Papers in Evolutionary Economic Geography

20.46

The Emergence of Knowledge Production in New Places

Christopher R. Esposito

Utrecht University Human Geography and Planning

The Emergence of Knowledge Production in New Places

Christopher R. Esposito

Ph.D. Candidate
Department of Geography, UCLA
9/23/20

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

Key words: Regional development, innovation, knowledge transmission, agglomeration

JEL codes: O33, R12

Acknowledgements: I thank David Rigby, Michael Storper, Allen Scott, Jacob Foster, and participants at Seminars in Economic Geography e-seminar series for helpful comments in the development of this research.

1) Introduction

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

While San Jose's rise is striking, nearly every innovative city in the United States started off in a similar position as a location that produced few patentable ideas. Table 1 shows the year that the top-15 patent-producing metropolitan areas in the U.S. emerged as centers for knowledge production, defined as the first five-year period that they produced 1% of the U.S.'s count of utility patents. Of the 15 top-ranked metropolitan areas, 12 crossed the 1% threshold after 1835 when the data series begin. San Jose, San Diego, and Austin, TX, for example, all emerged as centers for innovation after the 1950s. And although "old" cities like Chicago and Detroit rose in the rankings much earlier, knowledge production in those cities too had a beginning.

Table 1: The Rise of New Cities as Centers of Innovation

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

Note: To reduce volatility patent counts and thresholds are aggregated by half-decades.

How do inventors commence knowledge production in new places? This question is difficult to resolve using the traditional explanations of agglomeration from the literature on the geography of innovation. According to that literature, innovative activities concentrate in space (Audrestch and Feldman, 1996; Balland et al., 2019) because inventors use existing ideas to create new ideas (Nelson and Winter, 1982; Romer, 1988), and because ideas tend to be transmitted between actors located in close physical proximity or in distant but well-connected regions with established inventive milieus (Jaffe et al., 1993; Bathelt et al., 2004; Breschi and Lissoni, 2009; Kwon et al., 2020). While these arguments explain why knowledge production concentrates in space, they leave us to puzzle over how inventors begin to produce knowledge in places that lack existing knowledge stocks or inter-regional networks to begin with.

One possible solution to this puzzle is provided by the theory of the Window of Locational Opportunity. The theory of the Window of Locational Opportunity argues that new innovative agglomerations are able to form when impactful and disruptive inventions erode the competitive advantages of incumbent regions (Storper and Walker, 1987; Boschma and Lambooy 1999; Boschma and Frenken, 2006). So long as the geography of knowledge flow is not mechanistically governed, idiosyncratic factors may transport impactful ideas to "new" regions where they are used to make yet more ideas, thereby inducing local knowledge production.

In this paper, I find considerable support that knowledge production begins in new regions through a similar process as the one described by the theory of the Window of Locational Opportunity. In particular, I show that inventors initiate knowledge production in new places by sourcing ideas from impactful non-local inventions about 50 years before their home regions emerge as innovative centers, that inventors use the impactful ideas sourced from other regions to introduce a large number of impactful ideas of their own, and that inventors leverage their own impactful ideas to create even more ideas locally.

I substantiate these claims through analysis of a novel dataset, created with U.S. patent records, that traces the flow of technological knowledge between patents and across time. The new data provide reliable records of knowledge flow back to 1850. By combining these records with historical information on inventors' place-of-residence (Petralia et al., 2016), I am able to study how knowledge production initiated in all but the oldest U.S. cities. To study how local knowledge production initiates, I decompose the sources of knowledge used by the inventors in regions based on their sources' geographical origins and their sources' level of technological impact. To explore the causal effect that the composition of knowledge sources has on the emergence and growth of local knowledge production, I compare the composition used by inventors in regions that succeed in developing an appreciable volume of knowledge production with the composition used in places that fail to.

The results of the study contribute to three literatures: the geography of knowledge flow, agglomeration theory and evolutionary economic geography, and the urban lifecycle. With respect to the first literature, the decomposition of knowledge sources conducted in this study shows how local and non-local knowledge flows materialize in knowledge production growth in proximate and distant locations (Jaffe et al., 1993; Breschi and Lissoni, 2009; Kwon et al., 2020). Through that exercise, the results also report how the overall propensity for knowledge to flow locally has changed dramatically between 1850 and 2010. With respect to agglomeration theory and evolutionary economic geography, the analysis reveals how spatial concentrations of knowledge production form in their earliest years, before they enter the purview of agglomeration theory and evolutionary economic geography as units of observation. Finally, the findings expand the concept of the urban lifecycle (Audrestch et al. 2008) by revealing how new centers for innovation are conceived using ideas developed in other places, and that even after innovation in a places declines, its ideas can flock to and flourish in new regions.

In the text that follows, I discuss how inventors create and transmit technological knowledge and the geographical implications thereof, I introduce the methods used to infer historical flows of knowledge between patents and to identify high-impact inventions, and I present the results of the analysis, beginning with a birds-eye-view of knowledge production growth and decline in regions and continuing on to a decomposition of the sources used by inventors as their regions initiate, expand, and decline in patent production. In addition to outlining areas that require additional research, in the final section I discuss the relationship between incremental and disruptive innovation, and I elaborate on the respective repercussions for the emergence, evolution, and resilience of knowledge production in cities.

2) The Production and Transmission of Technological Knowledge

Technological knowledge, defined as the ability to assemble tangible and intangible elements into functioning systems, is exceedingly difficult to generate. Each element in a technology operates by interacting with other elements in the same system. Because of the high degree of interdependence between the constitutive components, inventors struggle to anticipate how their technologies will function before they assemble them (Fleming and Sorenson, 2001). Models and prototypes help inventors to simulate the interactions between components, but they are costly to create and time-consuming to administer (Usher, 1929; Arrow, 1962; Adler and Clark, 1991; Von Hippel and Tyre, 1995). These costs multiply when inventors design complex technologies with many elements arranged in irregular ways (Broekel, 2019).

To ease the process of designing complex technologies, inventors rely on prior knowledge (Fleming, 2001). Inventors that already know how an assembly of components functions can focus on integrating it into other known assemblies rather than developing it anew (Foster and Evans, 2019). The ability for inventors to build on prior knowledge is limited by the breadth of their individual accumulated knowledge assets. Because inventors have highly specialized areas of expertise, they often need to source ideas from other inventors and scientists (Wuchty et al., 2007).

Sourcing knowledge has its own challenges. In its native format, technical knowledge is a list of the experiences an inventor accumulates while developing a technology (Arrow, 1962).

For all but the simplest devices, that list is too detailed for an inventor to recollect let alone communicate (Polanyi, 1966), so inventors compress knowledge by recoding it into diagrams and metaphors (Nonaka and Takeuchi, 1995). These project-oriented coding schemas, however, can only be transmitted using supportive communication technologies. For most of the United States' industrial history, face-to-face communication has held an absolute advantage in communicating messages encoded in such schemas. Face-to-face communication allows for the use of visual clues such as body language and hand gestures to convey complex points, as well as the manipulation of vocal tone to stress key aspects of a message (Storper and Venables, 2004). The interactive nature of face-to-face communication allows speakers to notice misunderstandings and to correct their presentations to improve comprehensibility (Nohria and Eccles 1992), and to create norms, routines, and rhetorical devices that are specifically designed for the technical issues at hand (Powell et al., 1996; Kogut and Zander 1992; Gertler 2003).

Because close spatial proximity is a necessary condition for face-to-face communication, the ability for inventors to source knowledge is influenced by their socio-spatial environments. While inventors that are co-located with many other inventors are able to source a wide range of ideas face-to-face, inventors in isolated regions are at a severe competitive disadvantage. Empirical research shows that the frequency of technical knowledge transmission (Jaffe et al. 1993; Kwon et al., 2020) and the frequency of collaboration (Balland 2012; Van der Wouden 2020) between inventors decline as spatial distance increases. The disadvantage of isolation can be momentarily relieved through travel, but the logistical and economic costs of travel also pose constraints (Torre 2008). Inventors are unlikely to travel for work unless they or their organizations have strong incentives to undertake travel (Morrison et al. 2013). This incentive is a function of the quantity and quality of the knowledge they expect to gather through travel, or the expected value of a resulting product or invention (Cowan and Jonard 2004). The incentive to travel to places with few knowledgeable inventors is therefore small and most non-local flows of knowledge span between regions that already have dynamic inventive milieus (Bathelt et al. 2004, Wolfe and Gertler 2004). Travel to temporary face-to-face meetings is not an exception to this rule because the returns to attending tradeshows and conferences are greatest for inventors and firms based in regions with robust local inventive milieus (Bathelt and Henn; 2014; Esposito and Rigby 2018).

