**Cycles of Regional Innovative Growth**

Christopher R. Esposito
Mansueto Institute for Urban Innovation
Department of Sociology
University of Chicago

Accepted for publication in *Journal of Economic Geography*
Decisioned May 14, 2022

**Abstract:** For academics and policymakers invested in regional economic development, two pertinent questions are how innovative city-regions rise and whether it is inevitable that innovative city-regions will fall. Using data from 8 million patents granted to U.S.-based inventors between 1850 and 1999, this study describes a general process that city-regions undergo as innovation begins, expands, declines, and (sometimes) resurges in regions. The results of the study show that inventors experiment with a small number of promising, diverse, and non-local ideas in the years before innovation in their home regions begins to grow, that inventors build on early locally-introduced ideas as innovation in their home regions expands, and that inventors experiment with relatively homogeneous sets of ideas shortly before innovation in their home regions declines. The results also show that declining U.S. city-regions rarely experience second waves of local innovative growth. However, when they do experience second waves, those waves are anticipated by changes in the knowledge sourcing strategies of local inventors. In particular, the years leading up to second cycles of regional innovative growth, local inventors experiment with promising, diverse, and non-local ideas.

**Key words:** Regional development, innovation, regional lifecycle, agglomeration

**JEL codes:** O33, R12

**Acknowledgements:** I thank David Rigby, Michael Storper, Allen Scott, Jacob Foster, Dieter Kogler, participants at Seminars in Economic Geography seminar series, and the editor and three anonymous reviewers at *Journal of Economic Geography* for helpful comments and feedback.

## 1) Introduction

Prosperous city-regions derive their fortunes from the innovative activities that agglomerate in their vicinities. Agglomeration increases local productivity and attracts inflows of skilled labor and knowledge, which gives the prosperity found in regions with strong agglomerations a sense of permanence, as though they have always been and always will be centers for innovation. Yet innovative city-regions rise and fall. The San Jose–Sunnyvale–Santa Clara Metropolitan Area, for example, is the economic core of California's Silicon Valley and is a remarkable case of innovative growth. The region's inventors produced less than 1% of the United States' patents in 1950 but over 8% in 2000. The Detroit Metropolitan Area, on the other hand, exemplifies innovative collapse. Metropolitan Detroit's inventors produced 4.5% of U.S. patents in 1940 but only 2.5% in 2000.

Although the economies of Silicon Valley and Metropolitan Detroit are currently heading in opposite directions, the two regions may be undergoing the same general process of economic development. As an economic geography literature argues, innovative growth in regions may advance through a three-stage process called the regional lifecycle. In the regional lifecycle model, innovation first begins in regions, then expands, and eventually declines (Audretsch and Fritsch, 2002; Falk and Heblich, 2008; Audretsch et al., 2008). From the perspective of the classical regional lifecycle model, meaningful differences in the economic conditions of regions are mostly results of the different stages those regions are at in their growth cycles. For example, innovation in Silicon Valley may collapse, as it did in Detroit. In a critique of the classical lifecycle model, Martin and Sunley (2011) argue that regional development is not a deterministic process, but one shaped by interdependencies that occur within regions, within the industries in which regions specialize, and at the region-industry nexus. Martin and Sunley (2011) propose an alternative "adaptive cycle model" in which regional innovative growth can unfold in a variety of directions as local industries experience renewal, are replaced by new innovative local industries, or decline without replacement by new innovative local industries.

The main purpose of this paper is to test the validity of the classical and adaptive regional lifecycle models. To carry out this test, I examine the dynamics of invention in regions and the inventors that reside in those regions. At the regional level, I study the average growth trajectory of the inventive output of regions and the extent to which regions deviate from the average growth trajectory. In addition, I subset the regions that significantly deviate from the average trajectory and carefully study the dynamics of these cases. At the inventor level, I investigate the micro-level behavior of the inventors in that reside in inventive regions. I carry out this inventor-level analysis because the classical and adaptive regional lifecycle models argue that the initiation, expansion, and decline of innovation in regions result from micro-level behavioral changes on the part of inventors (Porter, 1996; Audretsch et al., 2008; Martin and Sunley, 2011).

The paper has a second related purpose: to generate and test an explanation for how regional lifecycles begin. Currently, neither the classical nor the adaptive lifecycle model provides a satisfactory explanation for the beginning of regional lifecycles. To be sure, it is difficult to explain how inventors initiate knowledge production in new places. Inventors create new ideas by recombining existing ones (Nelson and Winter, 1982; Romer, 1990), and inventors are better able to source ideas from other actors when they are located in close physical proximity or in distant but well-connected regions with established inventive milieus (Jaffe et al., 1993; Bathelt et al., 2004; Kwon et al., 2020). In this paper, I attempt to resolve the puzzle of how invention can commence in places that lack local knowledge stocks by drawing on a theoretical literature on the window of locational opportunity (Scott and Storper, 1987; Storper and Walker, 1991; Brezis and Krugman, 1997; Boschma and Lambooy 1999; Crespo, 2011). The theory of the window of locational opportunity states that invention may commence in new locations following the introduction of a disruptive invention. In the empirical section of this paper, I present empirical evidence that is consistent with the window of locational opportunity theory.

I carry out the empirical study using a long-run panel dataset of innovative outputs geocoded to U.S. metropolitan areas that stretches from 1836 to 2014. This panel dataset allows me to identify commonalities in the growth trajectories of individual regions. Unlike cross-sectional data, the panel records allow me rule out deviations from general patterns that are driven by local idiosyncrasies (c.f. Audretsch, et al. 2008). I created the main dataset by combining two more basic sets of patent data. The first is the set of all utility patents granted by the U.S. Patent and Trademark Office between 1836 and 2014, geocoded to metropolitan areas using HistPat (Petralia, et al. 2016).[1] I use this first dataset to explore the extent to which the innovative output of regions follows and deviates from the pattern predicted by the classical regional lifecycle model. The second dataset contains records of the flows of technological knowledge between individual patents. This dataset covers all USPTO utility patent granted between 1836 and 2014, a coverage range which greatly exceeds that of traditional records of knowledge flows (patent citations) that are only available for patents granted after 1947 (Akcigit, et al. 2017). I created the second dataset for the purpose of this study, and I use it to investigate how the knowledge sourcing behaviors of inventors evolve as knowledge production begins, expands, and declines in inventors' home regions.

In the following section of this paper, I describe how geographical qualities of the process of knowledge production produce cyclical patterns of regional innovative growth and how those same qualities create windows of locational opportunity for inventors to initiate knowledge production in new places. From there, I describe the construction and dimensions of the dataset used in the empirical analysis. In the empirical analysis, I document the general pattern of innovative growth in regions and the deviations from the general pattern, I analyze the micro-level mechanisms beneath regional-level growth cycles, and I investigate the dynamics of regions that experience a resurgence of local innovative growth. In the final section, I reflect on the empirical results and relate the findings to the literatures on regional lifecycles and regional diversification.

[1] Public download links: https://patentsview.org/download/data-download-tables and https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/BPC15W

## 2) Invention and the Life Cycle of Regions

Regional innovative development is a process that contains recognizable stages involving the initiation of local innovation, the expansion of local innovation, and (potentially) its eventual decline. Each of those stages result from the nature of the process of invention as inventors navigate the technological opportunities and constraints that they encounter.

Invention entails the recombination of existing technological ideas into new configurations (Romer, 1990; Weitzman, 1998). These new configurations are exceedingly difficult to generate because each element in a technology operates by interacting with other elements in the same system through complex interfaces (Arrow, 1962; Adler and Clark, 1991; Von Hippel and Tyre, 1995; Fleming and Sorenson, 2001; Broekel, 2019). In response to that complexity, inventors tend to rely on prior knowledge to ease the process of creating new inventions (Fleming, 2001). However, inventors' ability to rely on prior knowledge is limited by the extent of knowledge that they know. Because inventors tend to have narrow areas of expertise, they often source ideas from other inventors and scientists (Wuchty et al. 2007).

The high level of detail to technological knowledge also constrains the ability for inventors to learn from others. Detailed knowledge overwhelms the bandwidth of the communication technologies inventors use to source and share it. Face-to-face communication stands out as an exception because face-to-face communication supports knowledge transmission along multiple channels including verbal language, body language, and the manipulation of vocal tone (Storper and Venables, 2004). Additionally, face-to-face communication allows for interactive feedback (Nohria and Eccles, 1992) and facilitates the production of norms and routines (Kogut and Zander, 1992; Powell et al., 1996; Gertler, 2003). Because close spatial proximity is a necessary condition for face-to-face communication, the ability for inventors to develop collaborative relationships, integrate into scientific and technological communities, and source technological knowledge increases with geographical proximity (Jaffe, et al. 1993). These advantages of face-to-face interaction allow knowledge production to expand in places where it has already commenced, but these same advantages make it very difficult for inventors to commence knowledge production in places that lack histories of local knowledge production, where few ideas can be sourced face-to-face.

One way that inventors may circumvent the challenges of starting knowledge production in new places is by exploiting discontinuities in the process of knowledge creation. The renderings of knowledge growth advanced by Kuhn (1956), Nelson and Winter (1982), and Arthur (2007) all emphasize how ideas exist within larger knowledge domains. Knowledge domains contain the sets of ideas that share material or institutional properties. By virtue of those shared properties, ideas found within the same domain are more adept for recombination than ideas found across domains. Most inventions advance the state of knowledge along well-defined trajectories of knowledge generation through the recombination of ideas that have strong recombinant complementarities (Frenken, et al. 2007; Hidalgo, et al. 2018). However, inventors periodically create new knowledge domains. Because the ideas created in new domains do not share strong recombinant complementarities with the ideas found in existing domains, the location where new domains will ultimately produce innovative growth is initially indeterminate (Scott, 2020, pp. 46). Therefore, the advent of a new knowledge domain opens a window of opportunity for inventors to commence knowledge production in new places (Scott and Storper, 1987; Walker and Storper, 1991; Brezis and Krugman, 1997; Boschma and Lambooy, 1999; Crespo, 2011).

While the introduction of new knowledge domains are often taken as exogenous or random events (Arthur, 1989; Brezis and Krugman, 1997), other researchers have argued that new domains are created out of existing ones, usually by recombining ideas found across existing domains in radically creative and original ways (Kuhn, 1956; Jacobs, 1969; Castaldi, et al. 2015). Ideas that were not thought to be complements become complements as a new domain emerges (Garud and Karnoe, 2001; Shi and Evans, 2020). Statistically, the diversity of the knowledge sources used by inventors is high during the creation of new knowledge domains, in part because heterogeneous knowledge inputs are truly useful for the creation of novelty (Glaeser, et al. 1992; Duranton and Puga, 2001; Berkes and Gaetani, 2020), but also because language and coding schemes are not updated after new domains form.

After a new domain forms, the usefulness the combinations of the ideas found within the domain becomes apparent to other inventors and economic actors. New entrants of inventors and firms, often working in relatively dispersed locations across geographical space (Scott and Storper, 1987), begin to experiment with the highly fertile ideas found in the new domain, producing more ideas within it (Dosi, 1982; Klepper, 1996). The nascent knowledge domain begins to grow, slowly at first, but later exponentially, so long as the ideas within it face sufficient demand (Dosi, 1982). Inventors rely on face-to-face communication to source and recombine knowledge within the domain and thus start to accumulate knowledge in some geographical regions but not in others. Eventually, knowledge production in the domain begins to concentrate in space, potentially in locations that were not major producers of technological knowledge before the advent of the new domain (Scott and Storper, 1987).

These mechanisms for the initiation of invention in new places have yet to be integrated into the regional lifecycle literature. For example, Audretsch, et al. (2008) begin their description of the classical regional lifecycle later in the lifecycle process, after a stock of knowledge has already accumulated in a new region and spills over between the local firms: they describe the first phase of the regional lifecycle as a "first entrepreneurial phase, during which both incumbent [firms] and entrepreneurs benefit from inter-industry [knowledge] spillovers" (Audrestch et al., 2008 pp. 7). Belussi and Sedita (2009) also emphasize that prior local knowledge, in the form of ancient craft traditions, initiates regional lifecycles. Their argument begs the question of how the ideas used in ancient crafts originated. Martin and Sunley (2011) are similarly vague on the initiation of regional growth cycles; they describe how regional lifecycles begin through local "resource accumulation" (pp. 1306), but they do not specify how those resources are accumulated.

In contrast to the question of how regional lifecycles begin, both the classical and adaptive regional lifecycle models provide clear answers for how and why innovation grows and declines in regions where it has commenced. According to the classical regional lifecycle model, as time passes knowledge production in a once-novel domain falters, often when a dominant design is introduced that performs a core function with sufficient efficiency (Utterback and Abernathy, 1975; Klepper, 1996). Subsequent innovation in the domain becomes incremental, draws from a narrow body of knowledge, and has minimal impact on subsequent invention (Dosi, 1982; Klepper, 1996). As innovation in a domain declines, so too does innovation in the regions that specialized in the domain (Audretsch et al., 2008).

