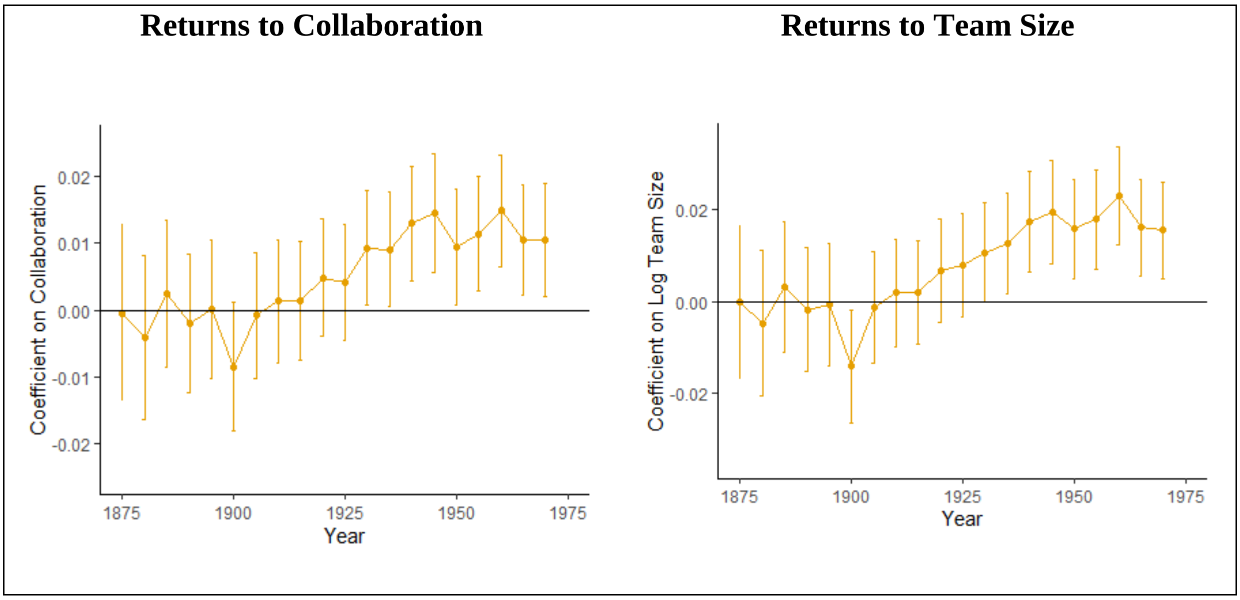

Figure 5.18 shows the mean probability that inexperienced inventors, experienced inventors, individuals, and teams create a high-impact patent, broken out by the growth rate of knowledge in a technology class.

Figure 5.18: Probability of High-Impact Patenting in Fast and Slow Classes, by Maturity and Collaboration

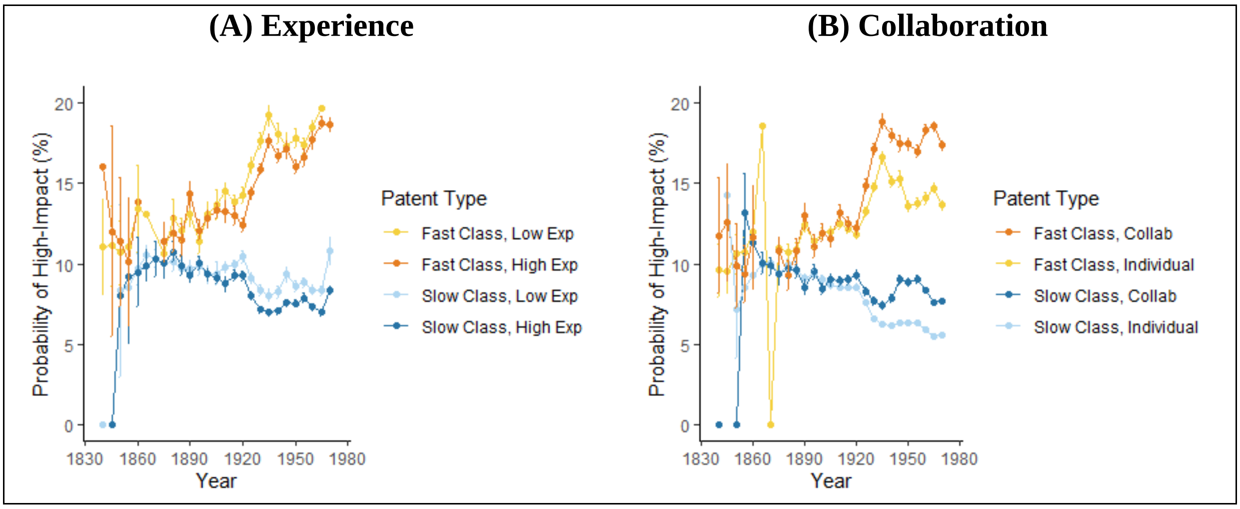

Panel A contains only inventors with careers that last 5+ years

Table 5.3 shows the regression results of the relationship between experience and the probability of creating a high-impact patent. The regression model is the same as Equation 5.15 in the main text. The results in Table 5.3 are slightly different from those in Table 5.1 (which is the same model but uses the 5% threshold to identify high-impact inventions).

