MANAGEMENT SCIENCE
Vol. 53, No. 7, July 2007, pp. 1113–1126
ISSN 0025-1909 | EISSN 1526-5501 | 07 | 5307 | 1113

DOI 10.1287/mnsc.1060.0624

Interfirm Collaboration Networks: The Impact of Large-Scale Network Structure on Firm Innovation

Melissa A. Schilling
Stern School of Business, New York University, 40 West Fourth Street, New York, New York 10012, mschilli@stern.nyu.edu

Corey C. Phelps
Department of Management and Organization, University of Washington, Box 353200, Seattle, Washington 98195, cphelps@u.washington.edu

The structure of alliance networks influences their potential for knowledge creation. Dense local clustering provides information transmission capacity in the network by fostering communication and cooperation. Nonredundant connections contract the distance between firms and give the network greater reach by tapping a wider range of knowledge resources. We propose that firms embedded in alliance networks that exhibit both high clustering and high reach (short average path lengths to a wide range of firms) will have greater innovative output than firms in networks that do not exhibit these characteristics. We find support for this proposition in a longitudinal study of the patent performance of 1,106 firms in 11 industry-level alliance networks.

Key words: alliances; networks; innovation; patents
History: Accepted by Brian Uzzi and Luis Amaral, special issue editors; received June 16, 2004. This paper was with the authors 5 months for 2 revisions.

Introduction

Although research has long recognized the importance of interfirm networks in firm innovation (see Freeman 1991 for a review), much of this work has treated the network concept as a metaphor. Only recently have researchers begun to assess the formal structural properties of alliance networks and their impact on firm innovation. This research has focused on a firm's position within a broader network of relationships or the structure of its immediate network neighborhood rather than the structure of the overall network. Studies have examined a firm's centrality (Smith-Doerr et al. 1999), number of alliances (Shan et al. 1994), and local network structure (Ahuja 2000, Baum et al. 2000). To our knowledge, empirical research has not yet examined the impact of the structure of industry-level[1] alliance networks on member firm innovation. In a related study, however, Uzzi and Spiro (2005) examined the network structure of the creative artists who made broadway musicals from 1945 to 1989, and concluded that the large-scale structure of the artists' collaboration network significantly influenced their creativity, and the financial and artistic performance of their musicals. This raises the following questions: Does the structure of an industry-level interfirm network influence the rate of knowledge creation among firms in the network? If so, what structural properties will enhance firm innovation?

To address these questions, we examine the impact of two key large-scale network properties, clustering and reach, on the innovative output of members of the network. The dense connectivity of clusters creates *transmission capacity* in a network (Burt 2001), enabling large amounts of information to rapidly diffuse, while reach (i.e., short path lengths to a wide range of firms) ensures that diverse information sources can be tapped. We argue that networks with both high clustering and high reach will significantly enhance the creative output of member firms. We test this hypothesis using longitudinal data on the innovative performance of a large panel of firms operating in 11 industry-level alliance networks.

This research offers several important contributions for understanding knowledge creation in interfirm networks. First, we find empirical support for our argument that the combination of clustering and reach increases member firm innovation. To our knowledge, no other study has attempted to assess the effect of industry-level interfirm networks on the innovation performance of member firms. Although recent studies have examined the structure of large-scale interfirm networks and the possible causes of these structures (Baum et al. 2003, Kogut and Walker

[1] An industry-level network is a specific type of whole or "large-scale" network. Wellman (1988, p. 26) defined a whole network as the relationships that exist among members of a population.

2001), little research has examined the *consequences* of large-scale network structure in an industrial setting (Uzzi and Spiro 2005 is a recent exception). Second, while most studies of network structure have examined a single industry, our study uses longitudinal data on 11 industries, which strengthens the generalizability of our findings.

We begin by describing two key structural characteristics of interfirm networks and their effect on information diffusion in the network. From this we develop a hypothesis about how the structure of interfirm networks will influence the innovative output of member firms. We test the hypothesis on a large, unbalanced panel of firms embedded in 11 industry-level alliance networks.

