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

Christopher Esposito

University of Chicago

Abstract

The geography of breakthrough invention in the U.S. – defined as the spatial distribution of the production of patents that are both novel and impactful – underwent three broad changes during the 20th century. At the start of the century, breakthrough invention was concentrated in populous metropolitan areas with high levels of local knowledge variety. 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 while their invention was once again concentrated in major metropolitan areas with high knowledge variety, they frequently involved long-distance collaboration. In this article, I document these changes and propose a theory to interpret why they occurred. The theory emphasizes how changes in inventors' institutional and communication technology environments influence the geographical locations that are advantageous for breakthrough invention. In support of the model, I find that the disruptiveness of the regime of technological change, the knowledge intensity of breakthroughs, the distance-based frictions incurred by collaboration technologies, and the distance-based frictions incurred by knowledge-sourcing technologies help to predict the spatial distribution of breakthrough invention. To conclude the article, I discuss lessons that the 20th century's geography of breakthrough innovation provide for anticipating 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, participants at the NARSC annual meeting, and the editor and four anonymous reviewers at Research Policy for their thoughtful feedback. In addition, I thank Frank van der Wouden for providing pre-1975 inventor location data.

1) Introduction

Breakthrough inventions are critically important to the economy, to organizations, and to policymakers. Breakthroughs are leading drivers of technological change (Anderson and Tushman, 1990) and to the competitive advantage of firms (MacGrath, et al, 1996). Recently, governments have started to view breakthroughs as instrumental for the achievement of policy goals, including the preservation of technological supremacy¹ and the accomplishment of societal grand challenges, from the combatting of infectious diseases to the mitigation of climate change (Mazzucato, 2021). A careful understanding of the types of institutional and geographical environments that are conducive to the creation of breakthroughs would help to advance the economic, organizational, and societal ambitions along all these fronts.

Economic geographers, urban economists, and innovation scientists have long sought to uncover the types of spatial environments that enhance creativity and promote invention. Past research has emphasized how agglomeration in regions with diverse stocks of circulating ideas increases the range of ideas that inventors can access and their propensity to invent (Jaffe, et al., 1993; Audrestsch and Feldman, 1996; Duranton and Puga, 2001; Mewes, 2019; Berkes and Gaetani, 2020; Moretti, 2021). Although important, this literature has two limitations. The first is that it faces an analytical puzzle: while inventive activity in the U.S. is currently concentrated in the country's large and dense metropolitan areas (Balland et al., 2020; Baum-Snow, et al. 2020), historically invention was prevalent in rural parts of the country (Perlman, 2016; Berkes and Nancka, 2021). Highly impactful inventions, such as airplane and the cotton gin, were developed in the countryside (Mokyr, 1992), and some inventions that are commonly regarded as urban creations, such as the Ford Model T, were developed through interactions between urban engineers and rural customer bases (Gordon, 2017). Even Silicon Valley, the leading cluster of invention in the United States, had a low population density during its formative – and potentially, most inventive – years (O'Mara, 2019). In addition to these historical fluctuations, the recent increase in the distance between patent co-inventors, which tripled between 1900 and 2015 (Van der Wouden, 2020; Clancy 2020), does not fit neatly into a vision of agglomeration as the primary driver of invention. The relationship between spatial density and invention is thus more complex than is often suggested.

The second limitation to the existing literature is that it focuses on invention in general, and not on breakthrough inventions in particular. The unique qualities of breakthrough inventions, as inventions that are both novel and highly impactful, suggest that their geography may deviate from the overall

¹ In 2022, the National Science Foundation in the U.S. established the Regional Innovation Engines program to "respond to the global competition for talent and leadership in science and technology". Source: https://beta.nsf.gov/funding/initiatives/regional-innovation-engines

geography of invention in important ways. The high quantity of knowledge needed to invent breakthroughs may make their invention more responsive to changes in the strength of long-distance communication technologies. In addition, the overall production of breakthroughs changes over time as the economy undergoes technological revolutions (Schumpeter, 1934). While overall invention may always have a geographical distribution, few breakthroughs are invented during non-revolutionary periods, so their production does not always assume a distinctive geography.

With these considerations in mind, this article has two objectives. The first is to systematically describe how the geography of breakthrough invention in the United States evolved over the 20th century. In this regard, the analysis benefits from two new datasets that make it possible to analyze the geography of breakthroughs over a long timeframe. The first dataset provides information on the residential location of pre-1975 U.S. patent inventors (van der Wouden, 2020). This dataset improves on existing historical inventor location records (Petralia, et al. 2017) by offering more precise information for multi-inventor patents. The second dataset are indicators on the impact of historical patents on subsequent invention (Esposito, 2022). Before 1947, patents did not make citations to prior art, so citation-based patent impact measures cannot be calculated for historical inventions (Akcigit, et al. 2017). The new dataset, introduced by Esposito (2022), uses subclasses on patents to trace knowledge flows between inventions and thus allows for the identification of breakthroughs invented before 1947.

In describing the changing geography of breakthrough invention, I study ordinal changes in the geographical distributions of breakthrough invention, which range from spatial concentration to spatial dispersion, as well as more complex geographies that can emerge from non-local collaboration. My analysis generates three findings. First, during the early 20th century, inventors residing in regions with large and varied stocks of local ideas were more likely than inventors residing in regions with smaller and less varied knowledge stocks to develop breakthroughs. Second, in the mid-20th century, inventors located in regions with large and varied knowledge stocks were no more likely than inventors in other regions to develop breakthroughs. Third, at the end of the 20th century, inventors that resided in regions with large and varied knowledge stocks *and* engaged in non-local collaborations were more likely than all other types of inventors to develop breakthroughs. These results demonstrate that the geography of breakthrough invention in the U.S. is complex and evolves over time.

The second objective of this paper is to propose an interpretation for why historical changes in the geography of breakthrough invention occurred. Toward this effort, I draw from three literatures. Two of those literatures, on technological regimes and the burden of knowledge, argue that inventors' knowledge environments influence how they invent and the types of inventions that they produce (Winter, 1984; Breschi, 2000; Breschi, Malerba, and Orsenigo, 2000; van Dijk, 2000; Jones, 2009; Diodato and Morrison, 2019; Fontana, Martinelli, and Nuvolari, 2021). The third literature, on localized

knowledge spillovers, argues that distance-based frictions associated with communicating complex and tacit knowledge cause inventive activity to concentrate spatially (Jaffe, et al. 1993; Audretsch and Feldman, 1996; Gertler, 2003; Storper and Venables, 2004). I bring together these sources to propose that four factors interacted to help produce the spatial distributions of breakthrough invention that emerged over the 20th century. These factors are disruptiveness of the regime of technological change, the knowledge intensity of breakthrough inventions, the distance-based frictions incurred by the technologies that inventors use to collaborate with colleagues, and the distance-based frictions incurred by the technologies that inventors use to source ideas from cities where they do not have active collaborators. The theory developed here is not intended to isolate causality. For example, some of the listed factors, such as the disruptiveness of a regime of technological change, are both causes and outcomes of the geography of breakthrough invention. Instead, the theory is intended to serve as an analytical tool that, when critically applied, can assist in interpreting the history of breakthrough invention in the U.S.

I begin in Section 2 by describing the process of invention and by showing how inventors' institutional and communication technology environments influence how they invent, their propensity to create breakthroughs, and the locations where they make breakthroughs. In Section 3, I introduce and explore the datasets used in the empirical analysis. In Section 4, I empirically examine how the geography of breakthrough invention changed in the U.S. over time. In Section 5, I study how inventors' institutional and communication technology environments changed over time, and I demonstrate how those changes are related to the evolution of the geography of breakthrough invention. Finally, in Section 6 I discuss how the findings of this analysis revise a common interpretation for why economic activities dispersed across space during the mid-20th century, and I share lessons that this historical revision generates for forecasting the future of the agglomeration of breakthrough innovation, including in the years after COVID-19.