Yet early in the history of every innovative region, there is a moment when local inventors overcome the constraints to sourcing knowledge and begin to produce patentable ideas. Given that the creation of technical knowledge is difficult and competitive, how do inventors in isolated regions accomplish this? A plausible answer is rooted in the interaction of two factors that are loosely conceptualized under the theory of the Window of Locational Opportunity: first, some ideas move to regions with underdeveloped knowledge bases for idiosyncratic reasons; second, there is immense heterogeneity in usefulness of ideas for the creation of new ones (Storper and Walker, 1988; Boschma and Lambooy, 1999). With regard to that heterogeneity, the count of the forward citations of patents, a record of the number of subsequent inventions that build on each patent as prior art, follow an extremely skewed distribution where the majority do not receive a single citation and a small fraction receive more than 100 citations (Hall et al., 2005). Over time, the skew in forward citations grows exponentially as high-impact ideas have more knowledge-based descendants after one generation of endogenous knowledge production, even more descendants after two generations, and so on (Martineli and Nomaler, 2014). Inventors in underdeveloped regions are not able to source many ideas face-to-face, but they nonetheless may be capable of commencing knowledge production if they manage to source a small number of highly-impactful ideas through idiosyncratic means.

The example of the semiconductor-based transistor shows how the transportation of a single promising idea to a "new" location with a minimal history of innovation can revolutionize the geography of knowledge production. The first two transistors, the point-touch transistor and the junction transistor, were invented at Bell Labs in the suburbs of New York City in 1947 and 1948. However, their development into useful tools took place primarily in Silicon Valley, the region of California centered on the San Jose Metropolitan Area. Today, transistor design is the main driver of knowledge production in Silicon Valley: the San Jose Metropolitan Area accounted for 27% of all U.S. semiconductor patents granted between 2001 and 2005.[1] While the design of new transistors is now a complex art around which large teams, firms, and agglomerations have organized (Balland et al., 2019), their design has not always been equally knowledge-intensive. When Bell Labs' star engineer William Shockley, the inventor of the junction transistor, relocated from the New York City area to Silicon Valley in 1956, he brought just one colleague from Bell Labs with him (Gertner, 2012 p. 181). Evidently, the knowledge base needed to design new transistors was simple enough for two people to collectively master at that time (cf. Wuchty et al., 2007). Because early inventions in semiconductors drew from a relatively small body of existing knowledge and could be made by small teams of experts, idiosyncratic factors such as Shockley's preference to be near his ailing mother were able to overcome the structural forces that tend to keep knowledge flows within established milieus.

In fact, in the decades after Shockley invented the junction transistor, the Window of Locational Opportunity in the design and manufacture of amplifiers was opened wide enough for the industry to disperse across the United States. By 1972, only 8.8% of U.S. employment in the SIC code 3674 (the code for semiconductors and related devices) remained in the states of New York and New Jersey, the vicinity of the initial transistor inventions. By that year, 18.3% of U.S. employment in SIC 3674 was in Texas, 5.8% was in Arizona, and 21% was in California (Scott and Storper, 1987). Although these data record employment counts and not technological knowledge per se, they suggest that the knowledge required to design and manufacture semiconductors was not particularly complex and thus was not spatially-sticky in the industry's early years.

Semiconductor design eventually consolidated in Silicon Valley, California in the early 1980s as California's share of U.S. employment in SIC 3674 reached 28.7% in 1982 while the

[1] Semiconductor patents defined as those with USPC 257, 438, or 716 listed as their primary class.

employment share in each of the leading competing states of Texas, Arizona, New Jersey, New York, and Pennsylvania declined (Scott and Storper, 1987). The reason Silicon Valley became the center of semiconductor design is an open debate. Saxenian (1994) and Storper (2016) argue that Silicon Valley's informal institutions encouraged the transmission of ideas and the superior generation of new knowledge. Along those lines, Fleming and Frenken (2007) analyze collaborations between inventors in Silicon Valley and find that principal firms helped the region's inventors to network. Other researchers have emphasized government defense spending as the stimulus for demand for transistors in California (Heinrich, 2002), though US Department of Defense contracts were awarded to several states during that time, including Texas and Arizona where semiconductor production and design ultimately faltered (Scott and Storper, 1987).

Another cause of the concentration of semiconductor design in Silicon Valley in the 1980s is that Silicon Valley's inventors made more impactful inventions than those in the competing regions. The first microprocessor (the Intel 4004), for example, was introduced in Silicon Valley in 1971, which made personal computers and hand-held computer devices possible. Other groundbreaking inventions, such as the first dynamic RAM chip (the Intel 1003, invented in 1970) were also made in Silicon Valley, as were hardware innovations such as the computer mouse in 1964 and the first commercialized computer monitor in 1973. The common aspect of these inventions is that they either used or extended the capabilities of the semiconductor introduced by William Shockley in 1948. In so doing, they increased the number of inventions that the semiconductor made possible.

The geographical history of the semiconductor transistor thus suggests two conditions that must be met for inventors to commence knowledge production in new places. The first is that a promising invention needs to be created. This invention does not need to be made in the same place where it ultimately produces growth, so long as it is simple enough to be transported across space, either through the movement of people or messages (cf. Kerr 2010). The second condition is that the inventors in an emerging innovative center need to out-invent their competitors in other regions. This second condition is critical because if inventors in one city receive a promising idea from afar, it is likely that inventors in other places will receive that idea as well.

3.1) Methods Overview

The study of how inventors initiate local knowledge production has been held back by a dearth of reliable, harmonized, and long-running records of the sources of knowledge that inventors use to make new ideas. Patent citation records have been used by researchers to study knowledge sourcing, but citations carry two limitations. First, because many citations are added by examiners and the attorneys of patent applicants, the extent to which 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 1975, patent citation records are unreliable before this date.

There are, however, implicit historical records of knowledge flow between patents hidden in the classification 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 class codes, at which level the codes describe the individual components that are contained in the patented invention (Fleming and Sorenson, 2001). Because technological knowledge is the ability to assemble components into functioning systems (Fleming, 2001; Arthur, 2009), the detailed classification codes listed on a patent indicate the technical know-how embedded in a technology. An illustrative example is the patent granted to Thomas Edison for the incandescent light bulb (USPTO patent number 223898). Edison's bright idea was that a vacuum chamber slows the combustion of a carbon filament. The physical components Edison used to build his bulb – vacuum-tight joints to seal the bulb and a carbon filament – appear on his patent with the classification codes H01J5/24 and H01K1/14. The USPTO defines these codes as "vacuum-tight joints between insulating parts of vessel" and "incandescent bodies characterized by shape".

Because the classification codes listed on a patent indicate the knowledge embedded in the technology, when two patents share many of the same classification codes, it is reasonable to infer that they draw from the same body of knowledge. Therefore, the flow of knowledge between patents can be predicted by tracing the shared classification codes on patents across time.

3.2) Data Construction

More specifically, I generate predicted flows of technological knowledge between patents by exploiting the information provided by USPC classification 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 the raw public files of granted patents and USPC classifications available on PatentsView. The USPTO reclassifies patents using the USPC coding schema as new classes 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 classes. Most patents are assigned between 2 and 6 classification codes; however, a very small number of patents are assigned more than 100 codes. To make the dataset less cumbersome, I discard excess classification codes on patents by selecting only the first 8 codes from each patent. Selection of the first 8 codes on each patent ensures that I use the primary classification code for each patent which indicates its dominant class.

The tree-building algorithm begins by selecting the most recently granted patent and recording its components based on its USPC classification codes. I define technological knowledge as knowledge of components and the interactions of those components, so I generate all combinations of degree n of its components, where n is the number of components in a patent.[2] For example, if a focal patent (FP) contains the USPC subclassification codes A, B, C, the knowledge vector is generated as follows:

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

Each element in $Knowledge_{FP}$ denotes a single unit of knowledge; the length of $Knowledge_{FP}$ indicates the total quantity of knowledge embedded in the focal patent. The knowledge units in $Knowledge_{FP}$ are used to link the FP to its 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, based on the sequence of patent ID numbers which are ordered by patents' grant date. For each possible parent that fits this 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:

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

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

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

The length of the above $SKnowledge_{FP,PP}$ vector indicates that the focal patent FPsourced 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.

