The Geography of Breakthrough Innovation in the United States over the 20th Century

Christopher Esposito

Papers in Evolutionary Economic Geography

21.26

Utrecht University Human Geography and Planning

The Geography of Breakthrough Innovation in the United States over the 20th Century

Christopher Esposito
University of Chicago

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 changes. At the start of the 20th century, breakthrough innovation was concentrated in populous and knowledge-diverse metropolitan areas. By the 1930s, breakthroughs were created less frequently across the entire country and so their invention had a less distinct geography. The substantial creation of breakthroughs resumed in the 1960s and was once their invention was concentrated in large and knowledge metropolitan areas. However, during the latter part of the century the invention of breakthroughs also frequently involved long-distance collaborations between inventors. In this paper, I document these historical changes to the geography of breakthrough innovation and propose a model to explain why they occurred. The model suggests 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 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. I generate support for the model, and 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.

Acknowledgements: I thank David Rigby, Michael Storper, Allen Scott, Wolf Uhlbach, Anthony Frigon, Carolin Ioramashvili, Ryan Hynes, participants at the SEG seminar series, and participants at the NARSC annual meeting for their thoughtful feedback. In addition, I thank Frank van der Wouden for providing pre-1975 inventor location data.

1) 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 studies 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 the flows of diverse ideas that circulate in those regions (Duranton and Puga, 2001; Mewes, 2019; Berkes and Gaetani, 2020). This research effort, however, is challenged 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 almost 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 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 not a singular territorial form of economic activity that optimizes creativity and innovation in the absolute sense.

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 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 (Arzaghi and Henderson, 2008; Baum-Snow et al. 2020). However, during the late 20th century, non-local collaboration between patent inventors also became increasingly frequent, which suggests that a more complex geography of innovation emerged had started to emerge (Bathelt et al., 2004; Van der Wouden, 2019). With these caveats in mind, Table 1 synthesizes the above-listed sources in a timeline of the broad shifts in the spatial distribution of innovative activity across in the United States.

Table 1: Geographical Distribution of Innovative Activity 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.

The timeline in Table 1 is suggestive, but it is not a rigorous account of how the relationship between agglomeration and innovation evolved over time. The innovative activities summarized in Table 1 – patent production 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 may be paid to non-innovative work. The lack of a harmonized definition of innovative activity across the literature summarized in Table 1 motivates the first objective of this paper: to systematically describe how the relationship between agglomeration density and breakthrough innovation in the United States evolved over the 20th century.

The second objective of this paper is to propose a cogent explanation for why the changes to the geographical distribution of breakthrough innovation occurred. In this respect, I argue that four interacting factors collectively determine the geographical distribution of breakthrough innovation: 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. Resulting geographies of breakthrough innovation can be understood as outcomes of these factors. For example, breakthrough innovation concentrates in big and knowledge-diverse cities when the knowledge intensity of breakthroughs is high and knowledge-sourcing technologies incur strong distance-based frictions because inventors need to build on a large quantity of ideas to develop breakthroughs and because inventors are best able to source ideas that are nearby. I elaborate on these four factors in Section 5 of the paper, and I in Section 6 I empirically show that changes in these factors' conditions are correlated with changes in the geography of breakthrough innovation.

Sections 3 and 4 of the paper describe the steps used to carry out the main empirical analysis. In that analysis, I examine how the propensity for inventors residing in knowledge-diverse city-regions in the U.S. to create breakthrough inventions changed over time using data from 4 million patents granted between 1900 and 1999. I define knowledge-diverse city-regions as the Core-Based Statistical Areas (CBSAs) where local inventors patent in a wide array of patent classes.[1] Because the geography of breakthrough innovation may also be non-binary (i.e. not just urban vs. rural) or non-ordinal (i.e. not just a continuum of local knowledge density), I also examine changes in the propensity for inventors engaged in non-local collaborations to invent breakthroughs. The empirical 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* were engaged in non-local collaborations were more likely than all other types of inventors to develop breakthroughs.

Before moving to the empirical analysis, I begin in Section 2 by introducing a literature on breakthrough invention and its geography. The main empirical analysis, the description of the theoretical model, and an analysis of the model's feasibility follow. Finally, in Section 7 I discuss how the model proposed in this study revises a common interpretation for 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.

2) 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 result from the loss of information incurred by long-distance communication technologies when actors transmit messages using audio, visual, and physical

---

[1] CBSAs include all metropolitan and micropolitan areas in the United States. There are a total of 982 U.S. CBSAs. Empirically, 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.

channels across distance (Storper and Venables, 2004). Face-to-face communication, which incurs minimal information loss, is only possible between actors that are located in the same physical location, so inventors that are 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 location in knowledge-rich environments should be particular relevant for inventors working to develop breakthrough inventions. Very few of the possible ideas that inventors can combine qualify as breakthroughs, so inventors have to search widely amongst the set of combinatorial possibilities to identify the few that do (Youn et al., 2015). In this regard, there are two characteristics that distinguish breakthroughs from other inventions. First, breakthroughs are novel in that they combine existing ideas in dramatically imaginative ways (Uzzi et al. 2013; Mewes, 20019; Berkes and Gaetani, 2020). Second, Breakthroughs are highly-impactful in that they combine ideas with a high level of complementarity and so enable a large quantity of subsequent innovation (Fleming et al., 2001). The combinations of ideas that are both novel and complementary comprise a very small region of the technological search space (Fleming and Sorenson, 2001), so inventors must search for knowledge extensively, must source knowledge intensively, and more often than not must locate in knowledge-diverse environments in order to invent breakthroughs.

Duranton and Puga (2001) and Berkes and Gaetani (2020) develop formal models where distance-based search costs cause innovating actors to locate in knowledge-diverse environments. 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 knowledge diversity, 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 locally generated, but complementary combinations are universally generated. Therefore, neither Duranton and Puga's (2001) nor Berkes and Gaetani's (2020) model generates an explicit prediction for how distance-based frictions affect the search for combinations of ideas that are *both* novel and complementary. However, the two models collectively propose that the creation of novel and complementarity combinations are positive functions of the heterogeneity of ideas that circulate in local environs.