2) Invention, Breakthroughs, and Location

A strong relationship currently prevails between agglomeration and invention, which suggests that the two are linked by a causal arrow. However, a detailed and historical perspective indicates that the relationship between agglomeration and invention is more complex. During the 18th and 19th centuries, inventions were frequently made in the rural regions of the U.S. (Mokyr, 1990; Gordon, 2017; Mewes, 2019; Balland et al., 2020). Although a big-city advantage for invention emerged at the start of the 20th century, it might not have lasted: while patent records suggest that invention remained concentrated in large cities in the middle of the 20th century (Bettencourt et al., 2007; Mewes, 2019; Berkes and Gaetani, 2020; Balland et al., 2021), employment records indicate that employment in the most innovative occupations spread out across space (Desmet and Rossi-Hansberg, 2009; Kemeny and Storper, 2020).

The contemporary big-city advantage for invention arrived by the end of the 20th century (Baum-Snow et al. 2020; Kemeny and Storper, 2020). In addition to these historical fluctuations, the recent increase in non-local collaboration also complicates the relationship between agglomeration and invention. The average distance between the co-inventors of patents tripled between 1900 and 2019, suggesting that permanent geographical co-location is no longer a necessary condition for inventors to create and maintain collaborative relationships (Van der Wouden, 2020; Clancy, 2020).

Together, the historical fluctuations in the relationships between agglomeration, non-local collaboration, and invention beg a question: why does the geography of invention change over time? Presumably, changes in the geography of invention are related to more fundamental changes in the process of invention. Inventors create new technologies by combining existing ideas into new forms (Schumpeter, 1936; Weitzman, 1997; Fleming and Sorenson, 2001). Past research has suggested that the way in which inventors carry out this general process is shaped by two main factors: inventors' competitive environments or *technological regimes* (Schumpeter, 1934; Schumpeter, 1942; Dosi, 1982; Winter, 1984; Breschi, 2000; Breschi, Malerba, and Orsenigo, 2000; van Dijk, 2000; Fontana, Martinelli, and Nuvolari, 2021), and inventors' long-distance communication technologies (Jaffe, et al., 1993; Leamer and Storper, 2001; Storper and Venables, 2004; Sorenson, Rivkin, and Fleming, 2006). These broader conditions influence the market structure of invention (Lamoreaux and Sokoloff, 1996; Lamoreaux, Sokoloff, and Sutthisphisal, 2013), the ways in which inventors undertake technological search, the types of technologies that inventors produce, and the types of locations where inventions are made.

Changes in inventors' technological regimes, and inventors' communication technologies also have the potential reshape the geography of breakthrough invention. Breakthrough inventions are distinguished from other inventions along two key dimensions (Grashof et al., 2019; De Noni and Belussi, 2021). First, breakthroughs are novel in that that they deviate from existing knowledge bases in highly imaginative ways. Second, breakthroughs are impactful in that they stimulate a large quantity of subsequent invention. Because breakthroughs must fit both criteria, they are exceedingly difficult to generate. The high impact of breakthroughs implies that inventors need to combine highly synergistic ideas to develop breakthroughs (Fleming and Sorenson, 2001). A very small percentage of the potential combinations that inventors can create are sufficiently synergistic, so inventors must search extensively and develop detailed expertise in order to identify the few that are (Youn, et al. 2016). Likewise, the novelty criteria of breakthroughs implies that inventors need to access unconventional ideas to create novelty (Mewes, 2019; Berkes and Gaetani, 2021). These combinations are only found at the "adjacent possible", or at the knowledge frontier, which is a relatively small region of complete technological search space (Kauffman, 1996). Therefore, novelty is generated more often by inventors that have access to larger, more varied, and more complex stocks of knowledge (Feldman and Audretsch, 1999;

Nieto and Santamaria, 2007; Mewes, 2019; Balland, et al. 2020; Solheim, Boschma, and Herstad, 20020; Antonelli, Crespi, and Quatraro, 2000; Bahar, Rapoport, and Turati, 2020; Lo Turco and Maggioni, 2020; Berkes and Gaetani, 2021).

Because of these difficulties, inventors benefit strongly from location in supportive institutional and geographical environments when creating breakthroughs. The literatures on technological regimes, the burden of knowledge, and the localization of knowledge spillovers have developed strong understandings of how inventors' environments affect their inventiveness. The technological regimes literature argues that inventors' technological-institutional environments vary in four important ways (Schumpeter, 1934; Schumpeter, 1942; Dosi, 1982; Winter, 1984; Breschi, 2000; Breschi, Malerba, and Orsenigo, 2000; van Dijk, 2000). The first is the availability of technological opportunities in a technological regime. In some knowledge-based institutional environments, new ideas are more plentiful than in others, which increases the overall productivity of inventors. The second factor is the ability for firms and organizations to appropriate the returns to invention. When appropriability is high, incumbent firms retain the technological leadership of their industries over time. The third factor is the nature of the knowledge base in a technological regime. This abstract factor has been interpreted to include the tacitness and the industrial specificity of the knowledge in a technological regime (Breschi, 2000; Breschi, Malerba, and Orsenigo, 2000). The fourth factor is the cumulativeness of technological advances in a regime. In cumulative technological regimes, new technologies build directly on prior ones in well-defined trajectories. By contrast, in disruptive regimes, technological change advances through the introduction of novelty that displaces incumbent technologies.

Of the four potential factors raised by the technological regimes literature, the cumulativeness of a technological regime is directly related to inventors' overall propensity to create breakthrough inventions. Because breakthroughs are novel and stimulate a large quantity of technological change, they can only be created in technological regimes where technological change advances from novel inventions (Fontana, et al., 2012). Historical studies suggest that the cumulativeness or disruptiveness of the technological regime in the U.S. changed over time. Schumpeter's pivot from his early view of technological change as a disruptive process (Schumpeter, 1911) to his later view that technological change is cumulative (Schumpeter, 1934) was likely induced by a decline in disruptiveness between his writings (Breschi, Malerba, and Orsenigo, 2000). The second industrial revolution, which occurred during the late 19th and early 20th century was a period of relative stability (Gordon, 2016). This period of stability was followed by the ICT revolution during the late 20th century, which led to another period of disruptive technological change (Storper and Kemeny, 2020). Therefore, the technological regimes literature predicts that inventors were in more advantageous positions to create breakthroughs if they worked during the beginning or end of the 20th century.

While the cumulativeness of a technological regime is related to the overall propensity for inventors to create breakthroughs, three other factors help to determine the geographical regions that are advantageous for their invention. The first factor is the knowledge intensity of invention. While all inventions build on existing ideas, extensive research has shown that the quantity of knowledge that inventors need to create useful inventions increased over the 20th century (Lamoreux and Sokoloff, 1996; Wuchty, Jones, and Uzzi, 2007; Jones, 2009). The increase in knowledge intensity, often called the "burden of knowledge", caused invention to concentrate in organizational forms that were able to access and assemble larger bodies of knowledge, including larger teams (Wuchty, Jones, and Uzzi, 2007) and denser inventor networks (Powell, Koput, and Smith-Doerr, 1996). As van der Wouden (2020) shows, the average number of collaborators on USPTO patents increased from 1.2 to 1.55 inventors between 1900 and 1975. Such increases in the intensity of collaboration will cause invention to concentrate is organizational form to create and maintain collaborative ties (Breschi and Lissoni, 2009; Balland, Boschma, and Frenken, 2015).