Large-Scale Interfirm Networks and Firm Knowledge Creation

We adopt a recombinatory search perspective in explaining the process of innovation (Fleming 2001). Innovation is characterized as a problem-solving process in which solutions to problems are discovered via search (Dosi 1988). Prior research suggests that search processes that lead to the creation of new knowledge, embodied in artifacts such as patents and new products, most often involve the novel recombination of known elements of knowledge, problems, or solutions (Fleming 2001, Nelson and Winter 1982) or the reconfiguration of the ways in which knowledge elements are linked (Henderson and Clark 1990). Critical inputs into this process include access to and familiarity with a variety of knowledge elements (e.g., different technological components and the scientific and engineering know-how embedded in them), novel problems and insights into their resolution, failed recombination efforts, and successful solutions (Hargadon and Fanelli 2002). Firms that have greater access to and understanding of these recombinatory resources should be advantaged in their innovation efforts.

As firms form and maintain alliances with each other, they weave a network of direct and indirect relationships. As a result, firms embedded in these networks gain access to information and know-how of direct partners and that of others in the network to which they are indirectly connected (Ahuja 2000, Gulati and Gargiulo 1999). The network of alliance relationships constitutes a conduit that channels the flow of information and know-how among firms in the network (Ahuja 2000, Owen-Smith and Powell 2004), with each member firm acting as both a recipient and transmitter of information (Ahuja 2000). The structure of these networks greatly influences the dynamics of information diffusion within the networks. Large-sample studies have found that *direct* alliance relationships facilitate knowledge flows between partners (Gomes-Casseres et al. 2006, Mowery et al. 1996) and enhance the innovative performance of firms (e.g., Deeds and Hill 1996, Stuart 2000). Research also shows that the extent to which a firm is *indirectly* connected to other firms in an alliance network enhances its innovativeness (Ahuja 2000, Owen-Smith and Powell 2004, Soh 2003).

Given the role of direct and indirect ties as channels for the flow of information and know-how, we argue that the structure of the interfirm network will significantly influence the recombination process. Two structural characteristics that have a particularly important role in diffusion are *clustering* and *reach*.

Clustering

Alliance networks tend to be highly clustered: Some groups of firms will have more links connecting them to each other than to the other firms in the network. A firm's clustering coefficient can be calculated as the proportion of its partners that are themselves directly linked to each other. The clustering coefficient of the overall network is the average of this measure across all firms in the network. Several mechanisms lead to clustering in interfirm knowledge networks, but two of the most common are linking based on similarity or complementarity. Firms tend to interact more intensely or frequently with other firms with which they share some type of proximity or similarity, such as geography or technology (Baum et al. 2003, Rosenkopf and Almeida 2003). This tends to result in a high degree of clustering.

Clustering increases the information transmission capacity of a network. First, the dense connectivity of individual clusters ensures that information introduced into a cluster will quickly reach other firms in the cluster. The multiple pathways between firms also enhance the fidelity of the information received. Firms can compare the information received from multiple partners, helping them to identify ways in which it has been distorted or is incomplete. Second, clusters within networks are important structures for making information exchange meaningful and useful. The internal density of a cluster can increase the dissemination of alternative interpretations of problems and their potential solutions, deepening the collective's understanding and stimulating collective problem solving (Powell and Smith-Doerr 1994). The development of a shared understanding of problems and solutions greatly facilitates communication and further learning (Brown and Duguid 1991, Powell et al. 1996). Third, dense clustering can make firms more willing and able to exchange information (Ahuja 2000). Sociologists (e.g., Coleman 1988, Granovetter 1992) have suggested that densely clustered networks give rise to trust, reciprocity norms, and a shared identity, all of which lead to a high level of cooperation and can facilitate collaboration by providing self-enforcing informal governance mechanisms (Dyer and Singh 1998). In addition to stimulating greater transparency, trust and reciprocity exchanges facilitate intense interaction among personnel from partnered firms (Uzzi 1997), improving the transfer of tacit, embedded knowledge (Hansen 1999, Zander and Kogut 1995). Thus, clustering enables richer and greater amounts of information and knowledge to be exchanged and integrated more readily.