While there is widespread interest in the drivers of breakthrough innovation, its geography has not been extensively studied. Two exceptions are Grashof et al. (2019) and De Noni and Belussi (2021). Grashof et al. (2019) study the creation of novel and impactful patents in Germany and find that these patents are disproportionately created by firms that are located geographically inside innovative clusters but whose inventors are in the periphery of their clusters' collaborative networks.[2] From these results, the authors conclude that both local and non-local interactions between inventors are important for the creation of breakthroughs. De Noni and Belussi (2021) study the creation of novel and impactful patents in regions of the European Union between 2008 and 2014 and find that they are most frequently invented in regions with multiple but related industrial specializations. The authors interpret the benefits of specialization within co-agglomerated indutries as an outcome of the extensive knowledge heterogeneity that can be found within industries. Therefore, while De Noni and Belussi's (2021) conceptualization of the diversity of knowledge in regions is more nuanced than a simplistic aggregate measure of specialization or diversity, their results are consistent with the view that the propensity for inventors to create breakthroughs increases when they have access to a diverse array of atomistic knowledge units.

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 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.[3] 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 novel 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 overall relationship between the novelty of patents and the impact of patents, measured using patent forward citation counts. They find that novel patents in the U.S. are disproportionately created in counties with high population densities and that novel patents are on average more impactful than non-novel inventions in terms of spurring subsequent innovation. However, Berkes and Gaetani (2020) do not analyze whether patents which are both novel and impactful are more often created in high-density counties. Finally, Castaldi et al. (2015) examine the knowledge-based characteristics of U.S. states

[2] 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.

[3] 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.

that are more likely to produce high-impact patents, measured again using patent 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 patents. In addition, Castaldi et al.'s (2015) study is at the state level, within which population density and local knowledge diversity substantially varies. 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. In addition, the two studies that do analyze the geography of the production of patents that are both novel and impactful (Grashof, et al., 2019; De Noni and Belussi, 2021) are focused on European regions. 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 conditions 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 change over time. Schumpeter (1934; 1942) proposes that there are times when technological change advances incrementally and that there are times when it advances disruptively. During an incremental regime, few if any breakthrough inventions are introduced so the production of breakthroughs should not have a distinctive geography.

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) find 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.

3) Methods

3.1) 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 scientific articles 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 subclassification codes listed on patents, taking those subclass codes as indicators of the knowledge components in each invention. Berkes and Gaetani (2020) apply z-scores to the citations made by patents to a similar effect. Atypicality, measured at the pairwise level between all ideas combined in an invention, can be aggregated to the invention level to compute the overall novelty of an invention.

In this study, I identify novel inventions by assessing the atypicality between their internal elements. I calculate the atypicality of all combinations of knowledge units in each patented invention by calculating z-scores for all USPTO utility patents issued between 1900 and 1999.[4] I use the coarse-grained subclasses 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 1:

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

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

In Equation 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 as follows:

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

In Equation 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, $\sigma_{i,j}^2$, is given by Equation 3:

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

$Z_{i,j}$ 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 straightforward measure 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 that introduce one 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). Esposito (2020) develops a related approach by tracing the flow of knowledge between individual patents based on the co-occurrence of combinations of subclassification codes found on different patents. There are two advantages to the latter method. First, citation records are unavailable for patents granted before 1947 (Akcigit, et al. 2017) but the subclass codes used by Esposito's (2020) 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. The one exception is the small number of patents that are assigned to just one subclass. Because patents assigned to a single subclass do not contain any combinations of subclasses, their novelty cannot be measured and I must omit them from the study.

To compute the impact of individual patents on subsequent invention, I follow the method of Esposito (2020) to count the number of subsequent inventions that draw knowledge from each patent. I deviate slightly from the method of Esposito (2020) by using course-grained USPC subclassification codes instead of the most-disaggregated codes. I make this change because course-grained subclasses allow me to use the same classification scheme across the entire analysis. In keeping with the method of Esposito (2020), I reduce the computational intensity of the task by restricting my dataset to the first 8 subclasses on patents. While less than 5% of patents are affected by omitting these excess subclasses, taking this subset requires me to restrict the sample of patents that I analyze to patents with 7 of fewer subclasses. After using the method described in Esposito (2020) to predict the flow of knowledge between patents, I compute the impact of patents by counting the number of patents that draw knowledge from each focal patent within 10 years of the grant year of each focal patent. Because my dataset ends in 2010, the 10-year window for forward citations allows me to compute the impact of patents granted up to 1999 without right truncation.

Table 2 presents a typology of patents that vary in terms of impact and novelty. In the subsequent analyses I treat the number of knowledge-based descendents of a patent as a continuous variable. However, for simplicity in Table 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 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

Note: Dataset continues all USPTO utility patents assigned between 2 and 7 subclasses and granted between 1900 and 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%.

3.2) The Geography of Breakthroughs

After classifying each patent based on the typology in Table 3, 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.[5] 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

[5] van der Wouden's (2020) dataset provides better geographical coverage for historical patents invented by multiple co-inventors than the publicly available dataset, HistPat (Petralia, et al. 2016),

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 knowledge diversity criterion.[6]

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, I find that there is no between the production of novelty and the knowledge diversity of the regions where those patents are invented in Figure 1. I produce Figure 1 by aggregating the total number of patents produced, the number of novel patents produced, and the number of high-impact patents produced to the CBSA level. The figure 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 of Mewes (2019), that novel patenting is no more concentrated in knowledge-diverse cities than overall patenting is. On the other hand, Figure 1 shows that high-impact patenting (dotted line, defined as in Table 3) is more concentrated in knowledge-diverse cities than is overall patenting. Because these results show that novelty does not benefit from location in knowledge-diverse regions but the creation of high-impact inventions does, local knowledge may help inventors less for finding novel combinations of ideas than for finding complementary combinations of ideas. Therefore, in the proceeding empirical analysis, I take the geographical distribution of the production of novelty as a given, and examine how the impact of novel patents varies with local knowledge.

Figure 1: Percentage of Patents Produced in Knowledge-Diverse Cities by Patent Type

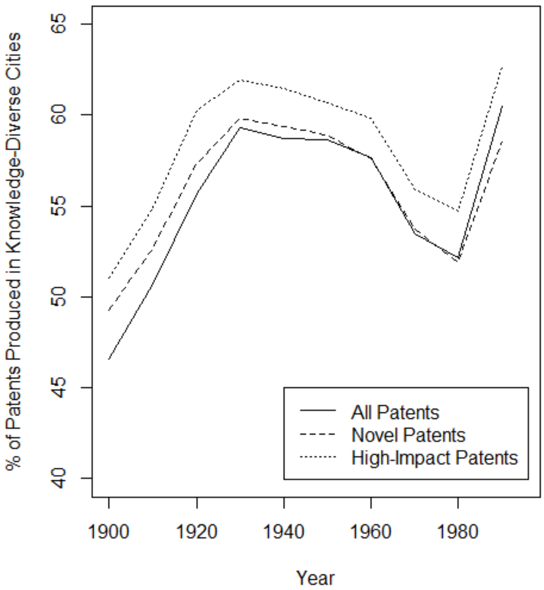

[6] In 1950, the knowledge-diverse CBSAs were (in descending order), New York, Chicago, Los Angeles, Philadelphia, Cleveland, Boston, Detroit, Pittsburgh, Cincinnati, San Francisco, Washington DC, Milwaukee, and Bridgeport CT.

