Organization Science

Vol. 18, No. 6, November–December 2007, pp. 938–954
ISSN 1047-7039 | EISSN 1526-5455 | 07 | 1806 | 0938

INFORMS

DOI 10.1287/orsc.1070.0289

Small Worlds and Regional Innovation

Lee Fleming
Harvard Business School, Morgan Hall 485, Boston, Massachusetts 02163, lfleming@hbs.edu

Charles King III
Greylock McKinnon Associates, One Memorial Drive, Suite 1410, Cambridge, Massachusetts 02142, and
Pleiades Consulting Group, Inc., P.O. Box 711, Lincoln, Massachusetts 01773, cking@gma-us.com

Adam I. Juda
Google, Inc., 76 Ninth Avenue, 4th Floor, New York, New York 10011, juda@goggle.com

Small-world networks have attracted much theoretical attention and are widely thought to enhance creativity. Yet empirical studies of their evolution and evidence of their benefits remain scarce. We develop and exploit a novel database on patent coauthorship to investigate the effects of collaboration networks on innovation. Our analysis reveals the existence of regional small-world structures and the emergence and disappearance of giant components in patent collaboration networks. Using statistical models, we test and fail to find evidence that small-world structure (cohesive clusters connected by occasional nonlocal ties) enhances innovative productivity within geographic regions. We do find that both shorter path lengths and larger connected components correlate with increased innovation. We discuss the implications of our findings for future social network research and theory as well as regional innovation policies.

Key words: small-world networks; innovation; regional advantage

Why are some regions more innovative than others? This question merits study because innovation drives productivity and, ultimately, economic growth (Solow 1957). Many answers have been put forward. Finance researchers focus on the importance of venture capital to innovative risk-taking (Gompers and Lerner 1996, Davila et al. 2003). Labor researchers cite the importance of skilled labor to innovation (Florida 2004) and the creation of a regional infrastructure to support new firm formation (Feldman 2001). Legal scholars and sociologists (Gilson 1999, Hyde 2003, Stuart and Sorenson 2003) propose that lax enforcement of laws designed to restrict the flow of people and ideas (including nondisclosure and noncompetition agreements) contributes to job mobility and higher levels of innovation. Many descriptions of innovative regions cite social networks as the crux of innovation (Marshall 1919, Piore and Sabel 1984, Almeida and Kogut 1999, Singh 2005). Silicon Valley's success, for example, has been attributed to its informal networks of friendship and collaboration (Saxenian 1994).

Despite the acknowledged importance of social networks to innovation, little research has systematically investigated the relationship between network properties and innovation within geographic regions.

The development and application of formal models of macro structure facilitates our understanding of how social networks influence regional productivity. Recent theoretical research on macro networks has focused on the properties of "small-world" networks (Watts and Strogatz 1998, Watts 1999). Small-world networks are defined as clusters of locally dense interaction connected via a few bridging ties. Empirical research has begun to investigate how small-world structure influences economic and sociological outcomes. The predominant hypothesis is that small-world networks should enhance innovative creativity (Watts 1999; Hargadon 2003; Cowan and Jonard 2003, 2004; Baum et al. 2003; Verspagen and Duysters 2003; Schilling and Phelps 2007; Uzzi and Spiro 2006). More innovation is argued to occur because small-world networks enable dense and clustered relationships to coexist with distant and more diverse relationships. The dense and clustered relationships enable trust and close collaboration, while distant ties bring fresh and nonredundant information to the cluster. These attractive hypotheses remain relatively untested (for exceptions, see Uzzi and Spiro 2006 on Broadway musicals and Schilling and Phelps 2007 on strategic alliances and patenting).

In this paper, we first review the nearly universal predictions that small worlds should enhance innovative productivity, explain our differences with those predictions, and develop our hypotheses. Next, we develop and exploit a novel database on patent coauthorship to investigate the effects of small world and collaboration networks on innovation. Because our data on patent coauthorship ties have not been widely used or characterized in the prior literature, we report our field studies on the nature, strength, and longevity of patent coauthorship ties. Econometric tests of regional network structure on subsequent patenting generally support our arguments: (1) Clustering demonstrates no statistically significant influence; (2) shorter path length exhibits a positive and statistically significant influence; and (3) the small-world interaction measure fails to demonstrate a statistically significant effect. A simple measure of the degree of connectedness between regional inventors demonstrates a stronger correlation in significance and magnitude with subsequent patenting than any small-world structure. We discuss the importance of these results for further theoretical and empirical research and for regional innovation policy.

The Small Worlds of Inventors

The idea of small worlds first arose from the finding that seemingly unrelated people are surprisingly close in social space (Milgram 1967). Milgram randomly sampled inhabitants of two small midwestern towns and asked each study participant to forward a letter through personal connections to a Boston address. Conditional on the letter reaching the target, only six personal contacts (on average) separated the participant and the target. Milgram's result passed into urban folklore until Watts and Strogatz (1998) offered a minimal model that reproduced the macro features of the phenomenon.

Watts and Strogatz proposed that small-world networks exhibit tight clusters of local interaction linked by occasional nonlocal interactions whereby any node in the network could still easily reach any other node. Network scholars generally agree that small-world networks facilitate information flow, the spread of epidemics, and the surprisingly short distances between linked sites on the World Wide Web. Watts and Strogatz's paper prompted a flurry of theoretical modeling by physicists and social scientists (Watts 1999; Newman and Watts 1999; Barabasi and Albert 1999; Amaral et al. 2000; Cowan and Jonard 2003, 2004). Empirical studies quickly followed and have found small-world properties in a variety of social network contexts, including German corporate ownership (Kogut and Walker 2001), American corporate boards (Davis et al. 2003), strategic alliances (Verspagen and Duysters 2003), Canadian investment bank syndicates (Baum et al. 2003), e-mail networks (Dodds et al. 2003), Italian scientific and academic collaboration networks (Balconi et al. 2004), and invisible scientific colleges (Goyal et al. 2004).

By Watts and Strogatz's definition (1998), small-world networks simultaneously exhibit high clustering and low path length. In the context of inventor collaboration, clustering increases when two inventors are more likely to patent together, if both have patented with the same third inventor. As a visual example, the boxed areas in Figures 1 and 2 contain highly clustered inventors. Highly clustered networks are less vulnerable to the removal of a single inventor from the structure. The concept is similar to cohesion (Uzzi and Spiro 2006) and opposite to brokerage (Burt 2004). Path length measures the social distance between any two inventors as the minimum number of collaborative links between them. For example, if Tom worked with Dick, Dick worked with Harry, and Tom did not work with Harry, then the path length between Tom and Harry would be two.

Watts and Strogatz (1998) integrated the ideas of clustering and path length by considering the extremes of *regular* and *random* graphs. *Regular* graphs, where each node is connected to its k nearest neighbors, exhibit high clustering and long path length. For example, for $k = 4$, your immediate neighbors would be directly connected, both to you and one another. Distant neighbors, however, would be connected through a large number of indirect ties. In contrast, *random* graphs, where nodes are randomly connected, exhibit low clustering and short path length. In a purely random graph, you are as likely to be connected to your immediate neighbors as to distant neighbors. Thus, local neighbors can be isolated and distant neighbors connected through only a few indirect ties. Between these two extremes are intermediate small-world regimes—essentially regular graphs with a small number of random connections. In these graphs, high clustering (relative to a random graph) and low path length (relative to a regular graph) coexist simultaneously. Figure 3 illustrates the extremes and the intermediate small-world regime.

An important limitation of small-world measures is that the network must be fully connected (i.e., there must exist a path between any two nodes). Real social networks often include isolates. We followed the conventions of most research in this area (Newman 2000, Kogut and Walker 2001, Davis et al. 2003, Verspagen and Duysters 2003, Baum et al. 2003, Uzzi and Spiro 2006, for an innovative exception, see Schilling and Phelps 2007 for development of a harmonic weighting method) and focus on the largest fully connected component within each region. The largest connected component is the largest set of inventors in a region who can trace a direct or indirect collaborative path to one another. Watts (1999) argues: "The graph must be connected in the sense that any vertex can be reached from any other vertex by traversing a finite number of edges.... Disconnected graphs pose a problem because they necessarily have [path length] $L = \text{infinity}$" (p. 499). Figures 1 and 2 illustrate the largest connected components of patented inventors in Silicon Valley and Boston from 1986 to 1990. They provide an empirical illustration of inventive small worlds that consist of clusters of cohesive interaction, linked together by occasional bridging connections.

These patent coauthorship networks provide a rich opportunity to study how small worlds influence creativity, because these networks represent a primary conduit

Figure 1 Inventors of Silicon Valley's Largest Component in 1986–1990 by Assignee and Importance of Inventions