When dense clusters are sparsely connected to each other, they become important structures for creating and preserving the requisite variety of knowledge in the broader network that enables knowledge creation. The internal cohesion of a cluster can cause much of the information and knowledge shared within a cluster to become homogeneous and redundant (Burt 1992, Granovetter 1973). The dense links provide many redundant paths to the same actors, and thus the same sources of information and knowledge. Cohesion can also lead to norms of adhering to established standards and conventions, which can potentially stifle experimentation and creativity (Uzzi and Spiro 2005). This limits innovation. Clusters of firms will, however, tend to be heterogeneous across a network in the knowledge they possess and produce due to the different initial conditions and causes for each cluster to form. The diversity of knowledge distributed across clusters in the network provides the requisite variety for recombination.

Clustering thus offers both local and global advantages. Firms benefit from having redundant connectivity among their immediate neighbors because it enhances the speed and likelihood of information access, and the depth of information interpretation. Firms also benefit from being embedded within a larger network that is clustered because the information a firm receives from partners that are embedded in other clusters is likely to be more complete and richly understood than information received from partners not embedded in clusters, and because information received from different clusters is likely to be diverse, enabling a wider range of recombinatorial possibilities.

Reach

The size of a network and its average path length (i.e., the average number of links that separates each pair of firms in the network) also impacts information diffusion and novel recombination. The more firms that can be reached by any path from a given firm, the more knowledge that firm can potentially access. However, the likelihood, speed, and integrity of knowledge transfer between two firms are directly related to the path length separating those two firms. The diffusion of information and knowledge occurs more rapidly and with more integrity in networks with short average path lengths than in networks with longer paths (Watts 1999). A firm that is connected to a large number of firms by a short average path can reach more information, and can do so quickly and with less risk of information distortion than a firm that is connected to fewer firms or by longer paths. To capture this we use *distance-weighted reach*.

A firm's distance-weighted reach is the sum of the reciprocal distances to every firm that is reachable from a given firm, i.e., $\sum_j 1/d_{ij}$, where d_{ij} is defined as the minimum distance (geodesic), d, from a focal firm i to partner j, where $i \neq j$. A network's *average distance-weighted reach* is this measure averaged across all firms in the network, $(\sum_n \sum_j 1/d_{ij})/n$, where n is the number of firms in the network. Other things being equal, a very large connected network with a very short average path length (e.g., a completely connected network where there are many firms and every firm is directly connected to every other firm, or a star graph with many firms all connected to the same central "hub" firm) will have the greatest average distance-weighted reach. Longer path lengths, smaller network size, or disconnects that fragment the network into multiple components all decrease average distance-weighted reach.

The preceding reveals one of the key benefits of using distance-weighted reach: It provides a meaningful measure of the overall size and connectivity of a network, even when that network has multiple components, and/or component structure is changing over time. It avoids the infinite path length problem typically associated with disconnected networks by measuring only the path length between connected pairs of nodes, and it provides a more meaningful measure than the simple average path length between connected pairs by factoring in the size of connected components.[2]

Because forming alliances is costly and constrained, there appears to be a trade-off between forming dense clusters to facilitate rapid exchange and integration of knowledge, versus forging links to create short paths to a wider range of firms. However, recent research has shown that even sparse, highly clustered networks can have high reach if there are a few links creating bridges between clusters (Watts 1999, Hansen 2002, Hargadon 1998). Bridges between clusters of firms provide member firms access to diverse information that exists beyond their local cluster, enabling

[2] The authors are grateful to Steve Borgatti for pointing this out. They are also grateful to Mark Newman for numerous discussions about how to handle the infinite path length consideration in our networks.

new combinations with their existing knowledge sets, while preserving the information transmission advantages of clusters. As Uzzi and Spiro (2005) note, bridges between clusters increase the likelihood that different ideas and routines will come into contact, enabling recombinations that incorporate both previous conventions and novel approaches. The combination of clustering and reach thus enables a wide range of information to be exchanged and integrated rapidly, leading to greater knowledge creation. In sum, we predict a multiplicative interaction between clustering and reach in their effect on firm knowledge creation. Consistent with the symmetrical nature of such interactions (Jaccard and Turrisi 2003), we have argued and expect that the effect of clustering on firm knowledge creation will be increasingly positive as reach increases, while the effect of reach on knowledge creation will be increasingly positive as clustering increases.