The adaptive cycle model rejects the deterministic decline of innovation in regions. Instead, it contends that innovation in regions may resurge because of interdependencies that occur within regions, within industries, or in the particularities of regional-industry pairings (Martin and Sunley, 2011). The empirical literature on regional innovative growth suggests significant heterogeneity across regions in their growth trajectories. Empirical evidence affirms this point: some regions achieve a greater quantity innovative growth, and innovative growth in some regions persists for longer periods of time, often because of the infusion of non-local ideas that re-spark novelty in a region (Nefke, et al. 2017). These differences are present even when regions specialized in the same industries are compared (Saxenian, 1996; Menzel and Fornhal, 2010; Crespo, 2011; Storper, et al. 2015). Additionally, the economies of some regions are more resilient than others (Balland et al. 2015), and some regional economies have experienced multiple cycles of innovative growth through a localized industrial branching process (Glaeser, 2003; Belussi and Sedita, 2009; Boschma and Capone, 2015; Neffke, et al. 2017).

In Table 1, I summarize a reconstructed regional lifecycle model. The model depicted in Table 1 includes the theory of the window of locational opportunity to describe how regional lifecycles may begin. As noted in Table 1, the classical and adaptive lifecycle models disagree on the dynamics of regions during the fourth stage of the lifecycle.

< Table 1 about here >

## 3) Data and Methods

### 3.1) Identifying Inventive Regions and Their Lifecycles

Of the nearly 1,000 metropolitan and micropolitan areas (CBSAs) in the United States, only a small number have emerged as innovative centers. The 20 CBSAs that emerged as innovative centers during the period for which I have reliable data that is not subject to severe truncation, 1850-1999, are the main observation units of this study. I define a city-region as an inventive center if its local inventors produce 1% or more of the U.S. total flow of patents in a given time period. While I use a 1% patenting threshold definition for inventive centers in all analyses in the main text, in Online Appendix 4 I show that the main results of the paper are robust to altering the 1% threshold value.

In Table 2, I provide descriptive statistics on the 20 CBSAs that crossed the 1% threshold during the 1850-1999 study period. These statistics include the half-decade that each of the 20

CBSAs crossed the 1% threshold, the number of patents their inventors produced in the half-decade that each CBSA crossed the threshold, the leading CBSA from which the inventors in each CBSA sourced their knowledge during the half-decade their home CBSA crossed the 1% threshold, the half-decade when local knowledge production peaked in each CBSA, and the technological specialization of each CBSA during the half-decade when local patent production peaked (defined as the modal primary USPC patent class in which their local inventors produced patents). Some innovative U.S. city-regions, including the New York and Boston Metropolitan Areas, are not shown in Table 2 because their local inventors produced more than 1% of all U.S. patents before 1850. Because these cities' early histories are significantly truncated in the data, I do not these cities the main analysis.

These descriptive indicators in Table 2 generate several insights. First, city-regions have emerged as innovative centers across a long historical period, with some cities breaking the 1% threshold during the first few years of the 1850-1999 study period and others breaking the threshold during the final years of the study period. Second, New York City's knowledge base has been the most important for the emergence of new innovative centers; 18 of the 20 CBSAs listed in Table 2 sourced more knowledge from New York City than from any other metropolitan area during their formative years. Third, from the specialization in the "sewing" technological class in Bridgeport, CT, to the specialization in "railway rolling stock" in St. Louis, MO, to the specialization in "multiplex communications" in San Jose, CA, new innovative centers have emerged by developing specializations across a diverse array of technological fields.

<Table 2 about here>

To study the lifecycles these 20 inventive regions, it is necessary to juxtapose their histories of inventive growth. To do so, I calculate the $CBSA\ Age_{c,t}$ of each city-region $c$ in each time period $t$. This variable measures the number of years before or after a city-region's local inventors first produce 1% or more of the U.S. flow of all patents in the same time period. For example, using the information in Table 1, San Jose, CA had a $CBSA\ Age$ value of 0 in 1965, while Detroit had a $CBSA\ Age$ value of 0 in 1865. Formally, $CBSA\ Age_{c,t}$ is calculated by subtracting the half-decade an observation is taken from the half-decade the CBSA first crosses the 1% patenting threshold, as in Equation 1:

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

$CBSA\ Age_{c,t}$ is negative for the decades leading up to the 5-year period a CBSA crosses the 1% patenting threshold and is positive thereafter. After aligning the time series of patent production in each CBSA based on $CBSA\ Age_{c,t}$, I compute the average innovative output of CBSAs at each $CBSA\ Age_{c,t}$ value by taking the mean percent of U.S. patents in each half-decade by $CBSA\ Age_{c,t}$. I plot the resulting values in Figure 1 using a Loess regression with a 100% search distance. Loess regression is a smoothing function that computes rolling means across a sample; in the case of Figure 1, the Loess regression shows the average percentage of patents produced in CBSAs at each value of CBSA Age. I use the same Loess regression to overlay 95% confidence intervals in the chart.[2]

<Figure 1 about here>

Figure 1 generates two observations. The first is that there is a general pattern of knowledge production growth in inventive city-regions. Local patent production, as a percentage of the U.S. total, is minimal during the earliest years in their inventive growth, such as when CBSA Age is less than -50. However, patent production expands at a fast rate when CBSA Age is between -50 and 50. Patent production continues to increase when CBSA Age is between 50 and 100, albeit at a slower pace. Patent production reaches its maximum value when CBSA Age is 100. As CBSA Age increases from 100 to 150, the mean patent production across CBSAs as a percent of the U.S. total declines slightly, though this decline is not statistically significant.

The second observation generated by Figure 1 is that invention in most city-regions follows the general pattern described above with some deviation. This is demonstrated by the 95% confidence interval bands plotted in Figure 1, which are narrow when knowledge production begins and expands in CBSAs. The confidence intervals widen somewhat as CBSA Age passes 100. This increase in the size of the confidence intervals has two possible causes. First, a smaller number of city-regions reach a CBSA Age of 100 during the course of the study period, which increases the size of the standard errors. Second, city-regions might experience a more diverse range of outcomes during their later years.

To summarize, Figure 1 presents strong evidence that knowledge production commences and expands in U.S. city-regions following a well-defined pattern. However, Figure 1 does not generate a strong conclusion as to whether innovative city-regions regularly enter sustained periods of innovative decline in their later years of development. To push this analysis further, the final empirical section of the paper (Section 6) focusses on the prevalence and potential drivers of regional innovative resurgence. Before that analysis, however, I analye how changes the knowledge-sourcing behaviors of inventors over the course of regional lifecycles.

**3.2) Overview of Construction of Knowledge Sourcing Data**

The cyclical patterns of regional innovative growth documented in the previous section may arise from characteristics inherent to the process of knowledge creation. Because inventors create new knowledge by building on existing ideas, any tendency for regional innovative growth to be cyclical is likely to be associated with changes in the ways that inventors source knowledge.

[2] Loess regression is a local weighted regression method that estimates a moving average fit line. The 100% search distance I use in the Loess model means that observation points factor into the estimate for patent production at each CBSA Age value; however, Loess weights nearby points more strongly in its estimation, so not all points factor in equally.

To study how inventors' knowledge-sourcing behavior changes over regional growth cycles, it is necessary to analyze harmonized, long-run records of the knowledge sources that inventors use to make new inventions. Such data have not been generally available because the common data source used to study how inventors source technological knowledge, patent citation records, carry two major limitations. First, because many patent citations are added by patent examiners and attorneys, the extent to which patent citations represent knowledge spillovers is debated (Arora et al., 2019). Second, because the United States Patent and Trademark Office (USPTO) did not require patents to cite prior art before 1947, patent citation records are unreliable before this year (Akcigit, et al. 2017).

There are, however, implicit historical records of knowledge flow between patents hidden in the subclassification codes that the USPTO assigns to patents. The USPTO classifies all utility patents using a highly detailed classification scheme. At the highest level of granularity, the USPC classification scheme contains over 160,000 unique subclass codes which describe the individual components contained in each patented invention (Fleming and Sorenson, 2001). Therefore, subclassification codes listed on a patent indicate the recombinant know-how embedded in a technology (Fleming, 2001; Arthur, 2009). This argument is illustrated by the patent granted to Thomas Edison for the incandescent light bulb (USPTO patent number 223898). Edison's bright idea was that a vacuum chamber slows the combustion of a carbon filament. These elements Edison used to build his bulb – a vacuum chamber and a carbon filament – appear on his patent with the subclassification codes 201/035000 and 313/333000. The USPTO defines these codes as "carbonizing under pneumatic pressure or vacuum" and "filament or wire shield or electrode".

Because the subclassification codes listed on a patent indicate the recombinant knowledge embedded in the technology, when two timestamped patents share many of the same subclassification codes, it is plausible that the more recent patent sourced knowledge from the one that was granted before it (Foster and Evans 2019). The method I develop to predict flows of recombinant knowledge between inventions works by identifying shared subclassification codes between temporally-sorted patents.[3] I elaborate on this method in Online Appendix 1, and I provide a validation exercise of the produced data in Online Appendix 2. Using the methods described in these online appendices, I produce directed a-cyclical graph of the knowledge flows that connect 8.7 million USPTO patents granted between 1836 and 2014. In the graph, edges point from "parent" patents to their "children" patents.

[3] Conceptually, the method developed in this study is an inversion of the knowledge persistence method developed by Martinelli and Nomaler (2014). Martinelli and Nomaler's knowledge persistence approach uses observed knowledge flows (in the form of patent citations) to infer the knowledge embedded in patents. The method I develop uses observed records of knowledge embedded in patents to infer flows of knowledge between them. Although these methods are closely related, they make use of different data inputs and serve different purposes.

### 3.2) Measurements of Inventor Behavior

The literature discussed in Section 2 emphasized that inventors may change their propensity to source knowledge that is local, diverse, and high-impact along the regional growth cycle because each of those types of knowledge sources has unique advantages for initiating and expanding knowledge production in regions. To analyze how the knowledge sources used by inventors change over the regional growth cycle, I leverage the graph of knowledge flows described in the previous section. The dataset records each prior patent from which each focal patent draws knowledge. By linking these data to historical records of the residential location of each inventor (Petralia, et al., 2015), the USPC classification codes assigned to each patent, and a novel measure of the impact of each patent on subsequent invention (as described later in this section), I am able to identify the location, knowledge diversity, and impact of the knowledge sourced by inventors.

To measure the propensity for inventors in each CBSA and in each half-decade to source local knowledge, I compute the percentage of all knowledge sources used by inventors in a CBSA*Half-Decade unit that are local. To do so, I first identify all patents granted to inventors that resided in a focal CBSA in a focal half-decade. Next, I use my graph of knowledge flows to identify the parent patents of the focal CBSA*Half-Decade unit. Finally, I compute the percentage of all parent patents that were invented in the same CBSA as the focal patent. Occasionally, parent patents are invented by co-inventors that reside in multiple CBSAs, in which case I define a parent patent as a local knowledge source if one or more of its co-inventors resided in the CBSA of the focal patent.

To measure the diversity of knowledge sources used by inventors in cities, I again use the graph of knowledge flows to identify the parent patents of the focal patents invented in each CBSA*Half-Decade unit. Next, I compute a diversity indicator of the parent patents by recording the aggregate primary classification codes (at the 438 unique code level) of the parent patents and computing the GINI coefficient across these USPC codes. The GINI coefficient of knowledge sources is a measure of inequality in the knowledge sources across technology classes. Theoretically, it can range from 0 to 1, with higher values indicating greater inequality and thus narrower technological search.

To measure the propensity for inventors to source knowledge from high-impact inventions, I first develop a novel indicator of patent impact. Traditional impact indictors such as the number of forward citations received by patents are not available for patents granted before 1947 and therefore are unsuitable for my historical analysis (Akcigit, et al. 2017). Therefore, I leverage my graph of knowledge flows to compute the number of patents that draw knowledge from each patent. This produces a raw (continuous) measure the impact of each patent on subsequent invention. Next, I normalize the raw values by half-decade time periods by creating a binary variable, *High-Impact*, which equals 1 for patents that are in the top-decile of the same half-decade in terms of their raw impact.[4] I define all other patents as low-impact patents.

After computing the binary impact indicator of each patent, I use the graph of knowledge flows to identify the parent patents of the patents invented in each CBSA*Half-Decade observation unit, and I calculate the percentage of those parent patents that are high-impact and low-impact. During this calculation, I also record whether the parent patents are local or non-local, because the dynamics of the regional lifecycle suggest that inventors may vary the extent to which they source high-impact knowledge from local and non-local environments over the course of the lifecycle. Therefore, I distinguish between four types of knowledge sources: sources that are non-local and high-impact (NL.High), non-local and low-impact (NL.Low), local and high-impact (L.High), and local and low-impact (L.Low).

