Papers in Evolutionary Economic Geography

12.18

The Geography of Knowledge Relatedness and Technological Diversification in U.S. Cities

David L. Rigby

Utrecht University
Urban & Regional research centre Utrecht

The Geography of Knowledge Relatedness and Technological Diversification in U.S. Cities

DAVID L. RIGBY
Department of Geography
UCLA

ABSTRACT

U.S. patent data and patent citations are used to build a measure of knowledge relatedness between all pairs of 438 major patent classes in the USPTO. The knowledge relatedness measures, constructed as the probability that a patent in class j will cite a patent in class i, form the links of a patent network. Changes in this U.S. knowledge network are examined for the period 1975 to 2005. Combining the knowledge network with patent data for each of the CBSAs in the United States permits analysis of the evolution of the patent knowledge base within metropolitan areas. Measures of knowledge relatedness are employed to explain technological diversification and abandonment in U.S. cities.

knowledge relatedness technological diversification patents citations

JEL classifications: O18, O31, R12

The Geography of Knowledge Relatedness and Technological Diversification in U.S. Cities

INTRODUCTION

Regardless of whether the theory of growth is guided by political economy or more orthodox economic frameworks, the accumulation of knowledge and technological change are viewed as central to economic performance (SCHUMPETER, 1950; SOLOW, 1956; MARX, 1970; NELSON and WINTER, 1982; ROMER, 1986; LUCAS, 1988). For economic geographers, concern with the creation of competitive advantage at the regional level focused attention on the ability of place-based economic agents to generate, or otherwise acquire, economically relevant knowledge, and on their capacity to use that knowledge effectively (VON HIPPEL, 1988; COHEN and LEVINTHAL, 1990). While for developing regions of the global economy that capacity might be manifest in the production of large volumes of standardized outputs at low unit prices, within industrialized regions it frequently rests on the invention of new product or process technologies and the creation of monopoly rents (LUNDVALL, 1992; STORPER, 1997; MASKELL and MALMBERG, 1999).

At least since the work of MARSHALL (1920), SCHUMPETER (1950) and PERROUX (1955), economic growth has been considered “lumpy”, or discontinuous in both a spatial and a temporal sense. While geographical lumpiness might once have been seen as the outcome of localized natural resources, sunk capital investments and market exchanges between a region’s firms and industries, it is increasingly viewed as a product of the stickiness of tacit knowledge and the difficulty of developing the creative *terroir* between individuals, firms, and the panoply of supporting institutions and ‘untraded interdependencies’ from which knowledge is born (POLANYI, 1966; LUNDVALL, 1988; STORPER, 1995; MASKELL and MALMBERG, 1999). These ideas have spawned a massive literature on regional innovation systems (FREEMAN, 1985; COOKE *et al.*, 1997), learning regions (MORGAN, 1997; LUNDVALL and JOHNSON, 1994) and localized knowledge economies more generally. MACKINNON *et al.* (2002) and ASHEIM and GERTLER (2005) offer reviews.

While considerable effort has been directed towards uncovering what it takes to be a learning region or a knowledge economy, much less attention has been given to the character of knowledge produced within such regions. What does the technology or the knowledge base of a region look like? How do the knowledge cores of regions vary over space and time? How might we measure that variance to gauge the degree of specialization or the relatedness of a region’s knowledge base, and how is that variance connected to the evolution of technology within the region? These are key questions for economic geographers who seek to understand regional differences in knowledge production and economic growth.

This paper proposes a way of using information contained within patents to answer some of the questions just posed. The arguments of the paper are laid out in five sections that follow. Section 2 provides additional motivation for this work and a brief review of related literature. Section 3 outlines a method of using patent citation data to measure the proximity of the different technology fields into which patents are placed. These measures form the links of a network that represents a U.S. knowledge space. Visualizations of that space are shown for the period 1975 to

2005. In Section 4, measures of proximity between patent technology fields are combined with metropolitan patent data to calculate the relatedness of invention, an index of knowledge specialization, in U.S. cities. In Section 5, the relatedness of urban invention is employed to model technological diversification and technological abandonment, key components of the evolution of urban knowledge cores. Section 6 concludes the paper.

LITERATURE REVIEW

At any moment of time, a regional economy may be conceived as a collection of economic agents, firms and workers, embedded within a set of organizational and institutional structures that guide behavior to a greater or lesser extent. As such, regions are repositories of accumulated knowledge, both codified and tacit, that is to varying degrees place-bound, locked within capital stocks of different vintage, within firm routines, within workers of varying skill, education and experience, and within the broader social capital of the region. We typically imagine the region’s knowledge base to comprise familiarity with the production of particular commodities (industry mix) and of specific techniques used in their production. However, the “softer capital” associated with behavioral conventions that regulate activities within and between firms, as well as the political and economic environments within which firms and workers compete, are increasingly understood to shape the ways in which knowledge is produced and distributed over space (STORPER, 1995; STORPER 1997).

The knowledge base of regions change over time through deliberate processes of search, through invention, learning and imitation that involves multiple actors engaged in competition. We know that the pace of such change is uneven, with some regions locked into particular technological regimes that yield diminishing returns (GRABHER, 1993) while other regions seem better able to maintain their capacity to innovate (SAXENIAN, 1994). Cloning the hard and soft capital of fast growth, high technology regions does not appear to be a viable competitive strategy, in part because some technologies do not travel well (GERTLER, 1995; 2003). Thus, the technological trajectories of most regions are relatively stable (RIGBY and ESSLETZBICHLER, 1997). The costs of technological diversification and technological abandonment assure that this is so.

From DAVID (1975), NELSON and WINTER (1982) and DOSI (1982), we hypothesize that firms accumulate knowledge about technology largely through experimentation with existing techniques. Localized search is thought to be conditioned by bounded rationality (SIMON 1959) and by sharply declining returns to investment in research and development efforts that are relatively dissimilar to existing technology, and by the costs of knowledge acquisition that rise steeply around the boundaries of existing knowledge bases (ATKINSON and STIGLITZ, 1969; WEBBER *et al*., 1992; ANTONELLI, 1995). Much of the literature on learning economies has extended these arguments, deepening our understanding of the different dimensions of proximity (BOSCHMA, 2005) and their importance to the regional evolution of technology and competitive advantage (KIRAT and LUNG, 1999). Yet, empirical research on knowledge production and its evolution over space, how regions diversify and add to their technological repertoire, and what kinds of technologies they abandon, remains underdeveloped.

Some clues are offered by research on the history of technological development within the firm. In early work on technological spillovers, JAFFE (1986) provides a measure of technological distance between the knowledge portfolios of different firms. A related idea is advanced by TEECE *et al*. (1994) who examine the coherence of knowledge developed by individual firms and show that technological diversification within firms is closely linked to their existing knowledge base. Similar measures of technological or knowledge proximity are used by ENGELSMAN and VAN RAAN (1994) to “map” technological fields. In more recent work, BRESCHI *et al*. (2003) and LETEN *et al*. (2007) extend earlier arguments, offering different ways of calculating knowledge relatedness to explore the nature of technological diversification and firm performance.

In related, and more explicitly geographical, research VERTOVA (1999) and CANTWELL and VERTOVA (2004) use patent data to illustrate the cumulative nature of technological development at the country-level, and the long-run relationship between country size and level of technological diversification. HAUSMAN and KLINGER (2007) and HIDALGO *et al*. (2007)) employ detailed trade data to measure the relatedness of products through patterns of co-exporting. They argue that specialization in the production of particular commodities provides countries with a set of capabilities that constrains diversification to related products. Thus, countries specializing in goods that are located in densely populated parts of “product space” can transition relatively easily among different product-sets. Countries specializing in products that are relatively isolated in product space have more narrowly-defined sets of capabilities that hinder diversification.

BOSCHMA *et al*. (2012) extend the work of HAUSMANN and KLINGER (2007) at the sub-national level exploring how different regions in Spain diversify into industrial sectors that are related to their existing product-based capabilities. Similar ideas are exploited by NEFFKE *et al*. (2011) to reveal path-dependence in the evolution of the industrial landscape in Sweden, and the more general concept of related variety is developed by FRENKEN *et al*. (2007) to explain differences in regional employment growth in the Netherlands. QUATRARO (2010) shows that knowledge variety and coherence play a significant role in productivity growth across Italian regions, while BOSCHMA and IAMMARINO (2009) demonstrate the importance of related variety to regional growth in Italy, alongside the role that trade plays in spilling related knowledge over space.

This paper builds on some of the ideas above to explore the relatedness of technological fields identified in U.S. patent data. A U.S. knowledge space is mapped and the evolution of this space is tracked since 1975. Knowledge relatedness within U.S. metropolitan areas is employed to explain patterns of technological diversification and abandonment.

THE U.S. KNOWLEDGE SPACE

One of the primary reasons we know so little about the spatial demography of knowledge is that we lack precise measures of knowledge and technology (PAVITT, 1982). Consequently, researchers have long made use of proxies such as the “high tech” industry mix of a region’s economy (HALL and MARKUSEN, 1985), or the distribution of “knowledge workers”, equated with particular sectors, occupations, and creative potential (FLORIDA, 2002; FESER, 2003). Unfortunately, these proxies are noisy and they do not tell us much about the nature of the knowledge or technology created in different places. We can do better, perhaps, by developing more sophisticated, multi-dimensional measures of the technologies employed by individual plants/firms within a region, and examining how far those technologies might be from an industry average or from a technological frontier. We can also examine the range of techniques available within different economic spaces and link technological heterogeneity to the expected rate and direction of invention and innovation (WEBBER *et al*., 1992), or to the different mechanisms by which techniques of production are altered (RIGBY and ESSLETZBICHLER, 2006). However, the data requirements of these measures are relatively high, often demanding access to establishment-level micro-data. And, outside manufacturing, these possibilities are typically unavailable.

Increasingly, we have turned to various measures of the inputs and outputs of invention and innovation to track knowledge production. On the input side, research and development spending has been shown to be closely correlated with counts of patents and innovations that typically form our indicators of knowledge outputs (FELDMAN, 1994) Many of these linkages are exploited by work on knowledge production functions and spatial derivatives thereof (GRILICHES, 1979; ACS *et al*., 2002). Patent data have become a staple for those interested in the geography and history of knowledge production (LAMOREAUX and SOKOLOFF, 1996; O'HUALLACHAIN, 1999; JAFFE and TRAJTENBERG, 2002; O'HUALLACHAIN and LEE, 2010), in inventors and inventor networks (BRESCHI and LISSONI, 2001; SINGH, 2005), in knowledge flows or spillovers (JAFFE *et al*., 1993; SONN and STORPER, 2008), in geographical and cognitive proximity (FISCHER *et al*., 2006), and on the types of knowledge produced (HALL *et al*., 2001; STRUMSKY *et al*., 2010).

The popularity of patent data is related to their availability and to the wealth of information that they provide. At the same time, the disadvantages of patent statistics as overall indicators of economic and inventive activity are legion (PAVITT, 1985; GRILICHES, 1990), and we are becoming increasingly aware of the difficulties of extracting increasingly sophisticated information from patent and citation records (BRESCHI and LISSONI, 2004; THOMPSON and FOX-KEANE 2005; ALCACER and GITTELMAN 2006). Bearing these difficulties in mind, and recognizing that patents do not represent all knowledge production, the primary aim of this section of the paper is to develop a measure of technological relatedness between patent classes, to use this measure to examine the technological specialization, or coherence, of knowledge production in different regions of the U.S., and to explore how this has changed over time. Patent technology class data and patent citations form the core of this work. These data are drawn from the United States Patent and Trademark Office (USPTO).