Hypothesis. *Firms participating in alliance networks that combine a high degree of clustering and reach will exhibit more knowledge creation than firms in networks that do not exhibit these characteristics.*

Methods

To test our hypothesis, we constructed a large, unbalanced panel of U.S. firms for the period 1990–2000. The panel includes all U.S. firms that were part of the alliance networks of 11 high-technology manufacturing industries: aerospace equipment (standard industrial classifications (SICs)): 3721, 3724, 3728, 3761, 3764, 3769; automotive bodies and parts (3711, 3713, 3714); chemicals (281-, 282-, 285-, 286-, 287-, 288-, 289-); computer and office equipment (3571, 3572, 3575, 3577); household audiovisual equipment (3651); medical equipment (3841, 3842, 3843, 3844, 3845); petroleum refining and products (2911, 2951, 2952, 2992, 2999); pharmaceuticals (2833, 2834, 2835, 2836); semiconductors (3674); telecommunications equipment (366-), and measuring and controlling devices (382-).

The choice of industries was particularly important for this study. The 11 industries selected have been designated as high technology in numerous Bureau of Labor Statistics studies (e.g., Hecker 1999).[3] These industries provide an excellent context for our study for three reasons. First, knowledge creation is fundamental to the pursuit of competitive advantage in high-technology industries. Second, firms in these industries actively use alliances in pursuit of their innovation activities (Vonortas 1997). Third, because we use patent data for our dependent variable, it is important to select industries that use patents. There is evidence that firms in these industries actively patent their inventions (Levin et al. 1987).

Alliance Networks

We chose to measure the network structure created by publicly reported strategic alliances for two reasons. First, there is a rich history of research on the importance of strategic alliances as a mechanism for knowledge sharing among firms (Freeman 1991, Gulati 1998, Powell et al. 1996). Second, alliances are used by a wide range of firms (both public and private) in a wide range of industries, and are often used explicitly for the exchange and joint creation of knowledge.

Social network research has identified three procedural tactics for establishing network boundaries for empirical research: attributes of actors that rely on membership criteria, such as membership in an industry; types of relations between actors, such as participation in strategic alliances; and participation in a set of common events (Laumann et al. 1983). Accordingly, we employed two rules to guide our construction of the 11 industry networks used in this study. First, each alliance included at least one firm that was a member of the target industry (indicated by its primary four-digit SIC). Second, each alliance had to operate in the target industry, as indicated by its primary four-digit SIC of activity.

Alliance data were gathered using Thomson Corp.'s SDC Platinum database. The SDC data have been used in a number of empirical studies on strategic alliances (e.g., Anand and Khanna 2000, Sampson 2004). For each industry, alliances were collected that were announced between 1990 and 1997. We chose 1990 as the initial year for our sample because information on alliances formed prior to 1990 is very sparse in the SDC database (Anand and Khanna 2000, p. 300). Separate alliance networks were created for each industry according to the alliance's primary SIC. Both public and private firms were included. We use data on only U.S. firms because the SDC alliance data are much more complete for U.S. firms than for non-U.S. firms (Phelps 2003). All alliances were aggregated to the parent corporation.

The resulting data set includes 1,106 firms involved in 3,663 alliances. Many of the alliances included more than two participating firms, so the number of dyads is greater, totaling 5,306. Because any type of alliance may provide a path for knowledge diffusion, and because prior studies indicate that the breadth of an alliance's true activity is often much greater than what is formally reported (Powell et al. 1996), we include all alliance types in our analysis. We do, however,

[3] We omitted high-tech manufacturing industries that rarely use alliances: special industry machinery (355), electrical industrial apparatus (362), search and navigation equipment (381), and photographic equipment and supplies (386).

control for the proportion of alliances in each network formed for the explicit purpose of technology exchange or development.

Alliances typically last for more than one year, but alliance termination dates are rarely reported. This required us to make an assumption about alliance duration. We took a conservative approach and assumed that alliance relationships last for three years, consistent with recent empirical work on average alliance duration (Phelps 2003). Other research has taken a similar approach, using windows ranging from one to five years (e.g., Gulati and Gargiulo 1999, Stuart 2000). We created alliance networks based on three-year windows (i.e., 1990–1992, 1991–1993, ... 1995–1997), resulting in six snapshots of network structure for each industry, for a total of 66 alliance network snapshots. Each network snapshot was constructed as an undirected binary adjacency matrix (Wasserman and Faust 1994).[4] Multiple alliances between the same pair of firms in a time window were treated as one link. UCINET 6 was used to obtain measures on these networks, as described below (Borgatti et al. 2002).