Finally, I identify high-impact patents based on the number of subsequent patents that draw knowledge from a focal patent. I define high-impact patents as those in the top decile of their half-decade cohort in terms of the count of subsequent patents that draw knowledge from them. I then aggregate patents to the metropolitan area level based on the home address of the inventors of that patent. When patents have inventors living in two or more metropolitan areas, I fractionally assign those patents to each metropolitan area.

---

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

4.1) Results Overview

The growth of knowledge production in U.S. metropolitan and micropolitan areas tends to follow a general pattern in which regions begin knowledge production by producing a small number of ideas, expand their production of ideas over time, reach a peak in knowledge production, and thereafter tend to enter a period of decline. Using black dots in Figure 1, I plot the production of patents by half-decade for four representative U.S. metropolitan areas centered on Detroit, Cleveland OH, San Jose CA, and Austin TX. I selected these cities because they are or have been major centers for innovation and because they initiated patent production growth during the time period for which I have reliable data, starting in 1850. To improve the comparison of patent production across years, I express the patents produced by a city in a given half-decade as a percentage of the U.S. total for that half-decade. I also plot the number of high-impact patents produced by that city using plus-signs, with high-impact patents defined as those in the top-decile in terms of the number of their number of knowledge-based progeny, as described in Section 3.2

Figure 1: High-Impact and Total Patents Produced in Four Representative Cities

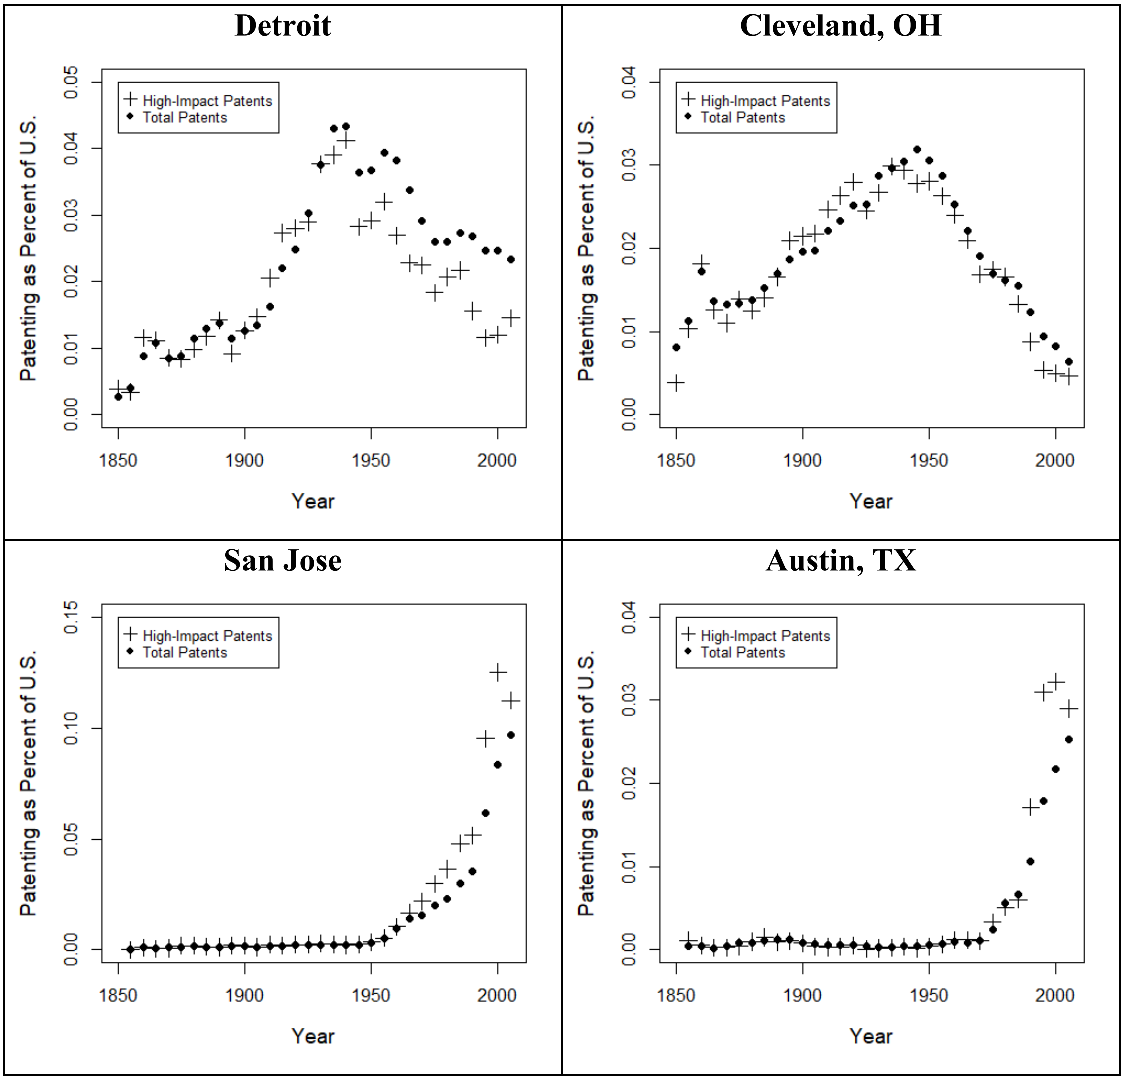

Figure 1 shows the years of 1865 in Detroit, 1855 in Cleveland, 1965 in San Jose, and 1990 in Austin all bear resemblance in terms of patent production growth: during these years, knowledge production in each city started to climb. Additionally, as those cities begin to increase their overall production of patents, they also increased their production of high-impact patents. Generally, their production of high-impact patents grew faster than their production of overall patents; when the plus-signs rise in Figure 1, the black dots rise even faster.

The examples of Detroit, Cleveland, San Jose and Austin in Figure 1 thus suggest two general patterns. First, overall patent production in cities experiences both a rise and a decline; second, the rise and decline of patenting in cities is preceded by the rise and decline of the production of high-impact patents. To test if these patterns are found generally across U.S. metropolitan and micropolitan areas, I compute the average percentage of U.S. total and high-impact patents that each city produces at each stage in its patenting growth. To compare cities that underwent knowledge growth during different periods of time (such as San Jose and Detroit), I align the time dimension of their patenting based on their city "age", defined as the first five-year period a city produces 1% or more of the U.S. total stock of patents. For example, I assume that San Jose in 1965 (when it first produced 1% of U.S. patents) was at the same stage of its growth trajectory as Detroit was in 1865. Extending that reasoning, 35 years later in 2000, San Jose was at the same stage of its growth process as Detroit was in 1900. Formally, I calculate the $Age_{c,t}$ of a city *c* in half-decade *t* by subtracting the observation year from the year it first crosses the 1% patenting threshold:

$$Age_{c,t} = Year_{c,t} - ThresholdYear_c \tag{4}$$

For all subsequent analyses, I show that the results are robust when $Age_{c,t}$ is recomputed using a 0.5%, 1%, 2.5%, or 5% threshold value. After aligning the curves of each city based on $Age_{c,t}$, I compute aggregate patent production curves by averaging the percent of U.S. overall and high-impact patents in a given half-decade that are produced in cities with a given $Age_{c,t}$ value, as in Figure 2.

Figure 2 generates three observations. First, patenting growth in cities appears to be a function of city age as defined in Equation 4. Second, the production of high-impact patents in cities generally increases before the production of overall patenting starts to climb; similarly, the production of high-impact patents starts to decline before overall patenting goes down. Third, the temporal order of the growth of high-impact and overall patent production documented in Figure 2 suggests that a causal relationship runs from the local production of high-impact patents to the increased local overall patenting in subsequent time periods. Moreover, when inventors make high-impact inventions locally, they expand the local knowledge base and enable more local inventions in the future. We will analyze this potential causational relationship in greater detail in Sections 4.4 and 4.5.