4. Results: The Geography of Breakthrough Innovation

For the reasons discussed above, I examine changes in the benefits of local knowledge diversity on the invention of breakthroughs by testing whether the average impact of novel patents varies with local knowledge diversity. I begin to perform this examination in Figure 2 by plotting changes in the average impact of four types of patents over time: 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, it is infeasibly to plot a scatterplot so instead I plot best-fit lines with 95% confidence intervals. The large number of observations also renders the most common moving-average fit line (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 2: Average Patent Impact by Novelty and Knowledge Diversity of City of Invention

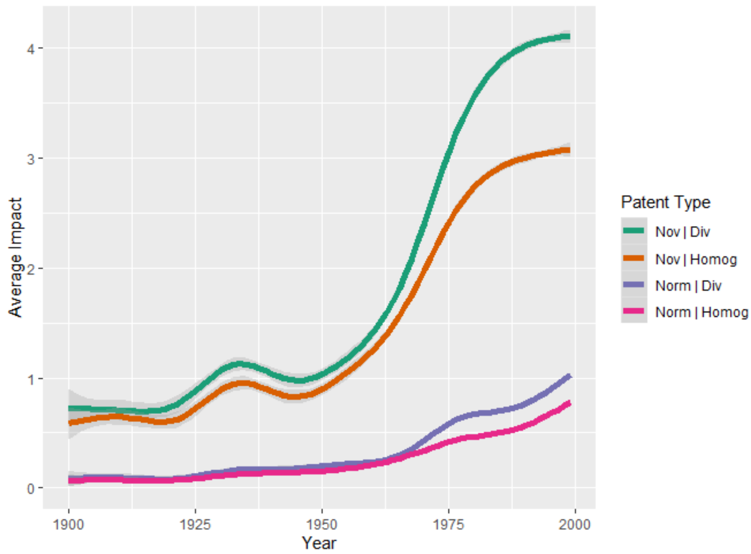

Figure 2 generates three inferences. First, across all years, novel patents invented in knowledge-diverse cities (Nov|Div) were on average the most impactful type of patents, followed by novel patents invented in knowledge-homogeneous cities (Nov|Homog) inventions. Second, the average impact of all types 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 (Esposito 2020). 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:

$$(4)\ Impact_p = Year_p * NrSubclasses_p + E_p$$

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 3.

Figure 3: Predicted Patent Impact by Novelty and Local Knowledge Diversity of CBSA of Invention

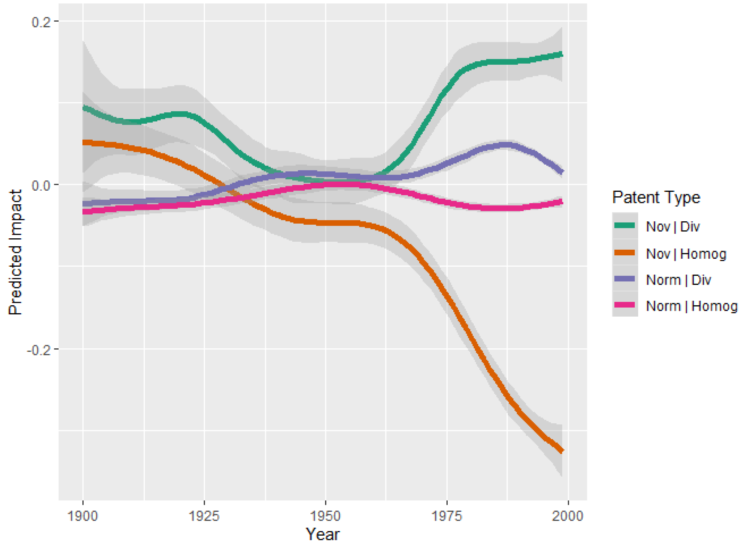

Note: The regression used to estimate predicated impact is given in Equation 4.

Figure 3 shows that the impact of patents with different levels of novelty and invented in cities with different levels of knowledge diversity had three distinct periods during the 20th century. The first period was 1900 to 1930. During this 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 was 1930 to 1965, during which 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 was 1965 to 1999, during which the predicted impact of novel inventions made in knowledge-diverse increased above that of normal patents. In addition, 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 knowledge-diverse cities.

While Figure 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 the study 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). Therefore, I examine the relationship between the engagement of inventors in non-local collaborations and the creation of breakthroughs in Figure 4. To do so by comparing the average impact of patents invented by inventor-teams located in single CBSAs and in multiple CBSAs. In addition, I decompose inventor-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 inventor-teams with teammates that resided in both knowledge-diverse and knowledge-homogeneous cities (I analyze these mixed teams in Appendix A). Finally, I omit all patents invented by lone inventors.

Figure 4: Average Patent Impact of Collaborative Patents by Type of Collaboration