### 3.4) Comparison Method

In Online Appendix 5, I present plots of the raw indicators of knowledge source types used by inventors by City Age. These raw indicators are informative, but they lack necessary controls for calendar-year shocks. To account for calendar-year shocks, in the main text I compare the types of knowledge sources used by inventors located in the 20 CBSAs that emerge as innovative centers during the study period (CBSAs that they break the 1% patenting threshold) with the knowledge sources used by inventors located in the 952 CBSAs that never break the 1% threshold.[5] To perform this comparison, I compute the frequency by which their inventors use each type of source at each CBSA age value. Unsuccessful cities do not have CBSA age values because they never break the patenting threshold. Therefore, I link unsuccessful cities to successful cities based on the calendar half-decade of observation.

To demonstrate the comparison method, assume that I seek to compare the knowledge sourcing strategies of inventors in San Jose when its CBSA Age was -15 with the knowledge sourcing behavior of inventors in all unsuccessful cities at that same moment in history. San Jose was age -15 in 1950, so I compute the difference between the reliance of San Jose's inventors on each knowledge type of knowledge in 1950 (local knowledge, diverse knowledge, and high-impact knowledge) and the reliance on each knowledge type of inventors in all unsuccessful cities on the three types of knowledge in 1950 using subtraction. I use patent-weighted means across all unsuccessful cities in the same time period when computing the knowledge source types used by inventors in unsuccessful CBSAs. Therefore, the knowledge sourcing behaviors of inventors in CBSAs with very limited patent production do not unduly influence the results.

[4] The top decile is a commonly-used decile to define high-impact inventions and I adopt it to keep with the convention in the literature (Uzzi, et al. 2013); however, the results in Figure 5 are robust to the use of a top-15% high-impact threshold

[5] Online Appendix 6 summarizes the number of patents produced in the CBSAs that break the 1% threshold and inventors in CBSAs that do not break the threshold.

## 5) Changes in Inventor Behavior over the Regional Lifecycle

### 5.1) Inventors' Propensity to Source Local Knowledge

In this section, I investigate how the propensity for inventors to source local knowledge changes as knowledge production initiates, expands, and declines in regions. There are two core mechanisms that may influence the propensity for inventors to source local knowledge. The first mechanism is the quantity effect: the number of local ideas that inventors may draw on changes over time as local knowledge production grows. Therefore, larger local knowledge bases may be positively associated with the propensity to source local knowledge. The second mechanism is the quality effect. Holding the size of local knowledge bases constant, inventors may prefer local to non-local ideas if the types of ideas available locally are more useful for creating inventions. I decompose the quantity and quality effects using a dartboard as developed by Ellison and Glaeser (1997) by randomly shuffling the patents produced across CBSAs. I do not change the grant years of patents during this random assignment, so each CBSA is assigned the same number of patents as its inventors produced during each year.

After randomizing patent production across CBSA, I compute inventors' reliance on local knowledge sources for each value of CBSA age. As described in Section 3.4, I calculate the difference between inventors' reliance on local knowledge sources in cities that break the 1% threshold and inventors' reliance on local knowledge sources in cities that do not break the 1% threshold. In Figure 2, I fit and plot a Loess regression with a 100% search distance to the randomized data using a dotted line. The distance between this dotted line and the horizontal axis shows the quantity effect on inventors' propensity to source knowledge locally.

In Figure 2, I add a solid line which is plot a Loess regression of inventors' reliance on local knowledge sources, measured using the observed geographical distribution of patent production. The distance between the solid line and the dotted line shows the quality effect on the propensity for inventors to source knowledge locally. The distance between the solid line and the horizontal axis shows the combination of the two effects.

< Figure 2 about here>

Figure 2 generate several insights. First, when CBSA Age is very low (<-50), the quantity effect is 0 (the dotted line is on the horizontal axis), which implies that the size of local knowledge bases does not have a positive effect on the propensity for inventors to source knowledge locally. This first result is expected because city-regions with very small local knowledge bases do not have a sufficient number of local ideas to spur invention. The second observation from Figure 2 is that when CBSA Age is very low (<50), the quality effect is negative (the solid line is significantly below the dotted line). This result shows that inventors in emerging regions strategically seek out non-local ideas. The third observation from Figure 2 is that as CBSA Age increases above 50, the quantity effect becomes positive (the dotted line is significantly above the horizontal axis). This shows that inventors in regions with expanding local knowledge production source local knowledge more intensively as the local knowledge pool grows. The fourth and final observation from Figure 2 is that as CBSA Age increases above 50, the quality effect is positive (the solid line is significantly above the dotted line). This fourth observation shows that regions with expanding knowledge production develop a technological niche that is well-suited to spur subsequent local innovation. Moreover, inventors in emerging innovative places benefit not just from the quantity of ideas that are locally accessible, but also from the quality of those ideas in terms of furthering local knowledge production.

### 5.2) Inventors' Propensity to Source Diverse Knowledge

In this section, I test whether inventors change the diversity of the knowledge sources that they use as local innovation initiates, grows, and plateaus in their home regions. Figure 3 plots a Loess regression fit line using a 100% search distance with 95% confidence intervals of the difference in the GINI coefficient of knowledge sources used by inventors in cities that break the 1% patenting threshold and inventors in city-regions that do not break the threshold. The figure generates three important results. First, inventors in city-regions that eventually break the 1% patenting threshold source a more diverse set of ideas in the years before their CBSAs break the 1% threshold than do inventors in city-regions that never break the 1% threshold: when CBSA Age is very low (less than -50), the difference between the GINI coefficients is negative. The second finding is that inventors in city-regions that break the 1% threshold source an increasingly narrow set of ideas as CBSA Age increases. The third result is that knowledge sourcing in city-regions tends to closely follow this general pattern: the 95% confidence intervals in the right panel are very narrow, suggesting that the knowledge sourcing strategies of inventors rarely deviates from this dominant pattern. These results again are consistent with the view that inventors initiate knowledge production in new places by creating new knowledge domains out of an initially diverse set of knowledge-based inputs.

<Figure 3 about here>

### 5.3) Inventors' Propensity to source High-Impact Knowledge

In this section, I examine how inventors source knowledge with varying levels of impact as knowledge production commences, expands, and declines in their home regions. In the analysis, I also decompose the types of knowledge sources based on their geography (local or non-local) in order to test if the geography of high and low-impact knowledge sourcing also evolves across the regional growth cycle.

In Figure 4, I plot Loess fit lines with a 100% search distance of the percentage of knowledge sources that used by inventors that are non-local and high-impact (NL.High), non-local and low-impact (NL.Low), local and high-impact (L.High), and local and low-impact (L.Low). As in the previous charts, the plotted values are the differences between the reliance on each of these four types of sources used by inventors in cities that break the 1% patenting threshold and their reliance in cities that do not break the 1% patenting threshold.
<Figure 4 about here>

Figure 4 shows that inventors in successful CBSAs rely more heavily on high-impact knowledge sources throughout all stages of the regional growth cycle than do inventors in unsuccessful CBSAs. However, the geographical origins of these high-impact knowledge sources evolve as local knowledge production takes hold. Early in their CBSA's innovative growth, when CBSA age is between -100 and -50, inventors in successful cities source about 2.5% more of their knowledge from NL.High patents than do inventors in unsuccessful cities, as indicated by the dotted grey line above the x-axis. The shaded confidence interval indicates that this difference is statistically significant at the 95% level. Between age -50 and 0, inventors in successful cities increase a growing share of their knowledge from local high-impact inventions, as indicated by dotted black line that rises well-above the x-axis. Eventually, when city age passes 50 years, inventors in successful cities source a growing share of their knowledge from local low-impact inventions. However, the dotted black line remains far above the solid black line for the duration of the chart, indicating that inventors in successful cities rely much more heavily on local high-impact knowledge than do inventors in unsuccessful cities for the full duration of their region's innovative growth and decline.

**6) The Renewal of Regional Lifecycles**

A proposition of the adaptive cycle model of regional innovative development is that regions may experience a resurgence of innovation if local industries are renewed or replaced. There are four U.S. city-regions that sustained innovation across multiple growth cycles, defined as CBSAs that broke the 1% patenting threshold, experienced a decline in patent output over multiple decades, and thereafter experienced a second increase in patent output. In this section, I describe the dynamics of patent production and knowledge sourcing in these four city-regions in order to better understand how their resurgence of innovation occurred. These four city-regions are the New York, Boston, San Francisco, and Seattle Metropolitan Areas. I did not include the New York City and Boston Metropolitan Areas in the prior analyses because their inventors produced more than 1% of U.S. patents before the coverage of my data begin in 1850, which makes the calculation of CBSA Age as defined in Section 3.1 impossible.

Figure 5 plots the number of patents produced, as a percentage of the U.S. total, by inventors in the New York City (NYC), Boston, San Francisco, and Seattle Metropolitan Areas over time. For each of these CBSAs, I define a year during which knowledge production reached a turning point, defined as the final year before which knowledge production in each CBSA started its final ascent. Those turnaround years are 2005 in NYC, 1955 in Boston, 1980 in San Francisco, and 1960 in Seattle. I mark each of these turnaround points with a star in Figure 5.

<Figure 5 about here>

I use the turnaround points of the four city-regions as benchmark events to study how the knowledge sourcing strategies of their inventors changed in the years leading up to the turnarounds. Specifically, I explore changes in the propensity for inventors to source knowledge that is local, diverse, and high-impact in the time leading up to the turnaround point of each CBSA by the number of years before or after the city-region undergoes a turnaround of its innovative output. As in Section 3.4, I account for system-wide changes in the knowledge sourcing strategies of inventors across time by comparing the types of knowledge sources used by inventors in these four city-regions with the types of knowledge sources used by inventors in city-regions that never cross the 1% patenting threshold.

I begin by analyzing whether the turnaround of innovation in city-regions is associated with the sourcing of non-local knowledge. As in Section 5.1, I decompose the quantity and quality effects on the propensity to source local knowledge. I plot the resulting propensities using Loess regressions with 95% confidence intervals in Figure 6.

<Figure 6 about here>

Figure 6 shows that inventors in city-regions that experience a resurgence of local innovation tend to source a relatively greater share of their knowledge non-locally in the years leading up to the turnaround. Part of this shift to non-local knowledge sources is driven by the quantity effect as the size of local knowledge bases decline. However, the larger share of this shift is driven by the decrease of the quality effect, which becomes statistically insignificant in the years leading up to the turnaround. The erosion of the quality effect indicates that the types of ideas that inventors can source locally in declining regions has no effect on their propensity to source local ideas. Moreover, residual knowledge left over in regions from previous growth cycles is not uniquely useful for spurring a second cycle of local innovation.

I next study how the propensity for inventors to source diverse knowledge changes as their city-regions begin to experience a renewal of local innovation. Figure 7 plots the technological narrowness of the knowledge sources used by inventors (the GINI coefficient) for the four city-regions that experience a turnaround in the years leading up to and following their turnaround point. Again, I compute the difference between the Gini coefficient for inventors in the four city-regions and the GINI coefficient for inventors in the city-regions that never break the 1% patenting threshold. Figure 7 shows that the inventors in the four city-regions that experienced a turnaround increasingly source diverse sets of ideas in the years leading up to the turnaround point.

<Figure 7 about here>

Finally, I examine how the propensity for inventors to source high-impact knowledge changes as city-regions approach a turnaround of local innovation. Figure 8 shows how the propensity for inventors in city-regions that experience a renewal of local innovation to source non-local high-impact (NL.High), non-local low-impact (NL.Low), local high-impact (L.High), and local low-impact (L.Low) knowledge, relative to those propensities for inventors in city-regions that do not cross the 1% patenting threshold, changes as the turnaround of local innovation nears.

<Figure 8 about here>

Figure 8 shows that the impact and location of the knowledge sources used by inventors in regions that experience a renewal of local innovation changes as the turnaround point nears. In particular, inventors in regions that experience a turnaround rely less heavily on local high-impact sources as the turnaround point approaches and instead source a larger percentage of their knowledge from non-local sources.

## 7) Discussion and Conclusions

This paper has used data on knowledge production and knowledge sourcing to study how innovation initiates, expands, declines, and resurges in regions. Regional knowledge production growth was found to generally conform to the patterns predicted by the classical regional lifecycle model. In particular, the share of technologies produced in each region tends to expand in regions in which it has already taken hold, eventually plateaus or declines, and in a few exceptional cases, may resume growth after long-term decline. An important caveat to these results is that city-regions that recently emerged as centers for innovation have not yet reached the age at which we would anticipate their decline. Therefore, it is possible that newer U.S. city-regions may break from these historical patterns.