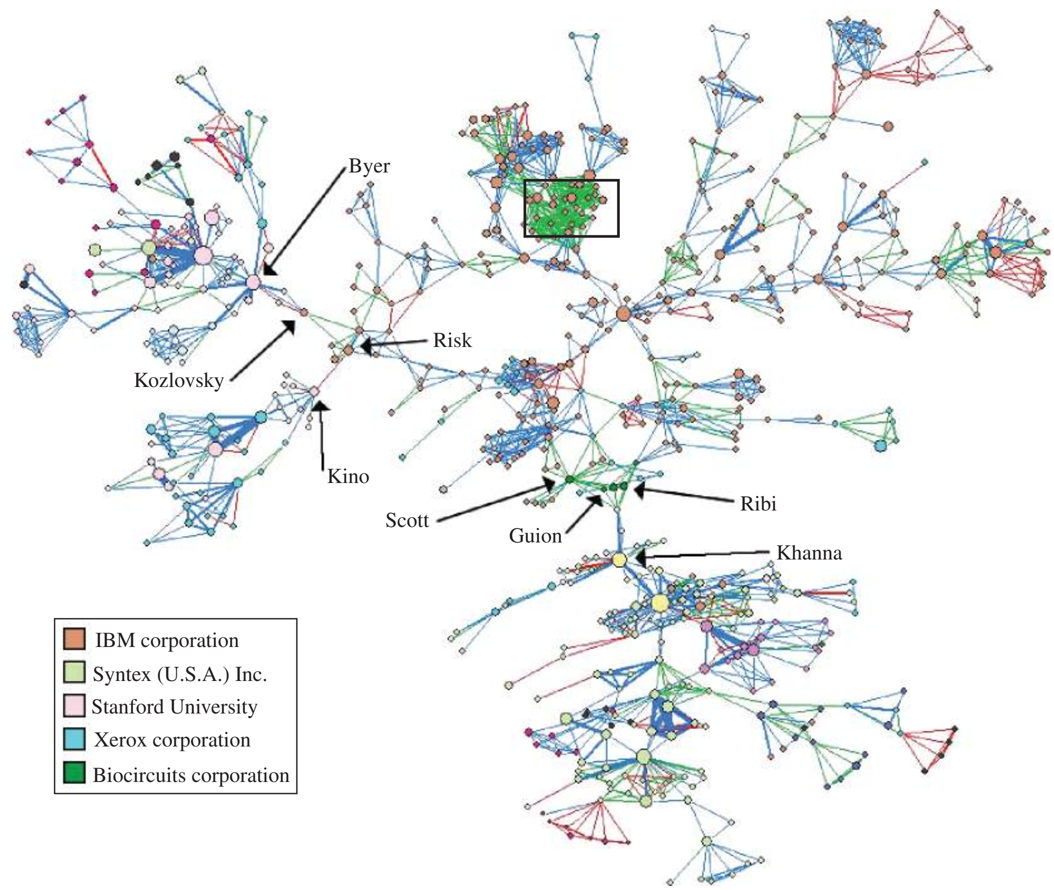

Notes. Node sizes reflect the number of future prior art cites to an inventor, normalized by the number of collaborators (future prior art cites correlate with value, *see* Albert et al. 1991). Tie width indicates number of collaborations, tie color indicates age of tie (red is five years prior, blue is two to four years prior, and green is prior year), and colors indicate assignee. Boxed area provides example of highly clustered inventors. Note that the figures do not illustrate the thousands of other (by definition) smaller components in each region; inventors need not connect to any extant component – or even another node. They can connect to small components, such as dyads or triads, or work their entire careers in complete isolation. Graphed in Pajek with Kamada-Kawai/Free algorithm (Batagelj and Mrvar 1998). Adapted from Fleming and Marx (2006).

of information for inventors. Although inventors do use nonsocial sources of novel information, detailed studies indicate that they rely heavily on social sources (Allen 1977). Social sources are particularly important for the transfer of tacit information.[1] Inventors are less likely to read documentation, textbooks, or scientific literature and more likely to approach a friend or colleague who has appropriate experience or does read the scientific and technical literature. Because asking for help by definition requires an admission of need, engineers are careful about whom they approach (Borgatti and Cross 2003). They tend to ask colleagues who reciprocate their requests or those who have little effect on their career evaluation—for example, outside suppliers and friends at other firms (Allen 1977). An engineer at the design firm IDEO remarked, "Where I worked for (*sic*) before, you just didn't ask for help. It was a sign of weakness" (Sutton and Hargadon 1996). Our fieldwork corroborates these observations and reveals that prior coauthors are prime candidates for information. This evidence implies that an inventor's past collaboration network will strongly influence subsequent productivity and, therefore, provide a powerful context in which to investigate how small worlds influence creativity and innovation.

Small Worlds and Innovation

Organizational researchers have adopted the formal models of small-world structure and argued that small worlds improve creativity and innovation. These arguments can be organized into the influences of clustering, path length, and their interaction. Uzzi and Spiro (2006) focus on clustering and argue that it improves creativity in musical productions "because clustering promotes collaboration, resource pooling, and risk sharing." These beneficial effects result from the increased trust that occurs within closed and embedded social contexts (Granovetter 1985, Uzzi and Spiro 2006, Obstfeld

Figure 2 Inventors of Boston's Largest Component in 1986–1990 by Assignee and Importance of Inventions

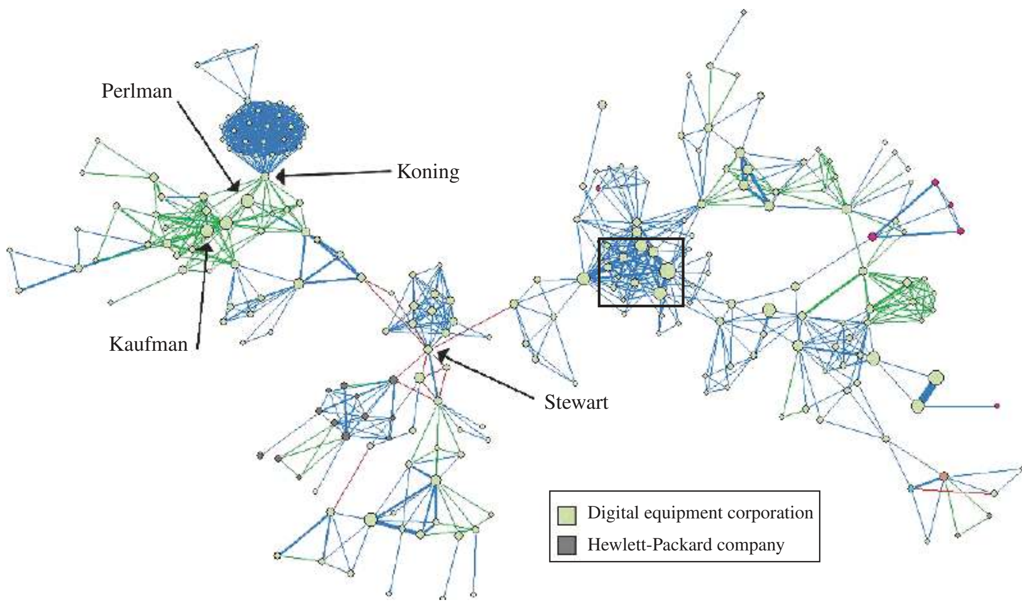

Notes. Node sizes reflect the number of future prior art cites to an inventor, normalized by the number of collaborators. Tie width indicates number of collaborations, tie color indicates age of tie (red is five years prior, blue is two to four years prior, and green is prior year), and colors indicate assignee. Boxed area provides example of highly clustered inventors. Graphed in Pajek with Kamada-Kawai/Free algorithm (Batagelj and Mrvar 1998). Adapted from Fleming and Marx (2006).

2005). Schilling and Phelps (2007) make similar arguments for clustering and also propose that once information crosses between clusters (made possible in a small world because the clusters are connected), it flows more easily within clusters. Uzzi and Spiro (2006) add that the effect of clustering is nonmonotonic because extreme clustering promotes recirculation of redundant information. All the work that we reviewed on small worlds and innovation (Cowan and Joward 2004, Uzzi and Spiro 2006, Schilling and Phelps 2007, Hargadon 2003, Verspagen and Duysters 2003) argued that decreased path length should improve innovation because of easier and improved information transfer.

Figure 3 Idealized Examples of Regular, Small-World, and Random Networks (From Watts and Strogatz 1998)

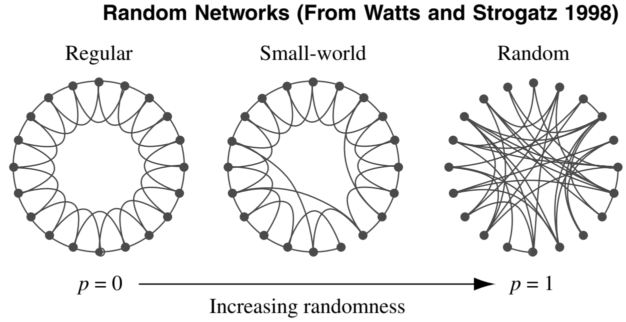

Notes. Only the middle configuration represents a small-world network that simultaneously exhibits high clustering and short path length. It is essentially a regular network with a few randomly rewired connections. Reprinted by permission from Macmillan Publishers Ltd: *Nature*, 393, 440–442 (1998).

We agree with the path length argument and expect that decreased path length will improve innovative productivity. Inventors will almost always profit from exposure to new information, although at some extreme they may face cognitive overload and be better off if they limited or filtered their exposure. Short path lengths expose inventors to new information because they connect them with different sources and nonlocal perspectives. Short path lengths in a network indicate that distant information—where distance can be technological, organizational, or geographical—is surprisingly close in social space (Singh 2005). Supporting these arguments, Cowan and Jonard (2004) develop an agent-based model that demonstrates the benefit of decreased path length for the diffusion of innovations. Without this exposure to new information and perspectives from others, inventors will become insular and less creative. Faster diffusion of innovations and the juxtaposition of diverse knowledge flows should increase the subsequent innovative productivity of regions whose largest components exhibit short average path length.

Hypothesis 1 (H1). *Decreased path length within a region's largest connected component will correlate positively with increased future patenting in the region.*

Independent of the path length within components, the formation of connections across components should also improve regional innovative productivity. Similar to the path length argument, the aggregation[2] of inventor components increases subsequent innovativeness because it enhances information flow and knowledge spillovers.

Regions whose inventors stay isolated will lack a large connected component. Isolates and small clusters will be left without access to new ideas and results. Because new results will remain unknown outside local contexts where the breakthrough occurred, opportunities to apply them will remain unexploited. A promising new combination will not occur because knowledge of a previous combination will not diffuse into a new and potentially fertile context. Network aggregation also enables greater opportunities for technological brokerage between previously disconnected technological communities (Stuart and Podolny 1999, Hargadon 2003, Burt 2004). The connecting bridges encourage cross-disciplinary fertilization, as is illustrated in Figure 1. As with decreased path length, and for similar reasons, we expect that the connection and aggregation of isolated components will correlate positively with subsequent patenting in the region.