Figure 2: Average Production of Total and High-Impact Patents in U.S. Cities by City Age

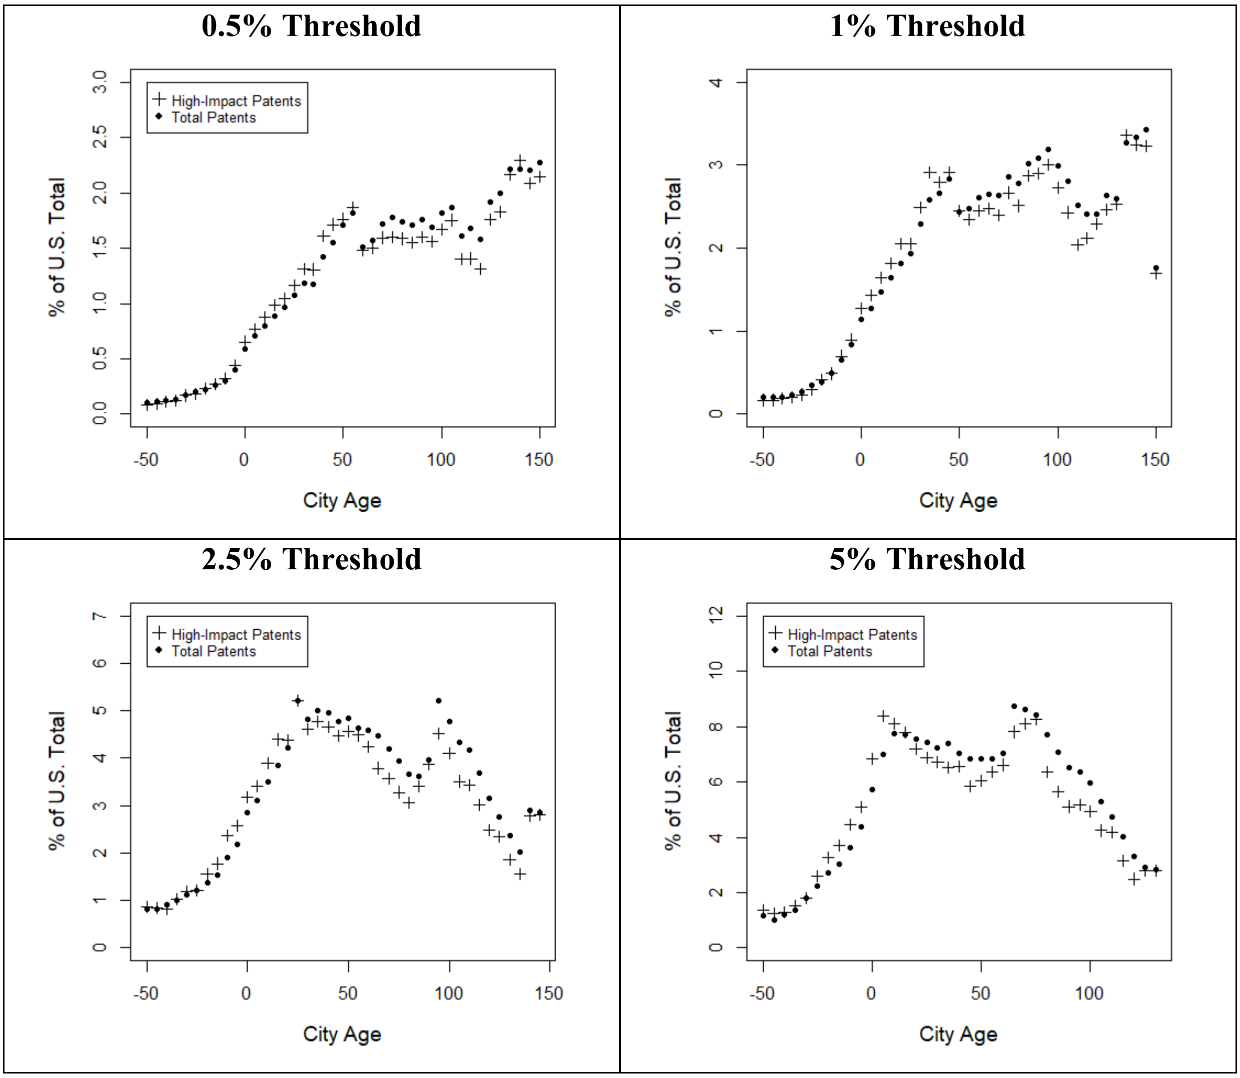

Note: Only cities that exceed the threshold value at one point in their history are included in the analysis. There are 33 cities using the 0.5% threshold, 19 cities using 1%, 10 cities using 2.5% and 4 cities using 5%. Hence, higher thresholds have more noise.

Before addressing causality, it is necessary to address the open question of how inventors access the ideas they use to develop their cities' initial high-impact inventions. Because inventors in cities with nascent knowledge production (when Age < 0) can access few technological ideas locally, the knowledge used to produce local high-impact inventions must come from other places. Section 4.2 explores whether inventors initiate local knowledge production by sourcing non-local ideas.

4.2: The Viability of using Non-Local Inventions to Initiate Local Knowledge Production

While it is difficult for inventors to source big or complex ideas non-locally, inventors in distant places may be able to source promising ideas before those ideas are elaborated or made complex. Figure 3 shows that there is significant heterogeneity in the usefulness of ideas for the creation of new ideas. Until 1930, high-impact patents had on average 4 first-generation descendants, while the average patent had on average one first-generation descendant. After 1930, the average number of first-generation descendants of high-impact patents began to climb rapidly, reaching nearly 15 descendants per patent by the start of the 21st century, while the increase for overall patents was smaller. The heterogeneity in the impact of ideas is even greater across multiple generations of endogenous knowledge production; in a sample of patents from the 1950s, my data show that high-impact patents had on average 6.5 first-generation descendants, each of which had on average 5 descendants of their own. By comparison, low-impact patents had 2.2 first-generation descendants, each of which had 2.3 descendants of their own.

Figure 3: Average Impact of High-Impact Patents and All Patents by Year

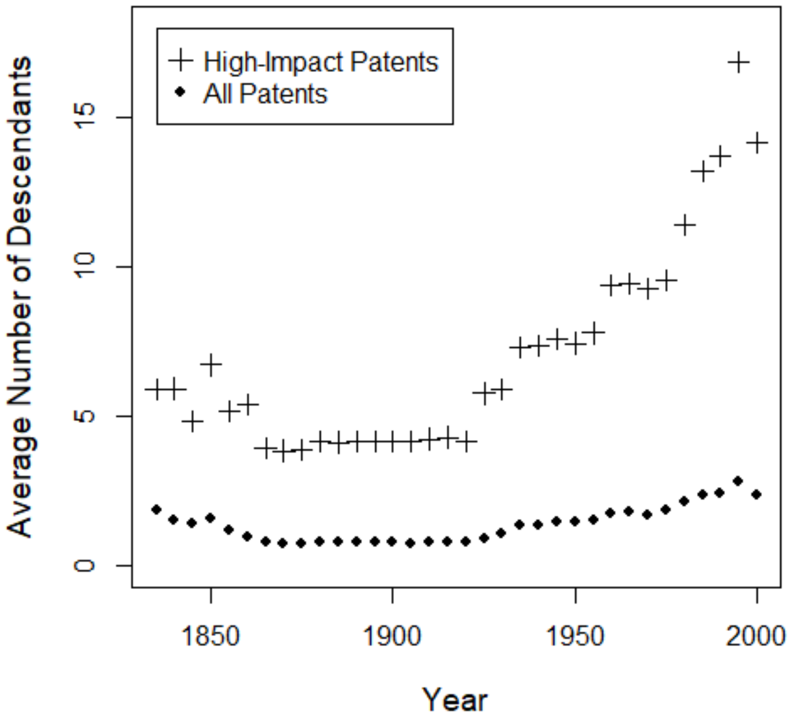

While high-impact inventions eventually are used to create complex technological systems, before their potential is fully realized high-impact ideas are not any more difficult to transport across space than low-impact ones. Figure 4 shows the average percentage of patents' descendants that are produced in the same city as their parent patents, broken out by impact level. If a larger share of a patent's descendants is produced locally, its knowledge travels between regions with greater difficulty. Figure 4 shows that the difficulty to transport knowledge across regional boundaries was highest during three time periods: before 1900, between 1920 and1950, and after 1980. These periods of low knowledge mobility correspond to periods of spatial economic concentration induced by the second and third industrial revolutions (Balland et al., 2019; Kemeny and Storper, 2020). The historical records of the geography of knowledge flow in Figure 4 also show that the recent increase in the importance of spatial proximity for knowledge transmission documented by Sonn and Storper (2005) and Kwon et al. (2019) is a recent phenomenon. More pertinent for our analysis, however, is the postive relationship between the impact of a patent and the percentage of its descendants that that are made in other regions: a larger share of the descendants of high-impact inventions are made by non-local inventors.