Upon review, individual patents are placed into one or more distinct technology classes that are supposed to reflect the technological characteristics of the underlying knowledge base that they embody. By the end of 2009, there were 438 such classes of utility patents in use by the USPTO. It is important to note that these technological classes do not remain constant over time. Through its bi-monthly "classification orders" the USPTO redefines classes, it adds new classes and, though rare, removes obsolete ones. Fortunately, the USPTO also reclassifies patents, providing the researcher with a consistent set of technology classes into which patents are placed for specific periods of time. Patents may be placed into a number of different technology classes, consistent with the range of knowledge that they introduce, though each granted patent is also allocated a primary technology class on the basis of the extent of the novelty generated across different technology fields. The research below focuses upon these primary technology classes.

Combining data on the location of inventors, the date of application of a patent and technology class information allows researchers to develop a technological profile of different places and to track changes in those profiles over time. Thus, ARCHIBUGI and PIANTA (1992) examine the links between country size and specialization in patenting, and CANTWELL and VERTOVA (2004) show that countries of a given size have become more technologically focused over time. At the sub-national level, KOO (2005) makes use of a concordance between patent technology classes and the standard industrial classification to identify geographical clusters of industries in the U.S. that are linked by knowledge flows, while O'HUALLACHAIN and LEE (2010) examine specialization and diversity in knowledge production across U.S. metropolitan areas.

Measurement of technological specialization is a significant problem in much of this research. In most cases, calculation of specialization/diversity is based on ordinal measures of inequality such as the Gini, Theil or Herfindahl indices. While these metrics are convenient, they say nothing about the nature of the association between technology (or other forms of economic) categories. Hence the "technological distance" between two primary patent technology classes in "organic compounds" is regarded as equal to the distance between the classes of "nanotechnology" and "boot and shoe making". This makes little sense. More useful characterizations of the knowledge stocks, or the technological capabilities, of different regions would be based on measures of the relatedness or coherence between the components of those stocks.

Attempts to generate measures of the relationships between patent technology classes can be traced back to JAFFE (1986), though SCHERER (1982) and TEECE *et al*. (1994) follow similar lines of argument in their measures of inter-industry R&D relatedness and corporate coherence, respectively. Three different methods have been used to measure the "distance" between patent technology classes. In the first of these, probabilities of the links between technology classes, and thus inter-class distances, are derived from the joint classification of individual patents across different technology categories (JAFFE 1986; VERSPAGEN 1997; BRESCHI *et al*., 2003). A second approach uses patent citations and the primary technology classes into which citing and cited patents are placed in order to measure distances between technology classes (LETEN *et al*., 2007). A third approach combines patent class data and geography. In this case measures of revealed comparative advantage in patent technology class codes are derived for different locations. Conditional probabilities linking all pairs of patent classes are estimated across all locations and measures of technological relatedness between classes so derived (VAN DER WOUDEN, 2012). This method builds on the product-space work of HIDALGO *et al*. (2007).

In the analysis below, the second of the methods just discussed is used to identify a U.S. knowledge/technology network and provide measures of technological relatedness (the inverse of distance) between each pair of U.S. patent classes. Before outlining how patent citation and technology class data are used to identify technological or knowledge relatedness between patent classes, some comments on patent citations, a crucial element of this method, is required.
U.S. patent law requires that an invention be novel and non-trivial in order for a patent to be granted. In addition, U.S., patent applicants are legally required to identify the prior knowledge upon which their inventions are based. This prior art is typically referenced through citations provided by patent applicants (inventors or their lawyers) and patent examiners. Patent

examiners are responsible for about 60% of the citations placed on U.S. patents. This has led many to question what these citations represent (BRESCHI and LISSONI, 2001; ALCACER and GITTELMAN, 2006; ALCACER *et al*., 2009). Do they measure real knowledge flows or not? JAFFE *et al*., (2000) provide evidence in the affirmative, though the jury is clearly still out. It is surely difficult to make the case that citations track knowledge flows if inventors have no knowledge of the patents cited by examiners. It is less dangerous, perhaps, to think of citations as measures of related sets of technologies or knowledge, whether added by patent examiners or inventors. This is the spirit in which citation data are used in the methods outlined below. No claim is made that knowledge only flows through citations.

To construct a U.S. knowledge space the location of individual patents has to be determined. For single inventor patents that is straightforward. For patents with multiple inventors, the country of the first listed inventor is taken as the location of the patent, and so only those co-inventor patents where the first inventor is located in the U.S. are designated as U.S. patents. Using the location of the first inventor, in cases of co-invention, assumes that the order of inventors reflects their relative weight in the invention process. In the citation analysis below, only citations to and from U.S. patents are considered.

The period of analysis runs from 1975 to 2005. The choice of the first year is easy to defend. 1975 is the first year for which patents are electronically linked to citations in the USPTO database. Patent data are available through 2011, though analysis here ends in 2005 largely because of the right censoring of patent applications. Table 1 provides some general information on numbers of patents and citations for selected years over the period 1975-2005. Note that these data are not averaged across a number of years: they are not smoothed in a temporal sense. There are significant swings in patent numbers and citations across individual years, but the years shown are certainly not outliers. Patents are aged by date of application rather than year of granting, for the usual reasons. Note that U.S. knowledge networks were built using both year of application and year of granting. There is little difference to the overall results. Self-citations are not removed in construction of the knowledge network for the analysis here is not attempting to identify spillovers. Rather, the focus is on the technological linkages across all patent classes. Investigation focuses on citations to patents that are no more than twelve years old. The twelve year cutoff is employed to ensure a long period of investigation while also capturing the majority of the citations that most patents generate (see HALL *et al*., 2001).

Table 1 shows the number of domestic patent applications to the USPTO for 1975, 1985, 1995 and 2005, along with the number of citations to existing U.S. patents recorded on those applications that are not more than twelve years old. For all these patents, citing and cited, their primary technological class is known. Clearly the number of U.S. patents has increased markedly over time. The apparent decline between 1995 and 2005 is, in large part, the result of right censoring, or truncation in the data: many patents applied for in 2005 have yet to be granted. What is also clear from Table 1 is the rapid growth in the number of total citations, and in the average number of citations per patent. Indeed, the number of citations on each patent has climbed from an average of 2.8 in 1975 to 9.6 in 2005. The extent to which this represents citation inflation or the increasingly derivative nature of new knowledge claims is unclear.

Table 1: Application Year Patents and Grant Year Patents, US Only

YEAR	PATENTS APPLCTN. YEAR	CITATIONS (12-YEARS)	PATENTS GRANT YEAR	CITATIONS (12-YEARS)
1975	41,385	118,742 (2.817)[2]	45,692	116,237 (2.544)
1985	36,996	135,866 (3.604)	37,772	122,364 (3.156)
1995	80,460	549,062 (6.708)	54,670	276,495 (5.058)
2005[1]	52,975	517,004 (9.649)	46,183	492,597 (10.666)

Notes: 1. 2005 application year patents are right censored. The patents identified here were granted by the end of 2009. 2. Ratio of citations to patents in parentheses.

The citation and technology class information on the patents described in Table 1 are used to derive a measure of the technological or knowledge relatedness between all pairs of patent classes. In a general sense, two technology classes are considered related if patents in one of these classes cite patents in the other class. There are 438 primary utility patent technology classes currently used by the USPTO. All patents in the database are located in one of those primary classes.

A method for calculating technological relatedness between patent classes is outlined next. All (granted) patent applications for a given year, say 1975, are recorded along with all citations on those citing patents that extend back for twelve years. This generates a database of cited patents that extend back to 1963 (for the digital records). The primary technology class of all citing and cited patents are recorded and arranged in the following matrix

$$\boldsymbol{C}_{ij}^{t} = \begin{bmatrix} c_{1\,1} & \cdots & c_{1\,438} \\ \vdots & \vdots & \vdots \\ c_{438\,1} & \cdots & c_{438\,438} \end{bmatrix}$$

where $\boldsymbol{C}_{ij}^{t}$ is a 438x438 matrix, the elements of which record the number of citations made by citing patents in technology class *j* to cited patents in class *i* in a given year *t*. Dividing each element of $\boldsymbol{C}_{ij}^{t}$ by the number of patent applications (granted) in the element's column class yields a matrix of the relative frequency that a patent in technology class *j* in a specified year will cite a patent in technology class *i*

$$\boldsymbol{P}_{ij}^{t} = \frac{c_{ij}^{t}}{N_j^t}$$

where N_j^t is the number of patents in technology class *j* in a given year.

$\boldsymbol{P}_{ij}^{t}$ provides a measure of the technological relatedness or knowledge relatedness between patents in technology classes *i* and *j*. The elements of $\boldsymbol{P}_{ij}^{t}$ take the value 0 when patents in class *j* do not cite patents in class *i*. In this case there is no technological relatedness between class *i* and class *j*. To be more concrete, in 1985, technology class 331 (oscillators) did not cite technology class 236 (automatic temperature and humidity regulation). However, note that technology class 236 did cite technology class 331 in 1985. Thus, the matrix of relatedness between patent classes is asymmetric. The values of technological relatedness are not bounded to the right. Patents in the same technological class, located on the principal diagonal of $\boldsymbol{P}_{ij}^{t}$, quite often exhibit relatedness values greater than 1. On average, technological relatedness should be greater for patents in the same technology class than for patents located in different classes. The values on the principal diagonal vary, perhaps, with the technological heterogeneity of patents found within individual classes. Technological relatedness values greater than 1 are rare off the principal diagonal of the matrix $\boldsymbol{P}_{ij}^{t}$. A similar method of identifying technological relatedness was applied to European patent and citation data by LETEN *et al*. (2007). That analysis focused only on 30 technology classes and thus the technological precision of the estimates developed here is considerably greater.

Individual technology classes reported in the USPTO have been aggregated into 30 intermediate classes and 6 broad technology classes by HALL *et al*. (2001). These aggregate groups afford a simple test of the efficacy of the technological relatedness measure just developed. It makes sense to assume that primary patent classes that are members of the same aggregate technology groups will cite one another with greater frequency than primary patent classes found in different aggregate technology groups. If we do not see this, then the meaning of our relatively frequency matrix $\boldsymbol{P}_{ij}^{t}$ is in doubt.

With the aid of UCINET (BORGATTI *et al*., 2002), the network of technological relatedness across the 438 primary patent classes is mapped. The technological relatedness network is generated with the Gower-scaling metric, itself derived to examine patterns of similarity across network nodes (GOWER, 1971). The nodes in the network correspond to each of the 438 distinct technological classes within the USPTO. The relative positions of the nodes are fixed by the frequencies of citation across each technology class pair ($\boldsymbol{P}_{ij}^{t}$). The principal diagonal plays no role in the relative locations of the nodes. Note that because of the asymmetry in the relatedness matrix, the links between network nodes *i* and *j* are the average of $\boldsymbol{P}_{ij}^{t}$ and $\boldsymbol{P}_{ji}^{t}$. The knowledge relatedness networks for 1975 and 2005 are shown below (see Figures 1 and 2). Note that the knowledge relatedness networks for 1985 and 1995 are quite similar to those displayed.

The node colors in the figures represent the aggregate technology (6 class) grouping of HALL *et al*. (2001): Black = Chemicals (1), Green = Computers & Communications (2), Yellow = Drugs & Medical (3), Red = Electronics (4), Blue = Mechanical (5), Grey = Miscellaneous (6). There is clear evidence of the clustering of individual patent categories within most of these classes, indicating that the relatedness measure is capturing what may be considered as a common knowledge base within these more aggregate technology groupings. All network links are not included for their density would render the network largely unreadable. The network links shown are illustrative of the total, representing the strongest links in the network at each time period. Because of the increase in the number of citations per patent throughout the study period, the density of links increases markedly between 1975 and 2005. The citation data used to build the knowledge networks are not adjusted for potential citation inflation. The size of each node illustrates the number of patents in that technology class in the given year. Node sizes have been scaled to allow comparison over time. To aid comparison, the size of the largest node in each year is indicated below each figure.