Hypothesis 2 (H2). *The size of a region's largest connected component will correlate positively with increased future patenting in the region.*

We agree with many of the arguments for the benefits of clustering but can cite alternative arguments for its detrimental (Burt 2004) and contingent (Nerkar and Paruchuri 2005, McEvily and Reagans 2005, Fleming et al. 2007a) effects as well. The main problem, readily acknowledged even by the proponents of clustering (Uzzi and Spiro 2006), is that clustering can lead to insularity, lack of exposure to new and diverse perspectives, and ultimately decreased creativity. Clustering still has benefits, however, and its optimal degree depends on a variety of technical, psychological, and social dynamics. Clustering probably improves productivity, for example, in groups with scientific training that generate original and pertinent knowledge as part of their inventive efforts. The group in this case will benefit more from internal communication and focus than external search and exploration. A similar situation would hold for nonscientific inventors who were ahead in the invention of a new-to-the-world technology. Because they would know more internally than any other group in the world, their marginal benefit from external connections would be less. They might be distracted, and/or lose their creative edge, if competitors gained access to their proprietary lead, or pressured the group into less creative processes. To the extent, however, that a group does not generate all the needed information internally, then clustering becomes less efficient.

Psychological and social dynamics further modify the optimal degree of clustering. Following the formation of a new team, inventors typically become more productive as they build team cohesion, learn from one another, and gain an appreciation of the location and abilities of expertise within the team. Clustering at this early stage probably improves inventive productivity. After some time, however, a group that stops importing new components, perspectives, or information (Katz 1982) will grow stale and productivity will stall. The effect will be contingent on the original diversity of training and experience within the group. Greater initial diversity would make clustering more helpful and delay the stall in productivity. In contrast, if management hires similar individuals, promotes in-group pressures for conformity (Hunt et al. 2003), or does not encourage contrary opinions and risk-taking, this will hasten the stall.

These examples illustrate how clustering can have a multitude of contingent effects, depending on the technical challenge, modes of knowledge generation and transfer, structural history, demographics, and norms of interaction. This implies that the benefit of clustering will vary greatly and conceivably—or even frequently—turn negative. As illustrated in Figures 1 and 2, a single firm or region might encompass a great variety of contingencies, with the result that the average value of clustering across the network would miss individual contingencies. Each organization in the figures works in a different technology (from pharmaceuticals to optics); employs different strategies and goals (from profit making to teaching and research); and provides different incentives for collaboration, resource pooling, and risk sharing. This locally contingent influence of clustering makes it difficult to predict the influence of average cohesion on the innovative productivity of the entire network. We therefore make no predictions about the independent influence of clustering on regional innovation.

Despite the difficulty of predicting the first-order effect of clustering, it remains a crucial component of small-world structure. The heart of the small-world and creativity argument lies not in the first-order effects but in the interaction of increased clustering and decreased path length. A small-world network should be more creative, to the extent that its clusters are tighter *and* the path lengths between the clusters are shorter. Dense clusters of collaboration within a larger network should become more creative as the path lengths of the larger network decrease. This occurs because inventors can rely on their close collaborators to collect and interpret increased amounts and diversity of external information. Alternately, if path lengths are held constant, the network should become more productive as clustering increases because increased information flow will be exploited more effectively. As a result, embedded clusters can maintain their productive focus without becoming distracted or stale. Independent of any main or additive effects of clustering or path length, an increase of one in the presence of the other should still create a positive interaction.

Hypothesis 3 (H3). *The interaction of decreased path length and increased clustering within a region's largest connected component will correlate positively with increased future patenting in the region.*

Figure 4 Regional Small-World Structures of U.S. Patented Inventors

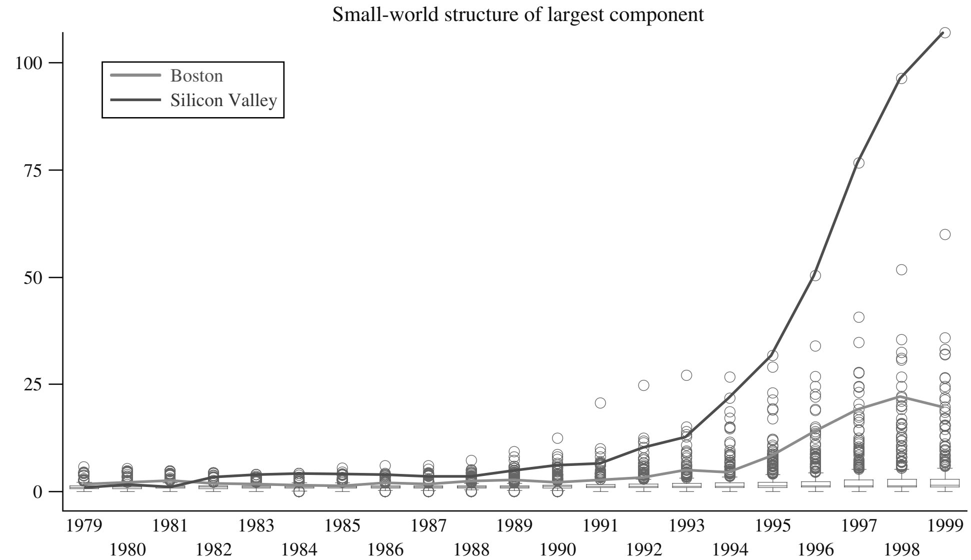

Notes. The small-world variable is calculated as clustering divided by path length for inventors within the largest component of each region. It is without units, because clustering and path length are both normalized. The *x* axis indicates the last year in a five-year moving window; the box plots illustrate quartile percentiles, upper and lower adjacency values, and outliers, for all 337 U.S. Metropolitan Statistical Areas.

Data and Methods

Each U.S. patent lists (1) inventors (also referred to as the authors), (2) assignee (i.e., the owner, typically a firm or university, but also individuals), (3) technological classes and subclasses, and (4) hometown(s) of the inventors. The patent does not, however, provide consistent listings of inventor names or unique identifiers for the inventors. Using a variety of conditional matching algorithms (described in the appendix), we identified 2,058,823 unique individual inventors and their patent coauthors from all U.S. patents granted from 1975 to 2002 (a total of 2,862,967 patents). These data enabled us to construct regional collaboration networks for moving five-year windows[3] in all 337 U.S. Metropolitan Statistical Areas (MSAs). An MSA is a geographic region with a large population, together with adjacent communities, that has a high degree of economic and social integration within that nucleus. Silicon Valley, for example, encompasses San Jose and the surrounding cities in Santa Clara County, California. Boston includes all of eastern Massachusetts and a small part of southeastern New Hampshire.

We identify small-world regimes following the literature's empirical conventions (Kogut and Walker 2001, Davis et al. 2003, Verspagen and Duysters 2003, Baum et al. 2003) and calculate normalized clustering divided by normalized path length for the largest connected component in each geographical region.[4] We normalize the measure to account for differing numbers of inventors in different connected components. Because clustering does not vary greatly according to component size, the main importance of normalization is for path length. Without normalization, a large network might appear to have long path lengths relative to a small but poorly connected network. For example, a few inventors connected in a line might have a similar path length to the inventors in Figures 1 and 2. We also restrict consideration to each region's largest component because small worlds remain undefined across disconnected components. Figure 4 shows how the small-world measure has been increasing—thus demonstrating that inventor collaboration structure has become more small-world for many MSAs, Silicon Valley in particular, since the early 1990s. Figure 5 graphs the aggregation measure over time for all MSAs.

The Emergence of Innovative Small Worlds

To investigate the nature, characteristics, ties, and dynamics of inventive small worlds, we interviewed inventors in Silicon Valley and Boston from the 1986–1990 timeframe as part of a larger qualitative study (Fleming et al. 2007b).[5] That study sampled inventors who linked—and did not link—smaller components into the emerging giant component of Silicon Valley (and based on his graphically compelling position in Figure 2, the person at the center of the disintegration of Digital Equipment Corporation (DEC) in 1990).[6] Researchers provided each inventor with an illustration of their collaborative network in 1990 and asked them if it was accurate (or whether it missed important collaborations that remained unpatented), the nature of the collaborative tie, whether information flowed across the tie, whether they had maintained personal or technical contact with their patent coauthors, and the details of their collaborative mobility.

Figure 5 Regional Largest Component of U.S. Patented Inventors

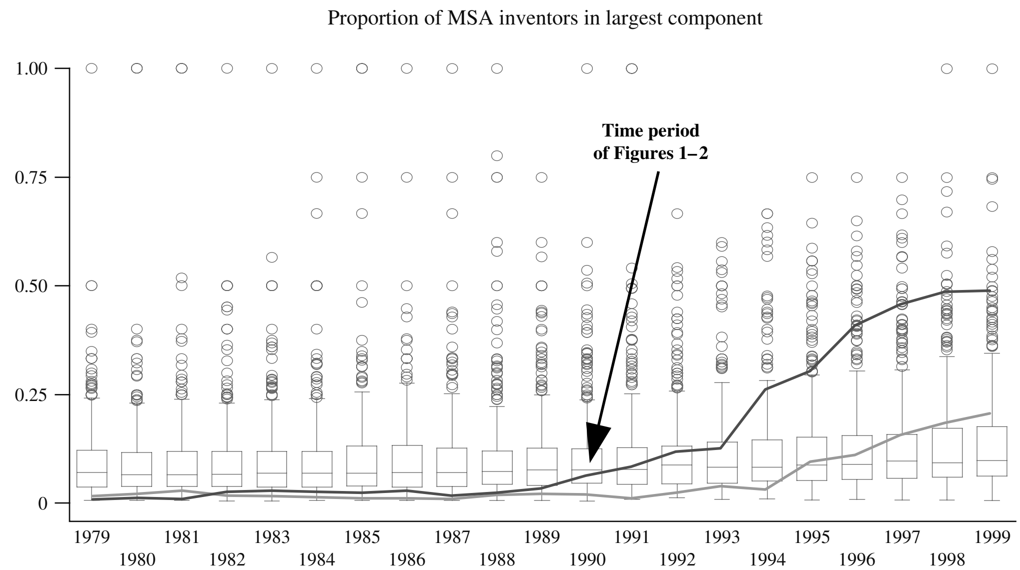

Notes. The measure is the largest proportion of the number of patented inventors in a region that can trace an indirect collaborative path to one another. The *x* axis indicates the last year in a five-year moving window; the box plots illustrate quartile percentiles, upper and lower adjacency values, and outliers, for all 337 U.S. Metropolitan Statistical Areas. Adapted from Fleming and Marx (2006).