Figure 4: Localization of Knowledge Flow by Parent Patent Impact Level and Year

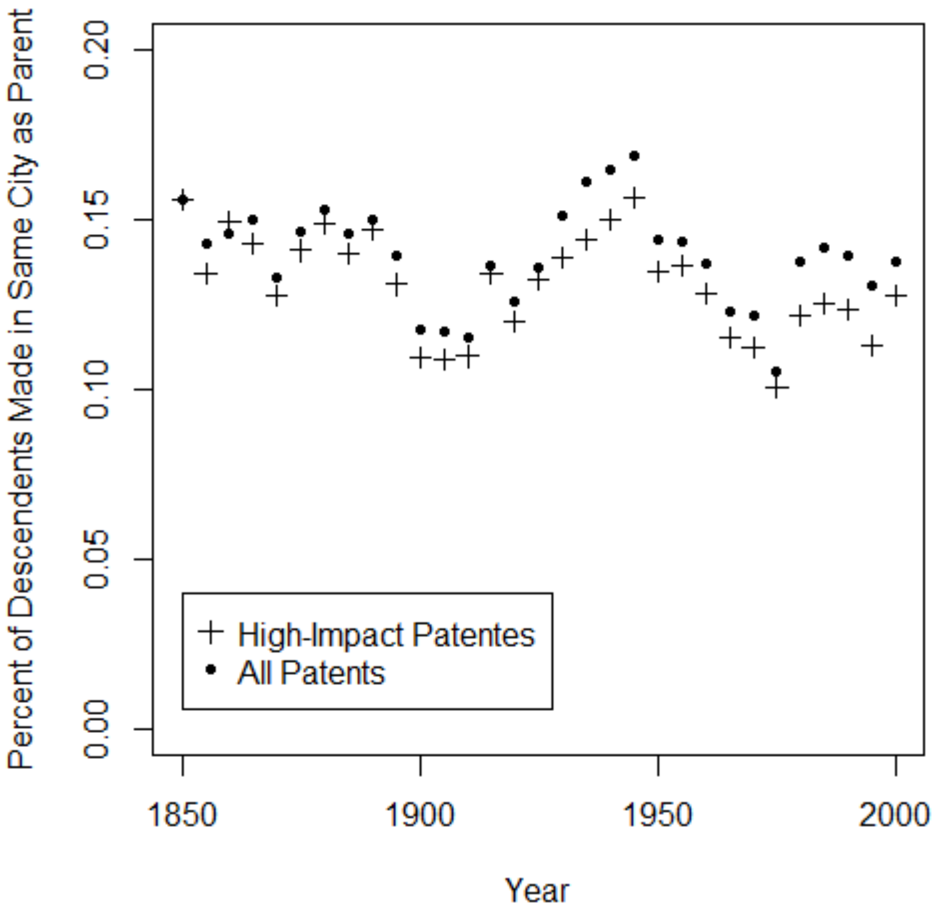

The higher propensity of high-impact inventions to stimulate knowledge production in other cities does not necessarily mean that high-impact inventions cause knowledge to diffuse. Inventors in the city where a high-impact invention is made are still likely to invent a greater share of its descendent inventions than inventors in any other city will. For example, if an inventor in Los Angeles introduces a high-impact invention, we expect from Figure 4 that about 10% of its descendants to be produced in Los Angeles. In the scenario of extreme dispersion, 90% of the focal inventions' descendants will be made in locations scattered across the nearly 1,000 metropolitan and micropolitan areas in the United States, yielding about 0.09% of its descendants in each location. Of course, dispersion is never this extreme; nonetheless, this example makes two related points: high-impact inventions do not cause knowledge production to mechanically diffuse across space; yet at the same time, inventors in far-away regions are not categorically excluded from accessing these promising ideas.

4.4) The Sources of Knowledge Production Growth in Cities

While the results in the previous section indicate that non-local high-impact ideas can be used to commence knowledge production in new places, the more important question is whether inventors actually use those ideas for that purpose. In this section, I empirically examine whether inventors source knowledge from impactful local and non-local sources as their cities initiate, expand, and decline knowledge production. To undertake this analysis, I develop a knowledge-source accounting framework that reveals the relative importance of various types of knowledge flows for instigating knowledge production. Most research on the geography of knowledge flow, such as Jaffe et al. (1993), Breschi and Lissoni (2007), and Arora et al. (2019) infer the geographical consequences of local and non-local knowledge flows based on the friction posed by distance. For reasons discussed in Section 4.3, even if geographical distance exerts strong frictions on the spreading of ideas across distances, non-local ideas can nonetheless be used to initiate knowledge production in new places. The source-accounting framework developed here overcome this analytical challenge by directly calculates the relative importance of local, non-local, high-impact, and low-impact knowledge sources for inventors to initiate local knowledge production.

Specifically, for each city and in each half-decade, I calculate the extent to which a metropolitan area's inventors source knowledge from local high-impact inventions (L.HIGH), local low-impact inventions (L.LOW), non-local high-impact inventions (NL.HIGH), and non-local low-impact inventions (NL.LOW) as knowledge production in the region begins, expands, and declines. I begin by defining the patent stock produced in city *c* in a half-decade t as $P_{c,t}$. Next, using records of knowledge flow between patents (see Section 3.2), I trace the patents in $P_{c,t}$ one generation back in time to identify their parent patents, defined as $PatParents_{c,t}$. Finally, I calculate the percentage of $PatParents$ that were invented in the same city as $P_{c,t}$, and the percentage of $PatParents_{c,t}$ that are high-impact inventions. I also calculate the cross-tabulations, yielding the measures L.HIGH, L.LOW, NL.HIGH, and NL.LOW. To illustrate these calculations, in Figure 5 I plot the composition of knowledge sources used by inventors in Detroit and San Jose by half-decade periods. To visualize how their knowledge sources evolve as their cities expand their patent production, I overlay the number of patents produced in the city (as a percentage of U.S. patent production in the same half-decade) using lines.

Figure 5 illustrates three important relationships. First, inventors in both Detroit and San Jose sourced a large percentage of knowledge from non-local high-impact (NL.HIGH) patents when their cities started to produce an appreciable amount of technological knowledge, at the turn of the 20th century in Detroit and during the second half of the 20th century in San Jose. The percentage of knowledge sourced from NL.HIGH patents, however, declined as their regional patent production grew: in 1905, Detroit's inventors sourced about 75% of the knowledge used to create their new inventions from NL.HIGH patents, but by the time Detroit's patent production peaked in 1940, that figure declined to about 60%. A similar albeit less sharp decline in knowledge sourcing from NL.HIGH patents occurred in San Jose

1950-2000. Second, as knowledge production grew in both cities, the inventors started to source knowledge more frequently from L.HIGH patents. In Detroit, L.HIGH patents accounted for about 1% of all knowledge sourced in 1905 but rose to about 10% by the 1920s. The increase was even larger in San Jose. Third, in both cities the sourcing of knowledge from L.LOW patents grew less quickly than the sourcing of knowledge from L.HIGH patents. Between 1905-1940 in Detroit and 1970 through the end of the data series in San Jose, inventors in each city sourced more knowledge from L.HIGH patents than from L.LOW patents.