Figure 1: U.S. Knowledge Space, 1975

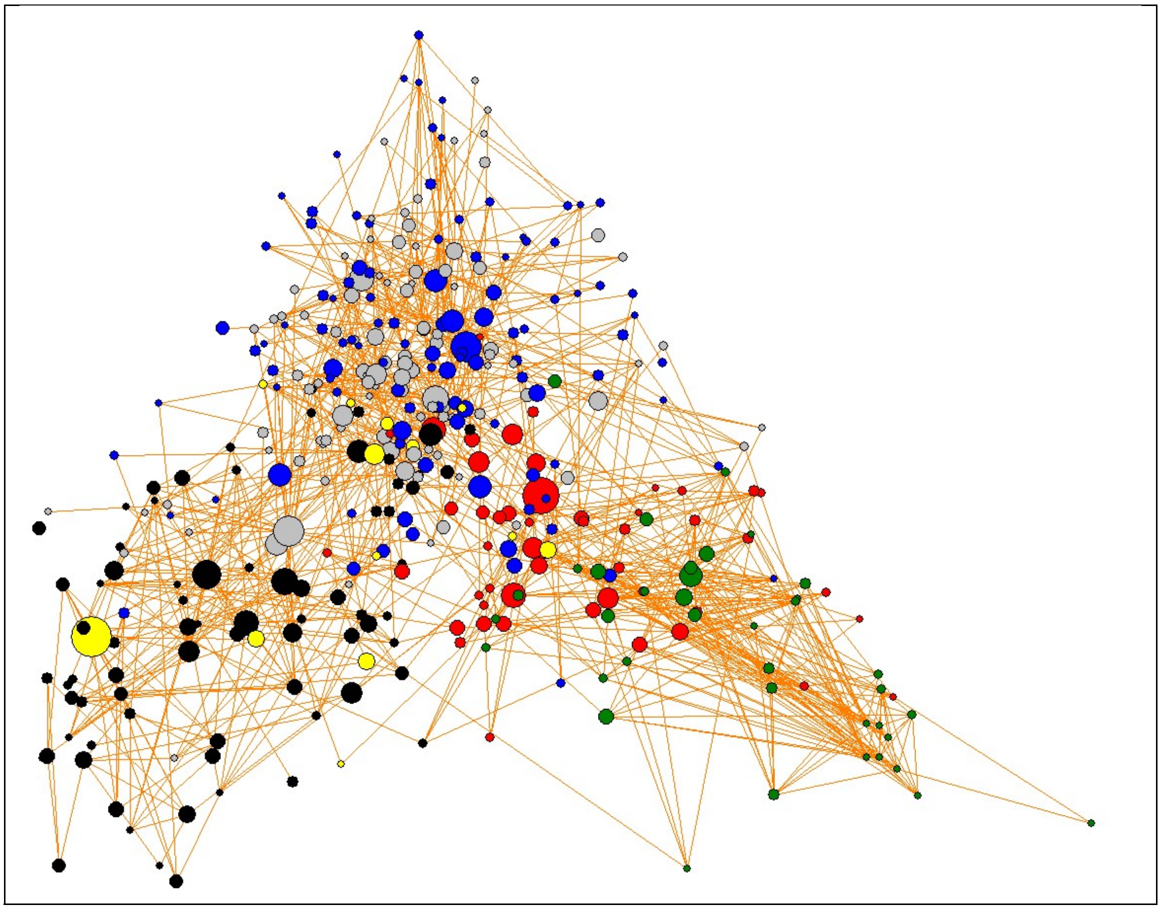

Notes: Black = Chemicals (1), Green = Computers & Communications (2), Yellow = Drugs & Medical (3), Red = Electronics (4), Blue = Mechanical (5), Grey = Miscellaneous (6). The largest node (514 = Drug: bio-affecting and body treating compositions) represents 746 patents.

The U.S. knowledge space in 1975 shows that patents appear to be reasonably evenly distributed across the six broad technology groupings. The shared knowledge cores of those groupings can be identified in 1975, though they become much more pronounced over time, especially in computers and communications, electronics, and to a somewhat lesser extent in drugs and medical patents. By 2005, the rapid expansion in the share of patents in those same three classes is evident. The computers and communications core becomes more closely aligned with the electronics core between 1975 and 2005 and the drugs and medical core moves away from the core of the chemicals cluster, perhaps overlapping increasingly with the electronics and computers clusters. The cores of the mechanical and miscellaneous groups appear to be reasonably closely correlated over the whole period examined.

Figure 2: U.S. Knowledge Space, 2005

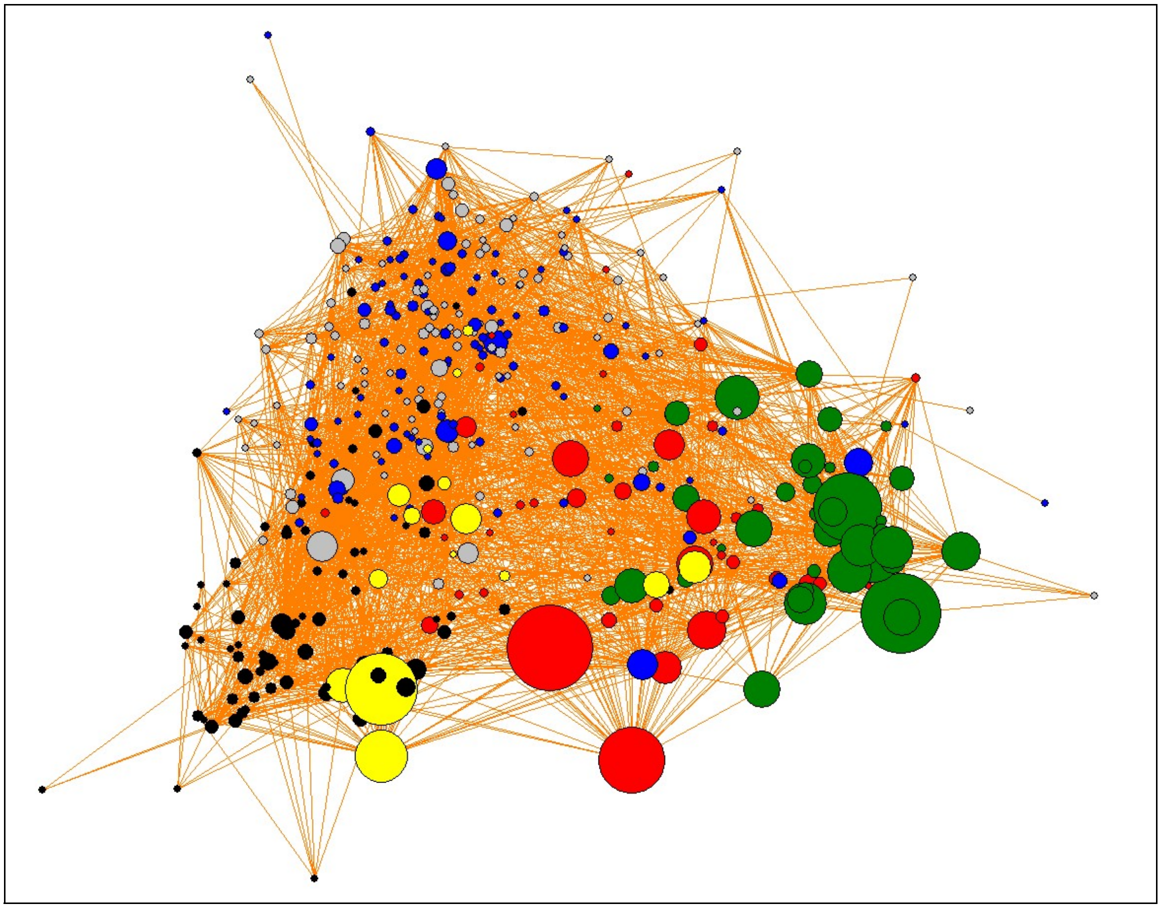

Notes: Black = Chemicals (1), Green = Computers & Communications (2), Yellow = Drugs & Medical (3), Red = Electronics (4), Blue = Mechanical (5), Grey = Miscellaneous (6). The largest node (438 = Semi-conductor device manufacturing) represents 1709 patents.

To measure the overall technological coherence of the patent network an average relatedness score is generated. Average relatedness measures the total "technological distance" between all pairs of patents divided by the number of such pairs. For a given number of patents, a higher average relatedness score indicates that patents are located in technology classes that are relatively close to one another in the U.S. knowledge network. These are patents found in classes that tend to cite each other with a relatively high frequency. A lower relatedness score would indicate that the patents are distributed over technology classes that are, on average, further apart from one another in technology space. Average relatedness measures the technological specialization or coherence of produced knowledge. Higher levels of relatedness indicate greater technological specialization. The average relatedness value for a region *r* in year *t* is calculated as:

$$AR^{t,r} = \frac{\sum_i \sum_j P_{ij}^{t} * D_{ij}^{t,r} + \sum_i P_{ii}^{t} * 2D_{ii}^{t,r}}{N^{t,r} * (N^{t,r} - 1)} \qquad for\ i \neq j$$

where P_{ij}^{t} measures the technological relatedness between patents in the 438 technology classes *i* and j, $N^{t,r}$ is a count of the total number of patents in region *r* in year *t*, and where $D_{ij}^{t,r}$ counts the number of pairs of patents that can be located in technology classes *i* and *j* in region *r* in year *t*. To clarify the meaning of $D_{ij}^{t,r}$, imagine a region with three patents, one in technology class 1 and two in technology class 2. Then, the pair counts $D_{ij}^{t,r}$ represent elements in the (438x438) symmetric matrix

$$\boldsymbol{D}^{t,r} = \begin{bmatrix} 0 & 2 & \ldots & 0 \\ 2 & 1 & \ldots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \ldots & 0 \end{bmatrix}$$

With three patents, there are 3x2 = 6 unique distance measures to calculate, the distance between the patent in class 1 and each of the patents in class 2, the distances from both patents in class 2 to the patent in class 1 and the distance between the two patents in class 2. Note that the latter distance is counted twice.

Tables 2 and 3 provide information on average technological relatedness between all patents in the U.S. knowledge network and between patents within each of the six aggregate technology classes. Table 2 reports that average relatedness increased by more than 94% between 1975 and 2005, even after adjusting for “citation inflation”. An increase in average relatedness indicates that more patents are being generated within technology classes that are closer to one another in technology space. This is consistent with the growth of technological specialization, an increase in the shared knowledge base that underpins invention. The rate of growth of specialization in U.S. patenting accelerated sharply after 1985, though it slowed somewhat between 1995 and 2005.

Table 2: Total and Average Knowledge Relatedness, US Total (All Patents)

		RELATEDNESS	
YEAR	PATENTS[1]	TOTAL	AVERAGE
1975	41,385	23,811,636 (23,811,636)[2]	0.01390 (0.01390)[2]
1985	36,996	28,041,166 (21,924,289)	0.02049 (0.01602)
1995	80,460	366,530,344 (153,939,666)	0.05662 (0.02378)
2005	52,975	259,254,666 (75,694,793)	0.09238 (0.02697)

Notes: 1- Number of (granted) patent applications in year indicated, with known CBSA. 2- Knowledge relatedness adjusted for citation inflation in parentheses.