The results of our interviews suggest that the aggregation of Silicon Valley's giant component in 1990 occurred primarily for three reasons. First, the IBM Almaden Valley Labs provided a stable structural backbone for the region (the tan-colored nodes of the upper right in Figure 1 denotes IBM inventors). At the time, IBM still invested heavily in research and hiring. Second, several Stanford University doctoral graduates took employment at IBM (Risk 2003). This is shown by the ties that connect the multicolored but predominantly pink Stanford University nodes on the upper left part of Figure 1. Specifically, William Risk connected Professor Gordon Kino's microscopy students and most of the Gintzon Applied Physics Lab when he began patenting at IBM in 1989 (Kino 2003). William Kozlovsky did the same a year later when he departed from Professor Robert Byer's optical technology lab (Kozlovsky 2003). The third linkage occurred across the now failed BioCircuits, a pioneering startup that attempted to integrate biology and engineering (Ribi 2003). The dark green nodes connected by a web of green ties in the lower right of Figure 1 illustrate the BioCircuits bridge between IBM's optical and computer technology and the pharmaceutical technology of Syntex. Glenda Choate provided the pharmaceutical abutment for the bridge when she moved to BioCircuits from Microgenics, a drug-testing technology firm. The IBM abutment resulted from a postdoctoral program that intentionally seeded local industry with scientists after a 1- or 2-year appointment at Almaden Valley Labs. Todd Guion, who worked for BioCircuits after his IBM postdoctoral program, recommended bringing in his former advisor, Campbell Scott, to solve a particular optical problem (interviewed by authors July 9, 2003). IBM reluctantly allowed Campbell Scott to act as a scientific advisor—and patent with BioCircuits—while maintaining his IBM employment.

For a variety of reasons, Boston failed to aggregate in the same period despite similar numbers of inventors, technologies, firms, and greater university patenting. First, Boston lacked a stable structural backbone. In Silicon Valley, IBM Almaden Valley Lab inventors became the basis of the largest component in 1989 and have since remained within the largest component. In Boston, dominance of the largest component passed back and forth between Massachusetts Institute of Technology (MIT) and DEC. DEC dominated in 1990 and 1992; MIT dominated in the prior and intervening years and permanently since 1993. The immediate cause of this instability is dramatically illustrated by the expiring red ties in the middle of Figure 2. Robert Stewart is the only inventor to integrate the three major subcomponents at DEC. Stewart (2004) said that his brokering role arose from his popularity as a design reviewer across different DEC product lines. The immediate cause of the upper tie disintegrations was that many related patents were filed just before product shipments (in this case, DEC's Nautilus project in early January 1986). The lower left tie had been one of several collaborations between Stewart and the research and development (R&D) and networking groups, and happened to expire at the same time. Second, graduate students (or at least those with patents) did not leave MIT to join DEC. MIT graduated many electrical engineers, but DEC maintained an official (if imperfect) hiring freeze during this period. Instead, MIT graduates took nonlocal employment, often with other academic institutions (Cohen 2003). Had they taken local employment, as the Stanford students did, Boston's largest components would have merged, similar to what occurred in Silicon Valley. Finally, both regions had little movement of established inventors from IBM or DEC to other large, established firm components in the region. Even though DEC was having problems during this period, all four DEC inventors we interviewed told us they were technically challenged and professionally content. They did not feel it necessary to change employers (Kaufman 2003, Koning 2003, Perlman 2003, Stewart 2004).

Although the disintegration of large components, such as occurred in Boston, has received little attention, the dramatic emergence of a largest or giant component such as happened in Silicon Valley (and happened in Boston three years later) has been modeled robustly, using a variety of assumptions and mechanisms. The phenomenon is consistent with critical transitions and percolation models (Amaral et al. 2000), the spread of fads or epidemics (Newman and Watts 1999), and the aggregation of isolated scientific collaborators (Newman et al. 2002). Such a dominant largest component size is not without precedent, for example, 50% of the entire network for economics collaboration networks (Goyal et al. 2004), and 90% and 100% for Canadian investment banks (Baum et al. 2003). In the current empirical context, we observe that as inventors work within larger components, the network reaches a critical point where the isolated components aggregate into a giant component. The significant and positive correlation between the measures of Figures 1 and 2 (the small-world interaction and size of the largest component, $r = 0.32$, $p < 0.001$) indicates that, at least for inventor collaboration networks, one cannot ignore the aggregation process when studying the dynamics of small-world networks.

Characterization of Patent Collaboration Ties

Coauthorship ties provide attractive illustrations and compelling stories of formation, but what do they represent, particularly with regard to information flow and the effect of that flow on subsequent inventive productivity? Singh (2005) reports significant information flow between patent coauthors, as measured by citations from future patents linked by direct—and even indirect—collaborative ties. The results hold even after econometrically controlling for the greater likelihood of a citation simply because the inventors work in similar technologies. Singh goes on to demonstrate that a large fraction of the geographical citation spillover (Jaffe et al. 1993) results from coauthorship networks. (Breschi and Lissoni 2004 find similar results for European inventors.)

All but three of our 16 interviewees reported some degree of technical interaction after a patent coauthorship (Fleming et al. 2007b). Most important, given our interest in spillovers and regional innovation, they reported these interactions even when they no longer worked at the same firm. Only three inventors denied any technical communication across firm boundaries in the years following a collaborative tie. Hans Ribi, CEO of the failed BioCircuits company that bridged IBM and the pharmaceutical component in the Valley, stated that patents exist to protect intellectual property (Ribi 2003). He argued that patents guarantee that information does not flow across organizational boundaries after inventors no longer work for the same firm. A prominent inventor within the pharmaceutical abutment, Pyare Khanna, kept his inventors in organizational silos to avoid such flows (Khanna 2003; he is now CEO of a biotech startup). Khanna's negative and strong managerial reaction can be interpreted as evidence that the potential for such flows is real. A third inventor, Salvador Umatoy from Applied Materials (not illustrated in the figures), also stated that such flows did not exist, at least across firm boundaries (although he has only patented with colleagues at Applied Materials and most of them remain with him at the firm; Umatoy 2003).

These were the most negative responses to our question about whether technical information flowed between past collaborators. Other responses varied from tepid to strong affirmation. University professors and their students, in particular, maintain close ties at conferences, and continue to visit each other long after their inventive collaboration (Gordon Kino and William Risk had met the week before the interview). Michael Froix (not illustrated) quit Raychem in the late 1980s and became a self-employed inventor within the Valley (Froix 2003). According to his account, Froix relied on his friends and former coauthors for infrastructural as well as informational support (for example, access to lab facilities after hours).[7] The first three DEC inventors at the core of Boston's 1990 largest component (Charles Kaufman, Paul Koning, and Radia Perlman, illustrated in Figure 2) still maintain close contact and have continued collaborating on patents and scientific papers across the firm boundaries of their subsequent employers; in fact, they had all communicated just before our independent interviews (Kaufman 2003, Koning 2003, Perlman 2003). Robert Stewart planned to see the collaborator to his lower left the following weekend (Stewart 2004).

These descriptions and the network graphs in Figures 1 and 2 indicating the number of dyadic collaborations by tie width imply a wide distribution of the strength of ties. This issue is important. Research correlates weak ties with nonredundant information (Granovetter 1973) and codified information transfer, strong ties with tacit information transfer (Kogut and Zander 1992, Szulanski 1996, Hansen 1999); cohesive and redundant ties with more accurate information transfer (Ahuja 2000, Sorenson and Stuart 2001, Reagans and McEvily 2003); and intraboundary ties with advantaged transfer of moderately complex information (Sorenson et al. 2006). Some ties would be extremely weak even during their initial formation, for example, if two inventors on a patent had hardly worked together. At the other end of the distribution, some collaborative ties would be exceptionally strong.

Whereas different types of small-world networks vary in their capability to transfer different types of knowledge, prior collaborative patenting ties are potentially effective for all types of information, including heterogeneous and difficult to transfer information. With redundant ties, Figures 1 and 2 indicate an impressive cohesion within organizations, but also obvious vulnerabilities across the boundaries, particularly if the connections narrow to a single cut-point.[8] We discussed this issue at length with Campbell Scott, the IBM scientist who took part in the BioCircuits bridge (interviewed by authors July 9, 2003). He said that he worked with every member of the startup, even though his name appears only with a subset of collaborators as patent coauthors. It is therefore likely that patent data overstate some ties and miss other redundant information paths. Robert Stewart (2004) also explained that although he was the integrating center of DEC's patent network, he had many nonpatent interactions with other inventors on the other side of each bridge.

In summary, we argue that the amount of technical information that flows along an observed collaborative tie in later years varies greatly—from none to a great deal. Given this distribution, we argue that, on average, the information flow between former patent coauthors is nonnegative and occasionally substantial. The distribution of tie strength varies from exceptionally weak to extremely strong, such that ties support a variety of different types of information flows. Because patent collaboration ties span a wide distribution of tie strengths and characteristics, we do not make any assumption about the content of information flow or the capability of a tie to transmit information. Rather, we introduce measures of information heterogeneity and alternate diffusion paths within our statistical models.

Statistical Models

To more systematically explore the relationship between structure and innovative productivity, we modeled regional patenting from 1980 to 2000 with fixed-effect conditional likelihood estimators (Hausman et al. 1984). Count models are appropriate because the dependent variable, regional patent counts by year, takes on whole number values. Using linear models for count data can result in biased and inconsistent coefficient estimates (Greene 2002). Negative binomial models are preferred because our data demonstrate overdispersion (rejection of Poisson model at $p < 0.0001$). Though it sacrifices efficiency, a fixed-effects model is preferred because it considers within-region variation only, i.e., it controls for time-invariant, regionally idiosyncratic characteristics. A Hausman test rejects random effects specification at $p < 0.0001$. An additional benefit of the Hausman model is that it allows dispersion to vary by region. We ran all analyses in STATA Version 7.