Figure 5: Composition of Knowledge Sources Used for Patent Production in Detroit and San Jose

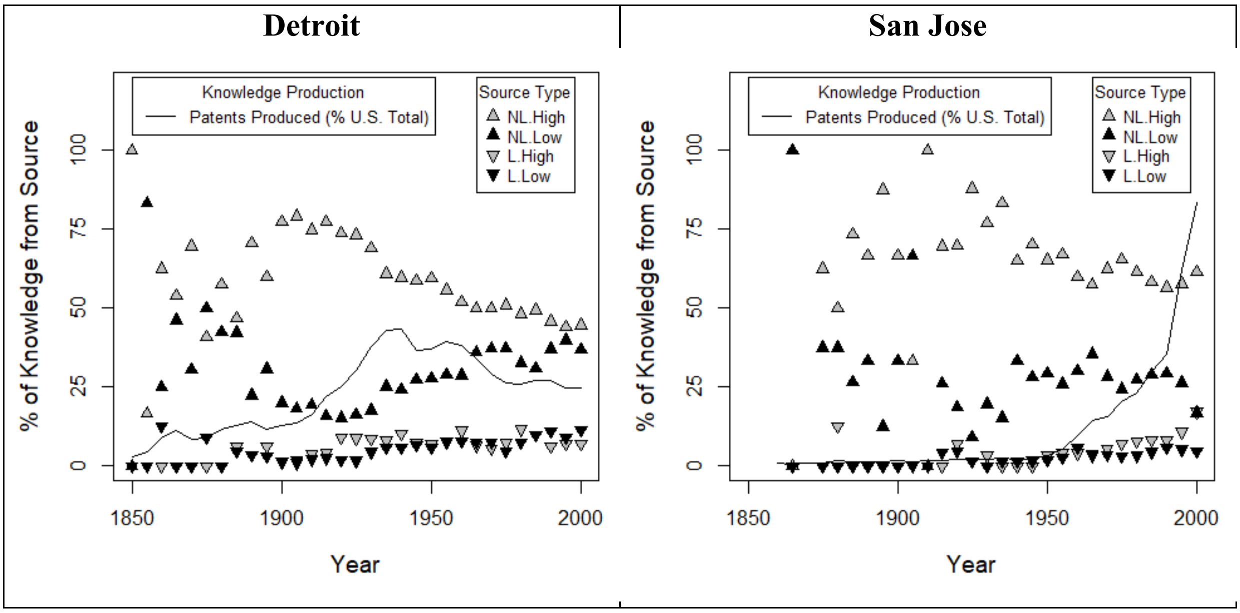

Note: Patent production (black squares) is not plotted to scale.

To test whether the relationships identified in Figure 5 are found more generally across cities in the U.S., I plot the averages of the decomposition of the sources of knowledge used by inventors across U.S. cities by their city "age" in Figure 6. As in Figure 2, I compute the age of a city as the number of years elapsed since each city first produces a threshold percentage of all U.S. patents. As before, I omit all cities that exceeded the threshold level of patent production before the start of my data series in 1850 and all cities that never break the patenting threshold.

Figure 6: Sources of Knowledge for Knowledge Production in Cities by City Age

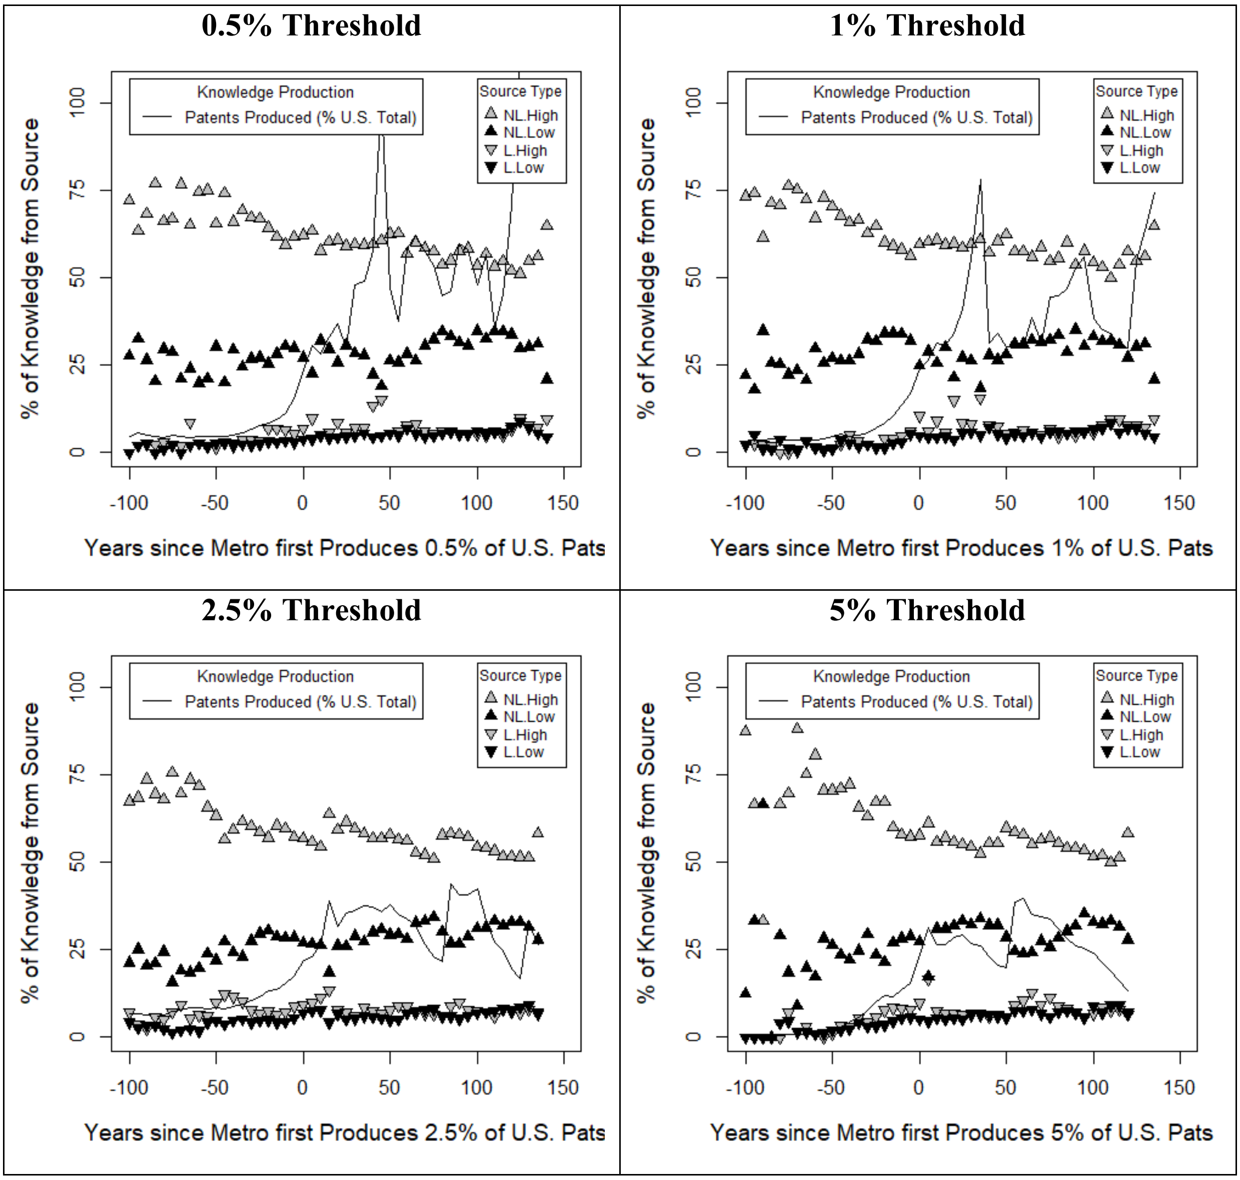

Note: Patent production (black squares) are not plotted to scale.

Figure 6 shows that the relationships identified in Detroit and San Jose are found more generally across U.S. cities. Inventors source knowledge from NL.HIGH inventions most frequently when their cities first commence patent production; they source an increasing share of knowledge from L.HIGH inventions as local patent production grows; and they source a larger share of their knowledge from L.LOW and NL.LOW inventions when local patent production reaches its peak and declines. As shown in Figure 6, these results are robust across the 4 different threshold values used to benchmark the age of the city. To provide a simpler presentation, Table 2 calculates the composition of sources used by inventors in three discrete ranges of city age. For brevity, Table 2 only reports values using the 1% threshold definition.

Table 2: Composition of Knowledge Sources by Source Type and by City Age

Age Range	Source Type NL.HIGH	NL.Low	L.High	L.Low
-100 to 0	67%	28%	3%	2%
0 to 50	60%	26%	9%	5%
50 to 150	56%	31%	7%	6%

Note: Table 2 uses the 1% patenting threshold

4.5) Results: Why Does Knowledge Production Consolidate in Certain Locations?

While Figure 6 and Table 2 show that knowledge production growth expands in cities as their inventors source ideas from NL.HIGH and L.HIGH sources, it remains unestablished whether these sources cause local knowledge production to expand. A full causal inference is beyond the scope of this paper; however, a meaningful counterfactual can be generated by comparing the knowledge sources used by inventors in cities that break the patenting threshold with the sources used by inventors in cities that fail to break that threshold.