The rising knowledge intensity of invention also produced changes in the market structure of invention, which further facilitated transformations in the geographical distribution of invention. As Lamoreaux and Sokoloff (1996; 2005) show, inventors' increasing reliance on specialized knowledge at the turn of the 20th century caused them to focus on the development of new technologies as opposed to the commercialization of their technologies. To profit from their inventions, inventors sold their patents to incumbent firms, and these sales were facilitated by networks of agents and lawyers that were concentrated in the nation's largest cities (Lamoreux, Sokoloff, and Sutthisphisal, 2003). Because independent inventor-entrepreneurs accounted for half of the patents granted by the USPTO until the 1930s (Nicholas, 1930), one would expect that invention would be spatially concentrated during the first decades of the 20th century, and to have deconcentrated thereafter.

More generally, the increase in then knowledge intensity of invention over time may interact with spatial frictions in the spillover of knowledge and cause inventive activity to concentrate in space (Jaffe, et al. 1993; Audrestsch and Feldman, 1996a; Audrestsch and Feldman, 1996b; Sorenson, Rivkin, and Fleming, 2006; Balland, et al. 2020). The friction that distance exerts on the flow of knowledge is shaped by the long-distance communication technologies that inventors have at their disposal to access and assemble knowledge (Storper and Leamer, 2003). The efficacy of long-distance communication technologies depends on the extent to which they substitute for face-to-face communication. Face-to-face communication is exceedingly efficient because it allows for the use of body language, the manipulation of vocal tone, and the use of physical contexts to enrich messages (Storper and Venables, 2004; Gertler, 2003). Certainly, early in U.S. history long-distance communication technologies were poor substitutes for face-to-face communication. While primitive communication tools such as the mail post diffused across the U.S. during the 19th century and increased local patenting rates (Acemoglu,

Moscona, and Robinson, 2016), the materials that can be sent via the post (text and diagrams) represent a small percentage of the information that can be conveyed through face-to-face interaction and the physical demonstrations of technologies.

When long-distance communication technologies do improve, they improve asymmetrically, and become more effective for transmitting certain types of information than for others. In this regard, it is useful to distinguish between long-distance collaboration technologies and long-distance knowledge-sourcing technologies. Long-distance collaboration technologies are the devices that inventors use to collaborate with non-local colleagues, such as letters, email, videoconferencing, and long-distance travel. An improvement in these tools allows inventors to collaborate more often and more effectively with non-local colleagues. Long-distance knowledge-sourcing technologies are the tools that inventors use to source ideas from regions where they do not have active collaborators. These tools include written or digitalized scientific articles and patent documents. For example, Lamoreaux and Sokoloff (1996) describe how trade magazines started to "print complete lists of patents issued on a weekly basis, and provided readers with copies of patent specifications for a small fee" during the early 20th century (pp. 12,687). Such publications allowed inventors to build on ideas known by non-local inventors in the absence of collaboration, similar to the service now performed by Google Patents.

The four factors of the disruptiveness of the regime of technological change, the knowledge intensity of breakthroughs, the effectiveness of long-distance learning technologies, and the effectiveness of long-distance collaboration technologies interact to influence the geographical distributions of breakthrough invention. In influencing the geography of breakthrough invention, the disruptiveness of the regime of technological change is an overriding factor. When the regime is not disruptive (or cumulativeness is high), few breakthrough inventions are created across the entire economy. Thus, the geography of breakthroughs is undefined. The remaining three factors are highly interactive. When the knowledge intensity of breakthroughs is low, changes in the effectiveness of long-distance communication technologies have minimal impact the geography of breakthrough invention, because inventors do not need to source much knowledge to create breakthroughs. Therefore, regardless of the strength of communication technologies, breakthrough invention will disperse across space when their knowledge intensity is low. These conditions may describe the 19th century in the U.S., when invention was less knowledge-intensive (Jones, 2009; Balland et al. 2020), inventors generally worked as individuals rather than in teams (van der Wouden, 2019), and relatively modest attempts to diffuse invention spatially, such as rail extensions and the building of public libraries in rural areas, had an observable effect on regional patenting rates (Pearlman, 2016; Berkes and Nencka, 2021). By contrast, when the knowledge intensity of breakthroughs is high, even a small decrease in the efficacy of longdistance communication vis-a-vis face-to-face communication will put geographically isolated inventors at a serious competitive disadvantage in the creation of breakthroughs.

The most interesting geography of breakthroughs occurs when the knowledge intensity of breakthroughs is high, and the technologies inventors use to access non-local knowledge improve asymmetrically. If long-distance knowledge sourcing technologies are strong but long-distance collaboration technologies are weak, innovation disperses across space because inventors can easily source non-local ideas. If long-distance knowledge sourcing technologies are weak but long-distance collaboration technologies are strong, the geography of breakthroughs becomes multi-nodal, with core hubs of inventive activity connected by long-distance collaborative networks. This outcome arises because it allows inventors to combine more varied sets of ideas. By concentrating spatially in regions with large and varied knowledge stocks, inventors can learn a wide range of ideas. These inventors do not lose access to non-local ideas, because they can collaborate with non-local colleagues (Bathelt, et al. 2004).

These arguments are synthesized in Table 1. For brevity, I omit from the table changes in the disruptiveness of the technological regime. I omit this factor because the spatial distribution of breakthroughs is undefined whenever disruptiveness low.

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

 Table 1: Predicted Geography of Breakthroughs in Disruptive Regimes of Technological

 Change

Existing research has explored the geography of invention, novel invention, and high-impact invention, but there are no existing large-scale studies across long historical periods in the United States of the geography of breakthrough invention. Grashof et al. (2019) studied the creation of novel and impactful patents in Germany and found 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) studied 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 industries 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 variety of knowledge in regions is somewhat different, their results are consistent with the view that the propensity for inventors to create breakthroughs increases when they have access to a variety elemental 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 that 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 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 five studies of U.S. invention suggest that agglomeration economies are important for overall patenting, novel patenting, and high-

 $^{^{2}}$ 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.

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 is likely to contain important variations across time. Two historical studies analyze the geographical concentration of innovation in the U.S. over an extended 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), using a slightly different measure of novelty Mewes (2019) finds no difference between the agglomeration of overall patenting and the agglomeration of novel patenting. Again, neither study examines changes in the geographical concentration of inventions that are novel *and* impactful.

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, 2019). While the growing prevalence of non-local collaborations is well-documented (Fitjar and Rodriguez-Pose, 2013; 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) Methods Overview

I define breakthroughs inventions are the subset of inventions that are both novel and highly impactful (Grashof, et al., 2019; De Noni and Belussi, 2021). A simple typology of inventions that vary in terms of their novelty and their impact is provided in Table 2. The table highlights how the measurement of patent's novelty and impact are combined to identify breakthrough inventions.

	Low impact	High impact
Low novelty	(1) Unsuccessful conservative experiments	(2) Incremental advancements
High novelty	(3) Unsuccessful radical experiments	(4) Breakthroughs

Table 2: Typology of Inventions by Novelty and Impact

3.2) Measuring Patent Novelty

Past studies have identified novel inventions based on the extent to which they introduce new ideas or recombine existing ideas in new ways. For example, Uzzi et al. (2013) computes the atypicality of the knowledge combinations in scientific articles using z-scores, which calculate the extent to which each combination of knowledge units in each invention deviates from the combinations inventors have made in the past. Kim et al. (2016) and Mewes (2019) apply this method to the subclassification codes listed on patents, taking subclass codes as indicators of the knowledge components in each invention. Berkes and Gaetani (2020) compute z-scores using the citations made by patents to a similar effect. After measuring the atypicality of component combinations at the pairwise level between all ideas combined in an invention, the novelty of patents can be computed by aggregating the atypicality of component combinations to the level of patents.