Dependent Variable

The models estimated the influence of small-world and largest component measures during moving five-year windows on the number of successful patent applications in the following year. The count included all utility patents granted to the region that had been applied for within the subsequent year, up to the end of the data in July 2002. Longer lags and different window sizes did not demonstrate substantively different results.

Explanatory Variables

Watts and Strogatz (1998) calculate clustering by considering the number of "triples" or the pairs of an inventor's collaborators who work with one another. We follow their approach (as has most subsequent empirical research) and calculate individual clustering as the number of actual triples for an inventor, divided by the number of potential triples, and average the inventor scores across the largest component (*Clustering of LC*). Inventors with one or zero ties receive a clustering score of zero. We normalized clustering by dividing the average clustering by the theoretical clustering of a fully connected regular graph (Watts and Strogatz 1998). Similarly, we calculated normalized path length (*Inverse Path Length of LC*) as the average path length divided by the theoretical path length of a comparable fully connected regular graph of size (N) and mean degree (z). We approximate the theoretical path length for a regular graph as $N/2z$ when $N > 2z$, and as $2 - (z/(N-1))$ when $N \leq 2z$. We determine N and z empirically for the largest component for each period and region.

Because the small-world measure requires path length in the denominator, we inverted the measure. After centering the cluster and inverse path length variables to facilitate interpretation of their effects (Friedrich 1982), we calculated the small-world measure as clustering multiplied by the inverse of path length (*Cluster∗Inverse Length*). To measure the extent of network aggregation, we calculated the proportion of inventors within each region who had a collaborative tie within the largest component in the region (*Size of LC*). Inventors within regions with small values of this variable remain more isolated. We tested the robustness of all proportion measures with Herfindahl indices and found similar results. The third author programmed network construction and variable measurement in C.

Control Variables

The models include the log of the number of inventors in the region (*Ln inventors*) to account for the number of people actively engaged in invention. We logged all count variables to account for their entry in exponential form into the models (and added 0.01 to variables with a zero minimum). *Ln extra-regional inventors* measures the number of inventors outside the region who coauthor with the region's inventors. Besides increasing the available resources, such collaborations provide additional learning and nonlocal insight for a region (Gittleman 2003). Given the importance of institutions in the formation of regional networks (particularly universities, see Owen-Smith and Powell 2004, McEvily and Zaheer 2004), the models include the number of patents in a region with a university assignee (*Ln university patents*). Labor mobility is higher during changes in employment levels (Angel 1989; measured as *Ln employment*; data from University of Virginia 2004). To control for general changes in wealth, the models include personal income per capita *PCPI (in thousands of dollars)*, deflated by Bureau of Labor Statistics' Consumer Price Index for 2002 (Bureau of Economic Analysis 2004). Minimum values of the variable tended to occur early in the time series within New Mexico and Arizona. Ideally, the models would include the amount of research and development spending in each region, but this remains impossible to measure. Even though R&D data are available for publicly traded firms, the data are not broken down by region (even within firms' internal accounting data), nor are private firm data available. Proxies such as technical professional employment would also be desirable, but are not available as a time series over the years of our study (Acs et al. 2002). Given that location patterns of R&D laboratories tend to be stable over time (Acs et al. 2002), the fixed effects models should account for much of the variation.

Because regions vary in their adoption and generation of new technologies, the models include the age of each region's technology (*Technology age*: measured by the average of the sequential prior art patent numbers cited by patents within a region). *Technology Herfindahl* controls for technological heterogeneity and diversity within each MSA by indexing the U.S. Patent Office technology classes by patent. The Patent Office divides all patents into approximately 400 technological classes, for example, class 437 (Semiconductor Device Manufacturing: Process) and 935 (Genetic Engineering: Recombinant DNA Technology, Hybrid). The models include the average number of assignees for each inventor in a region over the five-year period (*Assignees/inventor*). This is a direct measure of employee movement and its potential influence on network formation and creative efficacy. It also controls for the potentially greater fertility of ties that span organizations. An alternate control measure of the number of cross-organizational ties returned similar results. To account for the overall number of ties between inventors, the models include *Tie density*. We calculated density as the actual number of ties divided by the potential number of ties (actual ties/number of inventors in the region, choose 2). To measure the influence of alternate sources of collaborative norms and information diffusion mechanisms, we first determined the largest employer within each region's largest component and then calculated the proportion of inventors who worked for that employer (*Largest org in LC*). We also found no substantive difference with restricted models of only homogenous largest components (for example, regions where the networks are similar to Figure 2, where the largest component is predominantly DEC).

One empirical challenge with coauthorship data is that coauthors are clustered by definition (that is, all coinventors listed on any one patent will automatically be clustered with one another). This is common to all affiliation data—for example, musical collaborations, projects, or attendance at meetings—and creates a bipartite graph. As a result, the empirical measure of clustering can be biased. Newman (2001) develops a simulation model that estimates unipartite clustering from a bipartite graph based on the distribution of affiliation events (in this case, patents) and individuals (in this case, inventors) of the original graph. Uzzi and Spiro (2006) use Newman's simulations to estimate the clustering coefficient of a similarly distributed random graph, but we prefer a more direct approach of including controls for the number of patents and inventors within the graph (*Ln LC inventors* and *Ln LC patents*). Simple inclusion of controls allows more parsimonious modeling of the data and avoids the distributional and parametric assumptions of the Newman model. Although not shown, all models included categorical variables for each year with the exception of 1980 to control for systematic trends in patenting. 61 regions did not patent during the years of observation, and economic data were unavailable for 835 early observations. Models with all patent data demonstrate similar results (that is, including data points without economic data). Tables 1 and 2 list summary and correlation statistics, respectively.

Table 1 Summary Statistics ($n = 6{,}242$)

Variable	Mean	Std. dev.	Min	Max
Patents year $t+1$	231.00	528.49	0.00	9,162.00
Ln regional inventors	5.34	1.54	0.70	10.55
Ln extra-regional inventors	4.13	2.73	−4.61	9.82
Ln employment	11.63	1.14	9.28	15.12
PCPI (000s)	19.99	5.04	0.00	50.94
Ln university patents	−0.04	3.53	−4.61	7.41
Technology age	3.88	0.45	2.29	5.13
Technology Herfindahl	0.05	0.04	0.01	0.50
Ave assignees/inventor	1.16	0.09	1.00	2.00
Average tie density	0.01	0.02	0.00	0.42
Ln LC inventors	3.27	1.50	0.70	9.93
Ln LC patents	2.92	1.88	0.01	10.12
Largest org in LC	0.67	0.24	0.04	1.00
Clustering in LC	0.02	0.25	−0.98	0.35
Inverse path length in LC	−0.07	4.82	−1.90	130.73
Size of LC	0.11	0.10	0.01	0.75

Table 2 Correlation Statistics ($n = 6{,}242$)

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
(1) Patents year $t+1$															
(2) Ln regional inventors	0.65														
(3) Ln extra-regional inventors	0.38	0.45													
(4) Ln employment	0.61	0.85	0.46												
(5) PCPI (000s)	0.40	0.50	0.22	0.40											
(6) Ln university patents	0.48	0.74	0.40	0.63	0.37										
(7) Technology age	0.31	0.47	0.09	0.27	0.52	0.48									
(8) Technology Herfindahl	−0.22	−0.54	−0.33	−0.44	−0.22	−0.44	−0.20								
(9) Ave assignees/inventor	0.23	0.37	0.06	0.24	0.43	0.38	0.67	−0.16							
(10) Tie density	−0.23	−0.64	−0.34	−0.49	−0.28	−0.44	−0.22	0.58	−0.18						
(11) Ln LC inventors	0.65	0.84	0.42	0.65	0.51	0.63	0.60	−0.37	0.46	−0.42					
(12) Ln LC patents	0.59	0.81	0.41	0.62	0.45	0.57	0.52	−0.34	0.43	−0.45	0.93				
(13) Largest org in LC	−0.28	−0.33	−0.18	−0.27	−0.18	−0.29	−0.19	0.18	−0.37	0.21	−0.31	−0.34			
(14) Clustering in LC	−0.10	−0.16	−0.09	−0.14	0.03	−0.05	0.08	0.00	0.04	0.17	−0.13	−0.35	0.11		
(15) Inverse path length in LC	0.64	0.44	0.16	0.32	0.25	0.30	0.26	−0.14	0.21	−0.16	0.62	0.55	−0.18	−0.07	
(16) Size of LC	0.03	−0.20	−0.16	−0.27	−0.05	−0.19	0.15	0.36	0.08	0.45	0.28	0.23	0.06	0.00	0.40

Results

Table 3 presents the modeling results. The models include explanatory terms individually, then the small-world interaction, followed by a full model. Not surprisingly, the number of inventors, both within and outside the region, demonstrated large positive and statistically significant effects. General employment in the region exhibits no significant influence, and personal income generally exhibits negative significance. This is probably because some very poor regions patent a great deal (for example, New Mexico, where Pueblo Communities and National Research Labs colocate) and some very rich regions, such as New York City, do not. The number of university patents is not statistically significant, which might occur if graduate students leave the region for employment elsewhere. Regions with newer and less diverse technologies patent more, and both results are statistically significant. Regions with a higher number of assignees per inventor and tie density do not demonstrate significantly different rates of subsequent patenting. The presence of a single, large organization, however, demonstrates significantly increased rates of subsequent patenting. Of the bipartite controls (*Ln LC inventors* and *Ln LC patents*), only the number of patents

Table 3 Conditional Fixed-Effect Negative Binomial Models of Patenting in Year $t+1$, by Metropolitan Statistical Area 1980–2000^{+}