As expected, Table 3 shows that average relatedness values are much greater within each of the six aggregate technology classes than overall. This confirms expectations that technological information (a citation) is more likely to flow within a major technology grouping (such as the chemicals category) than it is to flow between such groupings. The drugs and medical group exhibits the highest average relatedness score of all major patent groups, indicating that knowledge, in the form of citations, circulates more frequently in this group than in others. On average, relatedness scores are also relatively high in the electronics and in the computers and communication patent groups. They tend to be lower in the chemicals and mechanical patents. Since 1975, before adjusting for citation inflation, average relatedness values have increased fastest in the electronic, the drugs and medical, and the miscellaneous patent groups.

Table 3: Average Knowledge Relatedness by Major Patent Class, US Total

PATENT GROUP	YEAR			
	1975	1985	1995	2005
CHEMICALS	0.05518 (0.05518)[1]	0.07951 (0.06216)	0.11638 (0.04888)	0.13150 (0.03839)
COMPUTERS & COMMNCTN.	0.14008 (0.14008)	0.14444 (0.11293)	0.24182 (0.10156)	0.26731 (0.07804)
DRUGS & MEDICAL	0.18804 (0.18804)	0.33333 (0.26061)	0.75794 (0.31833)	0.92738 (0.27074)
ELECTRONIC	0.07994 (0.07994)	0.10972 (0.08578)	0.19388 (0.08143)	0.47144 (0.13764)
MECHANICAL	0.03508 (0.03508)	0.03966 (0.03101)	0.08184 (0.03437)	0.17182 (0.05016)
MISCELLANEOUS	0.04221 (0.04221)	0.05753 (0.04498)	0.10743 (0.04512)	0.25746 (0.07516)

Notes: 1-Knowledge relatedness adjusted for citation inflation in parentheses.

THE KNOWLEDGE CORES AND AVERAGE RELATEDNESS OF U.S. METROPOLITAN AREAS

The patent knowledge cores of metropolitan areas in the U.S. can be identified by mapping the patents generated within individual cities on the U.S. knowledge space of a given year. Building a separate knowledge space for each city based upon localized patent citations would yield very sparse networks for most because of the tendency for the majority of knowledge flows to cross metropolitan boundaries. Once more the address of the inventor, or the first-named inventor on co-invented patents, is used to locate patents within the United States. There are 949 core based statistical areas (CBSAs) in the United States within which patents can be located. Note that the newest patent files provided by the USPTO provide zip codes for inventor addresses that are readily linked to CBSAs. For older files, linking the cities and counties within which inventors are located to CBSAs was performed using the geographical correspondence engine available through the Missouri Census Data Center (http://mcdc2. missouri.edu/websas/geocorr2k.html). The focus here is on the 366 metropolitan CBSAs that house well over 90% of all U.S. patents. Still with so many metropolitan areas, it is only possible to provide illustrations of the knowledge cores of a few. Three CBSAs highlighted - Boise, Idaho, Dayton, Ohio and Rochester, New York. These cities do not form a representative sample, they serve merely to illustrate quite different trajectories of knowledge production. Rather than map every patent within these three cities, only patent classes in which cities exhibit relative specialization, identified by location quotients, are displayed.

In 1975, Boise was home to inventors who developed 22 patents. 21 of these 22 patents were developed in different patent classes, with only one USPTO patent technology class accumulating two patents. Because of the relatively small number of patents generated in Boise, all 21 occupied patent classes in 1975 have patent location quotients greater than 1.0, indicating metropolitan specialization. Figure 3a shows the relative position of these 21 patent classes in the U.S. technology space for 1975. There is nothing in this figure that suggests an existing, or even a nascent, technology or knowledge core. The average relatedness score for Boise in 1975, the average technological distance between all pairs of patents (not just those in classes with a location quotient greater than 1) was a relatively low 0.0123, ranking Boise 275 out of 366 metropolitan CBSAs in terms of the coherence of its produced knowledge base. However, by 2005, Boise had become one of the leading centers of semi-conductor patenting in the United States (see Figure 3a). In that year, Boise inventors generated 577 patents, 68% of these in just three technology classes: 257 (active solid state devices), 365 (static information storage and retrieval) and 438 (semi-conductor device manufacturing). Of 74 metropolitan CBSAs in 2005 generating more than 100 patents, Boise ranked second highest in terms of average relatedness (1.413), indicating that most of its patents were located in technology classes closely related to one another. While a detailed history of invention in Boise is well beyond the scope of this paper, the interested reader might consult MAYER'S (2011) account of the evolution of invention in Treasure Valley.

A very different pattern of urban invention is offered by Dayton, Ohio that traces a long history of aircraft, computing and automobile related patenting, through the work of the Wright brothers, the National Cash Register Company (NCR) and Dayton Engineering Laboratories Company (DELCO), later to become the foundation of the General Motors research arm. One of the most productive U.S. cities for patenting at the end of the nineteenth century, by the 1970s the pace of invention in Dayton had already begun a slow decline. Nonetheless, in 1975 with 289 patents, Dayton ranked 30 out of all U.S. cities in terms of the number of patents produced. These patents were distributed across mechanical, chemical and electronic technologies with one pronounced cluster in refrigerants (see Figure 3b). In 1975 the average relatedness of Dayton's patents was 0.027, ranking 33 out of 60 metropolitan areas with more than 100 patents. As late as 2005, more than 100 patents were still produced in Dayton. However, these patents were widely scattered across the U.S. knowledge space. Indeed, the average relatedness score for Dayton's patents in 2005 was 0.0799, placing the city dead last in terms of technological coherence out of all U.S. metropolitan areas with more than 100 patents. The distributed character of Dayton's patents in 2005 is clear in Figure 3b.

Figure 3a: The Knowledge Core of Boise

1975 **2005**

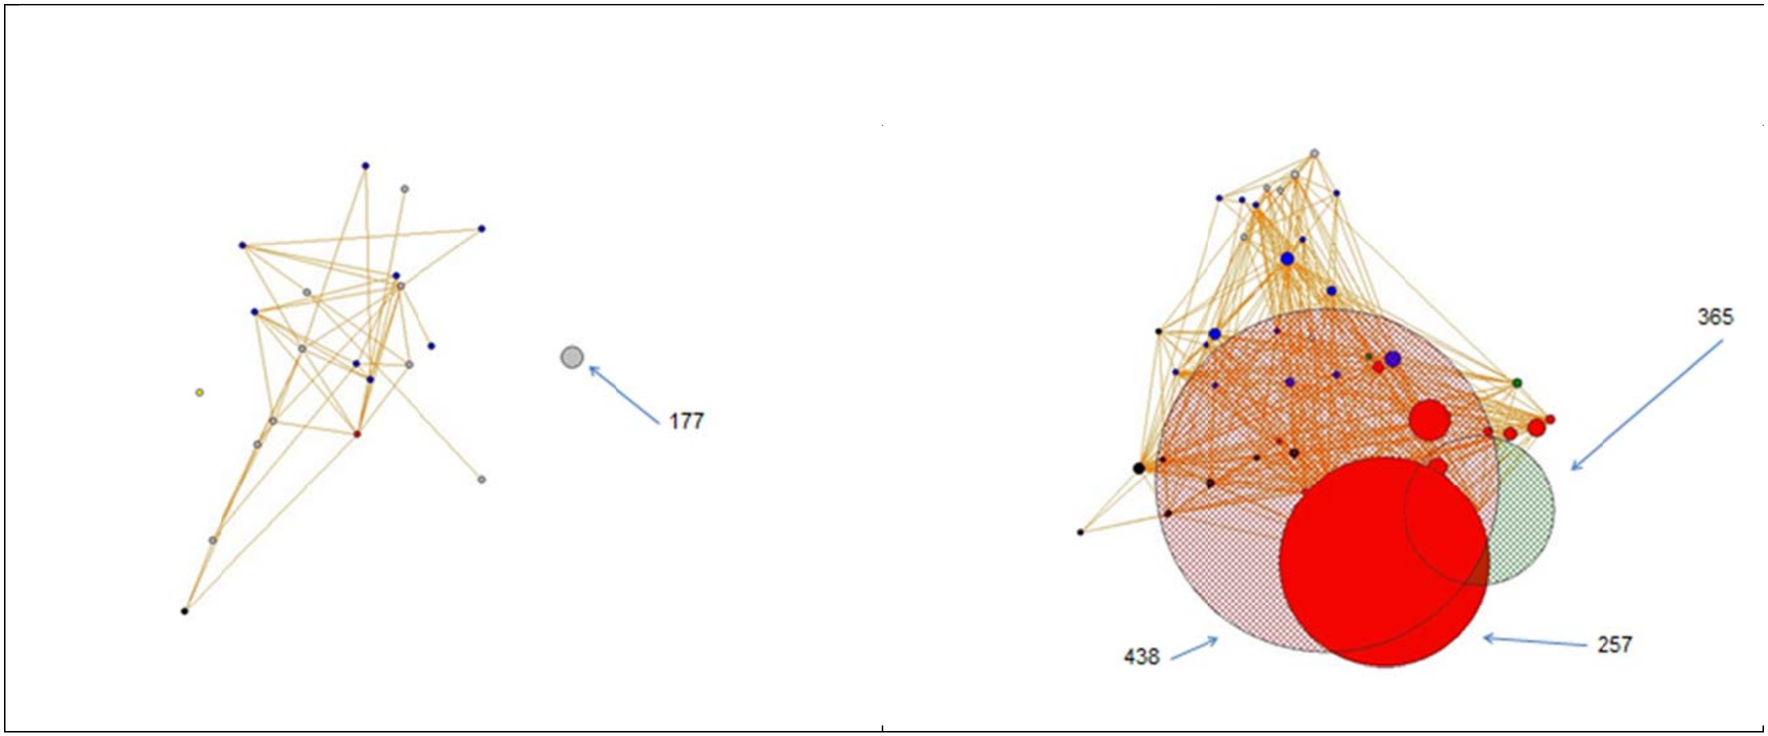

Notes: Technology class 177 is weighing scales, class 257 is active solid-state devices, class 365 is static information storage and retrieval and class 438 is semi-conductor device manufacturing.