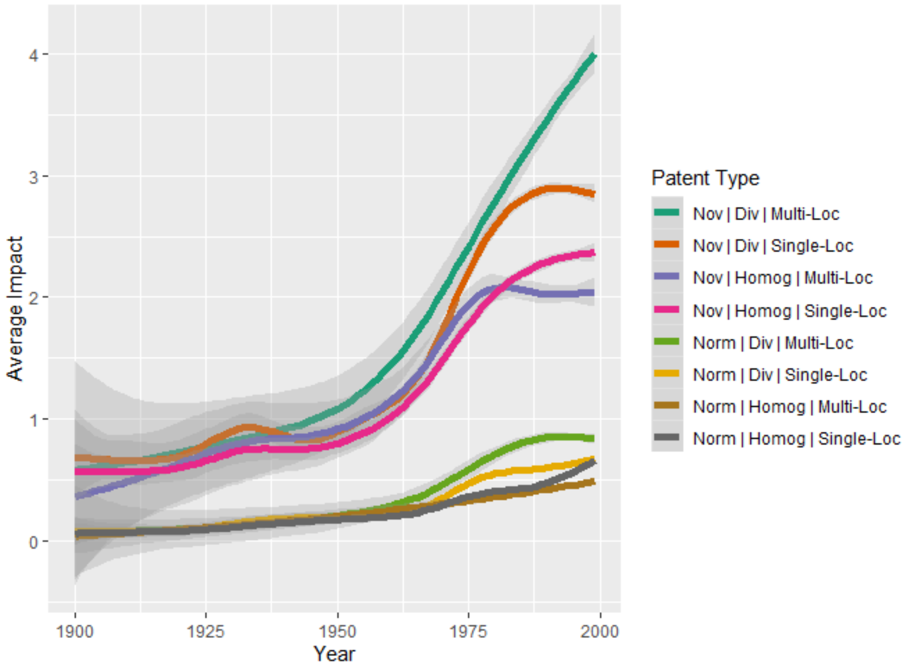

Figure 4 shows that the average impact of novel patents produced by teams in knowledge-diverse cities and in knowledge-homogeneous cities were statistically identical until 1960. In addition, before 1980 there was no significant difference in the average impact of novel patents produced by single-location 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 proposed in Section 5 will discuss how this spatial pattern can emerge when the state of collaborative and knowledge-sourcing technologies fit certain conditions.

Finally, in Figure 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 and aggregate the residuals by year and patent type.

Figure 5: Predicted Patent Impact of Collaborative Patents by Type of Collaboration

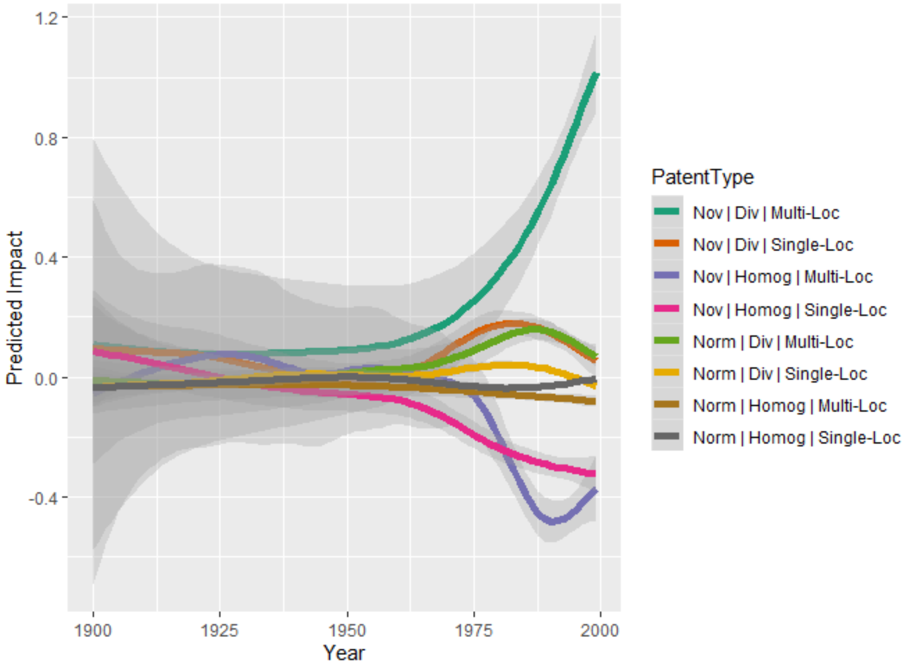

Note: Regression to estimate predicted impact is given in equation 4.

Figure 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 5 shows that the increasing concentration of breakthrough innovation in knowledge-diverse cities documented in Figure 3 was driven by inventors that collaborated with non-local teammates. Moreover, breakthrough innovation at the end of the 20th century was most common in large innovative clusters connected to other distant large innovative clusters.

5. 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 propose that the geography of breakthrough innovation is influenced by the conditions of four factors: (1) the disruptiveness of the regime of technological change, (2) the knowledge-intensity of breakthrough invention, (2) the state of long-distance collaboration technology, and (3) the state of long-distance knowledge-sourcing technology.

The *disruptiveness of the regime of technological change* captures the extent to which technological knowledge advances through the creation of novel inventions versus normal or incremental ones (Schumpeter, 1934; Schumpeter, 1942). 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 *knowledge intensity of breakthrough inventions* is the prevailing returns that sourcing a larger number of knowledge-based inputs has on the creation of high-impact novelty (Wuchty, et al. 2007; Jones, 2009; Balland et al., 2018; Bloom et al., 2020). 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.

Long-distance collaboration technology refers 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. Face-to-face collaboration suffers no minimal information loss but it is only readily viable for inventors that reside in the same region (Storper and Venables, 2004). One might imagine very robust long-distance collaboration technologies are perfect substitutes for face-to-face communication, which would allow inventors are able to create novelty through collaborations across distance with no loss in impact. However, it is uncertain if long-distance collaboration technologies will ever advance to that point.

Long-distance knowledge-sourcing technology refers to the items inventors use to source ideas from non-local regions in which they do not have active collaborators. These items include scientific articles, patent documents, and physical technological devices that can be reverse-engineered. The robustness of long-distance knowledge-sourcing technologies can be assessed based on the information lost when sourcing knowledge across long distances relative to the information lost when inventors source local knowledge. Inventors that source knowledge locally benefit from physical presence and embeddedness in an environment of shared norms, so the information lost in sourcing local knowledge is less than than the information lost when sourcing non-local knowledge. (Gertler, 2003). One might imagine fully robust long-distance knowledge-sourcing technologies that are perfect substitutes for local knowledge sourcing; however, such a state of knowledge sourcing technologies also has not yet come to pass.

The four factors, described above, interact to produce distinctive geographies of breakthrough innovation. To begin to understand how they do so, I first consider changes in the disruptiveness of the regime of technological change. Only when the regime of technological change is disruptive are a sufficient number of breakthroughs introduced to the economy for their invention to have a defined geography. Therefore, limit the subsequent discussion to the changes in the geography of breakthroughs under the condition that the regime of technological change is disruptive.