	Model 1	Model 2	Model 3	Model 4	Model 5
Ln regional inventors	0.4708*** (0.0193)	0.4623*** (0.0193)	0.4633*** (0.0193)	0.4717*** (0.0191)	0.4661*** (0.0192)
Ln extra-regional inventors	0.0663*** (0.0015)	0.0661*** (0.0015)	0.0661*** (0.0015)	0.0666*** (0.0015)	0.0663*** (0.0015)
Ln employment	0.0051 (0.0251)	0.0218 (0.0254)	0.0225 (0.0255)	0.0202 (0.0254)	0.0271 (0.0256)
PCPI (000s)	−0.0089*** (0.0015)	−0.0089*** (0.0015)	−0.0093*** (0.0016)	−0.0090*** (0.0015)	−0.0093*** (0.0016)
Ln university patents	−0.0029 (0.0023)	−0.0024 (0.0023)	−0.0024 (0.0023)	−0.0025 (0.0023)	−0.0022 (0.0023)
Technology age	0.2561*** (0.0363)	0.2501*** (0.0363)	0.2493*** (0.0363)	0.2678*** (0.0364)	0.2589*** (0.0366)
Technology Herfindahl	1.3974*** (0.2004)	1.3897*** (0.2004)	1.3981*** (0.2005)	1.3369*** (0.2008)	1.3653*** (0.2010)
Ave assignees/inventor	−0.0994 (0.0748)	−0.0866 (0.0749)	−0.0833 (0.0749)	−0.0998 (0.0744)	−0.0867 (0.0747)
Tie density	1.0436 (0.7074)	0.9565 (0.7076)	1.0173 (0.7077)	0.6905 (0.7126)	0.8329 (0.7134)
Ln LC inventors	−0.0026 (0.0091)	−0.0175* (0.0080)	−0.0124 (0.0096)	−0.0263** (0.0092)	−0.0213* (0.0105)
Ln LC patents	0.0371*** (0.0079)	0.0423*** (0.0061)	0.0374*** (0.0079)	0.0409*** (0.0061)	0.0363*** (0.0079)
Largest org in LC	0.0337* (0.0156)	0.0354* (0.0155)	0.0359* (0.0157)	0.0348* (0.0156)	0.0355* (0.0157)
Clustering in LC	−0.0240 (0.0241)		−0.0217 (0.0243)		−0.0237 (0.0243)
Inverse path length in LC		0.0016*** (0.0004)	0.0016*** (0.0004)		0.0012** (0.0004)
Clustering $*$ Inverse length			0.0027 (0.0030)		0.0024 (0.0030)
Size of LC				0.2191*** (0.0594)	0.1384* (0.0671)
Constant	−0.6305* (0.2719)	−0.7261** (0.2728)	−0.7356** (0.2729)	−0.8030** (0.2756)	−0.8172** (0.2756)
Max likelihood	−24,358.22	−24,350.52	−24,349.80	−24,351.94	−24,347.68
n	6,242	6,242	6,242	6,242	6,242

Notes. Dependent variables in year $t+1$; independent variables calculated from years $t-4$ to t. All models include yearly indicator variables with 1980 omitted.

$^{+}$: $p < 0.1$, *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, standard errors in parentheses.

in the largest component consistently correlates significantly (and positively) with future patenting. The lack of result for employee mobility (as proxied by the assignees per inventor variable) is surprising, as personnel movement across firms should arguably increase creativity. As demonstrated below, however, greater creativity may result from simple connection and aggregation of inventor networks.

Clustering correlates with a consistently negative but insignificant correlation with subsequent patenting. All of the count models demonstrated a positive and significant influence of decreased path length on the next year's patent counts. To correctly form the interaction of clustering divided by path length, we center and enter the inverse of normalized path length; a positive effect of the variable indicates a positive influence of shorter path length. The effect is robust but small: A one standard deviation decrease in path length corresponds to a 0.6% increase in patenting.[9] The small-world interaction never demonstrates even marginal significance. The size of the largest component variable demonstrates a large and positive correlation of 1.4% with subsequent patenting. We infer from these results that clustering and the small-world interaction are of much less consequence to subsequent patenting than decreased path length and the size of the largest component. These results are consistent with career models of patenting at the inventor level, which demonstrate a negative influence of clustering on generative creativity and a very small positive interaction of clustering and external ties (Fleming et al. 2007b).

These results support the first two hypotheses on path length and aggregation: (1) Shorter path length and (2) larger connected components correlate with an increase in subsequent patenting. The small-world effect is not observed in our patent data, and so the third hypothesis, on the small-world interaction of clustering divided by path length, fails to receive empirical support.

Discussion

These models should be interpreted cautiously for several reasons. First, if the influence of clustering on creativity depends strongly on local contingencies, it is difficult to generalize across a network of different clusters, let alone different creative endeavors. Contemporaneous evaluations of Broadway musicals (Uzzi and Spiro 2006) and patent counts in subsequent years (Schilling and Phelps 2007) represent different realms and make it difficult to directly compare results. Second, given that small worlds are connected by definition, we can only regress the largest component's structure against regional patenting. This may be a reasonable approximation in regions with small, connected inventor communities or large, aggregated communities, such as Boston or Silicon Valley because the largest component will comprise a large fraction of all inventors in a region. Still, as illustrated in Figure 5, most inventor communities remain fragmented, and the resultant decrease in the size of the largest component will increase the standard errors of the small-world structure estimates. Methods to develop a weighted average across disconnected components (Schilling and Phelps 2007) provide one alternative, but potentially conflate the influences of small-world structure and simple connection. Third, given that the small-world measure averages clustering and path lengths, the possibility of aggregation bias exists (Robinson 1950). Other biases due to unobserved linkages may also exist: For example, while one missing tie may not greatly bias the clustering measure, it might have a large impact on the path length and size of the largest component. While the regressions covered all regions, our deeper analyses covered only Silicon Valley and Boston. Other regions, such as San Francisco, California; Rochester, New York; and Pittsburgh, Pennsylvania, also underwent dramatic transitions. While San Francisco closely followed Silicon Valley, Rochester, and Pittsburgh split apart. Our understanding of phase transitions would benefit from qualitative study. Finally, given the difficulty of finding a natural experiment on small worlds or instrumental variables at the regional level of analysis, the results should be regarded as primarily descriptive.

These results motivate the exploration of several issues in future empirical research on small worlds. First, given that the regressions' strongest results come from the size of the largest component, future models of small worlds need to include it at least as a control variable. If the simple size measure continues to demonstrate stronger influence than the internal structural measures, it would appear that the current focus on small-world structure and creativity is misplaced. Supporting this possibility, the path length variable lost its significance in a robustness check that used weighted citations instead of patents as a dependent variable (all other results remained essentially unchanged). We interpret this result as indicating that the importance of path length rapidly decreases with longer paths. This is consistent with Singh (2005) who demonstrated that diffusion along collaborative paths drops off rapidly. One theoretical implication of the loss of path length significance is that individual gatekeepers may be of less importance than the redundant connections that begin to form as the size of the largest component increases. For example, Figure 1 illustrates an increasingly redundant core of connections, in which a single connecting gatekeeper becomes redundant. (Although not presented, graphs of bigger largest components illustrate even thicker and more redundant cores.)

Second, future work should proceed simultaneously at the individual and regional level to fully capture the complex dynamics of the process. Indeed, all small-world research might benefit from proceeding along micro and macro fronts simultaneously. Micro analysis would allow modeling of local contingencies and make the macro results less vulnerable to aggregation (Robinson 1950) and unipartite graph (Newman 2001) biases. To illustrate the value of such a combined approach, assume that researchers perform only macro studies. Further assume that clustering has a negative influence on creativity by itself, but that the marginal influence of clustering is positive when the participants bring diverse creative backgrounds to the collaboration. In other words, assume that clustering has not only a first-order negative effect, but also a positive interaction with a measure of the group's creative diversity (Fleming et al. 2007b demonstrate such effects). If the highly clustered groups in regions with high average clustering also have diverse creative backgrounds, and particularly if average diversity in the region remained unmeasured, then clustering might incorrectly appear to have a positive effect overall.

Micro studies would provide additional benefits. Consider the complex question about whether ties across institutions (universities and firms, for example) or industries provide more fertile inventions. Such a mechanism could provide many of the creative benefits of a small world and be masked by macro measures such as the size of the largest component and path length. An appropriate instrument for an individual would provide a natural experiment or control for the inherent endogeneity of such a mechanism (for example, did the inventor create the tie to take advantage of a previously identified opportunity, or did the tie enable the flow of fresh information that triggered more fertile inventions?). Consistent results at multiple levels of analysis would strengthen the inference and the foundation for subsequent theory building.

Third, future work should develop more nuanced and longitudinal measures of creativity. Although clustering may inhibit seminal creativity, it may aid in the diffusion of ideas (Fleming et al. 2007b). This simple observation might resolve the differing results on clustering and creativity. Building on this point, measures need to be carefully defined from a deep understanding of the contextual dynamics. To motivate this concern in the current research and to more fully understand its limitations, consider inventive creativity versus inventive productivity. Small worlds and network aggregation, because they make local inventors more aware of outside work, probably increase competitive pressures, particularly if they cross firm boundaries. While the influence of competition and pressure on creativity remains controversial (Amabile et al. 2002), it is possible that small worlds increase productivity, as measured by patent counts, at the expense of breakthrough creativity. Because most patents represent only incremental and relatively uncreative improvements, better measures of creativity versus productivity are needed to resolve this question.

In addition to the implications for research, the results have implications for policy and management. Figures 4 and 5 reveal large differences and rapid changes in the connectivity of regional collaboration networks. For example, even though Boston and Silicon Valley are similarly dynamic technological regions, our data support Saxenian's assertion that the Valley is much more connected (though the emergence of the giant component comes after the main focus of her study, see Saxenian 1994). Managers already take note of spillovers when they locate in a region (Feinberg and Gupta 2004). An understanding of the local collaborative networks may provide an additional benefit to entrepreneurs who found firms in their home region (Sorenson and Audia 2000). Location in a connected region means quicker access to better qualified personnel, a greater likelihood of university and scientific research, and most important, faster access to technological spillovers. But it also means a decision to live life in the fast lane. Personnel are probably more expensive, can leave more easily, and are probably better connected—even during employment—with externally employed professionals in the region.