The infrequency by which regions experience a second round of local innovative growth brings to question the analytical utility of the adaptive cycle model of regional innovation for the evolution of metropolitan areas in the U.S. (Martin and Sunley, 2011). There are several possible reasons for why the adaptive cycle may generally not apply to U.S. regions. First, the United States' economy is land-abundant which influences the behavior of its system of cities. In this regard, for most of its industrial history the U.S. economy had a frontier that was open to the western advance of new industries (Storper and Walker, 1991). This contrasts with Europe, where high levels of population density forced new rounds of economic growth to establish themselves in the same places where past growth occurred. The second possible reason is that the U.S. economy is strongly oriented toward disruptive innovation. Venture capital and supporting factors formed in the United States long before they did in other economies, such as the European Union or Japan. For this reason, the cognitive and spatial distances between existing old and new knowledge domains may be on average greater in the U.S. than in other economies. A final possible reason is the U.S. federal government's laissez-faire approach to regional economic development (Boschma and Capone, 2015). The federal government of the U.S. rarely engages in regional development. State governments, local governments, and industry associations are often left to coordinate regional development (Storper, et al. 2015). However, state and local governments are unable to distribute capital without the generation of tax revenue, and industry associations have minimal incentive to invest and organize in declining regions that have residual congestion, high tax rates, and organized labor forces (Scott and Storper, 1987; Brezis and Krugman, 1997; Storper and Walker, 1991).

Despite these conditions, four city-regions in the United States experienced significant resurgences of innovation between 1850 and 2010. These city-regions are the New York City, Boston, San Francisco, and Seattle Metropolitan Areas. In addition, this study found that inventors in resurgent city-regions change their knowledge-sourcing behaviors in the years leading up to moment of resurgence. In this regard, local inventors source an increasingly diverse and non-local stock of knowledge about 10 to 50 years before their home regions experienced a resurgence of local innovative growth. In addition, the decomposition of the quantity and quality effects on the propensity for inventors to source local knowledge in resurging city-regions also found that the quality of local knowledge has no bearing on the propensity for inventors in resurging city-regions to source local knowledge. Therefore, cities that resurge as innovative centers do so by virtue of the size of the local knowledge stock and not the type of ideas found in those places.

Together, the results of this study demonstrate that the process of innovative growth in U.S. regions as a cyclical one. Inventors initiate knowledge production in new regions by developing and advancing new domains of technological knowledge, and inventors initiate the resurgence of knowledge production in declining regions by again developing and advancing new domains of technological knowledge. Thus, the initiation and the resurgence of knowledge production are both iterations of a general process. In addition, these cycles of regional innovation – even when they occur in the same region – are generally distinct episodes, connected to other cycles by weak linkages of knowledge.

**Bibliography**

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

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

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

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

Arthur, B. (2007). *The Nature of Technology*. Free Press. New York.

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

Audretsch, D., and Fritsch, M. (2002). Growth Regimes over Time and Space. *Regional Studies* (36). 113-124.

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

Balland, P.A., Rigby, D., and Boschma, R. (2015). The Technological Resilience of U.S. Cities. *Cambridge Journal of Regions, Economy, and Society* (8). 167-184.

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

Belussi, F., and Sedita, S.R. (2009). Life Cycle vs. Multiple Path Dependency in Industrial Districts. *European Planning Studies* 17(4). 505-528.

Berkes, E., and Gaetani, R. (2020). The Geography of Unconventional Innovation. *The Economic Journal.* 131(636). 1466-1514.

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

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

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

Boschma, R., and Capone, G. (2015). Institutions and Diversification: Related versus Unrelated Diversification in a Varieties of Capitalism Framework. *Research Policy* 44(10). 1902-1914.

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.

Crespo, J. (2011). How Emergence Conditions of Technological Clusters Affect their Viability? Theoretical Perspectives on Cluster Life Cycles. *European Planning Studies* 19(12). 2025-2046.

Dosi (1982). Technological Paradigms and Technological Trajectories. *Research Policy*. 11(3).

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.

Ellison, G., and Glaeser, E. (1997). Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach. *Journal of Political Economy* 105(5). 889-927.

Falck, O., and Heblich, S. (2008). Modern Location Factors in Dynamic Regions. *European Planning Studies* 16(10). 1385-1403,

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).

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

Frenken, K., Van Oort, F., and Verburg, T. (2007). Related Variety, Unrelated Variety, and Regional Economic Growth. *Regional Studies* 41(5). 685-697.

Garud, R., and Karnøe, P. (2001). Path creation as a process of mindful deviation. In Garud, R. and Karnøe, P. (Eds), *Path Dependence and Path Creation*. Mahwah, NJ: Lawrence Earlbaum, 1–38.

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., Kallal, H., Scheinkman, J., and Shleifer, A. (1992). Growth in Cities. *Journal of Political Economy* 100(6). 1126-1152.

Glaeser, E. (2003). Reinventing Boston: 1640-2003. *NBER Working Paper #10166*.

Grant, R. (1996). Toward a Knowledge-Based Theory of the Firm. *Strategic Management* Journal 17: 109-122.

Hausmann, R., and Neffke, F. (2019). The Workforce of Pioneer Plants. *Research Policy* 48(3). 628-648.

Hidalgo, C., Balland, P.A., Boschma, R., Delgado, M., Feldman, M., Frenken, K., Glaeser, E. He, C., Kogler, D., Morrison, A., Neffke, F., Rigby, D., Stern, S., Zheng, S., and Zhu, S. (2018). The Principle of Relatedness. *International Conference on Complex Systems*. 451-457.

Jacobs, J. (1969). *The Economy of Cities.* Penguin Random House. New York.

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

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

Klepper, S. (1996). Entry, Exit, Growth, and Innovation over the Product Life Cycle. *American Economic Review* (86). 562-583.

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

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

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

Martin, R., and Sunley, P. (2011). Conceptualizing Cluster Evolution: Beyond the Life Cycle Model? *Regional Studies* 45(10). 1299-1318.

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

Menzel, M.P., and Fornahl, D. (2010). Cluster Life Cycles – Dimensions and Rationales of Cluster Evolution. *Industrial and Corporate Change* 19(1). 205-238.

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

Neffke, F., Hartog, M., Boschma, R., and Henning, M. (2018). Agents of Structural Change: The Role of Firms and Entrepreneurs in Regional Diversification. *Economic Geography* 94(1). 23-48.

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

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

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

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

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

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

Scott, A.J., and Storper, M. (1987). High Technology Industry and Reginal Development: A Theoretical Critique and Reconstruction. *International Social Science Journal* (112). 215-232.

Shi, F., and Evans, J. Science and Technology Advance through Surprise. arXiv:1910.09370.

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

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

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

Trippl, M., Grillitsch, M., and Isaksen, A. (2017). Exogenous Sources of Regional Industrial Change: Attraction and Absorption of Non-Local Knowledge for New Path Development. 42(5). 687-705.

Utterback, J., and Abernathy, W. (1975). A Dynamic Model of Process and Product Innovation. *OMEGA* 3(6). 639-656.

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

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

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

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

**List of Tables**

**Table 1: Reconstruction of the Regional Lifecycle with Associated Inventor Behavior**

| **Lifecycle Stage** | **Local Knowledge Production** | **Inventor Knowledge-Sourcing Behavior** | | |
|---|---|---|---|---|
| | | **Reliance on Local Knowledge** | **Reliance on Diverse Knowledge** | **Reliance on High-Impact Knowledge** |
| **1. Initiation** | Small but growing | Low | High | High |
| **2. Expansion** | Expanding rapidly | Moderate | Moderate | Moderate |
| **3. Decline** | Plateauing or declining | No prediction | Low | Low |
| **4. Resurgence (Only in Adaptive Cycle Model)** | Growing | Decreasing | Increasing | Increasing |

**Table 2: List of CBSAs (Metropolitan and Micropolitan Areas) that Emerged as Innovative Centers during the 1850-2000 Study Period**

| CBSA | Half-Decade CBSA broke 1% Patenting Threshold | Number of Patents Produced in CBSA during Half-Decade CBSA broke 1% Threshold | Leading CBSA of Knowledge Sources during Half-Decade CBSA broke 1% Threshold | Half-Decade Patent Production Peaked in CBSA as % of U.S. Total | Top USPC Technology Class in Half-Decade Patent Production Peaked in CBSA |
|---|---|---|---|---|---|
| Akron, OH | 1940 | 1,603 | NYC | 1945 | Synthetic Resins or Rubbers |
| Atlanta, GA | 1995 | 8,489 | NYC | Has yet to peak | Multiplex Communications |
| Austin, TX | 1990 | 5,010 | San Jose, CA | Has yet to peak | Multiple Computer or Process Coordinating |
| Bridgeport, CT | 1930 | 2,047 | NYC | 1940 | Sewing |
| Chicago, IL | 1855 | 169 | NYC | 1940 | Mineral Oils: Process and Products |
| Dallas, TX | 1970 | 2,656 | NYC | Has yet to peak | Multiplex Communications |
| Detroit, MI | 1865 | 494 | NYC | 1940 | Machine Element or Mechanism |
| Houston, TX | 1955 | 2,021 | NYC | 1985 | Wells (Oil) |
| Indianapolis, IN | 1880 | 756 | NYC | 1880 | Belt Power Transmission Systems |
| Los Angeles, CA | 1905 | 1,792 | NYC | 1970 | Fluid Handling |
| Miami, FL | 1980 | 2,791 | NYC | 1990 | Electrical Communications |
| Milwaukee, WI | 1890 | 942 | NYC | 1930 | Circuit Makers and Breakers |
| Minneapolis, MN | 1890 | 1,178 | NYC | Has yet to peak | Surgery: Light, Thermal, and |

| | | | | | |
|---|---|---|---|---|---|
| | | | | | Electrical Application |
| Phoenix, AZ | 1985 | 3,805 | NYC | 1995 | Solid-State Devices |
| Portland, OR | 1995 | 7,621 | San Jose, CA | Has yet to peak | Semiconductor Manufacturing |
| San Diego, CA | 1980 | 2,813 | NYC | Has yet to peak | Multiplex Communications |
| San Jose, CA | 1965 | 3,635 | NYC | Has yet to peak | Multiplex Communications |
| Seattle, WA | 1915 | 1,645 | NYC | Has yet to peak | Database and File Management |
| St. Louis, MO | 1855 | 170 | NYC | 1905 | Railway Rolling Stock |
| San Francisco, CA | 1865 | 532 | NYC | Has yet to peak | Database and File Management |

*Data source: Author's elaboration of USPTO utility patent records. In CBSAs where patent production as percentage of U.S. total has yet to peak, the top technology class is given using the half-decade ending in 2000.*

**List of Figures**

**Figure 1: Average Patent Production in U.S. CBSAs by CBSA Age**

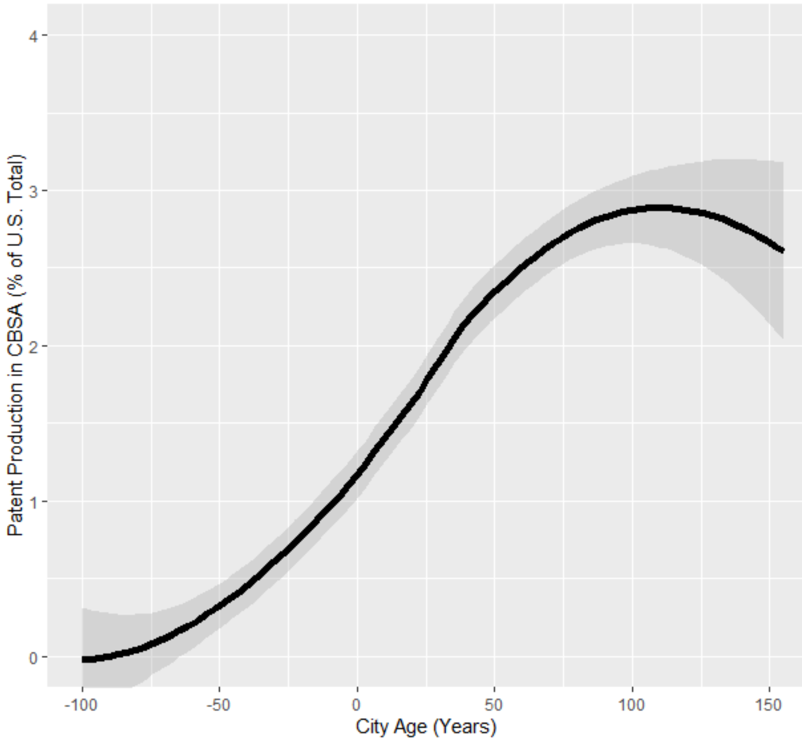


*Note: Fit lines estimated with Loess regressions with 100% search distance. Only CBSAs that produce at least 1% of U.S. patents during one or more 5-year periods of their history are included in the analysis (see Table 2)*

**Figure 2: Difference in Reliance on Local Knowledge Sources for inventors in CBSAs that break 1% Patenting Threshold and Inventors in CBSAs that do not break 1% Threshold, with Decomposition for Quantity and Quality Effects**