Figure 3b: The Knowledge Core of Dayton

1975 **2005**

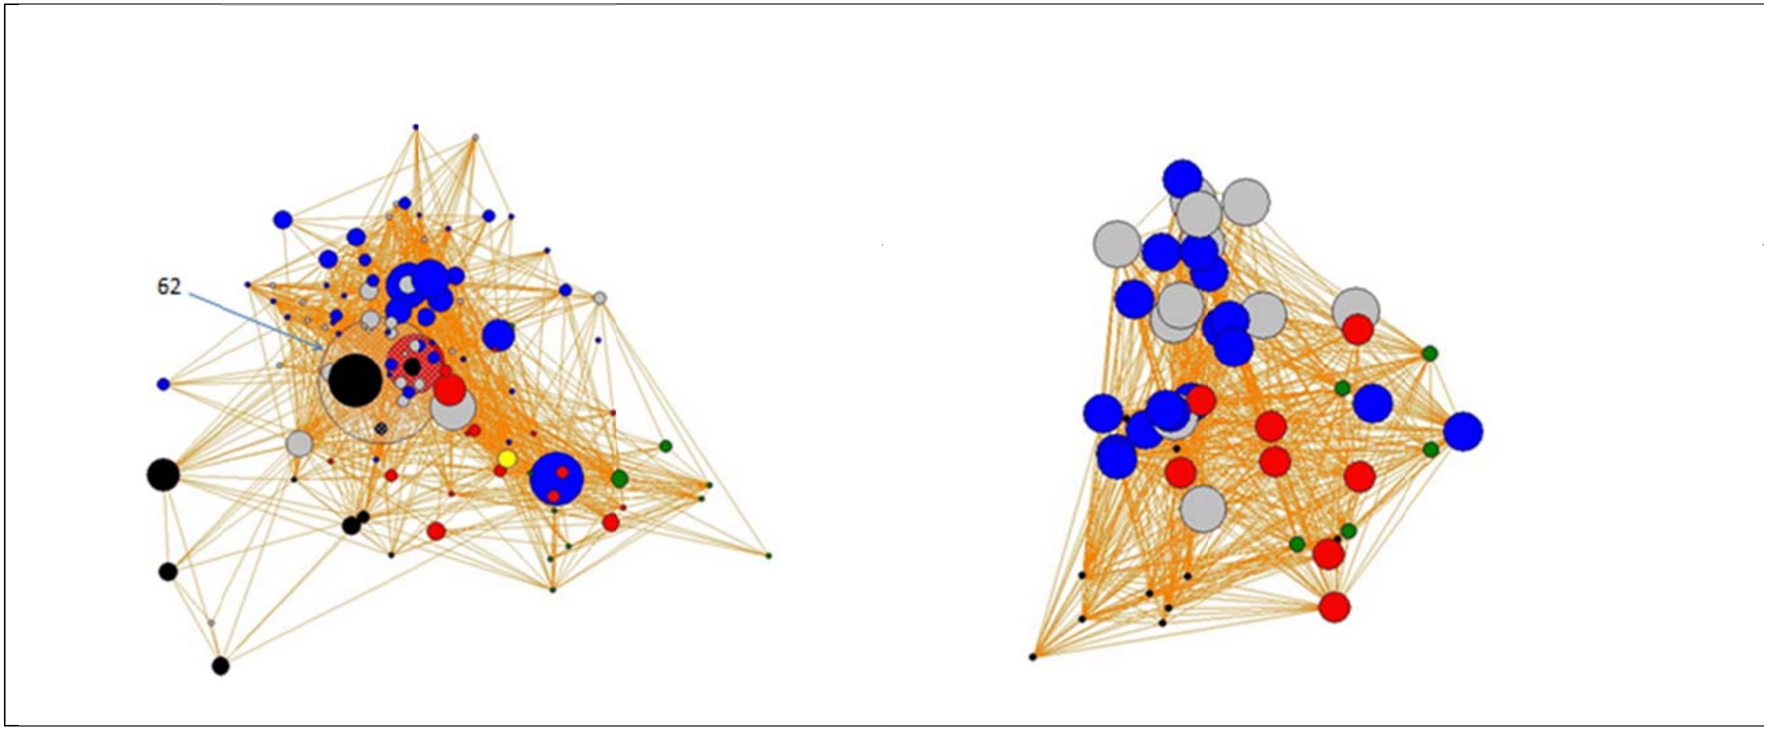

Notes: Technology class 62 is refrigeration.

Figure 3c: The Knowledge Core of Rochester

1975 **2005**

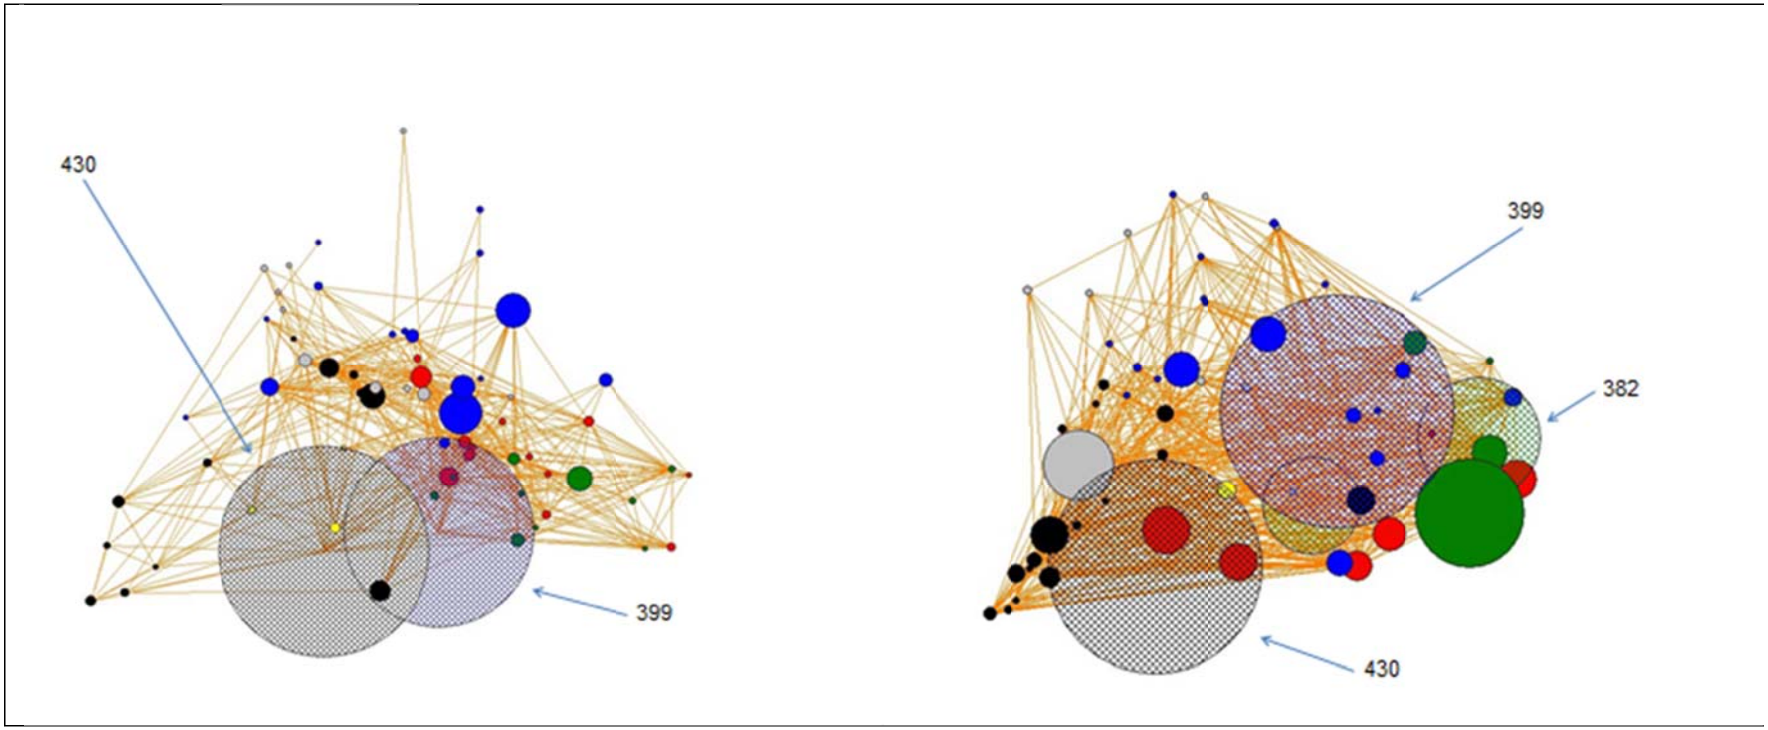

Notes: Technology class 382 is image analysis, class 399 is electrophotography and class 430 is radiation image chemistry.

Figure 3c maps the locations of patents developed in Rochester, New York in 1975 and 2005. Most noticeable from this figure is Rochester's long concentration in optics related invention, chiefly a function of the activities of the Kodak and Xerox Corporations. FELDMAN and LENDEL (2010) provide an overview of the emergence and development of the optics industry in the United States. In 1975, Rochester generated 654 patents, marking the city as the thirteenth most inventive in the U.S.. Even more remarkable is the concentration of these patents in technology classes strongly connected to the optics industry: 382 (image analysis), 399 (electrophotography) and 430 (radiation image chemistry). In 1975, Rochester had a higher average relatedness score (0.2495) than any other U.S. metropolitan area that housed more than 100 patents. By 2005, the average technological distance between Rochester's patents had not changed by much. In this year, the average relatedness of Rochester's patents measured 0.2731, ranking the city 14th most specialized out of the most inventive U.S. metropolitan CBSAs.

A more general overview of the technological coherence of patents produced across all U.S. CBSAs is provided in Figure 4 that maps the distribution of average relatedness values between 1975 and 2005. In general, average relatedness clusters at relatively low values, indicating that there is little coherence in the knowledge base of most CBSAs. This reflects the fact that in the majority of CBSAs the number of patents is relatively small and the technological links between those patents limited. The average relatedness values for each year exhibit a marked right skew as they are heavily influenced by a few CBSAs with considerable numbers of patents that cluster in a small number of technological fields. The most extreme examples of these are "company towns" like Midland Michigan, where patenting is dominated by Dow Chemical and Dow Corning, Bartelsville, Oklahoma where patenting is dominated by Philips Petroleum, and Duncan, Oklahoma where patenting is dominated by Haliburton. In general, average relatedness is negatively related to urban size, whether measured by patents or population: the largest cities tend to patent across many technology classes, some quite distant from one another in knowledge space. Over time, the average (unweighted) relatedness value for all CBSAs has increased from 0.0668 in 1975 to 0.3347 in 2005 (not adjusting for citation inflation).

Table 4 refocuses attention on metropolitan CBSAs with more than 100 patents each. These metropolitan areas are responsible for approximately 80% of all U.S. patents. The top and bottom ranked metropolitan CBSAs, in terms of average relatedness, are listed in this table. So in 1975, along with Rochester, New York, average relatedness or technological specialization was relatively high in Akron, Houston, Peoria and Springfield. In 2005, the metro CBSAs with the most specialized set of patents were Memphis, Boise, Burlington, Poughkeepsie and Houston. The cities with the least specialized patent knowledge base in 1975 tended to be those with the largest numbers of patents. This pattern had shifted by 2005, such that four of five metropolitan CBSAs with the least specialized patent set were the old, rust-belt cities of Dayton, Cleveland, Pittsburgh and Columbus. These cities still produce relatively large numbers of patents, but not necessarily in technological fields that the industrial histories of these cities would suggest.