To compare the types of knowledge sources in “successful cities” (SC; those that break the threshold) and “unsuccessful cities” (USC; those that fail to), I compute the frequency their inventors use each type of source at each age value. USCs do not have explicit age values because they never break the patenting threshold, so direct comparison is impossible. However, a robust comparison can still be generated by comparing the composition of knowledge sources used in SCs with the composition used in USCs in the same 5-year period of observation. To generate this within-time comparison, let the vector $Years_{Age}$ record the all 5-year periods in which SCs are observed at a given age value. For example, if a total of three SCs in reach age 10, the first in 1900 and the second and third in 1995, then

$$Years_{Age=10} = [1900, 1995,1995] \quad (5)$$

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

$$\%Parents_{USC,Type,Age} = \frac{\sum_{Years}\sum_{USC} PatParents_{Years,USC,Type}}{\sum_{Years}\sum_{USC} PatParents_{Years,USC}} \quad (6)$$

In Equation 6, the *Years* subscript is an index of the $Years_{Age}$ vector.[3] To provide an example of how Equation 6 is computed, let us make three assumptions:

(1) $Years_{Age=10} = [1900, 1995,1995]$, as in Equation 5
(2) $PatParents_{Years,USC} = [50,100,100]$, denoting that patents in in unsuccessful cities sourced knowledge from 50 parent patents in 1900 and 100 parents in 1995
(3) $PatParents_{Years,USC,Type=NL.HIGH} = [5,8,8]$, denoting that patents in in unsuccessful cities sourced knowledge from 5 high-impact parents in 1900 and 8 in 1995

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

$$\%Parents_{USC,Type=NL.HIGH,Age=10} = \frac{5+8+8}{50+100+100} = 8.4\% \tag{7}$$

Finally, to compare the composition of knowledge sources used by inventors in SCs with the composition used in USCs, the SC composition premium is taken as the difference between the compositions used in SCs and USCs:

$$SCCompositionPremium_{Type,Age} = \%Parents_{SC,Type,Age} - \%Parents_{USC,Type,Age} \tag{8}$$

In Figure 7, I create scatterplots of $SCCompositionPremium_{Type,Age}$ by plotting it against city age. I overlay Loess regression fit lines (search distance = 100%) to identify general trends in the data across age values.

Figure 7: Knowledge Sources used by Cities that Break the Patenting Threshold in Excess of Knowledge Sources used by Cities that Never Break the Patenting Threshold

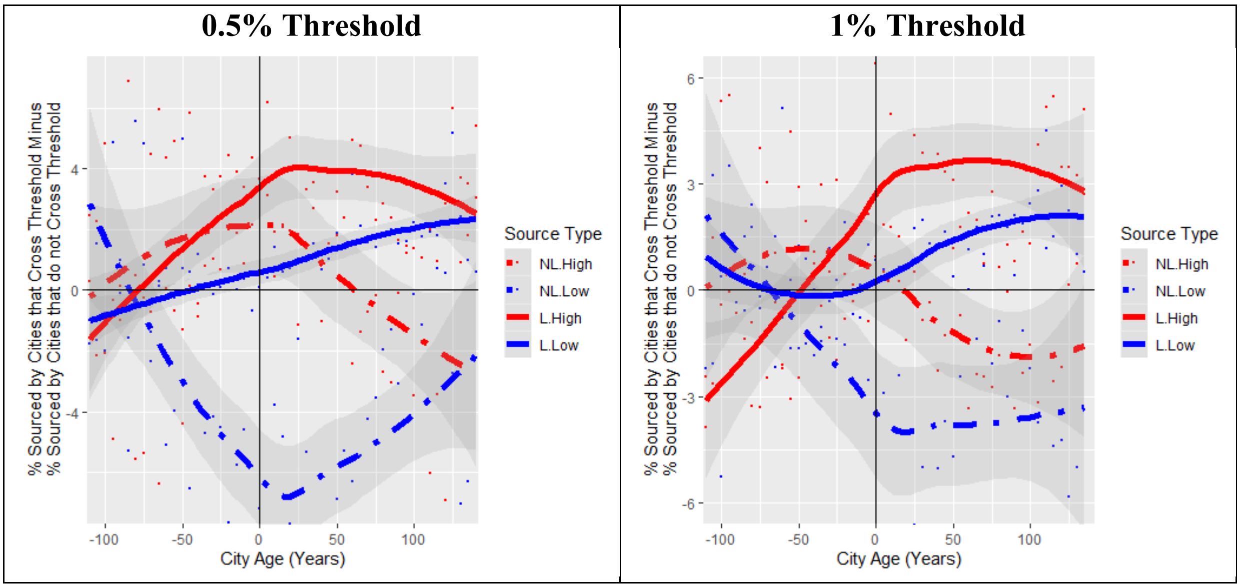

[3] The 1995 value is double-counted because it appears twice in the $Years_{Age=10}$ vector, which amounts to taking weighted means.

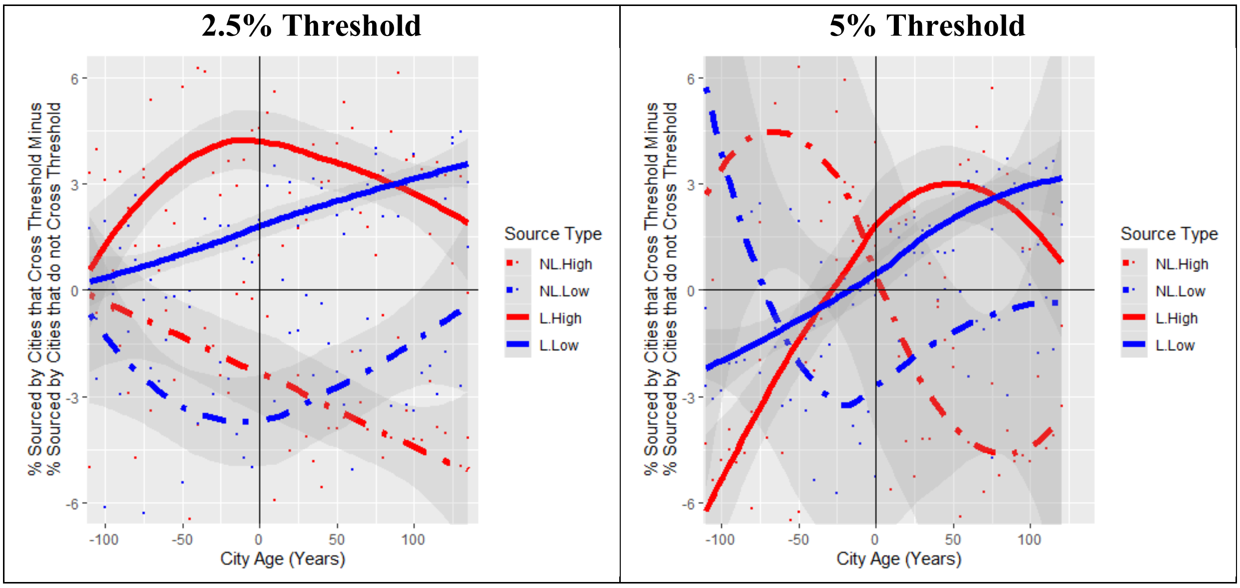

Figure 7 shows that inventors in SCs use a different composition of knowledge sources than inventors in USCs. These differences occur during and after SCs cities cross the patenting threshold. Early in their growth, when age is between -100 and 0, inventors in SCs source a larger share of knowledge from NL.HIGH patents, as indicated by the high dotted red lines using the 0.5%, 1%, and 5% threshold value definitions. The 2.5% definition did not yield this result, which may be due to the small number of SCs (10 cities) under that threshold definition. As SCs start to age, their inventors increasingly source knowledge from L.HIGH patents, as indicated by the solid red lines which peak when Age = 0. As SCs age further, their inventors increasingly source knowledge from local low-impact inventions, as indicated by the solid blue lines that rise over time. The percentage of knowledge that inventors in SCs source from L.LOW patents peaks after the percentage that sourced from NL.LOW does, indicating that L.LOW are the primary knowledge sources used by inventors as knowledge production in their home regions matures and declines.