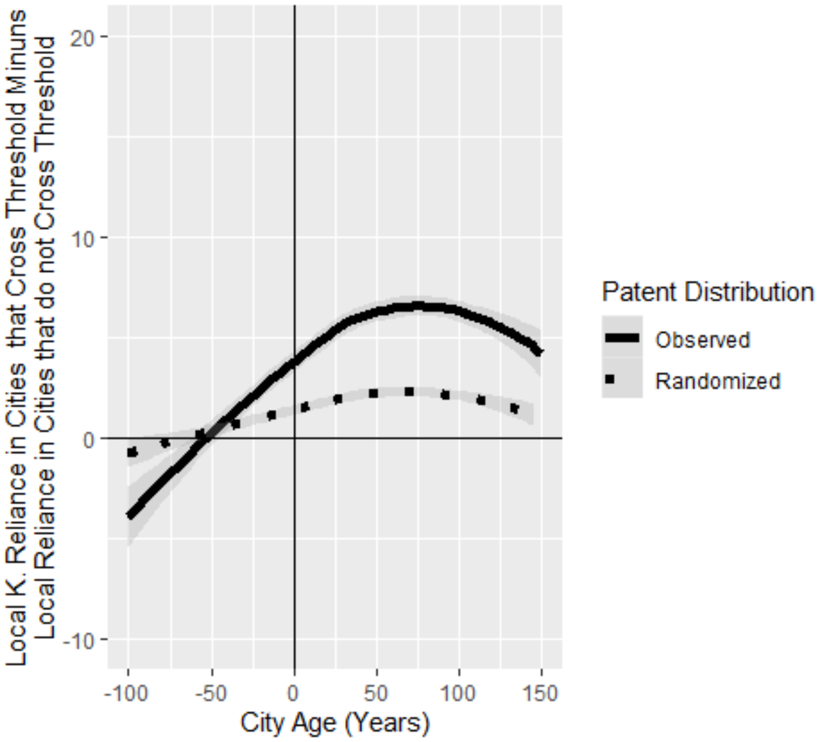


*Note: Dotted line is produced by randomizing the CBSAs in which patents are produced within each year. Patents are assigned the same number of patents as they produced in the same year. Solid line is computed taking the observed (real) distribution of patent production across CBSAs.*

**Figure 3: Difference in GINI Coefficient of Narrowness of Knowledge Sources for inventors in CBSAs that break 1% Patenting Threshold and Inventors in CBSAs that do not break 1% Threshold**

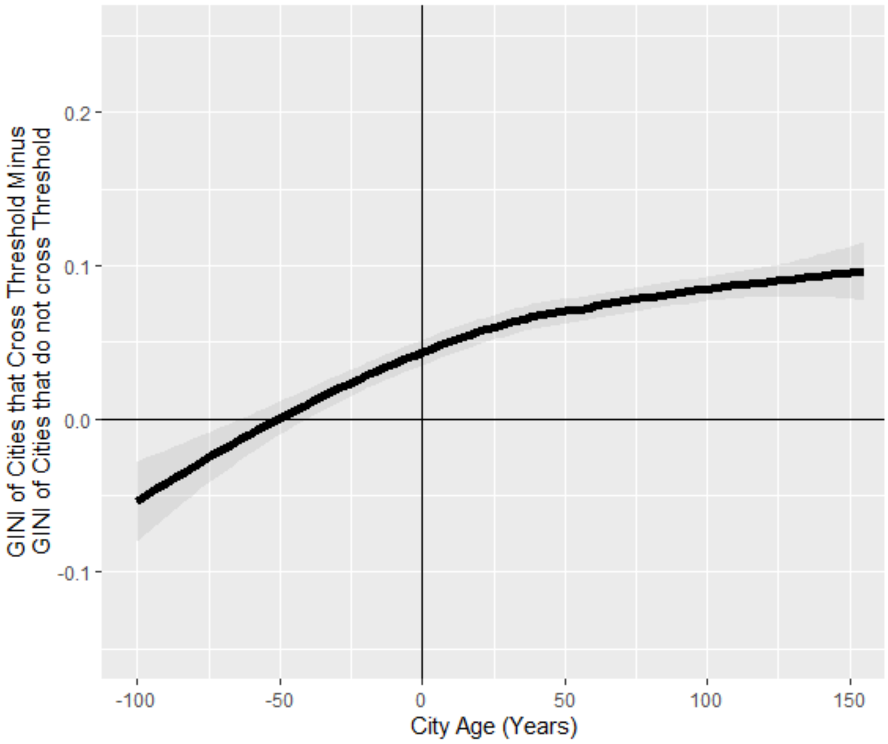


*Note: Higher y-values indicate a larger GINI coefficient (relative to the GINI in CBSAs that do not cross the 1% threshold) and thus more narrow/specialized technological search*

**Figure 4: Difference in Impact and Geography of Knowledge Sources for inventors in CBSAs that break 1% Patenting Threshold and Inventors in CBSAs that do not break 1% Threshold**

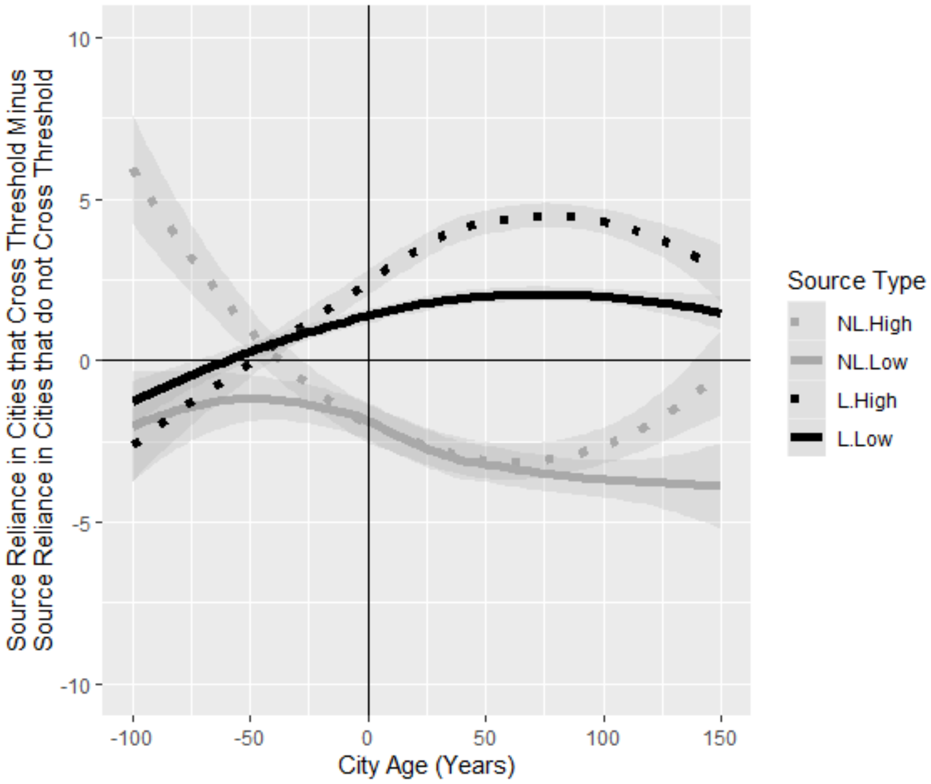

**Figure 5: Patent Production in Cities that Undergo a Renewal of Local Innovation**

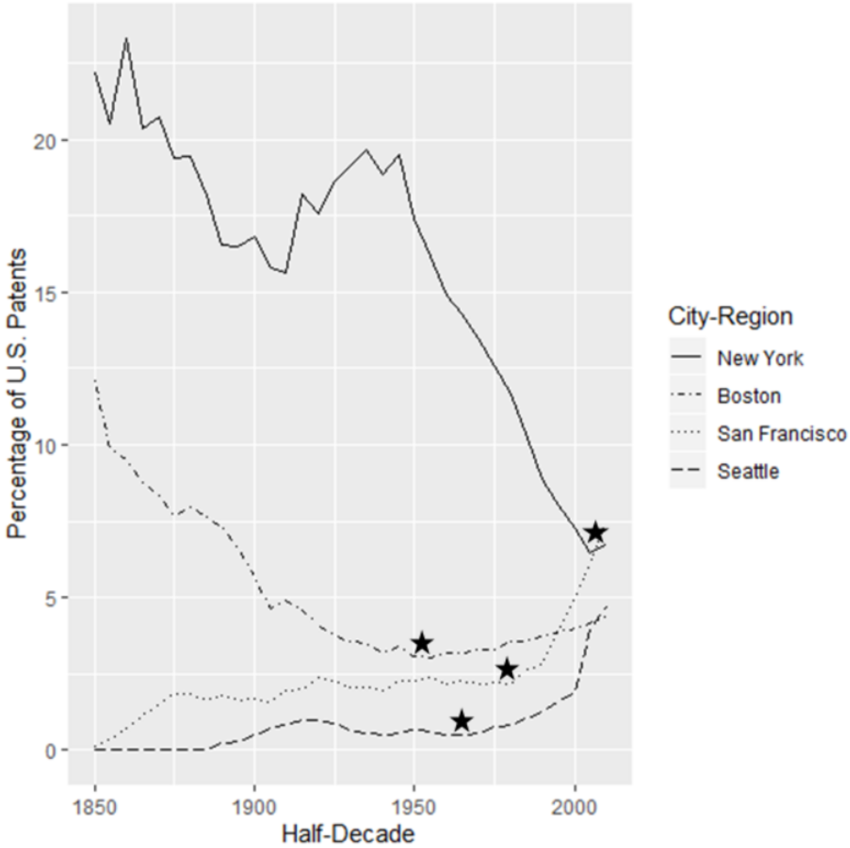


*Note: stars indicate the final year where patent production in each CBSA begins a monotonic increase over time that continues through the end of the dataset.*

**Figure 6: Difference in Narrowness of Knowledge Sources for inventors in CBSAs that Experience a Renewal of Local Innovation and Inventors in CBSAs that do not break 1% Patenting Threshold**

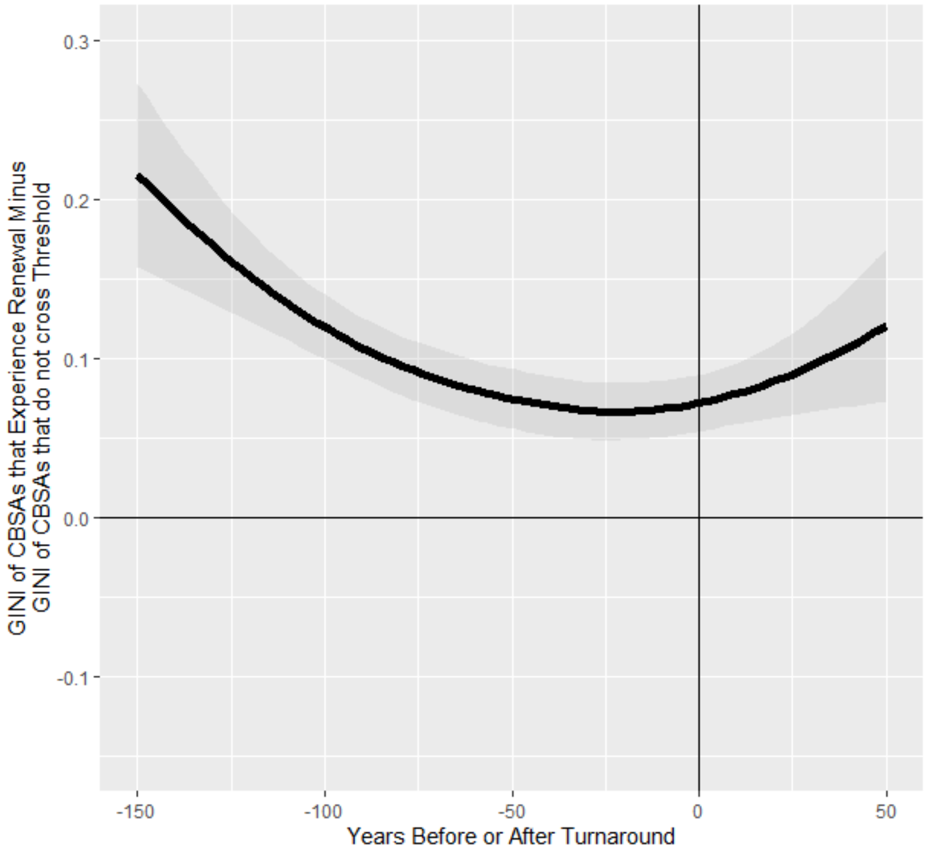


*Note: Higher y-values indicate a larger GINI coefficient (relative to the GINI in CBSAs that do not cross the 1% threshold) and thus more narrow/specialized technological search*

**Figure 7: Difference in Percentage of Knowledge Sources that are Local for inventors in CBSAs that Experience a Renewal of Local Innovation and Inventors in CBSAs that do not break 1% Patenting Threshold**