Figure 4: Distribution of Average Relatedness in U.S. Metropolitan CBSAs

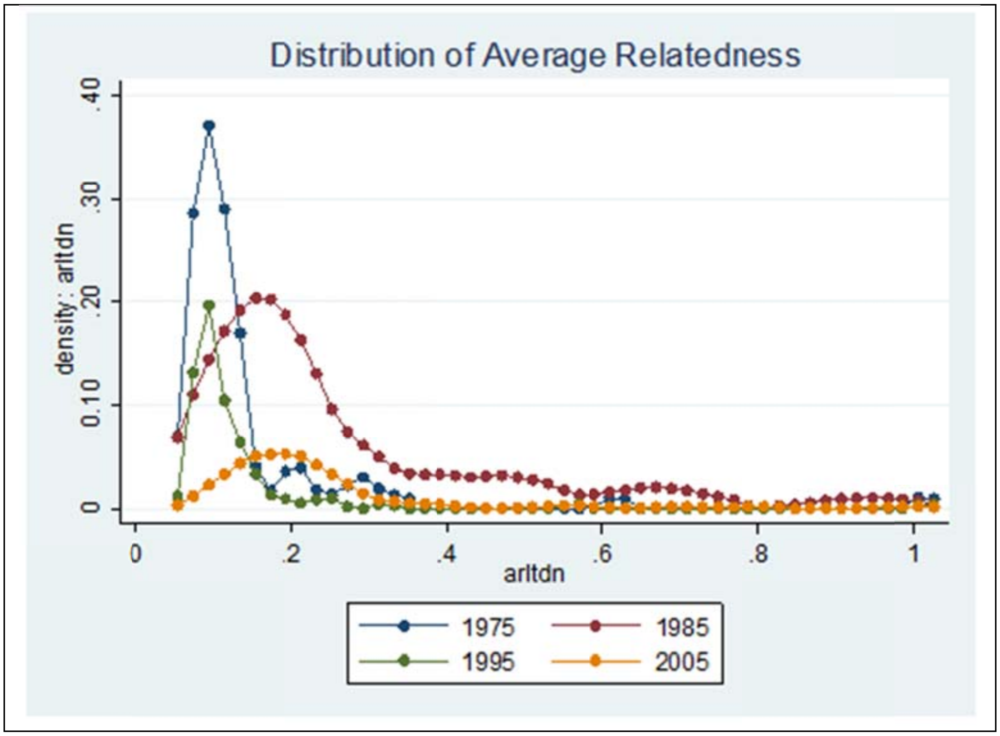

Table 4: Average Relatedness and Rank for U.S. Metropolitan CBSAs with >100 Patents

CBSA 1975	CBSA 1985	CBSA 1995	CBSA 2005
1 Rochester, NY 0.2495	1 Harrisburg, PA 0.8429	1 Ogden, UT 0.7686	1 Memphis, TN 2.4225
2 Akron, OH 0.0772	2 Ann Arbor, MI 0.1119	2 Harrisburg, PA 0.6411	2 Boise, ID 1.4130
3 Houston, TX 0.0719	3 Baton Rouge, LA 0.1073	3 Memphis, TN 0.04355	3 Burlington, VT 0.9615
4 Peoria, IL 0.0569	4 Houston, TX 0.1043	4 Indianapolis, IN 0.4243	4 Poughkeepsie, NY 0.8421
5 Springfield, MA 0.0538	5 Rochester, NY 0.1024	5 Appleton, WI 0.4197	5 Houston, TX 0.7637
57 Los Angeles, CA 0.0167	57 Milwaukee, WI 0.0264	105 Cleveland, OH 0.0508	70 Columbus, OH 0.0881
58 San Diego, CA 0.0164	58 Buffalo, NY 0.0234	106 Rockford, IL 0.0493	71 Riverside, CA 0.0878
59 Portland, OR 0.0161	59 Los Angeles, CA 0.0232	107 Sarasota, FL 0.0477	72 Pittsburgh, PA 0.0859
60 New York, NY 0.0146	60 Chicago, IL 0.0232	108 Orlando, FL 0.0477	73 Cleveland, OH 0.0805
61 San Francisco, CA 0.0146	61 Riverside, CA 0.0183	109 Milwaukee, WI 0.0472	74 Dayton, OH 0.0799

The knowledge cores and levels of technological specialization across U.S. cities exhibit considerable heterogeneity both in terms of the direction of inventive activity and in terms of how focused that activity is (see also O'HUALLACHAIN and LEE, 2010). Patterns of urban invention also vary markedly over time. Some places have maintained their capacity to produce specialized sets of knowledge closely aligned with their existing technology cores, others have seemingly lost their inventive focus and/or their capacity to patent in general, while yet others have been able to build a knowledge base and develop a localized trajectory of invention around that core. What accounts for the long-run inventive success of some cities over others? Technology development tends to be cumulative in nature and thus new knowledge production rests heavily on accumulated knowledge sets. This begs the question of whether we can use information on the characteristics of local knowledge cores along with measures of relatedness between different technology classes to account for local histories of technological diversification and technological abandonment? This question is examined in the next section.

TECHNOLOGICAL DIVERSIFICATION AND TECHNOLOGICAL ABANDONMENT IN U.S. METROPOLITAN AREAS

Knowledge accumulates over time through search, including learning, and diffusion. The productivity of search is uneven. Much of the time the process of invention proceeds incrementally, moving along familiar trajectories. Occasionally those trajectories are abandoned as knowledge breakthroughs occur, pushing invention and innovation along novel pathways, some of which become well-used routes to future discovery (DOSI, 1982; NELSON and WINTER, 1982).

The accumulation of knowledge proceeds unevenly over space as well as time. Space remains a barrier to the flow of particular kinds of knowledge and to other commodities within which technological information is embedded. Characterized by different histories of resources and industrial development, different modes of organizing production and competing institutional arrangements, the knowledge bases and production technologies of individual regions exhibit a variety of forms (SAXENIAN, 1994; RIGBY and ESSLETZBICHLER, 1997; STORPER, 1997). These forms often give rise to distinctive techno-industrial clusters, regions with long histories of development in particular sets of skills, commodities and industries that we can readily identify. Within many such regions, knowledge production exhibits considerable inertia. The patents produced within individual regions over time are not a random draw from the set of technology classes that defines the global knowledge space. Rather, the technological distribution of patents within a region reflects the underlying knowledge base of that geographical area. This knowledge base is directed by the past experience of search and discovery within the area, and it is also influenced by the characteristics of the networks of invention and innovation within which individual inventors, firms and other research units in the region are embedded.

The primary aim of this section of the paper is to explore how existing configurations of technological capabilities within U.S. metropolitan areas shapes emerging technological/ knowledge trajectories. I explore this question using information on the existing knowledge cores and measures of technological relatedness within cities, forecasting the direction of invention from the technological gap or proximity between current invention and the range of inventive paths not yet taken within a city. The goal is not to explain the precise patterns of technology production within individual cities, but rather to relate the technological structure of a place at time t to the technological structure of that place in some future time $t+n$. This analysis builds on the work of HIDALGO *et al*. (2007) and HAUSMANN and KLINGER (2007) in modeling a global commodity space and the evolution of countries within that space. In this sense it is closely related to work in economic geography by NEFFKE *et al*. (2011) and BOSCHMA *et al*. (2012) who examine the process of industrial diversification within regions.

To make my task a little easier, I focus only on those patent classes (technologies) in which cities exhibit relative technological advantage. I render this in a quantitative sense with a binary valued location quotient. Figure 5 illustrates the problem. In the left image, the shaded circles indicate those technologies in which a city is specialized in year ***t***, and the unshaded circles indicate those technologies in which a city is not specialized. Which of the open circles is most likely to become shaded as we move from time *t* to time *t+1*? That is, how does the knowledge core of a region diversify over time? In the right image of Figure 5, only those technologies in which a city has relative technological advantage are shown. The problem to solve in this image is which technology is most likely to be abandoned as we move from time *t* to time *t+1*?

Figure 5: Specialization of Technology Nodes in a City

Diversification **Abandonment**

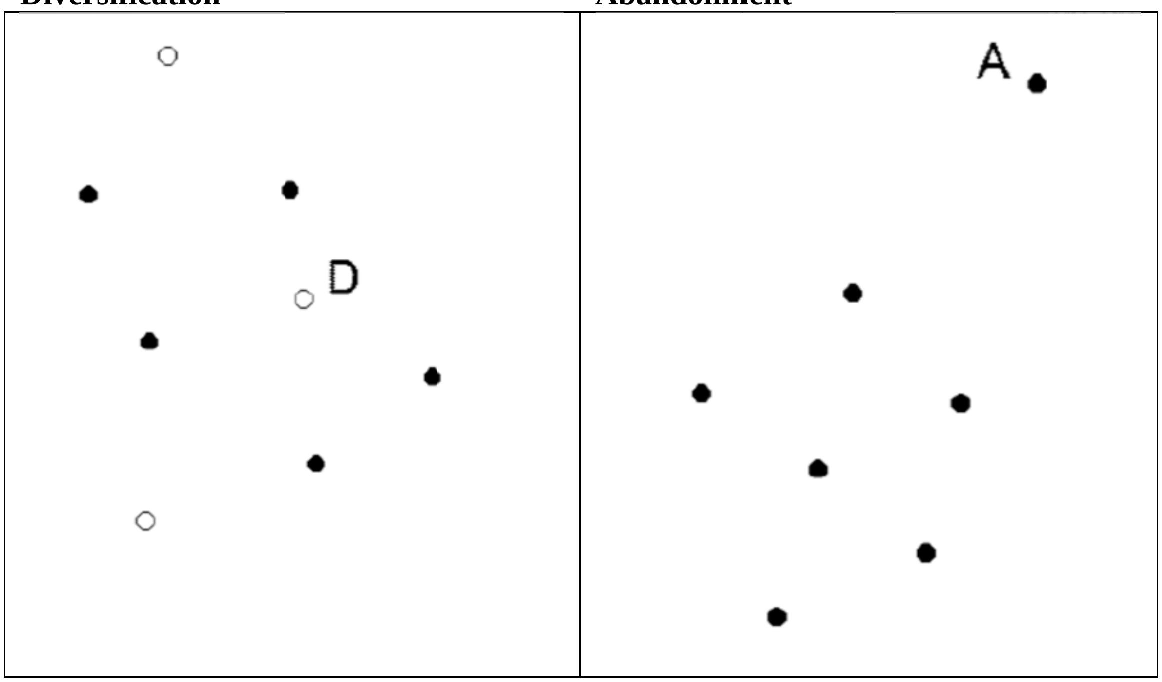

Notes: The shaded circles indicate the technologies in which a city is specialized.

A simple model assumes that technological diversification builds incrementally upon the existing knowledge base of the region. Thus, diversification to new technology classes, or gaining specialization in such classes, should be a function of their technological distance from the existing structure of knowledge within a region. In HAUSMANN and KLINGER (2007), diversification rests on the density of current practice within a product space and the value of that density around product classes that have yet to be exploited. I follow a similar logic and hypothesize that the probability of a city gaining specialization within a technology class is a positive function of the overall proximity (in knowledge space) of that class to all technology nodes in which the city is already specialized. Along the same lines, it follows that cities will be most likely to abandon those specialized technologies that are furthest from the core of their knowledge base.

In the left image of Figure 5, the city is most likely to diversify into knowledge class D because, out of the classes in which the city is not yet specialized, this class is closest to the set of techniques in which the city is already specialized. Turning to the case of technological abandonment, the right image of Figure 5 shows that of all technology nodes in which the city is specialized, node A is the most remote, or the furthest from the knowledge core of the city, and thus most likely to be abandoned. Note that in Figure 5 the set of all technology classes and their relative locations are fixed in U.S. knowledge space and are thus common to all cities. However, the pattern of technologies in which each city exhibits relative advantage (specialization) varies with the distribution of patents.

The discussion of technological diversification and abandonment to this point assumes that cities are independent spatial units that are not influenced by technological practices elsewhere. Yet, a long history of geographical scholarship suggests that there are strong linkages between urban areas (PRED, 1977). It seems reasonable therefore to extend the model of technological change to incorporate flows of knowledge between cities. I do this using co-inventor data. BRESCHI and LISSONI (2001) and SINGH (2005) provide strong evidence of the importance of co-inventor networks in understanding knowledge flows.

To capture co-inventor network effects, it is assumed that the probability of technological diversification in a city is influenced by the knowledge base (structure of knowledge) in other cities, together with an index of how closely those other cities are linked to the city in question. This is operationalized in the following way. For each year of data, a symmetric inter-city matrix (366x366) of co-inventor relations is multiplied by a (366x438) matrix of location quotients showing for each of the 366 CBSAs and 438 technology classes where each city has relative technological advantage (location quotient greater than 1). A given cell (i, j) in the resulting (366x438) matrix is a weighted average of the (binary) location quotient values in all cities k $(k \neq i)$ in technology class j where the weights specify the proximity of each (row) city to all other cities in the network. The distance between one city and itself is set to zero, so I ignore intra-city co-inventor relationships. It is hypothesized that spillovers of knowledge between cities, driven by co-inventor relationships, should exert a positive influence on the probability that a city adopts a new technology.