I define novel inventions as patents that contain one or more atypical combinations of components. I calculate the atypicality of the components in each patented invention by calculating z-scores between the pairs of USPC subclasses on each USPTO utility patent granted between 1900 and 1999.⁴ These USPC subclasses describe the components in each invention and thus allow me to identify inventions that combine components in unanticipated ways (Fleming, 2001). Before computing z-scores, I follow the approach of Kim et al. (2016) and coarse-grain the subclasses that the USPTO assigns to patents to the most aggregate level, at which level there are 12,986 unique coarse-grained subclasses that appear on patents granted between 1900 and 1999. I coarse-grain subclasses because z-scores require a sufficiently extensive pre-history of patenting to accurately measure the mean frequency of the combination of any two subclasses. Coarse-graining subclasses thus increases the frequency that each subclass is observed.

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

Following Uzzi, et al. (2013), Kim, et al. (2016) and Mewes (2019), I compute the z-score of the combination of subclass i with subclass j on a patent using Equation 1:

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

In Equation 1, $o_{i,j}$ is the number of past co-occurrences of coarse-grained subclasses *i* and *j* on all previously granted patents. The term $u_{i,j}$ gives the expected number of past co-occurrences of coarse-grained 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 coarse-grained 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 coarse-grained subclasses are combined more frequently than expected given a random process, and negative when two coarse-grained 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. To aggregate the atypicality of subclass pairs to the patent level, I define novel patents as those that contain one or more atypical combination of subclasses. I define all patents which do not introduce an atypical combination of subclasses as "normal" patents.

A significant share (34.9%) of patents granted between 1900 and 1999 are assigned to just one coarsegrained USPC subclass. Because these single-subclass patents do not contain any subclass combinations, their novelty cannot be measured and the patents must be omitted them from the main study. In Appendix D, I analyze these single-subclass patents in greater detail to anticipate how their omission influences the study's results. The general takeaway from that analysis is that the omission of single-subclass patents may cause the study to *understate* the concentration of breakthrough invention in cities with high knowledge variety and the propensity for breakthroughs to be invented by multilocational teams.

3.3) Measuring Patent Impact

The second criteria 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 (Hall, Jaffe, and Trajtenberg, 2001). Esposito (2022) 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 not available for patents granted before 1947 (Akcigit, et al. 2017) and are only publicly available via the PatentsView database from 1975 onward. In contrast, the subclass codes used by Esposito's (2020) method to trace knowledge flows are available for all USPTO utility patents starting in 1836. Therefore, the subclass method allows for longer historical studies. The second advantage is that the subclass method uses the same basic data input to identify high-impact inventions and to identify novel inventions. Using a common data input for both measures results in a more harmonized study.

To compute the impact of individual patents on subsequent invention, I follow the method of Esposito (2022) to count the number of subsequent inventions that draw knowledge from each patent. I deviate slightly from the method of Esposito (2022) by using course-grained USPC subclassification codes instead of raw subclasses. The course-grained subclass codes allow me to use the same classification scheme across the entire analysis, as discussed above. Applying the method described by Esposito (2022) creates a graph of predicted the flow of knowledge between patents, represented by the adjacency matrix **G**. After producing **G**, I compute the impact of each patent by counting the number of patents that draw knowledge from each focal patent. Patent impact is thus calculated by taking the outdegree (d) of each patent (p), using $d_p^{out} = \sum_i G_{i,p}$, where i indexes patents. To remove right-truncation from the measure of patent impact, I omit edges in **G** that span patents that were granted more than 10 years apart. Because the adjacency matrix **G** contains patents granted up to 2014, this allows me to compute the impact of patents granted through the end of the 20th century without right-truncation.

In keeping with the method of Esposito (2022), I reduce the computational burden of the knowledgetracing algorithm by restricting the dataset to patents that have 8 or fewer subclasses. This data truncation is necessary because the number of queries needed to identify the knowledge-based descendants for each patent increases at a rate of $2^n - 1$, where *n* is the number of subclasses on a patent. Only 0.6% of all 1900-1999 patents have more than 8 coarse-grained subclasses, so this truncation affects few patents. However, 9+ sublcass patents may represent a disproportionate share of the high-impact and novel patents. Therefore, in Appendix D, I examine whether patents with more than 8 subclasses are disproportionately produced in cities with high knowledge variety. The insight from that analysis is that the omission of patents with more than 8 subclasses causes me to *understate* the spatial concentration of high-impact invention in high-variety cities and the propensity for them to be invented by multi-locational teams.

3.4) Measuring the Variety of Knowledge in Metropolitan Areas

To measure the knowledge variety of metropolitan areas and micropolitan areas (CBSAs), I link each patent to the CBSA in which it is invented. To do so, I use place-of-residence data provided van der Wouden (2020) for all U.S. inventors between 1900 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 CBSA definitions for this purpose. I measure knowledge variety at the CBSA level by counting the number of unique USPC coarse-grained subclassification codes assigned to the patents produced by inventors that reside in each (CBSA) in each year. This measure captures the *absolute variety* of local knowledge production, meaning that it is the total number of unique ideas that an inventor can access using face-to-face communication and without the need for travel.

To simplify the empirical analyses, I transform raw local knowledge variety values into a binary variable by defining high-variety 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 the same year. All CBSAs that do not meet the variety criterion are labeled as low-variety CBSAs. For example, in 1950 the USPTO assigned patents using 7,454 unique course-grain subclass codes, so in 1950 high-variety CBSAs were those that produced patents with at least 745 unique subclasses. In 1950, 11 CBSAs met the knowledge variety criterion.⁵ In Appendix C, I repeat the main analysis using a more lenient 5% threshold to define high-variety CBSAs. The results do not change substantially when I use a 5% threshold.

3.5) Data Description

I plot key trends using the data in Figure 1. Figure 1A shows the median and top and bottom deciles of the number of coarse-grained USPC subclasses assigned to patents by year. Figure 1A uses the full sample of all 2,586,656 USPTO utility patents granted to U.S.-based inventors between 1900 and 1999. The objective of Figure 1A is to illustrate how the removal of patents with 1 subclass and with 9+ subclasses may bias the sample. Both the median and bottom decile are equal to one at the start of the

⁵ In 1950, the high-variety CBSAs were (in descending order), New York, Chicago, Philadelphia, Pittsburgh, Detroit, Los Angeles, Boston, Cleveland, San Francisco, Milwaukee, and Bridgeport CT.

century, and the median increases to two subclasses in the 1930s while the bottom decile remains at one subclass. Therefore, the dropping of single-subclass patents results in the loss of a larger number of patents during the early 20th century and a somewhat smaller loss of patents toward the end of the century. The dropping of patents with 9+ subclasses result in relatively few patents lost, particularly during the start of the century.

Figure 1B plots the percentage of patents in the subsample that are novel by year. The data subsample contains the 1,669,512 patents assigned to U.S.-based inventors between 1900 and 1999 that contain 2-8 coarse-grained subclasses. In 1900, 29.1% of the patents in the subsample were novel. By 1999, that figure declined to 17.8%. This negative trend is similar to the one identified by Kim et al. (1996). It is driven by the increased propensity for inventors to make conventional combinations of subclasses, given the prior distribution of subclass combinations. It is tempting to interpret the decrease in the novel share as evidence of increasingly conservative technological search. However, novel inventions could become more important for driving technological change even if they are created less frequently. Therefore, the decrease in the novel share does not necessarily imply that technological change has become less radical or disruptive over time.