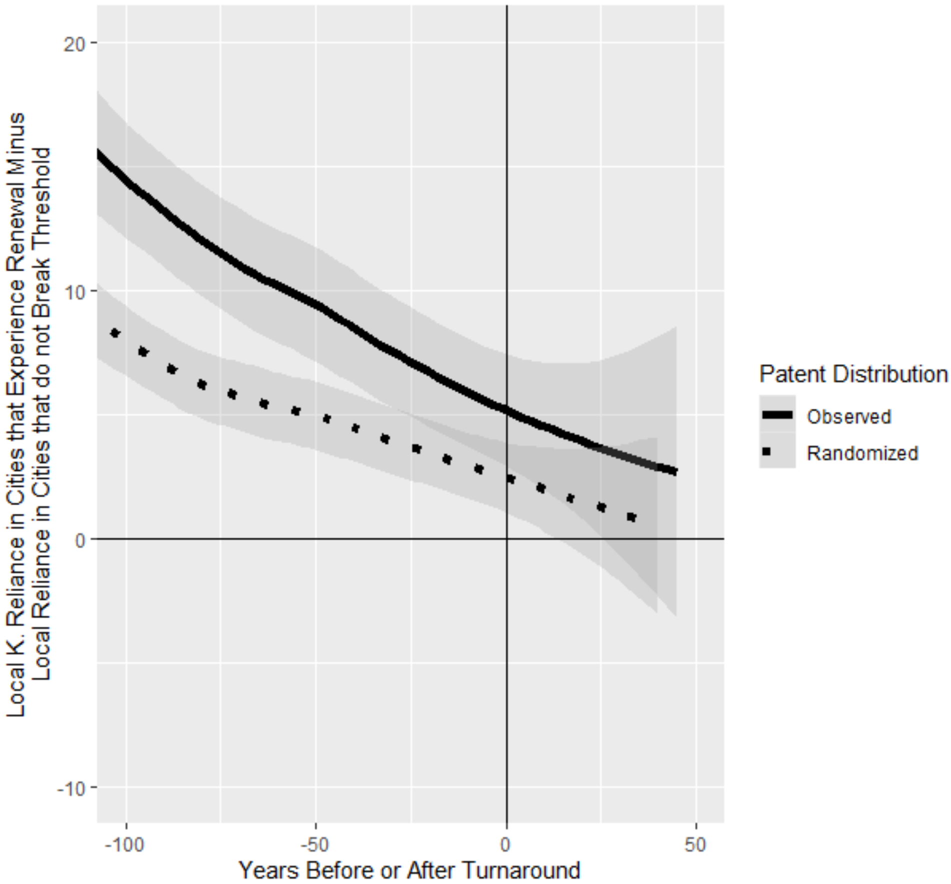

**Figure 8: Difference in Impact and Geography of Knowledge Sources for inventors in CBSAs that Experience a Renewal of Local Innovation and Inventors in CBSAs that do not break 1% Patenting Threshold**

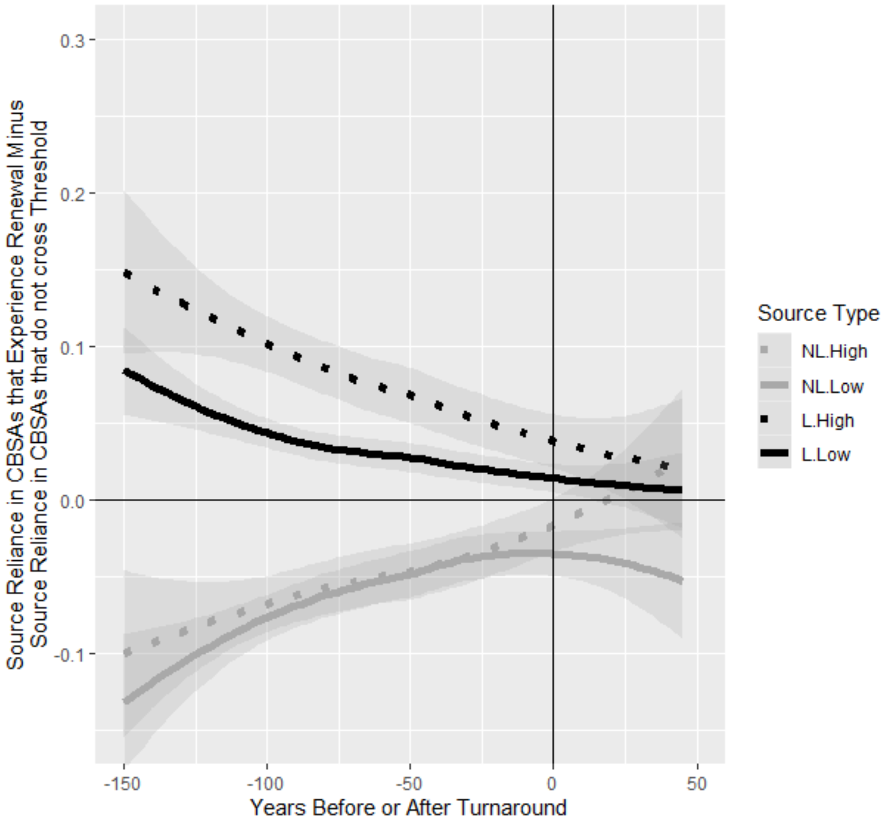

**Online Appendix 1) Specifics of Data Construction**

I infer flows of recombinant knowledge between patents by exploiting the information provided by USPC subclassification codes on all USPTO utility patents granted between 1836 and 2014. The resulting "tree of technology" is a directed a-cyclical graph that links each patent to its knowledge-based antecedents. To create the tree, I begin with raw public files of granted patents and USPC subclass assignments available on PatentsView. The USPTO reclassifies patents using the USPC coding schema as new subclasses are added over time, creating a harmonized, current system. I omit design patents but keep patents assigned to non-U.S. inventors, which leaves me with 8.7 million patents.

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

The tree-building algorithm begins by selecting the most recently granted patent and recording its components based on its USPC subclassification codes. I define recombinant knowledge as knowledge of components and the interactions of those components, so for each patent I generate all combinations of degree *n* of its components, where *n* is the number of components in a patent.[6] For example, if a focal patent (FP) contains the USPC subclassification codes A, B, C, the knowledge vector is generated as follows:

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

Each element in $Knowledge_{FP}$ denotes a single unit of knowledge; the length of $Knowledge_{FP}$ indicates the total quantity of knowledge units embedded in the focal patent. The knowledge units in $Knowledge_{FP}$ are used to link the FP to its knowledge-based predecessors, or "parent patents", based on the number of knowledge units that are found in both the focal patent and a possible parent patent. To identify the possible parents of a focal patent, I search for overlapping knowledge units in all patents that were granted before the focal patent was granted, based on the sequence of patent ID numbers which are numbered in order based on patents' grant date. I do not constrain the time window during which a parent patent can serve as a source of knowledge for a child patent because inventors often build on both old and new ideas (Mukherjee et al., 2017).[7] For each possible parent that fits the simple temporal criterion, I generate a shared knowledge vector (SKnowledge) to record the knowledge units that appear in both the focal patent and in the parent. For example, if a possible parents' knowledge vector, $Knowledge_{PP}$, is given by:

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

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

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

and the knowledge of the FP, $Knowledge_{FP}$, is given by Equation 1, the shared knowledge vector is taken as the intersection of the $Knowledge_{FP}$ vector and the $Knowledge_{PP}$ vector:

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

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

When an FP has multiple potential parents for an individual unit of knowledge, I assign a fractional weight to the edge based on the number of possible parents for that knowledge unit. For example, if two possible parents contain the component [ B ], I assume that the FP sources 0.5 units of knowledge from the [ B ] in the first possible parent and 0.5 units from the second. In practice, "simple" knowledge units of length 1 such as [ B ] tend to be found on many potential parent patents, while "complex" or lengthier knowledge units such as [ BC ] tend to be found on ~~far~~ fewer.[8] Therefore, the fractional assignment of edge weights, by virtue of its basic methodology, tends to create stronger ties between patents that share complex and thus non-ubiquitous combinations of knowledge. It is also important to note that two patents that share complex combinations of knowledge can be connected by stronger edge weights because complex combinations of knowledge also contain nested simpler combinations. For example, when two patents share the combination [BC], they also must share [B] and [C], which means that they share 3 total units of knowledge. Thus, the edge weight connecting the two patents could be as large as 3, depending on the fractional assignment.

I repeat the process described above for all 8.7 million USPTO patents granted between 1836 and 2014. This algorithm produces a directed a-cyclical graph where edges point from parent patents to their children patents.

**Online Appendix 2) Validation of the Patent Impact Measures**

This paper introduces a new dataset which records the flows of knowledge between individual patents granted between 1836 and 2010. In this online appendix, I validate the precision of these records of knowledge flow by testing whether my method generally agrees with external accounts of high-impact inventions.

I thus perform two validation checks to test whether the impact measures calculated using the out-degree of the technology tree corresponds to external sources. In the first validation check,

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

I test whether patents identified by technological historians as uniquely consequential inventions have a higher median out-degree than a comparison set of patents. I compare the median out-degree of these distributions instead of the means because of the extensive right-dispersion of the impact across patents.

I use two sets of historian-identified patents for this purpose. The first is provided by Rogers (2011), which lists over 100 important inventions made in the U.S. between 1840 and 1920.[9] The second is provided by the Computer History Museum in San Jose, California, which lists the patents issued for inventions that were milestones in the development of the silicon engine of modern computers.[10] For each set of historian-identified patents, I create three control groups of patents. The first control group consists of all patents that were granted in same year as the historians' patents. The second control group consists of all patents that were granted in the same year and assigned to the same primary class (at which level there are 438 unique classes) as the historians' patents. The third control group consists of all patents that were granted in the same year and assigned to the same primary subclass (at which level there are 160,000 unique subclasses) as the historians' patents. The median out-degree of the historian-identified patents and their reference group are given in Table OA1. In addition, I test whether the differences between the median out-degree of the historian-identified patents and the control group patents are statistically significant using a Fligner-Policello test. The results indicate that the out-degree of patents included in the historian-identified sets of important patents are significantly greater than all the control sets of patents.

**Table OA1: Mean Out-Degree of Historian-Identified Great Patents**

| | **Median Out-Degree in Genealogy** | | | |
|---|---|---|---|---|
| | **Patents in Historian List** | **Same-Year Control Group** | **Same Year and Class Control Group** | **Same Year and Subclass Control Group** |
| *Great American Patents* 1840-1920 Rogers (2011) | 1 | 0*** | 0*** | 0*** |
| *Milestones in the Silicon Engine* 1904-1983 Computer History Museum | 10 | 0*** | 1*** | 3*** |

**** Denotes p-value of difference with the median historian list is less than 0.01 (calculated using a Fligner-Policello test)*

[9] While Rogers (2011) lists impactful inventions back to 1750, his pre-1840 inventions cannot be linked to patents. Citation: Rogers, D. (2011). Inventions and their Inventors, 1750-1920. M-Y Books Limited. London.

[10] Source: https://www.computerhistory.org/siliconengine/timeline/

In a second validation exercise, I compare the mean out-degree of patents from the technology tree to the number of forward citations they receive from subsequent inventions. Because patent forward citations are generally broadly available before 1975, I perform this validation exercise using only patents granted starting in 1976. In addition, both patents' forward citation count and out-degree from the technology tree suffer right-truncation in recent years. Therefore, I do not include patents granted after 1990 in this exercise.

To test the association between patents' forward citations and out-degree, I run a regression model of the forward citations received by patent $p$ as a function of the out-degree of patent $p$. To compare similar types of technologies, I include a primary class of class fixed effect in the model. In addition, I include yearly fixed effects. The model is given by the equation below and its results are presented in Table OA2.

$$FwdCites_p = B_1 OutDegree_p + FE_{ClassCode} + FE_{Year}$$

Because the dependent variable, $FwdCites_p$, is a count variable with strong right-dispersion, I estimate the equation using a Quasi-Likelihood Poisson model. The estimates are given in Table OA2.

**Table OA2: Regression of 5-Year Forward Citations and 5-Year Out-Degree Impact of Patents, 1976-1990**

| | (1) | (2) | (3) |
|---|---|---|---|
| $B1$ | 0.107***<br>(0.00280) | 0.109***<br>(0.00349) | 0.0905***<br>(0.00295) |
| Fixed Effects | Year | Year and Primary Class | Year and Primary Subclass |
| Number of Fixed Effects | 15 | 389 | 16,917 |
| NOBS | 1,049,288 | 1,026,327 | 301,352 |

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

Table OA2 shows that for patents granted between 1976 and 1990, out-degree is positively associated with forward citation counts. Therefore, out-degree calculated using the technology tree is positively associated with three external records of patent impact: forward citations, inclusion in Rogers' (2011) list of great American patents, and inclusion in the Computer History Museum's list of milestone patents in the development of the silicon engine.

Next, I explore changes in the average impact of patents and the average number of subclasses assigned to patents over time. I plot these values in Figure OA1. In the figure, I also break out the average of high-impact patents (those in the top-10% of impact in the same grant-year cohort).

Figure OA1 shows that both the average impact of patents and the average number of subclasses on patents increased during the study period. These two trends may be related, because patents assigned more subclasses have a greater number of embedded knowledge units and thus can have more knowledge-based descendants. In the empirical analyses in the main text, I remove the temporal aspect of this relationship from the data by defining high-impact patents as those in the top-10% of impact in their same grant-year cohort. This normalization makes patents granted later in time (i.e. in the 1990s) no more likely to be high-impact than earlier-granted patents simply by virtue of the increase in the number of subclass codes on patents over time.