When the regime of technological change is disruptive, the geography of breakthrough innovation depends on how much knowledge inventors need to access in order to create breakthroughs and the distance-based frictions they incur when accessing that knowledge. Notably, when the knowledge intensity of breakthroughs is low, the effective cost of accessing knowledge across long distances also must be low. Therefore, whenever the knowledge intensity of breakthroughs is low, we might anticipate that breakthrough innovation will be dispersed across space. These conditions may describe U.S. invention during the 19th century, when innovation was less complex (Balland et al. 2020) and anecdotal evidence indicates that many breakthroughs were invented in the countryside (Mokyr, 1990). However, when the knowledge intensity of breakthroughs is high, the geography of their production is determined by the difficulties inventors face in accessing distant knowledge. When inventors face difficulty in accessing distant knowledge, they will agglomerate when they developing breakthroughs, resulting in a strong spatial concentration of breakthrough innovation.

The most interesting outcomes of the model occur when the technologies inventors use to access non-local knowledge, namely long-distance knowledge sourcing technologies and long-distance collaborative technologies, improve asymmetrically. When the knowledge intensity of breakthroughs is high, long-distance knowledge sourcing technologies are strong, and long-distance collaboration technologies are weak, innovation will disperse across space because inventors are able to source all ideas across distance with relative ease. When the knowledge intensity of breakthroughs is also high but long-distance knowledge sourcing technologies are weak and long-distance collaborative technologies are strong, the geography of breakthroughs takes a multi-nodal layout across space, with core innovative hubs connected by long-distance collaborative networks. This outcome arises because productive long-distance collaborative relationships require shared trust, norms, and expectations that are highly selective between individuals and require extensive investments of time and attention. These high costs to relationship building imply that inventors will only invest in building relationships with others that are located in regions with rich local knowledge environments. Moreover, when inventors use social proximity to overcome the distance-based frictions of sourcing knowledge (Boschma, 2005), the geography of the resulting inventions will be multi-nodal.

In Table 3, I summarize the geography of breakthrough innovation in a disruptive regime of technological change under different levels of knowledge-intensity and different strengths of long-distance collaborative and knowledge-sourcing technologies.

Table 3: Geography of Breakthroughs in Disruptive Regimes of technological Change

Long-Distance Collaborative Technologies	**Long-Distance Knowledge-Sourcing Technologies**	**Knowledge-Intensity of Breakthroughs**	
		Low	High
Weak	Weak	Dispersed	Perfectly Concentrated
	Strong	Dispersed	Dispersed
Strong	Weak	Dispersed	Multi-Nodal
	Strong	Dispersed	Dispersed

6. Empirical Assessment of the Theoretical Model

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 understand how these three factors that influence the geography of breakthrough innovated changed over the study period.

I begin by studying the evolution of the knowledge intensity of breakthrough innovation, which I define as the additional knowledge sources for helping inventors to create high-impact novelty. To perform this analysis, I 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. To measure the number of prior knowledge sources that each patent draws ideas from, I compute the in-degree of patents using the graph of knowledge flows described in Section 2. To simplify the analysis, 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 their grant year cohort in terms of the number of prior patents they draw knowledge from. I define patents as having "few knowledge-based parents" if they fall in the bottom 90% of their grant year cohort.

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. 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 6. A similar figure using raw impact values is presented in Appendix B.

Figure 6: Predicted Patent Impact by Novelty and Number of Patent Parents

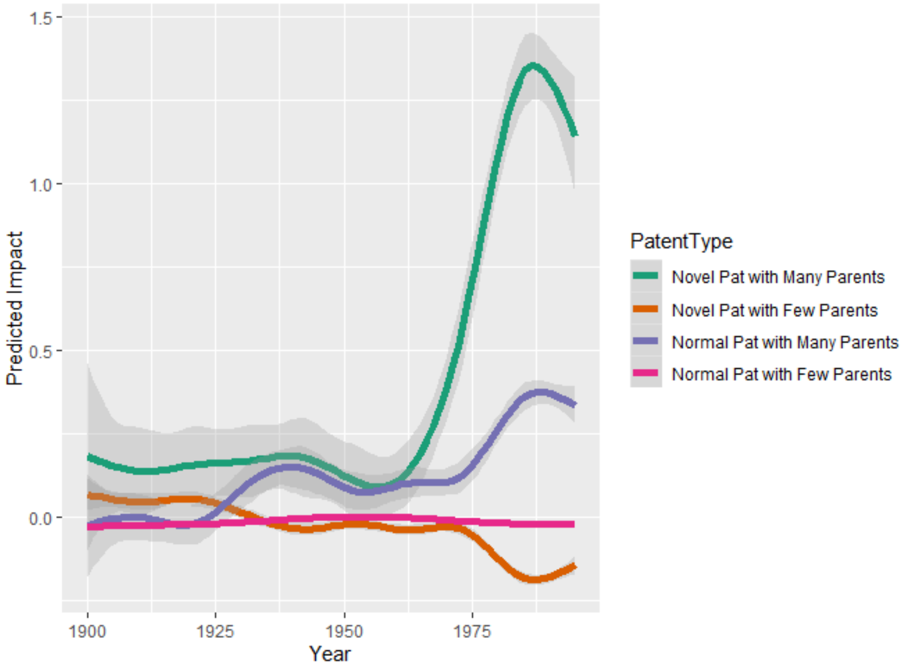

Note: Regression to estimate predicted impact is given in Equation 4.

In Figure 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. After that date, 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, I conclude that knowledge intensity of breakthroughs was moderate until 1965 but very high after 1965. The relationships identified in Figure 6 are also present within technological fields, as shown in Figure B2 in Appendix B 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. As discussed earlier, there are two types of long-distance communication technologies: long-distance collaboration technologies, and long-distance knowledge-sourcing technologies. 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 5 in Section 4 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 define = local knowledge sourcing as knowledge sourced from locations in which multi-locational teams have on-the-ground collaborators. Only one teammate needs to reside in a given CBSA for the knowledge sourced from that CBSA to be considered local. To aggregate this pairwise patent-patent level indicator to the individual patent level, I compute whether an above-average number of the knowledge sources used in a given patent were sourced from CBSAs in which the patent's inventors reside; I define patents that fit this condition as "patents that source knowledge with proximity".[7] Patents that source a below-average quantity of knowledge from CBSAs in which the patents' inventors reside are defined as "patents that source knowledge without 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 7.