Figure 1C plots the median, top decile, and bottom decile of the patent impact distribution for the subsample. While the median and bottom decile of the distribution are equal to 0 throughout the study period, the top decile begins at 1 and increases to 4 by the end of the study period. This right-skewed distribution is common in patent impact measures and needs to be accounted for when modeling the impact of patents. Often, researchers take the natural logarithm of patent impact measures to account for skew (Bakker, 2017). However, because of the large number of zero impact patents in the dataset, I use the inverse hyperbolic sine (IHS) transformation to adjust for the skew in the patent impact distribution in the subsequent analyses. The IHS transformation approximates the natural logarithm but is defined for values of zero (see Tubiana, Miguelez, and Moreno 2022 for a recent application).

Figure 1D brings together the measures of patent impact, novelty, and CBSA knowledge variety by plotting the percentage of overall patents, high-impact patents, and novel patents that are produced in CBSAs with high knowledge variety. Interestingly, the production of novel patents (red line) in high-variety CBSAs is nearly identical to the concentration of overall patenting (blue line) in high-variety CBSAs. This result is similar to the finding of Mewes (2019) and suggests that high local knowledge variety does not spur inventors to create novel inventions more frequently. On the other hand, Figure 1D shows that high impact patents (orange line) are disproportionately produced in CBSAs with high local knowledge variety. The difference in the concentration of novel and high-impact patenting in high-variety CBSAs suggests that there is a complex relationship between local knowledge variety and local breakthrough invention. Moreover, these divergent relationships documented in Figure 1D motivate a

more detailed examination of whether patents that are novel *and* impactful are disproportionately produced in CBSAs with high local knowledge variety.



Note: Figure 1A uses the full sample of all USPTO utility patents granted to U.S-based inventors. Figures 1B-1D restrict the full sample to patents assigned between 2 and 8 subclass codes.

4) Results: The Geography of Breakthrough Innovation4.1) Results for Breakthrough Concentration in High-Variety Cities

I examine changes in the relationship between local knowledge variety and the invention of breakthroughs by examining if the average impact of novel patents rises with local knowledge variety. I focus on the increase in the impact of novel patents in high-variety cities because, as Figure 1D showed, there is no relationship between the knowledge variety in a CBSA and the probability that a locally-invented patent is novel. By testing the relationship between local knowledge variety and the impact of locally-produced novel patents, I am able to examine whether the novel patents that are produced in high-variety CBSAs are highly impactful and thus breakthroughs.

I begin in Figure 2A, where I plot changes in the average impact of four types of patents over time. Those patent types are novel patents invented in high-variety CBSAs (Nov | High-Variety), novel patents invented in low-variety CBSAs (Nov | Low-Variety), normal patents invented in high-variety CBSAs (Norm | High-Variety), and normal patents invented in low-variety CBSAs (Norm | Low-Variety). As discussed in Section 3, I apply an inverse hyperbolic sine (IHS) transformation to the raw patent impact values to reduce skew in the variable's distribution. Because patents can be produced by multiple co-inventors, the observation unit of the data plotted this chart, and in all subsequent analyses, are patent-inventors. For example, a patent invented by two co-inventors appears twice in the dataset. There are 9,720,476 patent-inventor observations in the subsample. This large number of observations makes it infeasible to create scatterplot, so I plot best-fit lines with 95% confidence intervals. The large number of observations also renders infeasible the most common method used to compute rolling average fit lines (LOESS regression), so I produce 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 generates three inferences. First, across all years, novel patents invented in CBSAs with high knowledge variety (Nov|High-Variety) were the most impactful patent type, followed by novel patents invented in CBSAs with low knowledge variety (Nov|Low-Variety) inventions. Second, the average impact of all types of inventions increased over time. Third, the increases in average impact were larger in high variety CBSAs: the impact of Nov|High-Variety patents increased relative to Nov|Low-Variety, and the impact of Normal|High-Variety) increased relative to Norm|Low-Variety.

The increase in the average impact of Nov|High-Variety patents relative to Nov|Low-Variety patents suggests that the invention of breakthrough patents increasingly concentrated in high-variety CBSAs over time. However, Figure 2 should be interpreted cautiously for two reasons. First, the large increases in average impact for all types of patents 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 of patents 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, in Figure 3A I adjust the impact of patents based on the year a patent is granted and the number of subclasses assigned to it. To adjust the impact of patents, I run a linear regression model where patent impact is regressed against a Year*NumberOfSubclasses factor variable. Because there are 100 years in the dataset (1900-1999) and 7 counts of subclasses assigned to patents (patents can be assigned between 2 and 8 subclasses in my dataset), the factor variable has 100*7=700 unique values. I collect the residuals from the regression and plot them against the grant year of patents. As discussed before, I apply the inverse hyperbolic sine (HIS) transformation to the dependent variable to reduce its skew. The regression model used to generate the Year*NumberOfSubclasses adjusted impact values is given by Equation 4:

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

In Equation 4, Γ is a vector of coefficients that captures the mean of $IHS(Impact_p)$ by Year*NrSubclasses. The residual term E_p thus contains patent impact values that have been adjusted for the number of subclasses and the grant year of patents. I aggregate these adjusted impact values given by E_p by the novelty of the patent and the knowledge variety of the CBSAs in which they are invented and plot their values in Figure 3A.

Figure 3A shows that the impact of patents with different levels of novelty and invented in cities with different levels of knowledge variety changed over 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. The non-overlapping 95% confidence intervals around each of these lines indicates that this difference was statistically significant. In addition, starting in 1915 novel patents invented in high-variety CBSAs were significantly more impactful than novel patents invented in low-variety CBSAs. The second period was 1930 to 1965, during which the adjusted impact of novel inventions declined relative to normal inventions. During the third period, from 1965 to 1999, the adjusted impact of novel

inventions made in high-variety CBSAs increased significantly above that of normal patents. In addition, the adjusted impact of novel patents invented in low-variety cities did not increase. This latter result shows that by the end of the 20th century, breakthrough innovation was concentrated in cities with a high local level of knowledge variety.

The patterns in Figure 3A could potentially be driven by differences across industries or cities in the propensity to create high-impact patents that are independent of the level of local knowledge variety in a city. For example, the USPC subclassification scheme may be more detailed for some industries than for others. The impact of patents granted in industries with more detailed scheme could have higher or lower impact than patents granted in other industries, strictly because of the density of the classification scheme in the focal industry. In addition, cities with economies that specialize in industries with denser USPC patent classification schemes may appear to have higher local knowledge variety as an artifact of the classification scheme. To account for these potential confounding factors, in Figure 3B I plot 3-way adjusted patent impact values, where patent impact is adjusted for a patent's industry, the city it is invented in, the number of subclasses on the patent, and the year it is granted. I use the primary USPC class assigned to patents as a proxy of a patent's industry. There are 437 unique USPC primary classes in the dataset. To calculate adjusted patent impact, I regress the impact of a patent invented by inventor *I* against three sets of factor variables: *Year_p* * *NrSubclasses_p*, *Year_p* * *PrimaryClass_p*, and *Year_p* * *CBSA_i* factor variables. The regression model used to perform this adjustment is described by Equation 5:

(5) $IHS(Impact_p) = \Gamma Year_p * NrSubclasses_p + \Theta Year_p * PrimaryClass_p + \Phi Year_p * CBSA_i + E_p$

In Figure 3B, I extract the residual term from Equation 5 and aggregate it based on the novelty of patents and the knowledge variety of the CBSAs where those patents are invented.