The inter-city matrix of co-inventor relations was constructed for all 366 urban areas for every second year between 1975 and 2005. These measures were developed from co-inventor data. The USPTO does not track individual inventors. Fortunately, LAI and colleagues at Harvard University have produced a list of individual inventors and their co-inventors that can be linked to the individual patent records in the USPTO (LAI *et al*., 2009). From these data I take all patent applications in a given year that list co-inventors and I record the metropolitan areas within which co-inventors were located. If a co-inventor was located outside the United States, or in one of the micropolitan CBSAs, they were dropped from the analysis. I then construct a metropolitan co-invention matrix with 366 rows and columns, each corresponding to one of the metropolitan CBSAs. The matrix is initially populated with zero values in all cells. If two co-inventors on a patent are located in different metropolitan areas then the cells of the inter-city matrix of co-inventor relations corresponding to the cities where the co-inventors are located receive the value 1. The resulting matrix is symmetric and values on the principal diagonal are ignored. (Note that co-inventor counts along the principal diagonal provide a measure of co-inventors located in the same city. This information is not exploited below.) If there are three co-inventors on a patent, each living in a different metropolitan area then six cells in the matrix receive a count of 1 (three pairs of cities in the matrix are linked and the symmetry ensures a count of six). This process is repeated for all co-invented patents with the inter-city matrix counts building a representation of the interaction between co-inventors located in different U.S. metropolitan areas in a given year.

The processes of metropolitan technological diversification and abandonment are examined using a regression model. The observational units are the 438 technology classes within each city over time. The model is run for every second year from 1975 to 2005 inclusive. The values of the dependent variable are 0 or 1, so the regression model is predicting the probability that $Y = 1$, that a city exhibits technological specialization in a particular technology class in a given year. The binary nature of the dependent variable suggests use of a probit or logit model extended to panel form to take advantage of the time dimension in the data. This is not straightforward, for a probit model cannot be run with a fixed effects panel specification that is suggested by a simple HAUSMAN test as preferable to a random effects model. So I opt for a fixed effects panel version of the logit model. The incidental parameters issue (resulting from controlling for 438 fixed effects over 366 cities over 16 years and the need to estimate parameters for these effects) also raises problems of estimation. These are solved using a conditional logit specification.

The basic model to be estimated is

$$Y_{it}^{c} = \alpha + \beta_1 Y_{it-1}^{c} + \beta_2 proximity_{it-1}^{c} + \beta_3 cbsanet_{it-1}^{c} + z_t + city_c + tech_i + \varepsilon_{it}^{c}$$

where the binary dependent variable assumes the value 0 or 1, and represents the probability of city *c* in year *t* exhibiting relative technological specialization in technology class *i*. On the right hand-side of the equation, the lagged value of the dependent variable captures existing specialization by cities in particular technologies. Absent the fixed effects, this could be interpreted as a measure of persistence. With existing specialization controlled, the two key independent variables are the lagged value of the distance between each technology class where the city does not exhibit technological specialization and those technology classes where it does (*proximity*), and a lagged measure of technology in related cities (*cbsanet*), where related is proxied by the co-inventor relationship between cities. The z_t represent year fixed effects, $city_c$ and $tech_i$ are CBSA and technology (patent) class fixed effects.

The results are displayed in Table 5. Model 1 is produced for comparison and is similar in form to that offered in HAUSMANN and KLINGER (2007) and BOSCHMA *et al*. (2012). Model 2 adds the CBSA network variable. The proximity of technology classes where cities do not currently specialize to those in which cities do specialize has a positive and significant influence on technological diversification, or the probability that relative advantage will be developed in the future. A one unit increase in overall technological relatedness (knowledge proximity) means an increase in the probability that specialization will be developed in new technologies of about 5%. The city network effect is also positive and significant, but extremely small in size.

Clearly the linear probability models (Models 1 and 2) are not appropriate for estimating binary data as noted above. Estimated values might lie outside the unit interval and the errors cannot be homoscedastic. In Models 3 and 4 we take explicit advantage of the panel nature of the data estimating a fixed effect panel conditional logit model. This model handles omitted variable bias and deals with the incidental parameters problem. Still, some caution should be used in interpreting results because the lagged value of the dependent on the right hand side of the equation adds serial correlation (NICKELL, 1981), though with a relatively large number of years in the panel, this problem is dampened. Whether the techniques of ARELLANO and BOND (1991) provide a solution is unclear given the limited form of the dependent variable.

Model 3 is a fixed effects conditional logit panel model with three independent variables. Year specific fixed effects are incorporated in this model but not shown. Note that the effect of the lagged dependent variable is considerably smaller in this specification. The partial regression coefficient on the lagged technological proximity measure (*proximity*) is slightly larger in magnitude to that of Model 2, and it remains statistically significant with the correct positive sign. The coefficient on the lagged city network measure (*cbsanet*) is also significant with a positive sign, and is similar in value to Model 2. Note that we lose two-thirds of our total observations in the panel logit models that discard technology classes in cities where there is no change in the value of the dependent variable over time. Note also that the partial regression coefficients in this model represent the log odds ratio rather than marginal effects.

All the models examined so far do not distinguish between the factors that determine the probability of concentrating in a new technology class from the probability of abandoning a technology in which a city already exhibits specialization. To distinguish growth in specialization from decline in specialization, or from technological abandonment, in Model 4 of Table 5, a fixed effects panel version of the following model is estimated in conditional logit form

$$Y_{it}^{c} = \beta_1 Y_{it-1}^{c} + \beta_2\{(1 - Y_{it-1}^{c}) * proximity_{it-1}^{c}\} + \beta_3\{Y_{it-1}^{c} * xproximity_{it-1}^{c}\} + \beta_4 cbsanet_{it-1}^{c} + z_t + \varepsilon_{it}^{c}$$

where most terms are defined as above and where x*proximity* is the distance from each technology class where a city has existing relative specialization to all other technologies in the same city with relative specialization. The second term in this model will drop out of the regression when the lagged value of the dependent variable is 1 and then the partial regression coefficient β_3 captures the influence of knowledge proximity in resisting technological abandonment. When the lagged value of the dependent variable is 0, the parameter β_2 captures the log odds ratio of a city developing relative specialization in a new technology. The results show little change in the effects of the city network and the effect of technological proximity on technological diversity. The impact of proximity in resisting technological abandonment is an order of magnitude larger than the effect of proximity in technological diversification.

Overall the results in Table 5 provide strong support for evolutionary claims that the processes of technological diversification and technological abandonment exhibit significant path dependence. A region's existing knowledge core is a strong predictor of future relative technological specialization, at least over the short-run: the pattern of technological diversification within a city is positively conditioned by the proximity of new technological possibilities that are relatively close to the existing knowledge core. In similar fashion, technological abandonment is strongly resisted when technology classes are close, in terms of knowledge space, to the city's existing knowledge core. Inter-city knowledge flows aid technological diversification and help resist abandonment, though the effects are small.

Table 5: Regression Analysis of the Probability of Technological Diversification and Technological Abandonment

Independent Variables	Model 1 LPM	Model 2 LPM	Model 3 XT (Conditional Logit)	Model 4 XT (Conditional Logit)
Y^c_{it-1}	0.3747*** (0.0088)	0.3492*** (0.0078)	0.1043*** (0.0085)	0.0969*** (0.0085)
$proximity_{t-1}$	0.0471*** (0.0008)	0.0461*** (0.0008)	0.0501*** (0.0032)	
$cbsanet_{t-1}$		0.00003*** (4.07e-07)	0.00003*** (3.22e-06)	0.00003*** (3.21e-06)
$(Y^c_{it-1}) * xproximity^c_{it-1}$				0.5042*** (0.0474)
$(1 - Y^c_{it-1}) * proximity^c_{it-1}$				0.0522*** (0.0032)
Constant	-0.0025 (0.0012)	-0.0020*** (0.0034)		
	n=2,404,620 R^2=0.1513	n=2,404,620 R^2=0.1575	n=809,010 $\chi^2(17) = 6073.6$ p=0.000	n=809,010 $\chi^2(17) = 6191.0$ p=0.000

Notes: All independent variables are lagged one period (L). All regressions include city, technology class and year fixed effects. In Models 1 and 2 the errors are clustered by CBSA. Models 1 and 2 are fit by Ordinary Least Squares (LPM stands for linear probability model) and Model 3 and 4 (XT stands for cross-sectional and longitudinal) are fit by Maximum Likelihood Estimation and the chi-square test is based on the Likelihood Ratio. *** indicates significant at the 0.01 level.

CONCLUSION

Citation data were used to measure the proximity of different technology classes into which patents are placed. The resulting measures of knowledge relatedness formed the edges of a patent network that maps the U.S. knowledge space. The evolution of that space was traced between 1975 and 2005. Over that time period, average relatedness between U.S. patents, after adjusting for citation inflation, has just about doubled: patents are increasingly concentrating in fewer technology classes and the distance between those classes is shrinking. Since 1975, the share of patents in chemical and mechanical classes has been decreasing while the share in drugs and medical, electronics, and computing classes has been increasing. Average relatedness scores vary markedly between these broad patent groupings.

Knowledge relatedness also varies sharply between U.S. cities. In general there is a negative relationship between city-size and relatedness as larger cities typically patent across a broader range of technology classes. Since 1975 there has been considerable mixing in the technological specialization of cities. Many of the older, snowbelt cities such as Columbus, Dayton and Pittsburgh have seen the coherence of their knowledge cores decline as the industries on which their growth was based have become much less innovative. At the same time more specialized knowledge cities have grown rapidly, fuelled by their focus in new technologies where rates of invention are high. RIGBY and VAN DER WOUDEN (2012) extend the analysis of this paper to show that higher levels of average relatedness in cities, or greater specialization, increases the rate of invention.

While some cities transition extremely rapidly from one knowledge core to another, for most, the process of technological transition is relatively slow. Cities build competence around a range of related technologies over time and this competence shapes the knowledge trajectories that most cities tend to follow. Technological diversification in cities, an expansion of the knowledge core, depends upon current practice and the proximity of new technological possibilities to the set of existing specializations. Diversification is also influenced by information about knowledge production from other locations. Knowledge specialization exhibits considerable inertia and the same forces that guide diversification play an even stronger role in maintaining competence. Technological abandonment is most likely to occur at the frontier of the knowledge space occupied by a city.

Much more work remains to be done to define more precisely the knowledge cores of cities and how they evolve. Which technology classes are part of a knowledge core and which remain outside are important questions, along with identification of related knowledge sets that link different cores. What possible combinations of technologies are most productive, and how cities can efficiently transition from relatively barren parts of technology space to more fertile areas are key areas for future research.

REFERENCES

ACS Z., ANSELIN L. and VARGA A. (2002) Patents and innovation counts as measures of regional production of new knowledge, *Research Policy* **31**, 1069-1085.

ALCACER J. and GITTELMAN M. (2006) Patent citations as a measure of knowledge flows: the influence of examiner citations, *Review of Economics and Statistics* **88**, 774-779

ALCACER J., GITTELMAN M. and SAMPAT B. (2009) Applicant and examiner citations in U.S. patents: An overview and analysis, *Research Policy* **38**, 415-427.

ANTONELLI C. (1995) *The Economics of Localized Technological Change and Industrial Dynamics*, Kluwer, Dordrecht.

ARCHIBUGI D. and PIANTA M. (1992) Specialization and size of technological activities in industrial countries: The analysis of patent data, *Research Policy* **21**, 79-93.

ARELLANO M. and BOND S. (1991) Some tests of specification in panel data: Monte Carlo evidence and an application to employment equations, *Review of Economic Studies* **58**, 277-297.