**Figure OA1: Average Impact and Number of Classes of Patents by Year**

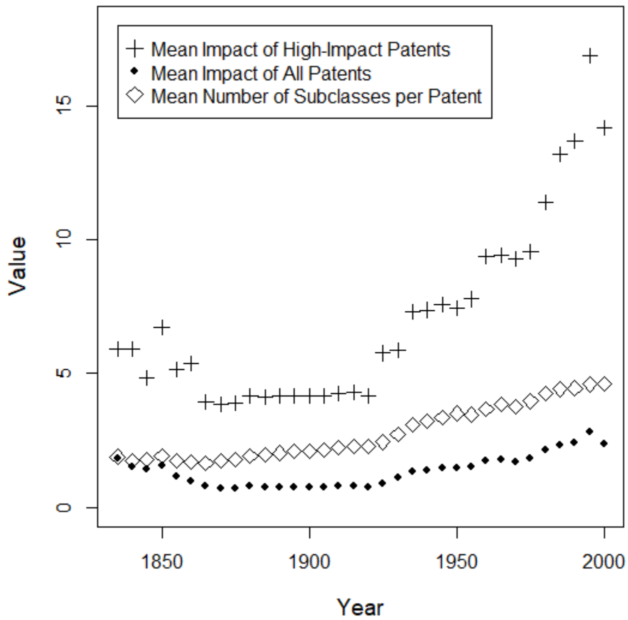


In the following analysis, I analyze the aggregate technological sectors in which high-impact patents were produced between 1850 and 1999. I use aggregated sectors at which there are 6 sectors in this analysis. The number of high-impact patents by sector are plotted in Figure OA2. The figure shows that most high-impact patents were created in the mechanical and miscellaneous sectors until 1940, after which the chemicals, computers, and drugs and medical sectors became the primary sectors in which high-impact patents were produced.

**Figure OA2: Count of High-Impact Patents by Technological Sector**

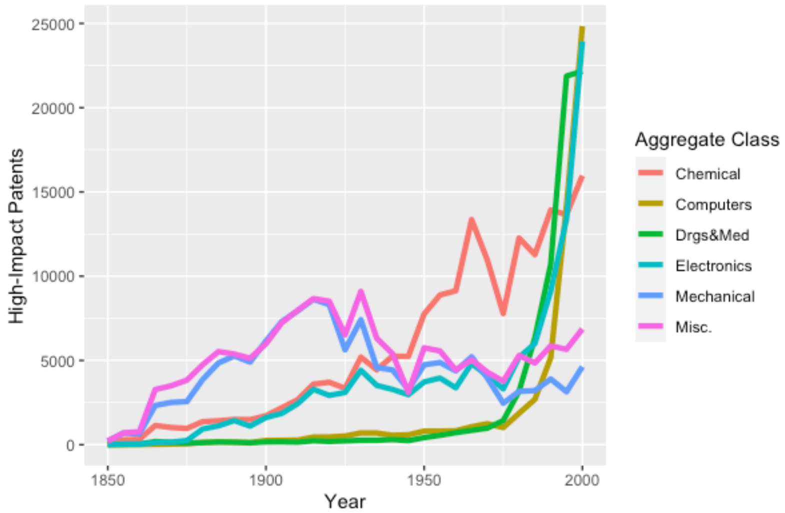


Finally, I examine the extent to which shared edges in the genealogical tree predicts citations between patents. Because this analysis is performed at the edge level, the full dataset is too computationally burdensome to analyze in full. Therefore, I take a random 10% sample of all backward citations made by patents granted in 1990 (there are 44,000 citations contained in this 10% sample). Next, I test whether the "genealogy-identified parent patents" of these sampled "child patents" are a significant predictor of their "citation-identified parent patents.

To perform this test, it is necessary to compare actual citations with citations that could plausibly have been made. The approach used by Jaffe, et al. (1993) involves creating two datasets: a case dataset and a control dataset. The case dataset contains citations that were observed on patents, while the control dataset contains citations that plausibly could have been made. I adopt a similar approach in this exercise. I create a case dataset that contains the random 10% sample of backward citations made by patents granted in 1990. Then, I create three control groups. In the first control group, I randomize the patent that receives each citation but keep the grant year of the receiving patent constant. Therefore, the first control group is a sample-year randomized-patent control set. In the second control group, I randomize the patent that receives each citation but keep the grant year and the primary aggregate technological class (at the 438 unique codes level) constant. Therefore, the citations made in the second control group have similar vintage and technological field. In the third control group, I randomize the patent that receives each citation but keep the grant year and the subclass (at the 160,000 unique code level) constant. In doing so, I omit citations that do not have a respective control patent granted in the same year and in the same subclass. Therefore, the third control group contains citations of the same vintage and of the same detailed subclass.

I combine the case dataset with one of the first control group and run a regression model in which the probability that a citation is a real citation is modeled as a function of whether the patent *i* and patent *j* are connected by an edge in the genealogical tree. After running this model using the case dataset and the first control group, I then combine the case dataset with the second control group and re-run the model. Finally, I run the model using the case dataset and the third control group. The regression model is given by the below equation:

$$Prob(RealCitation_{ij}) = B_1 GenealogicalEdge_{ij} + E_{ij}$$

I estimate the equation using a linear probability model and show the estimates of $B_1$ for the three datasets in Table OA3.

**Table OA3: Regression of $\boldsymbol{Prob(RealCitation_{ij})}$**

| | **(1)** | **(2)** | **(3)** |
|---|---|---|---|
| $B_1$ | 0.506***<br>(0.0153) | 0.494***<br>(0.0152) | 0.285***<br>(0.0136) |
| Intercept | 0.494***<br>(0.00173) | 0.494***<br>(0.00173) | 0.495***<br>(0.00174) |
| Control Parent Patents | Same vintage | Same vintage and aggregate class | Same vintage and detailed subclass |

The results in Table OA3 show that the connection between two patents with a genealogical edge is a significant predictor of whether two patents will be connected by a citation. These results are consistent regardless of whether the control group of "fake" cited patents are taken from the same vintage (column 1), the same vintage and aggregate class (column 2), or the same vintage and detailed subclass (column 3). Notably, the coefficient two patents sharing a genealogical edge is smaller when the granularity of the match with the control group increases (as in column 3). This result is expected because the control group in column three already is technologically very similar to the real citations. Therefore, the genealogical edges will not add as much explanatory power to the model.

**Online Appendix 3) Diversity of Patents in Innovative Centers**

The analysis in the main text shows that the knowledge sources used by inventors become increasingly specialized as their home regions emerge as innovative centers. In this section, I more generally describe the pattern of the increased specialization of knowledge produced by inventors in rising innovative places. I plot the Gini coefficient of the knowledge types produced by inventors in each rising innovative center in each given 5-year period. The patterns are plotted in Figure OA3. I use four separate charts to plot these patterns in order to make it easier to interpret the patterns.

**Figure OA3: Diversity of Knowledge Produced in Rising Innovative Centers**

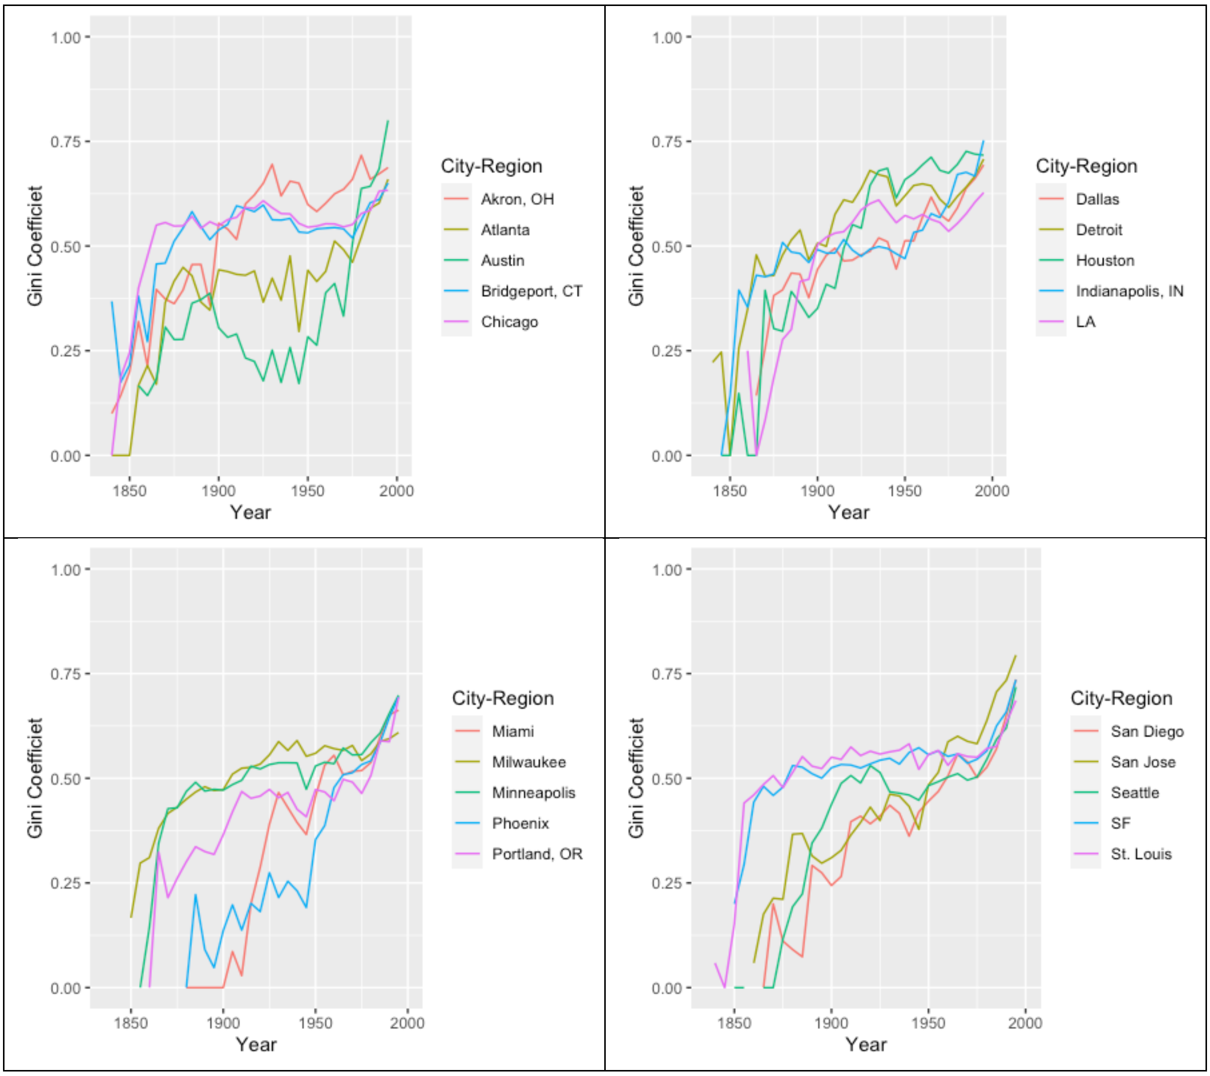


The patterns shown in Figure OA3 generally conform to those identified in the main text, in particular the main text's analysis of the diversity of knowledge sources used by inventors as their home regions rise as innovative centers. In particular, the knowledge produced by inventors in cities becomes increasingly specialized as their home regions emerge as innovative centers. This relationship is shown in examples such as Chicago, which both rose as an innovative center early on (in 1855, as per Table 2), and produced a relatively specialized set of knowledge early on (as shown in Figure OA3). By contrast, Austin TX did not produce a specialized set of knowledge until later in time, and it did not rise as an innovative center until the late 20$^{th}$ century.

**Online Appendix 4) Analysis using 0.5% and 5% Thresholds to Measure Emergence of Innovative City-Regions**

In the Online Appendix 4, I replicate the core analyses found in the main text but use 0.5% and 5% patenting thresholds to define the year that a CBSA emerges as an innovative center, instead of the 1% threshold used in the main text. The conclusions of the results are generally the same as those in the main text.