Figure 7: Predicted Patent Impact by Novelty and Extent to which a Patent Sources Knowledge with Proximity

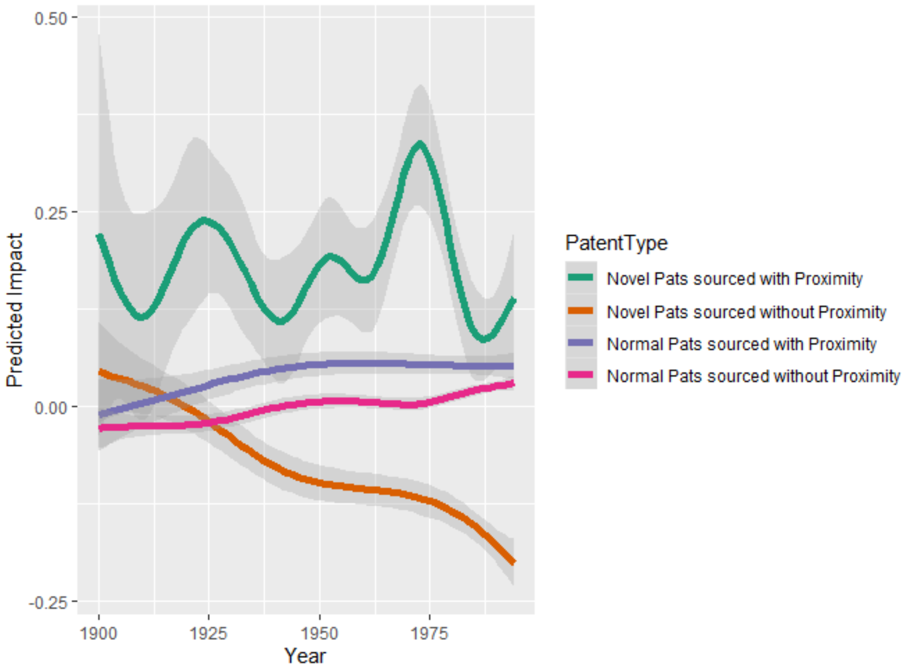

Note: Regression to estimate predicated impact is given in equation 4.

---

[7] I re-compute the average number of local knowledge sources on patents each each year, so in any given year half of all granted patents are defined as patents that source knowledge with proximity.

Figure 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 (Appendix C). 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 5's finding that breakthroughs were disproportionately produced by multi-locational teams toward the end of the 20th century, Figure 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 2. The predicted impact values, presented in Figure 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 8: Predicted Patent Impact by Novelty

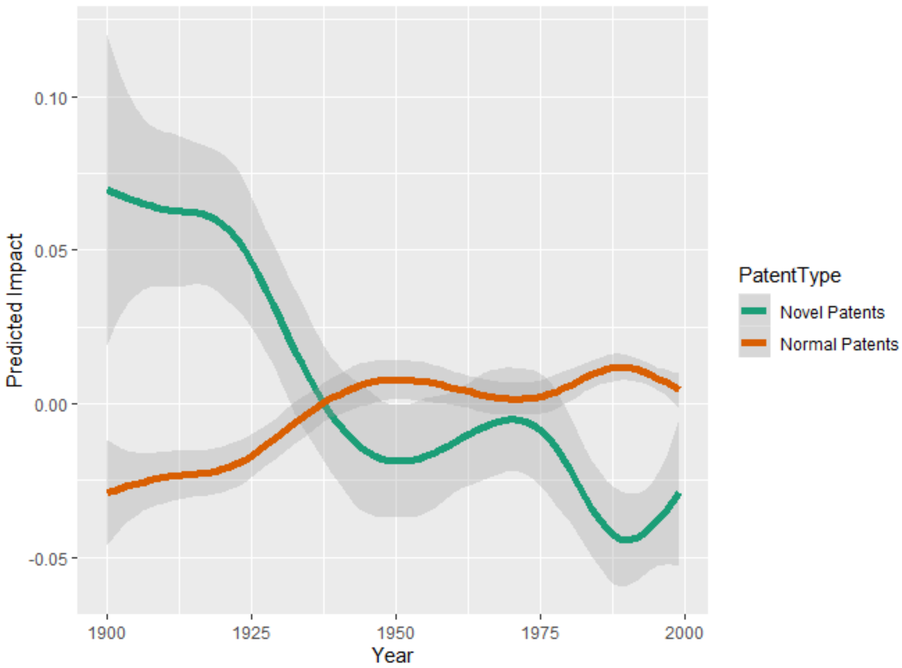

The decline in the impact of novel patents relative to normal patents between 1900 and 1950 indicates that technological change was less disruptive during the middle and end of the 20th century. The low level of disruptiveness at the end of the century seems at odds with the earlier finding that the average impact of novel patents invented in knowledge-diverse cities rebounded during the 1960s (Figure 5). One possible explanation to reconcile these two findings is that the novelty produced in knowledge-diverse at the end of the 20th century was qualitatively distinctive. Therefore, 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. In Figure 9, I plot the residuals of a model of patent impact that adds USPC primary class fixed effects to the normal control variables of Year*NrSubclasses fixed effects. The regression model is given by equation 7, where FE_{C438} designates fixed effects at the primary class level, at which scale there are 438 unique classes:

$$(7)\ Impact_p = Year_p * NrSubclasses_p + FE_{C438} + E_p$$

Figure 9: Predicted Patent Impact by Novelty with Aggregate Technology Class Fixed Effects

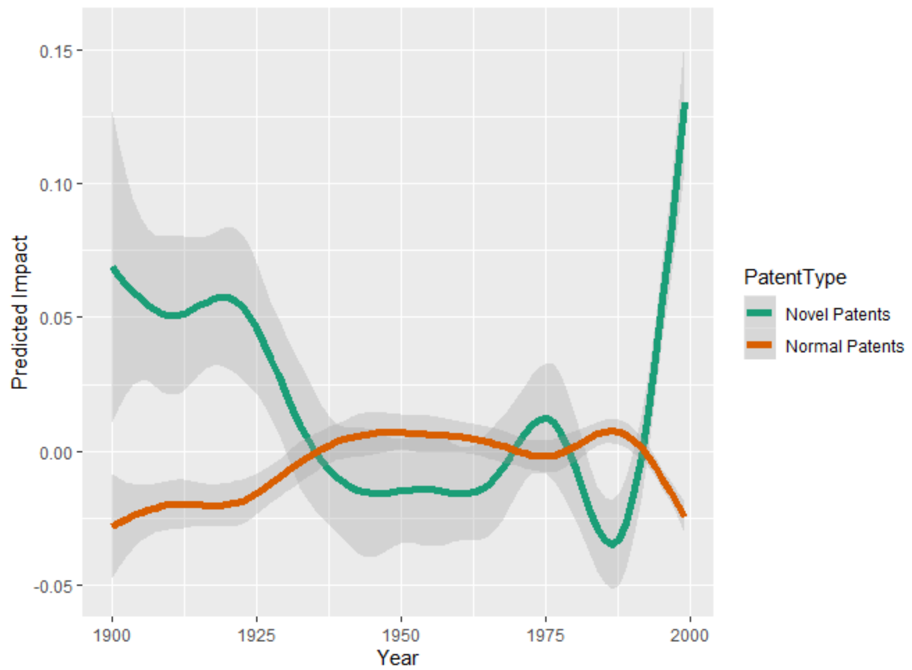

Note: The regression used to estimate predicated impact is given in Equation 7.