ASHEIM B. and GERTLER M. (2005) The geography of innovation: Regional innovation systems, in FAGERBERG J., MOWERY D. and NELSON R. (eds.) *The Oxford Handbook of Innovation*, Oxford University Press, Oxford (pp 291-317).

ATKINSON A. and STIGLITZ J. (1969) A new view of technical change, *Economic Journal* **79**, 573-578.

BORGATTI S. EVERETT M. and FREEMAN L. (2002) Ucinet for Windows: Software for Social Network Analysis. Analytic Technologies, Harvard, MA.

BOSCHMA R. (2005) Proximity and innovation: A critical assessment, *Regional Studies* **39**, 61-74.

BOSCHMA R. and IAMMARINO S. (2009) Related variety, trade linkages and regional growth in Italy, *Economic Geography* 85, 289-311.

BOSCHMA R., MINONDO A. and NAVARRO M. (2012) Related variety and regional growth in Spain, *Papers in Regional Science* **91**, 241-156.

BRESCHI S. and LISSONI F. (2001) Knowledge spillovers and local innovation systems: A critical survey, *Industrial and Corporate Change* **10**, 975-1005.

BRESCHI S. and LISSONI F. (2005) Knowledge networks from patent data: Methodological issues and research targets, in MOED H., GLANZEL W. and SCHMOCH U. (eds.) Handbook of Quantitative Science and Technology Research, Kluwer, Dordrcht (pp613-643).

BRESCHI S., LISSONI F. and MALERBA F. (2003) Knowledge-relatedness in firm technological diversification, *Research Policy* **32**, 69-87.

CANTWELL J. and VERTOVA G. (2004) Historical evolution of technological diversification, *Research Policy* **33**, 511-529.

COHEN W. and LEVINTHAL D. (1990) Absorptive capacity: A new perspective on learning and innovation, *Administrative Science Quarterly* **35**, 128-152.

COOKE P. URANGA M. and ETXEBARRIA G. (1997) Regional innovation systems: Institutional and organizational dimensions, *Research Policy* **26**, 475-491.

DAVID P. (1975) *Technical Choice, Innovation and Economic Growth*, Cambridge University Press, Cambridge.

DOSI G. (1982) Technological paradigms and technological trajectories: A suggested interpretation of the determinants and directions of technical change, *Research Policy* **11**, 147-162.

ENGELSMAN and VAN RAAN (1994) A patent-based cartography of technology, *Research Policy* **23**, 1-26.

FELDMAN M. (1994) *The Geography of Innovation*, Kluwer, Boston.

FELDMAN M. and LENDEL I. (2010) Under the lens: The geography of optical science as an emerging industry, *Economic Geography* **86**, 147-171.

FESER, E. (2003) What regions do rather than make: A proposed set of knowledge-based occupation clusters, *Urban Studies* **40**, 1937-1958.

FISCHER M., SCHERNGELL T. and JANSENBERGER E. (2006) The geography of knowledge spillovers between high-technology firms in Europe: Evidence from a spatial interaction modeling perspective, *Geographical Analysis* **38**, 288-309.

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

FREEMAN C. (1985) The ‘National System of Innovation’ in historical perspective, *Cambridge Journal of Economics*, **19**, 5-24.

FRENKEN K., VAN OORT F. and VERBURG T. (2007) Related variety, unrelated variety and regional economic growth, *Regional Studies* **41**, 685-697.

GERTLER M. (1995) Being there: Proximity, organization and culture in the development and adoption of advanced manufacturing technologies, *Economic Geography* **75**, 1-26.

GERTLER M. (2003) Tacit knowledge and the economic geography of context, or the undefinable tacitness of being (there), *Journal of Economic Geography* **3**, 75-99.

GOWER J. (1971) A general coefficient of similarity and some of its properties, *Biometrics* **27**, 857-871.

GRABHER G. (1993) The weakness of strong ties: The lock-in of regional development in the Ruhr area, in GRABHER G. (ed.) *The Embedded Firm: On the Socioeconomics of Industrial Networks*, Routledge, London (pp 255-277).

GRILICHES Z. (1979) Issues in assessing the contribution of R&D to productivity growth, *Bell Journal of Economics* **10**, 92-116.

GRILICHES Z. (1990) Patent statistics as economic indicators: a survey, *Journal of Economic Literature* **28**, 1661-1707.

HALL B., JAFFE A. and TRAJTENBERG M. (2001) The NBER patent citations data file: lessons, insights and methodological tools, National Bureau of Economic Research, Working Paper 8498.

HALL P. and MARKUSEN A. (1985) High technology and regional-urban policy, In HALL P. and MARKUSEN A. (eds.) *Silicon Landscapes*, Allen and Unwin, Boston (pp 144-151).

HAUSMANN R. and KLINGER B. (2007) The structure of the product space and the evolution of comparative advantage, Harvard University, Center for International Development, Working Paper No. 146.

HIDALGO C., KLINGER B., BARABASSI A. and HAUSMANN R. (2007) The product space conditions the development of nations, *Science* **27**, 482-487.

JAFFE A. (1986) Technological opportunity and spillovers of R&D, *American Economic Review* **76**, 984-1001.

JAFFE, A., TRAJTENBERG M. AND HENDERSON R. (1993) Geographical localization of knowledge spillovers as evidenced by patent citations, *Quarterly Journal of Economics* **108**, 577-598.

JAFFE A., TRAJTENBERG M. and FOGARTY M. (2000) The meaning of patent citations: Report on the NBER/Case-Western reserve survey of patentees, National Bureau of Economic research, Working Paper No. 8498.

KIRAT T. and LUNG Y. (1999) Innovation and proximity: Territories as loci of collective learning processes, *European Urban and Regional Studies* **6**, 27-38.

KOO J. (2005) Technology spillovers, agglomeration, and regional economic development, *Journal of Planning Literature* **20**, 99-115.

LAI R., D'AMOUR A. and FLEMING L. (2009) The careers and co-authorship networks of U.S. patent-holders, since 1975, Harvard Business School, Harvard Institute for Quantitative Social Science.

LAMOREAUX N. and SOKOLOFF, K. (1996) Long-term change in the organization of inventive activity. *Proceedings of the National Academy of Sciences* **93**, 12686-12692.

LETEN B., BELDERBOS R. and VAN LOOY B. (2007) Technological diversification, coherence, and performance of firms, *Journal of Product Innovation and Management* **24**, 567-579.

LUCAS R. (1988) On the mechanics of economic development, *Journal of Monetary Economics* **22**, 3-42.

LUNDVALL B. (1988) Innovation as an interactive process: From user-producer interaction to the national system of innovation. In: DOSI G., FREEMAN C., NELSON R., SILVERBERG G. and SOETE L. (eds.) *Technical Change and Economic Theory*, Pinter, London (pp. 349-369).

LUNDVALL B. (ed.) (1992) *National Systems of Innovation: Towards a Theory of Innovation and Interactive Learning*, Pinter, London.

LUNDVALL B. and JOHNSON B. (1994) *The Learning Economy*, Routledge, London.

MACKINNON D., CUMBERS A. and CHAPMAN K. (2002) Learning, innovation and regional development: A critical appraisal of recent debates, *Progress in Human Geography* **26**, 293-311.

MARSHALL A. (1920) *Principles of Economics*, Macmillan, London.

MARX K. (1970) *Capital* 3 vols., International, New York.

MASKELL and MALMBERG (1999) The competitiveness of firms and regions: 'ubiquitification' and the importance of localized learning, *European Urban and Regional Planning Studies* **6**, 9-25.

MAYER H. (2011) *Entrepreneurship and Innovation in Second Tier Regions*, Edward Elgar, Cheltenham.

MORGAN K. (1997) The learning region: Institutions, innovation and regional renewal, *Regional Studies* **31**, 491-503.

NEFFKE F., HENNING M. and BOSCHMA R.(2011) How do regions diversity over time? Industry relatedness and the development of new growth paths in regions, *Economic Geography* **87**, 237-265.

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

NICKELL S. (1981) Biases in dynamic models with fixed effects, *Econometrica* **49**, 1417-1426.

O'HUALLACHAIN B. (1999) Patent places: Size matters, *Journal of Regional Science* **39**, 613-636.

O'HUALLACHAIN B. and LEE D. (2010) Technological variety and specialization in urban invention, *Regional Studies* **45**, 67-88.

PAVITT K. (1982) R&D, patenting and innovative activities: A statistical exploration, *Research Policy* **11**, 33-51.

PAVITT K. (1985) Patent statistics as indicators of innovative activities: Possibilities and problems, *Scientometrics* **7**, 77-99.

PERROUX F. (1955) Note sur la notion de 'pole de croissance', *Èconomie Appliquèe* **7**, 307-320.

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

PRED A. (1977) *City Systems in Advanced Economies*, Hutchinson, London.

QUATRARO F. (2010) Knowledge coherence, variety and economic growth: Manufacturing evidence from Italian regions, *Research Policy* **39**, 1289-1302.

RIGBY D. and ESSLETZBICHLER J. (1997) Evolution, process variety, and regional trajectories of technological change, *Economic Geography* **73**, 269-284.

RIGBY D. and ESSLETZBICHLER J. (2006) Technological variety, technological change and a geography of production techniques, *Journal of Economic Geography* **6**, 45-70.

RIGBY D. and VAN DER WOUDEN F. (2012) Knowledge Relatedness and the Rate of Invention in U.S. Cities. Draft manuscript available from the authors.

ROMER P. (1986) Increasing returns and long-run growth, *Journal of Political Economy* **94**, 1002-1037.

SAXENIAN A. (1994) *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*, Harvard University Press, Cambridge, MA.

SCHERER F. (1982) Interindustry technology flows in the United States, *Research Policy* **11**, 227-245.

SCHUMPETER J. (1950) *Capitalism, Socialism and Democracy*, Harper, New York.

SIMON H. (1959) Theories of decision making in economics, *American Economic Review* **49**, 252-283.

SINGH J. (2005) Collaborative networks as determinants of knowledge diffusion patterns, *Management Science* **51**, 756-770.

SOLOW R. (1956) A contribution to the theory of economic growth, *Quarterly Journal of Economics* **66**, 65-94.

SONN J. and STORPER M. (2008) The increasing importance of geographical proximity in knowledge production: An analysis of U.S. patent citations, *Environment and Planning A* **40**, 1020-1039.

STORPER M. (1995) The resurgence of regional economies, ten years later: The region as a nexus of untraded interdependencies, *European Urban and Regional Studies* **2**, 191-221.

STORPER M. (1997) *The Regional World: Territorial Development in a Global Economy*, Guilford, New York.

STRUMSKY D. LOBO J. and S.VAN DER LEEUW (2010) Using patent technology codes to study technological change, unpublished manuscript, Department of Geography and Earth Sciences, University of North Carolina, Charlotte.

TEECE D., RUMELT R., DOSI G. and WINTER S. (1994) Understanding corporate coherence: theory and evidence, *Journal of Economic Behavior and Organization* **23**, 1-30.

THOMPSON P. and FOX-KEANE M. (2005) Patent citations and the geography of knowledge spillovers: A reassessment. *American Economic Review* **95**, 450-460.

VAN DER WOUDEN F. (2012) The role of relatedness in economic variety on knowledge production of U.S. cities between 1975 and 2005. MA thesis, Department of Economic Geography, Utrecht University.

VERSPAGEN B. (1997) Measuring intersectoral technology spillovers: Estimates from the European and US Patent Office Databases, *Economic Systems Research* **9**, 47-65.

VON HIPPEL (1988) *The Sources of Innovation*, Oxford University Press, Oxford.

WEBBER M., SHEPPARD, E. and RIGBY D. (1992) Forms of technical change. *Environment and Planning A* **24**, 1679-1709.