Notably, in Table 5.1, the coefficient for the interaction terms for patents granted 1920-1975 is negative and significant, while the coefficient on the interaction terms is statistically insignificant in Table 5.3.

Table 5.3: Regression Estimates of Effect of Years of Experience on Prob(HighImpact)

	Patents Granted 1836-1919		Patents Granted 1920-1975	
$Log(YearsOfExperience)$	-0.00326*** (0.000985)		-0.00637*** (0.000617)	
$FastClass$	0.0132*** (0.00183)	0.0157*** (0.00192)	0.0317*** (0.00125)	0.0310*** (0.00132)
$Log(YearsOfExperience) * FastClass$	0.00142 (0.000945)		0.000248 (0.000620)	
$Experienced\ (\geq 5\ years)$		0.00130 (0.00181)		-0.00737*** (0.00104)
$Experienced * FastClass$		0.000649 (0.00248)		0.00184 (0.00152)
Inventor and year fixed effects	Yes	Yes	Yes	Yes
R^2	0.329	0.290	0.304	
Inventor Subset	All Inventors	Inventors with careers 5+ years	All Inventors	Inventors with careers 5+ years
NOBS (unique inventors)	211,298	135,899	447,847	254,434

Standard errors clustered at inventor level in parenthesis.

There are a few possible reasons why the coefficient on the interaction term is not significant in Table 5.3 but significant and negative in Table 5.1. One possibility is that experience in a fast-growing class may be negatively associated with inventing very high-impact patents (top 5%) but less so with inventing slightly less high-impact patents (patents falling in the top 10% to top 6% of the impact distribution). This argument is supported by the more negative coefficient on log years of experience using the 10% threshold (-0.00637, Table 5.3) than using the 5% threshold (-0.00384, Table 5.1). The less negative coefficient using the 5% threshold indicates that the challenge of creating high-impact inventions is greater between fast and slow-growth fields under the 5% threshold. Moreover, a 10% threshold may be embracing enough to reduce the distinction between the field growth rates.

Finally, in Table 5.4 I plot regression coefficients for the effect of team size on the probability of creating a high-impact patent. These regression results are statistically identical to those I arrive at using a 5% threshold in the main text (Table 5.2).

Table 5.4: Regression Estimates of Effect of Team Size on Prob(HighImpact)

	Patents Granted 1836-1919		Patents Granted 1920-1975	
$Log(TeamSize)$	-0.000816 (0.00172)		0.0101*** (0.000811)	
$FastClass$	0.0150*** (0.00135)	0.0147*** (0.00136)	0.0281*** (0.000940)	0.0294*** (0.000952)
$Log(TeamSize) * FastClass$	0.00200 (0.00285)		0.00993*** (0.00141)	
$Collab$		-0.000872 (0.00135)		0.00704*** (0.000702)
$Collab * FastClass$		0.00252 (0.00227)		0.00593*** (0.00125)
Inventor and year fixed effects	Yes	Yes	Yes	Yes
R^2	0.329	0.329	0.304	0.304
Inventor Subset	All Inventors	All Inventors	All Inventors	All Inventors
NOBS (unique inventors)	211,298	135,813	447,847	254,434

Standard errors clustered at inventor level in parenthesis.

5.10) Bibliography

Akcigit, U., Grigby, J., and Nicholas, T. (2017). The Rise of American Ingenuity: Innovation and Inventors in the Golden Age. *HBS Working Paper 17-063.*

Arrow, K. (1962). The Economic Implications of Learning by Doing. *Review of Economic Studies* 29(3). 155-173.

Arthur, B. (2009). *The Nature of Technology: What It Is and How It Evolves*. Free Press, New York.

Balland, P.A., Jara-Figueroa, C., Petralia, S., Steijn, M., Rigby, D., and Hidalgo, C. (2020). Complex Economic Activities Concentrate in Large Cities. *Nature Human Behavior* (4).

Berkes, E., and Necka, P. (2020). 'Novel' Ideas: The Effects of Carnegie Libraries on Innovative Activities. Working Paper.

Bloom, N., Jones, C., Van Reenen, J., and Webb, M. (2020). Are Ideas Getting Harder to Find? American Economic Review 110(4). 1104-1144.