In Figure 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 9 is robust to the use of a more detailed course-grained subclass-specific fixed effect (Appendix D). Contrasting Figure 9 with Figure 8 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. Thus, we may conclude from Figure 9 that a narrow sector of the economy became disruptive between 1985 and 1999.

To conclude the analysis, in Table 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: Observed States of Factors of the Model and the Predicted Geography of Breakthrough Inventions

	Time Period		
Factor	**1900-1930**	**1930-1970**	**1970-1999**
Knowledge Intensity	Moderate	Moderate	High
Disruptiveness	High	Low	High *within* sectors
Long-Distance Collaboration Tech	Weak	Weak	Strong
Long-Distance Knowledge-Sourcing Tech	Weak	Weak	Weak
Predicted Geography of Breakthroughs by Model	Weakly concentrated	Undefined	Multi-Nuclei

Note: Empirical observation of the model factors are given in Section 6. The predicted geography of breakthrough inventions is presented in Table 3.

To explore the validity of the model, the predicted geographies of breakthrough innovation from Table 4 can be compared to the observed geographies documented in Figures 3 and Figure 5. Notably, the states of breakthrough innovation predicted in Table 4 closely correspond to the empirical distributions found in Figures 3 and 5. During the first third 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 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. The disruptiveness of the regime of technological change was low, the geography of breakthrough innovation was undefined. This proposition is confirmed in Figure 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 predicts a multi-nuclei geography of breakthrough innovation. The geography predicted by these parameters corresponds to the observed distribution described in Figure 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 by comparing Figure 8 with Figure 9, indicates that breakthroughs made in knowledge-diverse cities during the late 20th century were made were solved 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.

7. Discussion

The spatial concentration of innovation is not an inherent quality of density, agglomeration, or urbanization (c.f. Duranton and Puga, 2001; Bettencourt, et al. 2007; Mewes, 2019; Balland et al., 2020; Berkes and Gaetani, 2020). Instead, innovation organizes in concentrated, dispersed, or multi-nuclei spatial arrangements as a result of fundamental changes in institutions and communication technologies. These institutional and technological factors determine whether novelty will be rewarded by the economy, inventors' general need to interact in order to create impactful novelty, and the frictions involved in sustaining innovative 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 propose 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 factors and communication technologies 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 this literature, the spreading out of economic activities in the middle of the 20th century was in-part caused by improvements in communication technologies and decreases in transportation costs. 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 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. This reduction in the disruptiveness of the regime of technological change may have suppressed firms' demand for location in dense agglomerations, because 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; 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 (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. Therefore, any prediction 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 rents accrued by oligopolistic firms concentrate wealth in those their 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 Schumpeterian competition, meaning 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 factors is needed for researchers and policy makers to understand and rectify the new geographical and technological challenges that are bound to emerge.

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. "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. *Econ. Geog.* 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).

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

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

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(1). 31-56.

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*.

Baum-Snow, N., Gendrom-Carrier, N., and Pavan, R. (2020). Local Productivity Spillovers. Working Paper. <https://www.nicolasgendroncarrier.com/_pdf/localspillovers_20200506_paper_v2.pdf>

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 (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*.

Esposto, 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. Innovation in Cities: Science-Based Diversity, Spcialization, and Localized Competition (34). 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). 1019-1039.

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.

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.

Jones, B. (2009). The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder? *The Review of Economic Studies* 76(1). 283–317.

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. *PNAS* 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). 341-361.

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

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). 994-10017.

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(5827). 1036-1039

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

Appendix A: 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 their 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 A1: Average Impact of Multi-Locational Patents by Collaboration Type

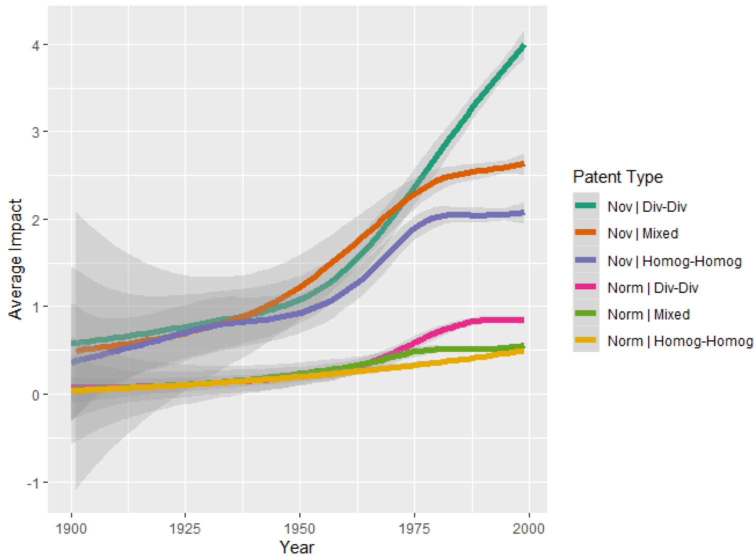

To compute the residual impact of inventions, I collect residuals from the following model and display them in Figure A2:

$$Impact_p = Year_p * NrSubClasses_p + E_p$$

Figure A2: Predicted Impact of Multi-Locational Patents by Collaboration Type,

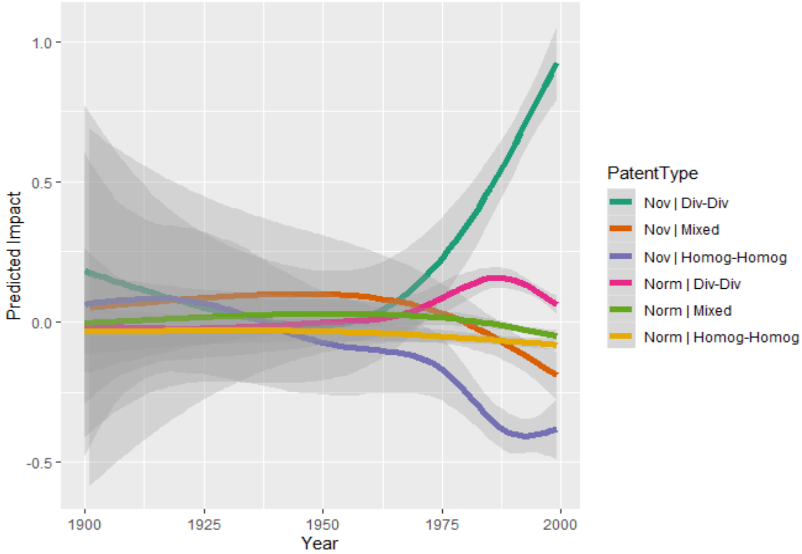

Appendix B: Knowledge Intensity and Impact of Patents

Figure B1 plots the raw patent impact of novel and normal patents with many and few patents by year.

Figure B1: Predicted Impact by Novelty and Knowledge Intensity

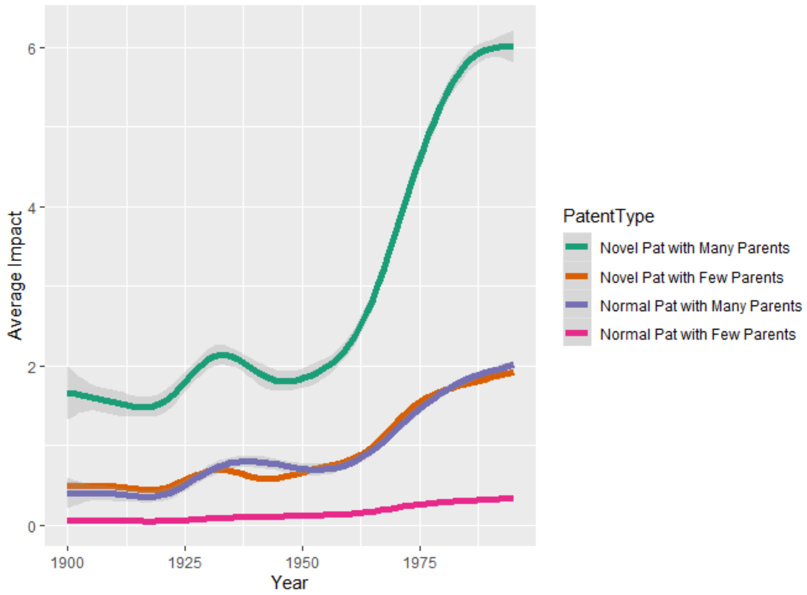

Figure B2 plots the residuals from the following model:

$$(4)\ Impact_p = Decade_p * NrSubClasses_p + PrimaryClassFE_p + E_p$$

Figure B2: Predicted Impact by Novelty and Knowledge Intensity

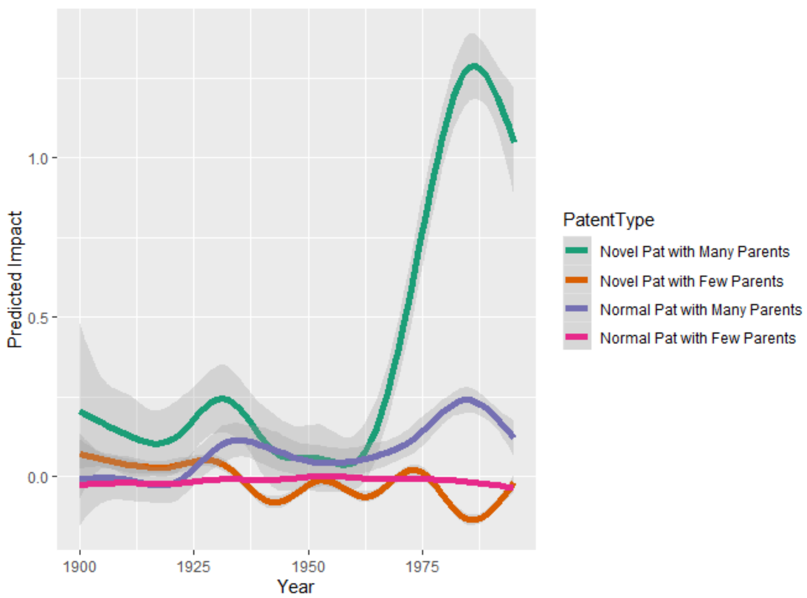

Appendix C: Residual Impact of Locally-Sourced and Non-Locally Sourced Patents,

Figure C 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 C: Residual Impact of Locally-Sourced Knowledge

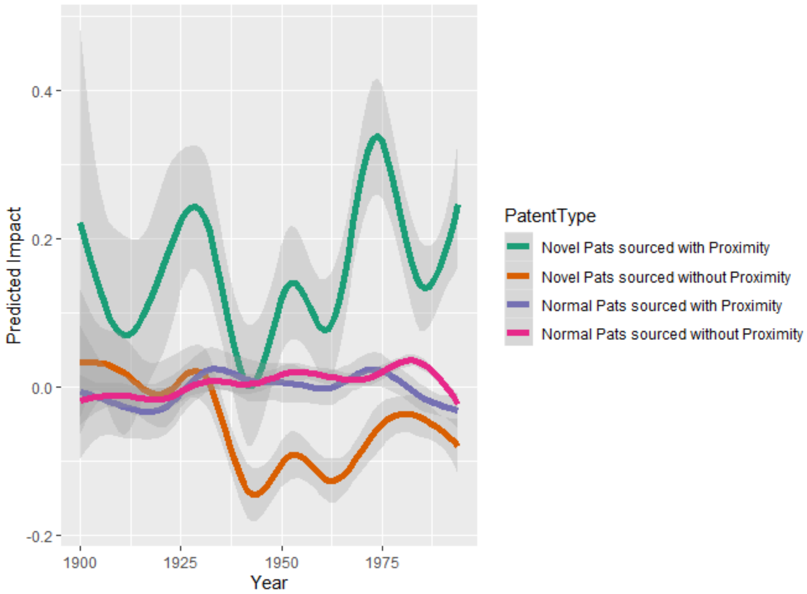

Appendix D: Disruptiveness of Regime of Technological Change

Figure D 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 D: Residual Impact of Novel and Normal Patents