From Figure 7 it is difficult to discern whether the differences in the composition of knowledge sources used by inventors in SCs and USCs are statistically significant. I thus develop two regression models. First, I develop a logit model of the probability that a city reaches the 1% patenting threshold in a regression framework as a function of the types of knowledge that its inventors used over the previous 50 year:

$$Prob\left(BreakThreshold_{c,t} = 1\right) = \\ B_1 * \%NL.HIGH_{c,tPrev} + B_2 * \%L.HIGH_{c,tPrev} + B_3 * \%NL_{c,tPrev} + FE_t + E_c \tag{9}$$

where the probability that a city breaks the threshold in the half-decade *t* is a function of the percentage of the knowledge its inventors source from NL.HIGH and L.HIGH sources during the previous 50 years. Moreover, the *tPrev* time index sums over the range of *t* values [t-50, t-45, . . . t-5]. To account for the small growth in the local knowledge stock that precedes the breaking of the 1% threshold, I include the percentage of parents that are non-local (NL) as a control variable in some model variants. Time fixed effects control for changes in the relationship between the various knowledge sources and the propensity for cities to break the threshold across time. Cities that break the patenting threshold in a prior time period are excluded from the model so that the model includes all observations of USCs but only the observations of SCs up to and including the 5-year period they break the threshold. Results are given in Table 3.

Table 3: Model of Prob(BreakThreshold =1)

	Breakthrough Threshold					
	1%		**0.5%**		**2.5%**	
% NL.HIGH	1.01*** (0.272)	1.33*** (0.318)	1.10*** (0.215)	1.28*** (0.255)	0.797** (0.384)	1.27** (0.567)
% L.HIGH	5.64*** (0.796)	1.43 (0.961)	6.47*** (0.919)	3.18*** (1.16)	7.79*** (1.29)	3.17** (1.45)
% NL		-4.41*** (0.561)		-3.27*** (0.827)		-4.91*** (.847)
FE_t	Y	Y	Y	Y	Y	Y
NOBS City*Time	14130	14404	13300	13738	14681	14854
NOBS where $BreakThreshold_{c,t} = 1$	19	19	33	33	10	10

Standard errors clustered at the CBSA level

The logit estimates in Table 3 indicate that cities are more likely to break the patenting threshold if their inventors source a larger share of their total knowledge from NL.HIGH or L.HIGH inventions over the previous 50 years. These results are robust to the threshold definition used, although the 2.5% threshold only leaves 10 cities that break the threshold in the dataset.

In addition to enabling cities to break the patenting threshold, sourcing knowledge from NL.HIGH and L.HIGH sources may also help inventors in those cities expand local production of patents in the following years. To test this proposition, I develop a similar regression where the dependent variable is replaced by $PatentProduction_{c,tfuture}$ which calculates the percentage of all U.S. patents that a city produces over the next 50 years. The model tests whether cities that source a large share of their knowledge from NL.HIGH and L.HIGH patents over the previous 50 years produce more patents over the following 50 years. The OLS model is given by:

$$PatentProduction_{c,tFuture} = B_1 * \%NL.HIGH_{c,tPrev} + B_2 * \%L.HIGH_{c,tPrev} + B_3 * \%NL_{c,tprev} + FE_t + E_c \tag{10}$$

Equation 10's estimates are given in Table 4. Again, cities that break the patenting threshold are only included in the model for the first time period during which they break that threshold.

Table 4: Model of $\boldsymbol{PatentProduction_{c,tFuture}}$

	Breakthrough Threshold					
	1%		0.5%		2.5%	
% NL.HIGH	0.0120*** (0.0000247)	0.0240*** (0.0000262)	0.0130*** (0.00157)	0.0157*** (0.00162)	0.0354*** (0.00507)	0.0446*** (0.00581)
% L.HIGH	0.412*** (0.000541)	0.208*** (0.000469)	0.228*** (0.0277)	0.0990*** (0.0262)	1.04*** (0.173)	0.581*** (0.128)
% NL		-0.191*** (0.0236)		-0.121*** (0.013)		-0.427*** (0.0692)
FE_t	Y	Y	Y	Y	Y	Y
R-Squared	0.094	0.104	0.108	0.116	0.0822	0.0902
NOBS City*Time	14130	14130	13738	13738	14854	14854

NB: Standard errors clustered at CBSA level.

Table 4 indicates that sourcing knowledge from NL.HIGH and L.HIGH inventions is positively associated with future patenting growth above and beyond the patent threshold level.

5) Discussion

This paper has used data on knowledge production and its sourcing to study how inventors initiate and expand knowledge production in new places. The creation of technological knowledge is a complex art that is aided by building on existing ideas. In regions without a history of knowledge production, the number of existing ideas that inventors can access through face-to-face communication is limited and so inventors tend to create new technologies in established milieus. Occasionally, new and impactful ideas are generated that are less reliant on existing stocks of knowledge (Kuhn, 1962; Dosi, 1982). Certainly, many of the subsequent inventions enabled by these breakthroughs are realized within immediate environs where the breakthroughs are initially made. However, as the example of the semiconductor transistor showed, it is possible for inventors to commence knowledge production in new places if they manage to source impactful ideas created far away *before* the impact of those ideas is realized. Ideas as promising as the junction transistor arise infrequently, but regions rarely emerge as innovative centers without seizing at least one such opportunity.

While a wealth of research examines how geographical proximity enhances the generation, transmission, and retention of technological knowledge, the environments in which these geographically-proximate interactions take place are created through the actions of economic actors over the course of time. The formation of these environments is an under-researched subject and there is significant room for further analysis. To begin, this study inferred the effect of high-impact knowledge sources on the creation of innovative environments through the temporal coincidence of sourcing high-impact knowledge and local patenting growth. A more rigorous identification would quantify the number of new local ideas that each previous local and non-local idea helped ferment. Secondly, as the examples of semiconductor employment in Phoenix, AZ and Dallas, TX show (see Section 2), inventors in some locations manage to source high-impact non-local ideas and yet fail to sustain long-run local innovation. The results of this study indicate that inventors also need to introduce high-impact ideas locally in order to capture and sustain knowledge production in the long run. This finding nonetheless begs an additional question: why do inventors in some regions introduce more high-impact inventions than in others? For Saxenian (1996) and Storper et al. (2015), competitive advantage in the creation of impactful new technologies is derived from malleable local institutions that allow regional actors to develop new methods for organization and coordination. Accordingly, the variation in the fluidity of inventor networks across regions and its association with the local creation and capture of high-impact inventions is a promising area of for further research.

Finally, while the study indicated that high-impact ideas provide opportunities for inventors to commence knowledge production in new places, too little is still known about what makes high-impact ideas impactful. In addition, further research is needed to reveal why some high-impact ideas stimulate knowledge production in far-away places while other impactful ideas stimulate knowledge production in their immediate environs. Both of these issues may be addressed by generating more accurate models how new ideas influence the usefulness of existing ones. Sometimes, inventors integrate new and impactful ideas into existing technological systems while in other cases they develop technologies that drive existing technological systems into obsolescence. These technological relationships have regional consequences as certain inventions promote regional diversification and resilience (Neffke et al., 2013; Rigby, 2015; Boschma 2015) while other inventions render regional knowledge bases obsolete and establish new outposts for innovation (Scott and Storper, 1987; Storper and Walker, 1989; Boschma and Lambooy 1999). Ecological models of symbiotic and adversarial relationships between inventions such as Foster et al. (2014) are encouraging starting points to unpack these conflicting sources of evolution and revolution in the geography of innovation.

Bibliography

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

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

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

Arthur, B. (2009). *The Nature of Technology*. Free Press. New York.

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

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

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

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

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

Bathelt, H., 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.

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

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

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

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

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

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

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

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

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

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

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.

Garfinkel, H. 1987. *Studies in Ethnomethodology*. Oxford: Blackwell.

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

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

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

Hall, B., Jaffe, A., and Trajtenberg, M. (2005). Market Value and Patent Citations. *The RAND Journal of Economics* 36(1). 16-38.

Heinrich, T. (2002). Cold War Armory: Military Contracting in Silicon Valley. *Enterprise and Society* 3(2). 247-284.

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

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

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

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

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

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

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

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

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

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

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

Neffke, F., Henning, M., and Boschma, R. (2013). How do regions diversify over time? Industry relatedness and the development of new growth paths in regions. *Economic Geography* 87(3). 237-265.

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

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

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

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

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

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

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

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

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

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. (1989). *The Capitalist Imperative: Territory, Technology, and Industrial Growth*. Basil Blackwell, Inc. New York.

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

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

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

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

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