As we focus on publicly reported contractual alliance agreements, we do not observe the numerous informal collaborative arrangements that exist between firms in our sample. Such informal arrangements often lead to the types of formal agreements that we observe (Powell et al. 1996, Rosenkopf et al. 2001). Thus, our analysis represents a conservative test of our diffusion argument because our data do not include informal relationships that promote knowledge transfer.

Dependent Variable: Patents

One way that knowledge creation is instantiated is in the form of inventions (Schmookler 1966). Knowledge embedded in artifacts such as inventions represents the "empirical knowledge" of organizations (Hargadon and Fanelli 2002). Inventions thus provide a trace of an organization's knowledge creation. Patents provide a measure of novel invention that is externally validated through the patent examination process (Griliches 1990). Patent counts have been shown to correlate well with new product introductions and invention counts (Basberg 1987). Trajtenberg (1987) concluded that patents are valid and robust indicators of knowledge creation. One of the challenges with using patents to measure innovation is that the propensity to patent may vary with industry, resulting in a potential source of bias (Levin et al. 1987). We addressed this potential bias in three ways. First, we sample only high-tech manufacturing industries, which helps to ensure a degree of commonality in the industries' emphasis on innovation. To further capture differences in emphasis on innovation, we control for industry-level R&D intensity. Third, to control for unobserved factors that influence the propensity to patent that are likely to be stable within industries, we control for industry fixed effects. The propensity to patent may also differ due to firm characteristics (Griliches 1990). We attempt to control for such sources of heterogeneity using covariate, *Presample Patents* (described below), and firm fixed and random effects in our estimations.

We measure the dependent variable, $Patents_{it}$, as the number of successful patent applications for firm i in year t. We used the Delphion database to collect yearly patent counts for each of the firms, aggregating subsidiary patents up to the ultimate parent level. Granted patents were counted in their year of application. Yearly patent counts were created for each firm for the period of 1993 to 2000, enabling us to assess different lag specifications between alliance network structure and patent output.

Independent Variables

Clustering Coefficient. To measure the clustering in each network for each time period, we used the weighted overall clustering coefficient measure (Borgatti et al. 2002, Newman et al. 2002):

$$Clustering_w = \frac{3 \times (\text{number of triangles in the graph})}{(\text{number of connected triples})},$$

where a *triangle* is a set of three nodes (e.g., i, j, k), each of which is connected to both of the others, and a *connected triple* is a set of three nodes in which at least one is connected to both the others (e.g., i is connected to j and k, but j and k need not be connected). This measure indicates the proportion of triples for which transitivity holds (i.e., if i is connected to j and k, then by transitivity, j and k are connected). The factor of three in the numerator ensures that the measure lies strictly in the range of zero and one because each triangle implies three connected triples. The weighted overall clustering coefficient represents the percentage of a firm's alliance partners that are also partnered with each other, weighted by the number of each

[4] Each matrix reflects the alliances maintained within the network as of the end of the focal year. Because alliances often endure longer than one year, constructing adjacency matrices using only alliances announced in the focal year would bias the connectivity of the observed networks downward. Consider the initial year of the panel for the network variables (1992): Using only alliances formed in 1992 would not capture the alliance relationships formed prior to, yet maintained through, 1992. Data on both presample alliance formation and alliance duration is needed to accurately assess network structure in each of the sample years. Moving three-year windows more accurately reflects the structure of an alliance network in the annual adjacency matrices. Robinson and Stuart (2007) use a similar approach in assessing alliance networks in the biotechnology industry.

firm's partners, averaged across all firms in the network. This variable can range from zero to one, with larger values indicating increasing clustering. While network density captures the density of the entire network, the clustering coefficient captures the degree to which the overall network contains localized pockets of dense connectivity. A network can be globally sparse and still have a high clustering coefficient.

Reach. To capture the reach of each network for each time period, we use a measure of average distance-weighted reach (Borgatti et al. 2002). This is a compound measure that takes into account both the number of firms that can be reached by any path from a given firm, and the path length it takes to reach them. This measure is calculated as

$$\textit{Average distance weighted reach} = \left[\sum_{n}\sum_{j} 1/d_{ij}\right] \Big/ n,$$

where n is the number of nodes in the network, and d_{ij} is defined as the minimum distance (geodesic), d, from a focal node i to partner j, where $i \neq j$. Average distance-weighted reach can range from 0–n, with larger values indicating higher reach.