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

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

Figure 3B shows that the concentration of high-impact novel patenting in high-variety CBSAs is robust to the inclusion of industry and city-specific fixed effect controls. The one notable difference between Figure 3A and Figure 3B is that the average impact of novel patents produced in low-variety CBSAs is somewhat higher toward the end of the 20th century once industry and CBSA controls are included in the model.

4.2) Results for Breakthroughs and Non-Local Collaboration

While Figures 3A and 3B show that breakthrough innovation concentrated in high-variety CBSAs 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 over time (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, I compare the average impact of patents created by inventor-teams located in single CBSAs and in multiple CBSAs. In addition, I decompose inventor-teams based on the knowledge variety of their home cities by differentiating between multi-locational teams that reside in low-variety and low-variety CBSAs. To ease interpretation, I momentarily omit all inventor-teams with teammates that resided in

both high-variety and low-variety CBSAs (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, up until the 1950s, the average impact of novel patents produced by multilocational teams was not significantly different from that of novel patents produced by single-locational teams. However, after 1950 the impact of novel patents produced by multi-locational teams in high-variety CBSAs increased significantly beyond those produced by single-locational teams or teams in low-variety CBSAs.

The patterns observed in the raw data in Figure 4 could be driven by confounding changes in the propensity for the USPC to assign more subclasses to patents over time. Therefore, in Figure 5A I adjust the impact of patents based on the number of subclasses on patents and the year each patent was granted. To perform this adjustment, I collect the residuals from the regression model described by Equation 4. In addition, in Figure 5B I adjust patent impact for the industry (measured using primary USPC classes) and the CBSA a patent is produced in, using the regression model described by Equation 5. The latter adjustment controls for the average patent impact of each primary technology class in each year, and the average impact of patents produced in each CBSA in each year.



Figure 5: Impact of Collaborative Patents by Type of Collaboration

Note: Regressions to estimate adjusted impact are given in equations 4 and 5.

Figures 5A and 5B both show there were no statistically significant differences in the adjusted impact across patent type until 1950, after which the impact of novel patents produced by multi-locational teams in high-variety CBSAs increased well beyond that of the other types of patents. 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) Empirical Assessment of the Theoretical Model

How are the changes in the geography of breakthrough invention documented in Section 4 related to changes in inventors' institutional and communication technology environments? In this section, I review evidence from patent records to understand how 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 changed over the study period.

I begin by analyzing the disruptiveness of the regime of technological change. I assess the disruptiveness of the technological change regime that prevailed during historical time periods by comparing the average impact of novel patents relative to that of normal patents. As in the prior analyses, I control for changes in the impact of patents across decades and across patents with different numbers of subclasses, and for patents in different industries and invented in different CBSAs. In Figure 6, I plot the adjusted impact of patents by novelty and knowledge intensity. In Figure 6A, I adjust for

the year and number of subclasses on a patent, as described by Equation 4, to conduct this adjustment. In Figure 6B, I additionally adjust for the primary class and the CBSA of a patent, as described by Equation 5. For brevity, I plot raw patent impact values, broken out by patents' novelty and knowledge intensity, in Appendix B.



Figure 6: Patent Impact by Novelty

Note: Patent impact is adjusted using Equations 4 and 5

Figure 6 shows that novel patents were more impactful than normal patents at the start and end of the 20th century, but were no more impactful than normal patents during the middle of the 20th century. The 95% confidence intervals around the mean values are narrow, indicating that the differences at the start and end of the 20th century were statistically significant. These patterns are found in both Figures 6A and 6B, indicating that the patterns are not driven by industry or city-specific factors. These results suggest that breakthrough inventions came in two waves during the 20th century. The first wave, which subsided during the 1920s, coincides with the second industrial revolution (Gordon, 2017). The second wave, which began in the 1970s, coincides with the IT revolution (Kemeny and Storper, 2020)

Next, I study 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 test whether the average impact of patents that draw knowledge from a larger number of knowledge-based parents increased relative to average impact of patents that draw knowledge from fewer knowledge-based parents. 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. This measure of knowledge intensity is the same as the "tree size" measure used by Jones

(2009), except that Jones (2009) uses citations to identify patent's knowledge-based parents while I use the graph of knowledge flows produced by tracing patent subclasses between patents.

To simplify the analysis, I transform the number of knowledge sources used to make 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 before, I use the regressions described by Equations 4 and 5 to make these adjustments. I plot adjusted patent impact by patent novelty and knowledge intensity in Figures 7A and 7B. Raw patent impact values, broken out by patent novelty, are shown in Appendix B.



Figure 7: Patent Impact by Novelty and Number of Patent Parents

Note: Patent impact is adjusted using Equations 4 and 5

Figure 7A shows that the adjusted impact of novel patents with many knowledge-based parents patents was greater than that of novel patents with few knowledge-based parents during the full study period. This difference in impact increased significantly starting in 1965. The 95% confidence interval is very narrow, indicating that the rise in the adjusted impact of patents with many parent patents at the end of the 20th century was statistically significant. These relationships are also found in Figure 7B, indicating that city-specific and industry-specific factors do not explain the difference in impact and the trend. Therefore, I conclude that knowledge intensity of breakthroughs was moderate until 1965 but very high after 1965.

Next, I investigate changes in the strength of long-distance communication technologies. As discussed earlier, long-distance communication technologies can be categorized into two groups: long-distance

collaboration technologies, and long-distance knowledge-sourcing technologies. I measure the strength of each type of long-distance communication technology by the revealed ability for inventors to create high-impact novelty while collaborating with distant teammates or while sourcing knowledge from non-local CBSAs. Figure 5 in Section 5 presented suggestive evidence that long-distance collaborative technology was weak before 1950 but grew stronger thereafter. In particular, the average impact of novel patents invented by multi-locational teams in knowledge diverse cities climbed significantly above that of novel patents invented by single-location teams in the second half of the 20th century.

Finally, 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. I define a focal patent as one that sources knowledge locally if an above-average share of its parent patents were developed in CBSAs in which the inventors of the focal patent reside.⁶ 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 described by Equation 4. I plot the adjusted values, aggregated by patent novelty and the localness of its knowledge sources in Figure 8A. I adjust for the industry of a patent and the CBSAs of its co-inventors using Equation 5 and plot those values in Figure 8B. In Appendix B, I present similar results using unadjusted patent impact values.



Figure 8: Patent Impact by Novelty and Local Knowledge Sourcing

Note: Patent impact is adjusted using Equations 4 and 5

⁶ I re-compute the average number of local knowledge sources on patents each year, so in any given year half of all granted patents are defined as patents that source knowledge with proximity.

Figure 8A 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, as indicated by their non-overlapping 95% intervals. Figure 8B shows that this relationship is robust to the inclusion of industry and CBSA fixed effects, although the adjusted impact of novel patents from non-local sources does trend upward toward the end of the 20th century in the latter figure. This increase at the end of the 20th century suggests that long-distance knowledge-sourcing technologies improved between 1965 and 1999, but fell short of becoming perfect substitutes for face-to-face communication. 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 part as a response to the difficulty for inventors to source knowledge from regions where they do not have collaborators.

To conclude the analysis, in Table 3 I summarize the observed state of the disruptiveness of the regime of technological change, the knowledge-intensity of breakthrough innovation, the state of collaborative technologies, and the state of knowledge-sourcing technologies to generate the predicted geographies of breakthrough innovation for the early, mid, and late 20th century.