We are also struck by the importance—completely unintended by its founders—of the IBM postdoctoral hiring program in creating the Silicon Valley structure. Campbell Scott and William Risk stated that IBM's main motivations for hiring postdoctorals included less expensive labor and new ideas (IBM rarely hired postdoctorals permanently unless their contributions were outstanding). IBM apparently also wished to seed the Valley with friendly professionals. The explanation makes more sense in light of the firm's dominant position, when the program was originally set up. The program had not been intended to provide spillovers, in either direction, across the firm's boundaries. Spillovers obviously occurred between IBM and BioCircuits, but these were of no competitive threat to the larger firm. Scott said that he did learn some biology while visiting at BioCircuits, and that he is applying this learning as IBM moves into life science technologies. In the spirit of unintended innovative benefits for all involved, managers might recreate these positive benefits with institutions, such as the IBM postdoctoral hiring program, while also taking care to avoid competitively harmful spillovers.

Our results should also be of interest to regional policymakers who seek to enhance social welfare. Even though the emergence of small worlds does not appear to greatly improve innovation, decreased path length and component aggregation correlate strongly with subsequent patenting. Simple prescriptions to encourage personnel movement and knowledge spillovers will not be popular with firms, however, and might dampen managerial enthusiasm to locate in the region. Such benefits, unfortunately, appear to come at the expense of private firms. The conflict between regional welfare and firms' protection of their intellectual property strikes us as an important topic for future research.

Conclusion

Despite attracting significant theoretical and more recently, empirical attention, our understanding of the dynamics and impact of small-world networks remains incomplete. Using the patent collaboration histories of over 2 million inventors, we showed that large, technologically dynamic regions undergo dramatic aggregations of isolated inventor components into giant components. Contrary to the growing consensus that small worlds should always improve innovative productivity, we argued that universal predictions remain problematic due mainly to the locally contingent effects of clustering. Count models of network structure and subsequent patenting support these arguments. Decreased path length and increased aggregation increased subsequent patenting. The small-world interaction of clustering divided by path length demonstrated no statistically significant influence on subsequent inventive productivity, once the region's network aggregation was taken into account. While we hesitate to make strong claims based on a null result, it appears that small-world structure is less important than the basic degree of connection within a region.

This research makes several theoretical, methodological, and empirical contributions to our understanding of small-world networks and regional innovation. Theoretical contributions include the separation of the components of the small-world argument and recognition that the structure of a network may have less influence on innovation than the fact that a connected network exists. Methodological contributions include the use of the new relationship measure of patent coauthorship; its characterization through interviews; and the use of large sample databases, visual analysis, and hierarchical clustering algorithms to identify representative inventors for field study. Empirical contributions include the identification and visualization of critical junctures in network evolution, and the first analysis of the effect of small worlds on whole network productivity (as opposed to the effects on firms or groups within a small-world network).

Our results suggest a possible tension between regional policy planners and managers. Ironically, while network aggregation improves regional productivity, it also makes it more likely that firms will suffer unwanted knowledge spillovers. Our discovery that collaboration networks are growing, along with the quantitative and qualitative evidence that spillovers occur quite easily along current and historical collaborative ties, implies that managers must pay increasing attention to the incentives, socialization, and collaborative opportunities of their primary inventors.

Acknowledgments

The authors thank Ivin Baker and Jonathan McPhie for their help in compiling data. The authors gratefully acknowledges support from the Harvard Business School Division of Research. The authors also thank Debbie Strumsky for her compilation of regional economic data, and two anonymous reviewers for their reviews and constructive comments; comments from Santiago Mingo, Melissa Schilling, Jasjit Singh, Olav Sorenson, Mark Szigety, Brian Uzzi, Giovanni Valentini, Woody Powell, Dave Waguespack, and seminar participants at Stanford University, Harvard University, and the Santa Fe Institute benefited our work greatly. Mistakes and errors remain the responsibility of the authors.

Appendix. Matching Algorithm

We extracted source data on all U.S. utility patents granted from 1975 to 2002 (U.S. Patent Office 2003) and MSA data for 2003 (ZIPList5 MSA 2003). Every patent includes the inventor's last name(s) (with varying degrees of first and middle names or initials), home town, detailed information about the technology in class and subclass references (over 100,000 subclasses exist), and the owner or assignee of the patent (if not owned by the inventor, generally a firm, or less often a university). The U.S. Patent Office indexes source data by patent number, not by inventor. We therefore devised an inventor-matching algorithm to determine each inventor's patents and other inventors with whom the focal inventor coauthored at least one patent. The database includes 2,058,823 inventors and 2,862,967 patents.

The matching algorithm refines previous approaches (Newman 2001). If last names match, first initials and middle initials (if given) then must match. Whole first names and whole middle names (if given) then are compared. If all comparisons are positive, the algorithm then requires an additional nonname similarity: hometown (and city and state), corporation (via assignee codes), or technology (via technology subclassifications). We also implemented a common name parameter that ignored the additional match requirement if the last name comprised less than 0.05% of the U.S. population, as determined by the 1995 U.S. census. For 26 randomly selected inventors, the algorithm correctly assigned 215 of 226 patents (as determined by resume searches and personal contact). The 11 incorrectly determined patents were assigned to four isolated nodes (that is, they did not create spurious cut-points). Given the sensitivity of the measures to cut-points, false negatives remain preferable to false positives or incorrectly matching two different inventors.

Endnotes

[1]Even though we study patent collaboration networks, which are based on the codified publication of patent records, much information remains tacit, private, and not communicated. This occurs because written documents can rarely capture all the richness of a technology and because inventors and their firms often withhold information contained in patents for strategic reasons.

[2]We define network aggregation as the process whereby previously isolated components become connected, in the current empirical case, through collaborative mobility. The term is intentionally dissimilar to the word agglomeration used in the economics literature, which refers to economies that firms gain

by clustering together and sharing pooled labor availability, infrastructure, suppliers, and other services.

[3]Alternate window sizes had little effect on illustrated processes (surprisingly, even the cut-point or crucial bridging inventors), substantive trends as illustrated in Figures 3 and 4, and econometric results. The illustrations include inventors both within and outside the MSA and the models control for each of these quantities independently.

[4]The small-world measure lacks units because the normalized measures of clustering and path length lack units.

[5]We chose the period 1986–1990 because it encompasses the initial differentiation that eventually turned into the runaway aggregation of Silicon Valley. Boston, in contrast, demonstrated no appreciable change until three to five years later. The regions also demonstrated remarkably similar trends in several factors, including patents, inventors, isolated inventors, corporate ownership of patents, number of collaborators, tie density, age and diversity of technologies, and number of assignees per inventor (a measure of personnel movement). There remain a few differences: Boston has more university patents and its inventors cite nonpatent references more heavily (Fleming et al. 2007b).

[6]The study focused on the connection dynamics of the six largest components in Silicon Valley and Boston in the 1986–1990 time-frame. The top six were chosen because the first, second, and sixth components aggregated in the Valley, compared with no aggregation of the top six components in Boston. The inventors that linked smaller components into a giant component were identified and interviewed, along with similar "counter-factual" nonlinking inventors. The nonlinking inventors were chosen by minimizing their distance to a linking inventor with a hierarchical clustering algorithm, along a variety of personal and network variables (see Fleming et al. 2007b for details). Despite our best efforts to find inventors with similar social structure, selection bias remains possible, since inventors that left firms and bridged components may differ from those that did not. They may, for example, be less risk averse and perceive the processes of invention and information flow differently.

[7]Froix stressed that this support was not financial; that is, he paid for any incremental expenses he caused during his use of laboratory facilities.

[8]A cut-point is an actor or node in a graph whose removal would separate the graph into two subcomponents (Wasserman and Faust 1994, pp. 112–114). In the case of Campbell Scott, it would separate the IBM and pharmaceutical components of the Silicon Valley graph.

[9]We estimate effect sizes from model 5 unless otherwise noted. For the path length example, effect size from model 5 is 0.0012 and the standard deviation of the measure is 4.82. Hence, $e^{(0.0012*4.82)} = 1.0058$, and $100*(1.0058-1) = 0.6\%$.

References

Acs, Z., L. Anselin, A. Varga. 2002. Patents and innovation counts as measures of regional production of new knowledge. *Res. Policy* **31** 1069–1085.

Ahuja, G. 2000. Collaboration networks, structural holes, and innovation: A longitudinal study. *Admin. Sci. Quart.* **45**(3) 425–457.

Albert, M. B., F. Narin, D. Avery, P. McAllister. 1991. Direct validation of citation counts as indicators of industrially important patents. *Res. Policy* **20** 251–259.

Allen, T. J. 1977. *Managing the Flow of Technology: Technology Transfer and the Dissemination of Technological Information Within the R&D Organization*. MIT Press, Cambridge, MA.

Almeida, P., B. Kogut. 1999. Localization of knowledge and the mobility of engineers in regional networks. *Management Sci.* **45**(7) 905–917.

Amabile, T., C. N. Hadley, S. J. Kramer. 2002. Creativity under the gun. Special issue on the innovative enterprise: Turning ideas into profits. *Harvard Bus. Rev.* **80**(8) 52–61.

Amaral, L., A. Scala, M. Barthelemy, H. Stanley. 2000. Classes of small-world networks. *Proc. National Acad. Sci.* **97**(21) 11149–11152.

Angel, D. 1989. The labor market for engineers in the U.S. semiconductor industry. *Econom. Geography* **65**(2) 99–112.

Balconi, M., S. Breschi, F. Lissoni. 2004. Networks of inventors and the role of academia: An exploration of Italian patent data. *Res. Policy* **33** 127–145.

Barabasi, A., R. Albert. 1999. Emergence of scaling in random networks. *Science* **286** 509–512.

Batagelj, V., A. P. Mrvar. 1998. A program for large network analysis. *Connections 21*, 47–57. Pajek v. 0.91, 2003. http://vlado.fmf.uni-lj.si/pub/networks/pajek/.

Baum, J. A. C., A. V. Shipilov, T. J. Rowley. 2003. Where do small worlds come from? *Indust. Corporate Change* **12** 697–725.

Borgatti, S. P., R. Cross. 2003. A relational view of information seeking and learning in social networks. *Management Sci.* **49**(4) 432–445.

Breschi, S., F. Lissoni. 2004. knowledge networks from patent data: Methodological issues and research targets. Working Paper 150, Centro di Ricerca sui Processi di Innovazione e Internazionalizzazione.