Clustering × Reach. We predict that the combination of clustering and reach will have a positive impact on member firm innovation, and thus include the interaction term, *Clustering × Reach*.

Firm-Level Control Variables

Presample Patents. To control for unobserved heterogeneity in firm patenting, we follow the presample information approach of Blundell et al. (1995) and calculate the variable *Presample Patents* as the sum of patents obtained by a firm in the five years prior to its entry into the sample.

Betweenness Centrality. We control for the possibility that firms that occupy more central positions in alliance networks may generate more innovations than more peripheral firms (e.g., Owen-Smith and Powell 2004, Soh 2003). We operationalize *Centrality* using Freeman's (1979) measure of "betweenness centrality," which captures the extent to which a firm is located on the shortest path (i.e., geodesic) between any two actors in its alliance network. Formally, betweenness centrality for firm i in year t is calculated as

$$\textit{Betweenness Centrality}_{it} = \sum_{j<k} g_{jk}(n_i)/g_{jk},$$

where $g_{jk}(n_i)$ refers to the number (n) of geodesics (i.e., shortest paths) linking firms j and k that contain focal firm i. The term $g_{jk}(n_i)/g_{jk}$ captures the probability that firm i is involved in the shortest path between j and k. Betweenness centrality is the sum of these estimated probabilities over all pairs of firms (excluding the ith firm) in the network. We use normalized betweenness centrality (i.e., betweenness divided by maximum possible betweenness, expressed as a percentage) to make the measure comparable across time and industry networks.

Local Efficiency. While studies have found that the extent to which a firm's partners are nonredundant enhances its knowledge creation (Baum et al. 2000), other research shows that redundant links improve innovation (Ahuja 2000). Although the empirical evidence is mixed, controlling for the effect of local structural holes is important if we wish to demonstrate that the global structure of the alliance network in which a firm is embedded has an independent and significant influence on its subsequent patenting. We control for the influence of a firm's local network structure using Burt's (1992) measure of efficiency. Efficiency captures the extent to which a firm's partners are nonredundant, indicating the presence of structural holes. Local efficiency for firm i in year t is computed as

$$\textit{Local Efficiency}_{it} = \left\lfloor \sum_{j} \left\lfloor 1 - \sum_{q} p_{iq} m_{iq} \right\rfloor \right\rfloor \Big/ N_i, \quad j \neq q,$$

where p_{iq} is the proportion of i's relations invested in the relationship with q, m_{jq} represents the marginal strength of the relationship between alter j and alter q (as we use binary data, values of m_{jq} are set to one if the relationship is present and zero otherwise), and N_i represents the number of unique alliance partners connected to firm i. This measure can range from zero to one, with higher values indicating greater efficiency.

Industry (Network) Control Variables

Network Density. We control for the overall density of the network with the variable *Network Density*, calculated for each industry network and time period. We do so because the rate and extent to which information diffuses increases with density (Yamaguchi 1994). This variable measures the ratio of existing links in the network to the number of possible pairwise combinations of firms, and may range from zero to one, with larger values indicating increasing density.

Centralization. The extent to which a network is centralized can also influence its diffusion properties. A highly centralized network is one in which all ties run through one or a few nodes, thus decreasing the distance between any pair of nodes (Wasserman and Faust 1994). To control for network centralization, we employ Freeman's (1979) index of group betweenness centralization, calculated for each industry network and time period. Group betweenness centralization for network j in year t is

$$\textit{Betweenness Centralization}_{jt} = 100 \times \left\{ \sum_{i=1}^{g} [C'_B(n^*) - C'_B(n_i)]/(g-1) \right\},$$

where $C'_B(n^*)$ is the largest realized normalized betweenness centrality for the set of firms in network j in year t, $C'_B(n_i)$ is the normalized betweenness centrality for firm i (in industry network j for year t), and g is the number of firms. This variable is expressed as a percentage and can range from zero, where all firms have the same individual betweenness centrality, to 100, where one firm connects all other firms.