	Time Period		
Factor	1900-1930	1930-1970	1970-1999
Disruptiveness of regime of technological change	High	Low	High
Knowledge intensity of breakthroughs	Moderate	Moderate	High
Effectiveness of long-distance collaboration technology	Weak	Weak	Strong
Effectiveness of long-distance knowledge-sourcing technology	Weak	Weak	Moderate
Predicted geography of breakthrough invention	Weakly concentrated	Undefined	Multi-Nuclei

 Table 3: Observed States of Factors of the Model and the Predicted Geography of

 Breakthrough Inventions

The predicted geographies of breakthrough innovation from Table 3 can be compared to the observed geographies documented in Figures 3 and Figure 5 to test the accuracy of the model. Notably, the states of breakthrough innovation predicted in Table 3 closely correspond to the empirical distributions found

in Figures 3 and 5. During the first part of the 20th century, the weak long-distance collaboration and knowledge-sourcing technologies, the high level of technological disruptiveness, and moderate knowledge intensity of breakthroughs predict a weakly concentrated geography of breakthroughs. Figure 3 bears out this prediction by showing that the adjusted impact of novel patents was slightly higher for patents invented in knowledge-diverse cities than for patents invented in knowledgehomogeneous cities. During the mid-20th century (approximately 1930-1970), long-distance collaboration and knowledge-sourcing technologies were weak 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 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 20^{th} 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 predicts a multi-nuclei geography of breakthrough innovation. The geography predicted by these parameters corresponds to the observed distribution described by Figure 5, where high-impact novelty was shown to be produced by multi-location teams with coinventors residing in multiple knowledge-diverse cities.

6. Discussion

Breakthrough invention is not an inherent outcome of agglomeration, but rather organizes in concentrated, dispersed, or multi-nuclei spatial arrangements because of broader changes to inventors' institutional and communication-technology environments. These factors influence whether technological change is driven by normal or novel inventions, the quantity of knowledge sources that inventors need to create novel and impactful inventions, and the frictions involved in sourcing knowledge and sustaining collaborations across distance.

The two goals of this paper were to document changes in the spatial distribution of breakthrough invention in the United States evolved over the 20th century and to propose an interpretation for why those changes occurred. To this end, the paper began by describing how the advantages of local knowledge variety and multi-locational collaboration changed over time. Thereafter, a theory was developed in which breakthroughs are generated in different technological regimes and with different distance-based frictions for communication. Finally, the paper showed how the geographical distribution of breakthrough invention predicted by the theory closely aligned with the observed distributions in the United States over the 20th century.

In addition to the geography of breakthrough invention, explicit recognition of the institutional factors and communication technologies that shape the process of invention can inform our understanding of how and why the geography of economic activity evolves over time. 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, 2008; Glaeser, 2008). According to these classical sources, the spreading out of economic activities in the middle of the 20th century was caused by a decline in transportation costs. While a decline in transportation costs may have helped to facilitate geographic dispersal, the results from this study suggest that an additional factor, a reduction in the disruptiveness of the regime of technological change, was also important. As documented in Figure 6, fewer breakthroughs were invented during the mid-20th century. This reduction in the disruptiveness 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; Lin, 2012; Frank, et al. 2018; 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 high-skilled 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). Thus, there is now widespread interests in the possibility that inventive activities will remain dispersed in the future. The findings of this study indicate that there is no historical precedent from the 20th century in which breakthrough invention was carried out in geographically-dispersed environments. The only time period that breakthrough invention was not concentrated in large CBSAs was during the mid-20th century, because fewer breakthroughs were invented altogether during that period. While it is theoretically possible that long-distance communication technologies will eventually become perfect substitutes for face-to-face communication, such extrapolations step outside the historical record.

Besides communication costs, this study emphasizes that the level of technological disruption is an important driver of the spatial concentration of breakthrough invention. Therefore, predictions of the post-COVID-19 geography of breakthroughs should pay careful attention to a possible decline in technological disruption. Notably, market concentration in firms in the United States has reached its highest level 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 disruption as the technologies and routines of incumbent firms are less frequently displaced by new product or process technologies. If technological change is increasingly advanced through incremental inventions, then technological uncertainty may also decline. As a result, economic activity could disperse across space

not because of a reduction in communication costs, but because of the reduction in technological disruption.

To conclude, 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 of inventors 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

Acemolgu, D., Moscona, J., and Robinson, J. (2016). State Capacity and American Technology: Evidence from the Nineteenth Century. *AER: Papers & Proceedings* 106(5). 61-67.

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

Akcigit, U., Grigsby, J., and Nicholas, T. (2017). The Rise of American Ingenuity: Innovation and Inventors of the Golden Age. NBER Working Paper #23047.

Anderson, P., and Tushman, M. (1990). Technological Discontinuities and Dominant Designs: A Cyclical Model of Technological Change. *Administrative Science Quarterly* 35(4). 604-633.

Antonelli, C., Crespi, F., and Quartraro, F. (2020). Knowledge Complexity and the Mechanisms of Knowledge Geneartion and Exploitation: The European Evidence. *Research Policy*. In press.

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

Audrestsch, D., and Feldman, M. (1996a). Innovative Clusters and the Industry Life Cycle. *Review of Industrial Organization* 11(2). 253-273.

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

Autor, D., Dorn, D., Katz, L., Patterson, C., and Van Reenen, J. (2017). Concentrating on the Fall of the Labor Share. *American Economic Review: Papers and Proceedings* 107(5). 180-185.

Bahar, D., Rapoport, H., and Turati, R. (2020). Birthplace Diversity and Economic Complexity: Cross-Country Evidence. *Research Policy*. In press.

Bakker, J. (2017). The Log-Linear Relation Between Patent Citations and Patent Value. *Scientometrics* 110(2). 879-892.

Balland, P.A., Boschma, R., and Frenken, K. (2015). Proximity and Innovation: From Statics to Dynamics. *Regional Studies* 49(6). 907-920.

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(1). 31-56. Baum-Snow, N., Gendron-Carrier, N., and Pavan, R. (2021). Local Productivity Spillovers. Working paper.

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.

Berkes, E., and Nencka, P. (2021). Knowledge Access: The Effects of Carnegie Libraries on Innovation. SSRN # 3629299.

Breschi, S. (2000). The Geography of Innovation: A Cross-Sector Analysis. *Regional Studies* 34(3). 213-229.

Breschi, S., Malerba, F., and Orsenigo, L. (2000). Technological Regimes and Schumpeterian Patterns of Innovation. *The Economic Journal* 110. 338-410.

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.

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

De Noni, I., and Belussi, F. (2021). Breakthrough Invention Performance of Multispecialized Clustered Regions in Europe. *Economic Geography* 97(2). 164-186.

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*.

Diodato, D., and Morrison, A. (2019). Technological Regimes and the Geography of Innovation: A Long-Run Perspective on U.S. Inventions. PEEG Working Paper #19.24.

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. (2022). The Emergence of Knowledge Production in New Places. Forthcoming at *Journal of Economic Geography*.

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.

Fitjar, R., and Rodriguez-Pose, A. (2011). Firm Collaboration and Modes of Innovation in Norway. *Research Policy* 42(1). 128-138.

Fontana, R., Nuvolari, A., Shimizu, H., and Vezzulli, A. (2012). Schumpeterian Patterns of Innovation and the Sources of Breakthrough Inventions: Evidence from a Data-Set of R&D Awards. *Jorunals of Evolutionary Economics* 22. 785-810.