Bell, A.M., Chetty, R., Jaravel, X., Petkova, N., & Van Reenen, J. (2019). Who Becomes an Inventor in America? The Importance of Exposure to Innovation. *Quarterly Journal of Economics* 134(2). 647-713.

Esposito, C. (2020). The Emergence of Knowledge Production in New Places. Working Paper. PEEG #20.46.

Fleming, L. (2001). Recombinant Uncertainty in Technological Search. Management Science 47(1). 117-132.

Gordon, R. (2016). *The Rise and Fall of American Growth*. Princeton University Press.

Haller, M. (2021). Firm Decline, the Patent Gap, and Inventor Performance. Working Paper.

Jones, B.F. (2010). Age and Great Invention. *Review of Economics and Statistics* 92(1). 1-14.

Jones, B.F., Wuchty, S., and Uzzi, B. (2008). Multi-University Research Teams: Shifting Impact, Geography, and Stratification in Science. *Science* 322(5905). 1259–1262.

Jung, T., and Ejermo, O. (2014). Demographic Patterns and Trends in Patenting: Gender, Age, and Education of Inventors. *Technological Forecasting and Social Change* 86(C). 110-124.

Kats, J., and Martin, B. (1997). What Is Research Collaboration? *Research Policy* 26(1). 1-18.

Lamoreaux, N.R., and Sokoloff, K.L. (1996). Long-term change in the organization of inventive activity. *Proceedings of the National Academy of Sciences of the United States of America*. 93(November), 12686–12692. https://doi.org/10.1073/pnas.93.23.12686

Lamoreaux, N.R., and Sokoloff, K.L. (2005). The Decline of the Independent Inventor: A Schumpterian Story? *National Bureau of Economic Research*. Working Paper 11654.

Lamoreaux, N.R., Sokoloff, K.L., and Sutthiphisal, D. (2011). The Reorganization of Inventive Activity in the United States During the Early Twentieth Century. In D.L. Costa and N.R. Lamoreaux (Eds.), *Understanding Long-Run Economic Growth: Geography, Institutions, and the Knowledge Economy* (235–274). University of Chicago Press.

Lamoreaux, N.R., Sokoloff, K.L., and Sutthiphisal, D. (2013). Patent Alchemy: The Market for Technology in US History. *Business History Review* 87(01). 3–38. https://doi.org/10.1017/S0007680513000123

Levin, D., and Cross, R. (2004). The Strength of Weak Ties You Can Trust: The Mediating Role of Trust in Effective Knowledge Transfer. *Management Science* 50(11). 1477-1490.

Liu, L., Wang, Y., Sinatra, R., Lee Giles, C., Song, C., and Wang, D. (2018). Hot Streaks in Artistic, Cultural, and Scientific Careers. *Nature* 559. 396-399.

Liu, L., Dehmamy, N., Chown, J., Lee Giles, C., and Wang, D. (2021). Understanding the Onset of Hot Streaks Across Artistic, Cultural, and Scientific Careers. Working Paper. arXiv: 2103.01256.

Mewes, L. (2019). Scaling of Atypical Knowledge Combinations in American Metropolitan Areas from 1836 to 2010. *Economic Geography* 95(4). 341-361.

Neffke, F., and Hartog, M. (2021). Organizing Complexity: The Changing Social Technology Behind Team Work, 1850-1940. Working Paper.

Petralia, S., Balland, P., and Rigby, D. (2016). Unveiling the Geography of Historical Patents in the United States from 1836 to 1975. Scientific Data (3).

Powell, W., Koput, K., and Smith-Doerr, L. (1996). Interorganizational Collaboration and the Locus of Innovation: Networks of Learning in Biotechnology. *Administrative Science Quarterly* 41(1). 116-145.

Sarada, S., Andrews, M., and Ziebarth, N. (2019). Changes in the Demographics of American Inventors, 1870-1940. *Explorations in Economic History* 74.

Sinatra, R., Wang, D., Deville, P., Song, C., and Barabasi, A-L. (2016). Quantifying the Evolution of Individual Scientific Impact. *Science* 354(6312).

van der Wouden, F. (2020). A History of Collaboration in US Invention: Changing Patterns of Co-invention, Complexity and Geography. *Industrial and Corporate Change* 29(3). 559-619.

Verhoeven, D. (2021). Knowledge Diversity in Teams and Innovation. Working Paper. http://ssrn.com/abstract= 3795476.

Weitzman, M. (1998). Recombinant Growth. *Quarterly Journal of Economics* 113(2). 331-360.

Wuchty, S., Jones, B.F., and Uzzi, B. (2007). The Increasing Dominance of Teams in Production of Knowledge. *Science* 316(5827). 1036–1039. https://doi.org/10.1126/science.1136099