Industry R&D Intensity. To control for differences in the emphasis on and costliness of innovation across industries, we employ a time-varying measure of industry-level R&D intensity (R&D expenditures/sales). We collected annual R&D expenditures and sales of firms in each industry from Compustat. We would have preferred to control for R&D intensity at the firm level; however, nearly 42% of our sample firms were privately owned during some portion of the sample, and R&D expenditures are not available for private firms. In investigating the robustness of our results, we utilize a control variable (stock of patents obtained in the past four years) that has been shown to be highly correlated with annual firm-level R&D expenses. Our results are unchanged when including this variable in our models.

Proportion of Alliances for R&D, Cross-Technology Transfer, or Licensing. Alliances that are established for the purpose of technology exchange or development may be more directly related to firm patenting. To examine this possibility, we include a time-varying measure of the percentage of alliance agreements in each network that were established explicitly for the purpose of joint research and development, cross-technology transfer, or technology licensing.

Model Specification

The dependent variable in this study, *Patents*, is a count variable and takes on only nonnegative integer values. The linear regression model is inadequate for modeling such variables because the distribution of residuals will be heteroscedastic nonnormal. A Poisson regression approach is appropriate to model count data (Hausman et al. 1984). However, the Poisson distribution contains the strong assumption that the mean and variance are equal. Patent data often exhibit overdispersion, where the variance exceeds the mean (Hausman et al. 1984). In the presence of overdispersion, coefficients will be estimated consistently, but their standard errors will generally be underestimated, leading to spuriously high levels of significance (Cameron and Trivedi 1986). Each model that we report, when estimated using the Poisson specification, exhibited significant overdispersion.

A commonly used alternative to the Poisson regression model is the negative binomial model. The negative binomial model is a generalization of the Poisson model and allows for overdispersion by incorporating an individual, unobserved effect into the conditional mean (Hausman et al. 1984). The panel data implementation of the negative binomial model accommodates explicit control of persistent individual unobserved effects through both fixed and random effects. In the present study, unobserved heterogeneity refers to the possibility that unmeasured (or unmeasurable) differences among observationally equivalent firms affects their patenting. Unobserved heterogeneity may also stem from unmeasured, systematic time period and industry effects. Failing to control for such unobserved heterogeneity, if present, can result in specification error (Heckman 1979).

We employ a number of strategies to control for these sources of unobserved heterogeneity. First, we include year fixed effects to control for systematic period effects such as differences in macroeconomic conditions that may affect all sampled firms' patent rates. Second, we employ individual firm effects to control for unobserved, temporally stable firm differences in patenting. We use both firm fixed and random effects in alternative estimations of our model. We use the Hausman et al. (1984) implementation of fixed effects in the context of a negative binomial model, which employs a conditional maximum-likelihood estimation procedure.[5]

[5] Allison and Waterman (2002) recently criticized the Hausman et al. (1984) conditional negative binomial fixed-effects model as not being a "true" fixed-effects method in that it does not control for all time-invariant covariates. Allison and Waterman (2002) developed an unconditional negative binomial model that uses dummy variables to represent fixed effects, which effectively controls for all stable individual effects. This procedure has been implemented in Limdep 8.0. However, estimates of β are inconsistent in negative binomial models when using such a dummy variable approach in short panels, due to the incidental parameters problem (Cameron and Trivedi 1998, p. 282). The number of unit-specific (e.g., firm) parameters (α_i) increases with the sample size, while the number of periods (T) stays fixed, resulting in a limited number of observations to estimate a large number of parameters. Contrary to linear regression models, the maximum-likelihood estimates for α_i and β are not independent for negative binomial models because the inconsistency of the estimates of α_i are transmitted into the maximum likelihood estimate of β. Given that this method is a true fixed-effects specification, it does not allow for time-invariant covariates such as *Presample Patents*. Thus, we chose not to employ Allison and Waterman's (2002) unconditional estimator. We report the results using the Hausman et al. (1984) conditional fixed-effects

We also use Hausman et al.'s random effects specification, which assumes that overdispersion due to unobserved heterogeneity is randomly distributed across firms. Because the random effects specification assumes that the unobserved firm-specific effect is uncorrelated with the regressors, we report the results from both fixed and random effects as a robustness check.