Fontana, R., Martinelli, A., and Nuvolari, A. (2021). Regimes Reloaded! A Reappraisal of Schumpeterian Patterns of Innovation, 1977-2011. *Journal of Evolutionary Economics* 31. 1495-1519.

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

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.

Frank, M., Sun, L., Cebrian, M., Youn, H., and Rahwan, I. (2018). Small Cities face Greater Impact from Automation. *Journal of the Royal Society Interface* 15(139).

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

Gordon, R. (2017). 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).

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

Lamoreaux, N., and Sokoloff, K. (2001). Market Trade in Patents and the Rise of a Class of Specialized Inventors in the 19th Century United States. *AEA Papers and Proceedings* 91(2). 39-44.

Lamoreaux, N., and Sokoloff, K. (2005). The Decline of the Independent Inventor: A Schumpetarian Story. NBER Working Paper 11654.

Lamoreaux, N., Sokoloff, K., and Sutthisphisal, D. (2013). Patent Alchemy: The Market for Technology in U.S. History. *Business History Review* Spring 87(1). 3-38.

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

Lo Turco, A., and Maggioni, D. (2020). The Knowledge and Skill Content of Production Complexity. *Research Policy*. In press.

Mazzucato, M. (2021). *Mission Economy: A Moonshot Guide to Changing Capitalism*. Harper Business. New York.

McGrath, R., Tsai, M-H., Venkataraman, S., and MacMillan, I.C. (1996). Innovation, Competitive Advantage, and Rent: A Model and Test. *Management Science* 42(3). 389-403.

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.

Moretti, E. (2021). The Effect of High-Tech Clusters on the Productivity of Top Inventors. *American Economic Review* 111(10). 3328-3375.

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

Perlman, E. (2016). *Connecting the Periphery: Three Papers on the Developments Caused by Spreading Transportation and Information Networks in the Nineteenth Century United States.* Doctoral Dissertation, Boston University. ProQuest Dissertations Publishing.

Petralia, S., Balland, P.A., and Rigby, D. (2016). Unveiling the Geography of Historical Patents in the United States from 1836 to 1975. *Scientific Data* 3(1).

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.

Roback, J. (1982). Wages, Rents, and the Quality of Life. *Journal of Political Economy* 90(6). 1257-1278.

Rosen, S. (1979). Wage-Based Indexes of Urban Quality of Life. In P. Mieszkowski and M. Straszheim, Eds., *Current Issues in Urban Economics*. John Hopkins University Press. Baltimore. 74-104.

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.

Solheim, M., Boschma, R., and Herstad, S. (2020). Collected Worker Experiences and the Novelty Content of Innovation. *Research Policy* 49(1). 1038-56.

Sorenson, O., Rivkin, J., and Fleming, L. (2006). Complexity, Networks, and Knowledge Flow. *Research Policy* 35(7). 994-10017.

Tubiana, M., Miguelez, E., and Moreno, R. (2022). In Knowledge we Trust: Learning-by-Interacting and the Productivity of Inventors. *Research Policy* 51(1). 1043-88.

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).

van Dijk, M. (2000). Technological Regimes and Industrial Dynamics: The Evidence from Dutch Manufacturing. *Industrial and Corporate Change* 9(2). 173-194.

Winter, S. (1984). Schumpeterian Competition in Alternative Technological Regimes. *Journal of Economic Behavior and Organization*. 5(3). 287-320.

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 Types

The following figure examines 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). To compute the residual impact of inventions, I collect residuals from the following regression model and display them in Figure A:

 $Impact_p = Year_p * NrSubClasses_p + E_p$



Figure A: Adjusted Impact of Multi-Locational Patents by Collaboration Type,

Appendix B: Analysis of Model Parameters using Unadjusted Patent Impact

The following figures replicate the analysis in Section 4 but use unadjusted patent impact.



Figure B1: Raw Patent Impact by Novelty

Figure B2: Raw Patent Impact by Novelty and Number of Patent Parents



Figure B3: Raw Patent Impact by Novelty and Proximity of Knowledge Sources



Appendix C: Analysis using 5% Threshold for High-Variety CBSAs

This analysis looks at the impact of patents produced in high-variety CBSAs, using a 5% cutoff value to define high knowledge variety. The plotted values are adjusted for the number of subclasses on patents and the year the patents are granted, using the following regression model:

 $IHS(Impact)_p = Year_p * NrSubClasses_p + E_p$



Figure C: Impact by Local Knowledge Variety using 5% Threshold

Appendix D: The Geography of Single-Subclass and 9+ Subclass Patenting

In this appendix, I test whether patents with one subclass or 9+ subclass are more likely to be invented in CBSAs with high knowledge variety or by multilocational teams. I perform this analysis because patents with 1 subclass and patents with 9+ subclasses were omitted from the main analysis due to data construction limitations discussed in Section 2. In addition, it is plausible that patents with 1 subclass are less novel and lower-impact than other patents, and that patents with 9+ subclasses are more novel and higher-impact than other patents.

To identify the geography of the invention of single-subclass and 9+ subclass invention, I administer two regression models. The first model is given by equation D1:

D1)
$$HighVarietyCBSA_p = B_0 + B_1OneSubclass_p + B_2GreaterThan8Subclasses_p + E_p$$

In the model, $HighVarietyCBSA_p$ is a binary variable that equals 1 if the CBSA in which the inventor *i* of patent *p* is located is above the 10% knowledge variety threshold, as described in Section 3, *Onesubclass_p* is a dummy variable that equals 1 if a patent has a single subclass, and *GreaterThan8Subclasses_p* if a patent has more than 8 subclasses. The model intercept, B_0 , represents the probability that an inventor that co-invents a patent with 2-8 subclasses is in a high-variety CBSA. I estimate the model using a logistic regression.

The second model I administer tests for a correlation between the number of subclasses on a patent and its propensity to be invented by a multilocational team of co-inventors. The regression model is given by equation D2:

D2) $MultiLocational_p = B_0 + B_1OneSubclass_p + B_2GreaterThan8Subclasses_p + E_p$

In the model, $MultiLocational_p$ is a dummy variable that equals 1 if the co-inventors of patent p reside in more than one CBSA. The model otherwise resembles Equation D1. As with Equation D1, I estimate D2 using a logit model.

In addition to the regressions represented by equations D1 and D2, I also estimate these models with the inclusion of year fixed effects. These yearly fixed effects are particularly important for the estimation of the model of multilocational teams, because the propensity for inventors to collaborate in multilocational teams and the average number of subclasses on patents both increased over time. Results for the regressions are presented in Table D.

	Dependent Variable: HighVarietyCBSA _p		Dependent Variable: MultiLocational _p	
Intercept	0.336*** (0.000710)		-1.64*** (0.000949)	
Single subclass	-0.04888*** (0.00191)	-0.0559*** (0.00196)	-0.391*** (0.00290)	-0.0630*** (0.00310)
Greater than 8 subclasses	0.152*** (0.00410)	0.165*** (0.00413)	0.181*** (0.00511)	0.177*** (0.177)
Yearly fixed effects		Yes		Yes

Table D: Regression Results for the Geography of Single-Subclass and 9+ Subclass Patenting

Note: 9,720,476 observations in each regression model

The results in Table D show that single subclass patents are less likely to be invented by inventors in CBSAs with high local knowledge variety and by multilocational inventor teams. On the other hand, patents with greater than 8 subclasses are more likely to be invented in high-variety CBSAs and by multi-locational teams. Because patent impact and novelty increase with the number of subclasses on patents, the results in Table D suggest that the omission of patents with one subclass and more than 8 subclasses in the main analysis may understate the extent to which breakthroughs are invented in high-variety CBSAs and by multilocational teams of inventors.