**Figure OA4: Production of High and Low-Impact Patents in CBSAs by CBSA Age using 0.5% and 5% Thresholds**

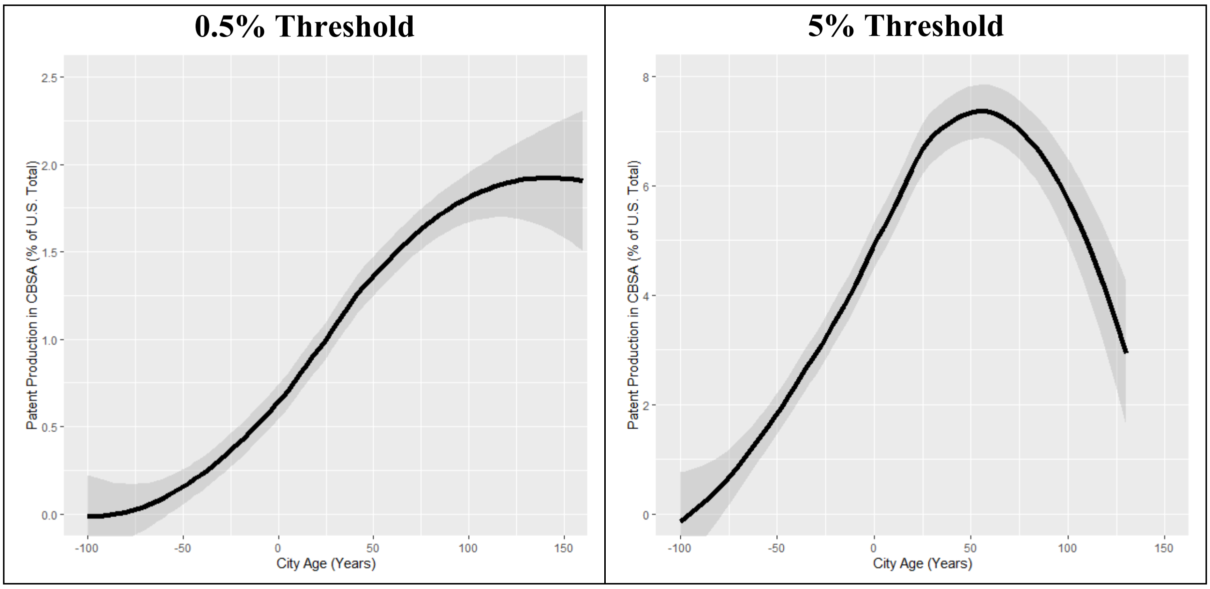


**Figure OA5: Geography of Knowledge Sources using 0.5% and 5% Thresholds**

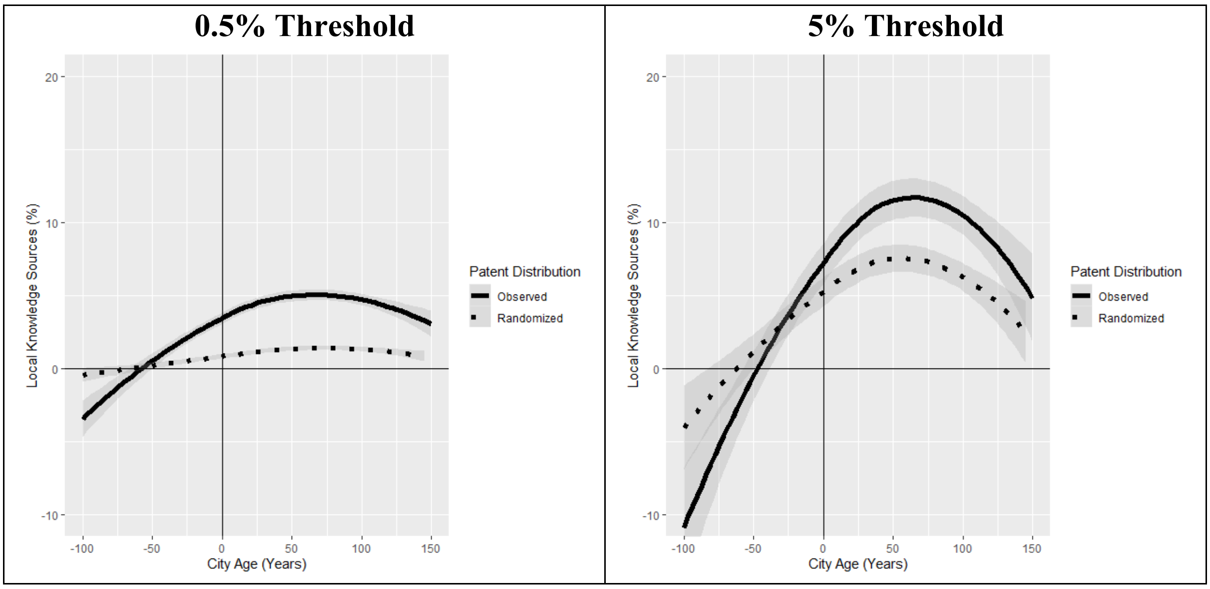

**Figure OA6: GINIs of Innovative CBSAs Relative to Non-Innovative CBSAs using 0.5% and 5% Thresholds**

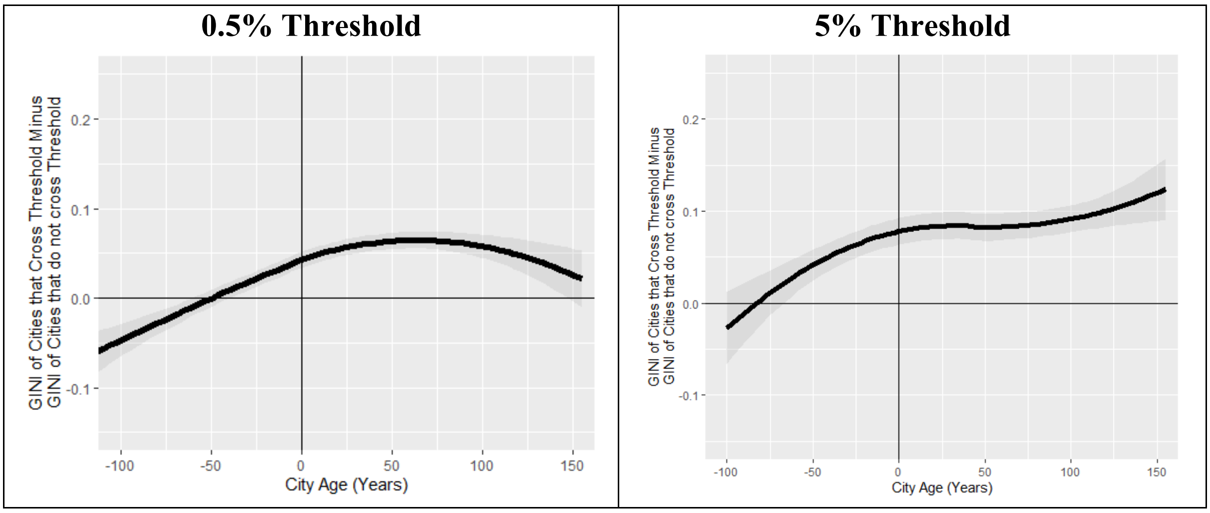


**Figure OA7: Impact and Geography of Knowledge Sources using 0.5% and 5% Thresholds**

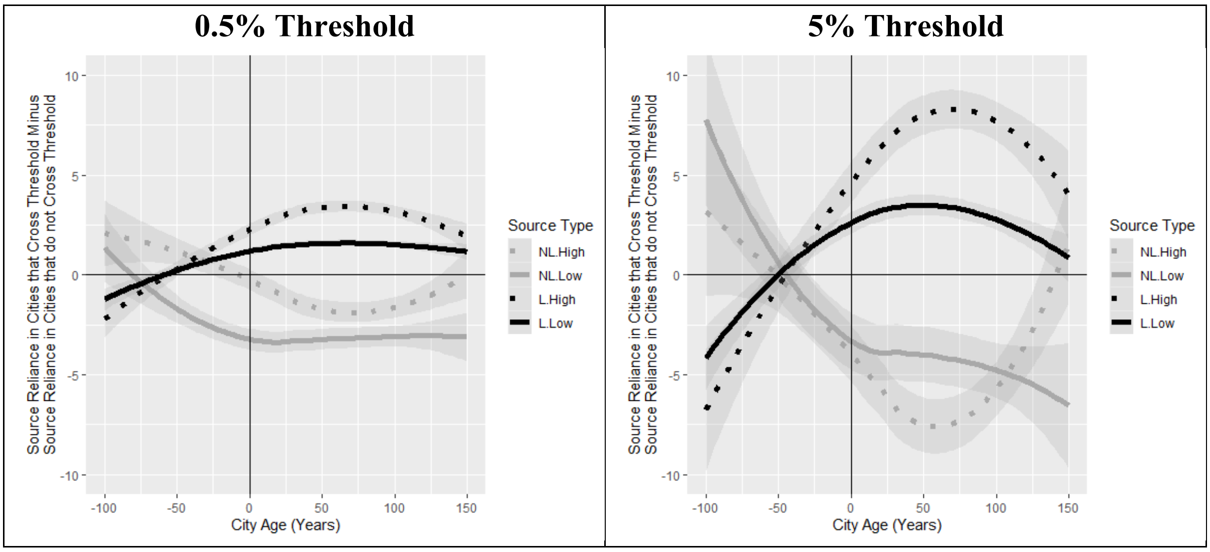

**Online Appendix 5) Analysis of Raw Knowledge Sourcing Propensities**

In the Online Appendix 5, I show the raw indicators of the types of knowledge sourced by inventors in CBSAs that cross the 1% patenting threshold. In contrast, the results in the main text show the difference in the types of knowledge sources used by inventors in CBSAs that cross the 1% threshold and the sources used by inventors in CBSAs that do not cross the 1% threshold.

**Figure OA8: Percentage of Knowledge Sources that are Local in CBSAs that break the 1% Threshold using Raw (Non-Differenced) Values**

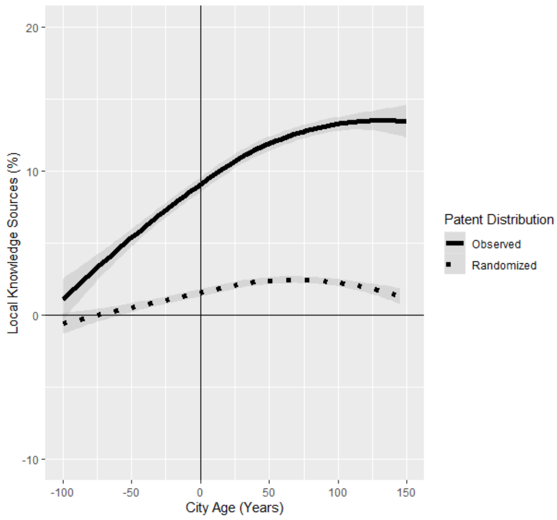


**Figure OA9: GINI Coefficient of Knowledge Sources in CBSAs that break the 1% Threshold using Raw (Non-Differenced) Values**

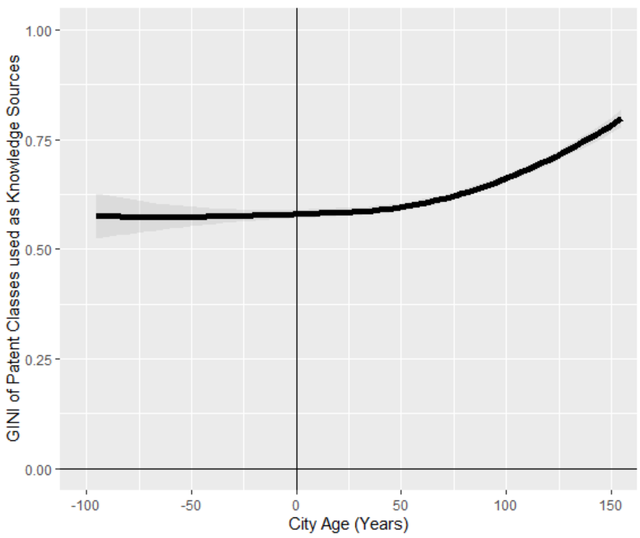

**Figure OA10: Impact and Geography of Knowledge Sources in CBSAs that break the 1% Threshold using Raw (Non-Differenced) Values**

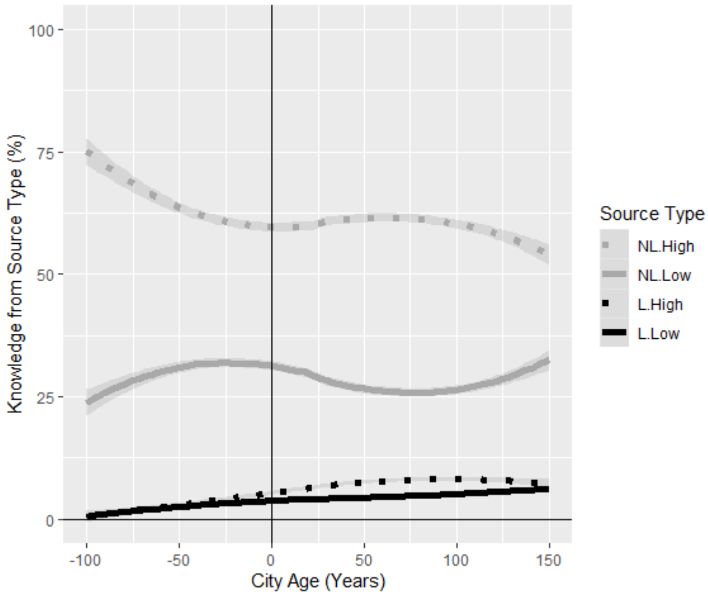

**Online Appendix 6) Patenting in CBSAs that break and do not break the 1% patenting threshold**

**Figure OA11: Log Patent Production in CBSAs that break the 1% Patenting Threshold and in CBSAs that do not break the 1% Patenting Threshold**