Bureau of Economic Analysis. 2004. Bureau of economic analysis, regional economic accounts local area personal income. http://www.bea.doc.gov/bea/regional/reis/. Deflation data http://www.census.gov/hhes/income/income02/cpiurs.html.

Burt, R. 2004. Structural holes and good ideas. *Amer. J. Sociol.* **110** 349–399.

Cohen, R. 2003. Interview by author. Cambridge, MA (November 12).

Cowan, R., N. Jonard. 2003. The dynamics of collective invention. *J. Econom. Behav. Organ.* **52** 513–532.

Cowan, R., N. Jonard. 2004. Network structure and the diffusion of knowledge. *J. Econom. Dynam. Control* **28** 1557–1575.

Davila, A., G. Foster, et al. (2003). Venture capital financing and the growth of startup firms. *J. Bus. Venturing* **18**(6) 689–708.

Davis, G. F., M. Yoo, W. E. Baker. 2003. The small world of the American corporate elite, 1982–2001. *Strategic Organ.* **1** 301–326.

Dodds, P., R. Muhamad, D. Watts. 2003. An experimental study of search in global social networks. *Science* **301** 827–829.

Feinberg, S. E., A. K. Gupta. 2004. Knowledge spillovers and the assignment of R&D responsibilities to foreign subsidiaries. *Strategic Management J.* **25** 823–845.

Feldman, M. 2001. The entrepreneurial event revisited: Firm formation in a regional context. *Indust. Corporate Change* **10**(4) 861–892.

Fleming, L., M. Marx. 2006. Managing creativity in small worlds. *California Management Rev.* **48**(4) 6–27.

Fleming, L., S. Mingo, D. Chen. 2007a. Collaborative brokerage, geuevative creativity, and creative success. *Admin. Sci. Quart.* **52** 443–475.

Fleming, L., L. Colfer, A. Marin, J. McPhie. 2007b. Why the valley went first: Aggregation and emergence in regional inventor networks. W. Powell, J. Padgett, eds. *Market Emergence and Transformation*. Santa Fe Institute, Santa Fe, NM. Forthcoming.

Florida, R. 2004. America's looming creativity crisis. *Harvard Bus. Rev.* **82**(10) 122–124, 126, 128.

Friedrich, R. 1982. In defense of multiplicative terms in multiple regression equations. *Amer. J. Political Sci.* **26**(4) 797–833.

Froix, M. 2003. Personal interview by author. Palo Alto, CA (July 8).

Gilson, R. J. 1999. The legal infrastructure of high technology industrial districts: Silicon Valley, Route 128, and covenants not to compete. *New York University Law Rev.* **74** 575–629.

Gittleman, M. 2003. Does geography matter for science based firms? Epistemic communities and the geography of patenting and research in biotechnology. Working paper, Stern School of Business.

Gompers, P., J. Lerner. 1996. The use of covenants: An empirical analysis of venture partnership agreements. *J. Law Econom.* **39** 463–498.

Goyal, S., M. Leij, J. Moraga-Gonzalez. 2004. Economics: An emerging small world? Tinbergen Institute Discussion Paper TI2004-001/1.

Granovetter, M. 1973. The strength of weak ties. *Amer. J. Sociol.* **78** 1360–1380.

Granovetter, M. 1985. Economic action and social structure: The problem of embeddedness. *Amer. J. Sociol.* **91**(3) 481–510.

Greene, W. 2002. *Econometric Analysis*. Prentice Hall, New York.

Hansen, M. 1999. The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Admin. Sci. Quart.* **44**(1) 82–112.

Hargadon, A. 2003. *How Breakthroughs Happen: The Surprising Truth About How Companies Innovate*. Harvard Business School Press, Boston, MA.

Hausman, J., B. Hall, Z. Griliches. 1984. Econometric models for count data with an application to the patents R&D relationship. *Econometrica* **52** 909–938.

Hunt, M., T. Ogden, M. Neale. 2003. Who's really sharing? Effects of social and expert status on knowledge exchange within groups. *Management Sci.* **49**(4) 464–477.

Hyde, A. 2003. *Working in Silicon Valley: Economics and Legal Analysis of a High-Velocity Labor Market.* M.E. Sharpe, Armonk, NY.

Jaffe, A. B., M. Trajtenberg, R. Hendeson. 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *Quart. J. Econom.* **434** 578–598.

Katz, A. 1982. The effects of group longevity on project communication and performance. *Admin. Sci. Quart.* **27** 81–104.

Kaufman, C. 2003. Interview by author. Westford, MA (August).

Khanna, P. 2003. Interview by author. Fremont, CA (July 10).

Kino, G. 2003. Interview by author. Stanford, CA (July 8).

Kogut, B., G. Walker. 2001. The small world of Germany and the durability of national networks. *Amer. Sociol. Rev.* **66** 317–335.

Kogut, B., U. Zander. 1992. Knowledge of the firm, combinative capabilities, and the replication of technology. *Organ. Sci.* **3** 383–397.

Koning, P. 2003. Interview by author. Cambridge, MA (August).

Kozlovsky, W. 2003. Interview by author. Stanford, CA (December 11).

Marshall, A. 1919. *Industry and Trade*. MacMillan, London.

McEvily, W., R. Reagans. 2005. Networks and knowledge transfer: The search-transfer problem reconsidered. Working paper, Tepper School of Business.

McEvily, W., A. Zaheer. 2004. Architects of trust: The role of network facilitators in geographical clusters. R. Kramer, K. Cook, eds. *Trust and Distrust in Organizations*. Rusell Sage Foundation, New York.

Milgram, S. 1967. The small world problem. *Psych. Today* **1** 60–67.

Nerkar, A., S. Paruchuri. 2005. Evolution of R&D capabilities: The role or knowledge networks within the firm. *Management Sci.* **51**(5) 771–785.

Newman, M. 2001. The structure of scientific collaboration networks. *Proc. National Acad. Sci. of the United States of America* **98** 404–409.

Newman, M., D. Watts. 1999. Scaling and percolation in the small-world network model. *Physical Rev. Part E* **60** 7332–7342.

Newman, M., D. Watts, S. Strogatz. 2002. Random graph models of social networks. *Proc. National Acad. Sci. of the United States of America* **99** 2566–2572.

Obstfeld, D. 2005. Social networks, the tertius iungens orientation, and involvement in innovation. *Admin. Sci. Quart.* **50** 100–130.

Owen-Smith, J., W. W. Powell. 2004. Knowledge networks as channels and conduits: The effects of spillovers in the Boston biotechnology community. *Organ. Sci.* **15**(1) 5–21.

Perlman, R. 2003. Interview with author. Cambridge, MA (August).

Piore, M., C. Sabel. 1984. *The Second Industrial Divide*: *Possibilities for Prosperity*. Basic Books, New York.

Reagans, R., B. McEvily. 2003. Network structure and knowledge transfer: The effects of cohesion and range. *Admin. Sci. Quart.* **48** 240–267.

Ribi, H. 2003. Interview by author. Cambridge, MA to Burlingame, CA (June 12).

Risk, W. 2003. Interview by author. San Jose, CA (July 9).

Robinson, W. 1950. Ecological correlations and the behavior of individuals. *Amer. Sociol. Rev.* **15**(3) 351–357.

Saxenian, A. 1994. *Regional Advantage*. Harvard University Press, Cambridge, MA.

Schilling, M. A., C. C. Phelps. 2007. Interfirm collaboration networks: The impact of large-scale network structure on firm innovation. *Management Sci.* **53** 1113–1126.

Singh, J. 2005. Collaborative networks as determinants of knowledge diffusion patterns. *Management Sci.* **51** 756–770.

Solow, R. 1957. Technical change and the aggregate production function. *Rev. Econom. Statist.* **39** 312–320.

Sorenson, O., P. Audia. 2000. The social structure of entrepreneurial activity: Geographic concentration of footwear production in the U.S., 1940–1989. *Amer. J. Sociol.* **106** 324–362.

Sorenson, O., T. Stuart. 2001. Syndication networks and spatial distribution of venture capital investments. *Amer. J. Sociol.* **106**(6) 1546–1588.

Sorenson, O., J. Rivkin, Fleming. 2006. Complexity, networks and knowledge flow. *Res. Policy* **35** 994–1017.

Stewart, R. 2004. Interview by author. Cambridge, MA (June 17).

Stuart, T., J. Podolny. 1996. Local search and the evolution of technological capabilities. *Strategic Management J.* (Summer).

Stuart, T., J. Podolny. 1999. Positional causes and correlates of strategic alliances in the semiconductor industry. S. Andrews, D. Knoke, eds. *Research in the Sociology of Organizations*. JAI Press, Greenwich, CT, 161–182.

Stuart, T., O. Sorenson. 2003. Liquidity events and the geographic distribution of entrepreneurial activity. *Admin. Sci. Quart.* **48** 175–201.

Sutton, R., A. Hargadon. 1996. Brainstorming groups in context: Effectiveness in a product design firm. *Admin. Sci. Quart.* **41**(4) 685–718.

Szulanski, G. 1996. Exploring internal stickiness: Impediments to the transfer of best transfer within the firm. *Strategic Management J.* **17** 27–43.

Umatoy, S. 2003. Interview by author. San Jose, CA, July 9.

University of Virginia. 2004. County business patterns online. Geospatial and Statistical Data Center, University of Virginia Library. http://fisher.lib.virginia.edu/collections/stats/cbp/.

U.S. Patent Office. 2003. Patents BIB (January).

Uzzi, B., J. Spiro. 2006. Collaboration and creativity: The small world problem. *Amer. J. Sociol.* **111**(2) 447–504.

Verspagen, B., G. Duysters. 2003. The small worlds of strategic technology alliances. *Technovation* **23**.

Wasserman, S., K. Faust. 1994. *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge, MA.

Watts, D. 1999. Networks, dynamics, and the small-world phenomenon. *Amer. J. Sociol.* **105** 493–527.

Watts, D., S. Strogatz. 1998. Collective dynamics of small-world networks. *Nature* **393** 440–442.

ZIPList5 MSA. 2003. CD Light, LLC. http://www.zipinfo.com/products/z5msa/z5msa.htm.