As an additional control for firm-level unobserved heterogeneity, we adopt the presample information approach of Blundell et al. (1995). Blundell et al. (1995) argued that because the main source of unobserved heterogeneity in models of innovation lies in the different knowledge stocks with which firms enter a sample, a variable that approximates the build-up of firm knowledge at the time of entering the sample is a particularly good control for unobserved heterogeneity. The *Presample Patents* variable described above serves as a control for unobserved differences in firm knowledge stocks upon their entry into the sample. Blundell et al. (1995) showed that the use of a presample patent entry stock measure virtually eliminated persistent serial correlation in their panel data models. We also include industry dummies in our models to control for unobserved industry effects that are not captured by the firm effects.

A final estimation issue concerns the appropriate lag structure of the independent variables. Based on prior research that investigates the relationship between interfirm alliances and innovation (e.g., Ahuja 2000, Sampson 2004, Stuart 2000), we employ alternative lags of our independent variables relative to our dependent variable. We estimate models using one-year, two-year, and three-year lags. We do so to explore the robustness of our findings across alternative specifications. All models were estimated with Limdep 8.0. The model we estimate takes the general form provided below (aerospace is the omitted industry and 1992 is the omitted year). Variables are indexed across firms (i), industry (j), and time (t):

$$\begin{aligned}&Patents_{it+1(2,3)}\\&\quad = f(Clustering_{jt}, Reach_{jt}, Clustering * Reach_{jt},\\&\qquad R\&D\ Alliance\ \%_{jt}, R\&D\ Intensity_{jt}, Centrality_{it},\\&\qquad Local\ Efficiency_{it}, Centralization_{jt}, Density_{jt},\\&\qquad Presample\_Patents_{it}, Automotive, Chemicals,\\&\qquad Computers, Audiovisual, Medical, Petroleum,\\&\qquad Pharmaceuticals, Semiconductors,\\&\qquad Telecommunications, Measuring, 1993, 1994, 1995,\\&\qquad 1996, 1997).\end{aligned}$$

approach. We point out that the results we obtained from both fixed- and random effects specifications are highly consistent (see the Results section). Studies that have employed both the Hausman et al. (1984) negative binomial fixed-effects approach and that of Allison and Waterman (2002) have found very similar results (e.g., Dee et al. 2005).

Results

A summary of the network size and component structure for each industry, averaged over time, is provided in Table 1. As shown, there is substantial variation across industries in the number of firms that participate in alliances. This is largely due to differences in industry size. The average number of alliances per firm within each industry exhibits much less variation. The next column provides the average number of firms in each network. This number includes firms from the industry and their partners, some of which are not in the target industry. The next column indicates the percentage of nodes in the network that are connected to the single largest ("main") component. This number varies significantly both across industry and over time (not shown). While researchers often study only the main component, in our study this would have yielded misleading results. Whereas in some industries there is a large main component that is relatively stable over time (e.g., pharmaceuticals), in other industries there are multiple large components, and those components merge and split apart over time. For example, between 1996 and 1997 in the computer industry, a large component broke away from the main component (see Figure 1). If we had focused only on the single largest component, we would have both understated the amount of

Table 1 Network Size and Component Structure, Averages over 1992–2000

Industry	Average number of firms from industry in alliances[a]	Average number of alliances per firm	Average network size (nodes)[b]	Percent in main component (%)
Aerospace	9	3.05	28	46
Automotive	15.67	3.43	53.2	37
Chemicals	45.17	2.97	199.8	11
Computers and office equipment	79.67	4.48	347	45
Household audiovisual equipment	9	1.5	28.3	10
Measuring and controlling	22.67	1.96	48.33	21
Medical equipment	66.17	1.66	172.33	7
Petroleum refining and products	5.3	2.65	24.83	18
Pharmaceuticals	218.33	2.54	510	64
Semiconductors	58.67	3.51	204	55
Telecommunication equipment	44.83	6.53	266.33	54

[a]This number includes only those firms with the designated primary SICs; it does not include partners in the network that are not in those SICs.
[b]Includes all U.S. firms in the network, including both those with the designated primary SICs and their alters, regardless of SIC.

Figure 1 Network Size and Component Structure (Common Shade of Gray Indicates Firms in Same Component)

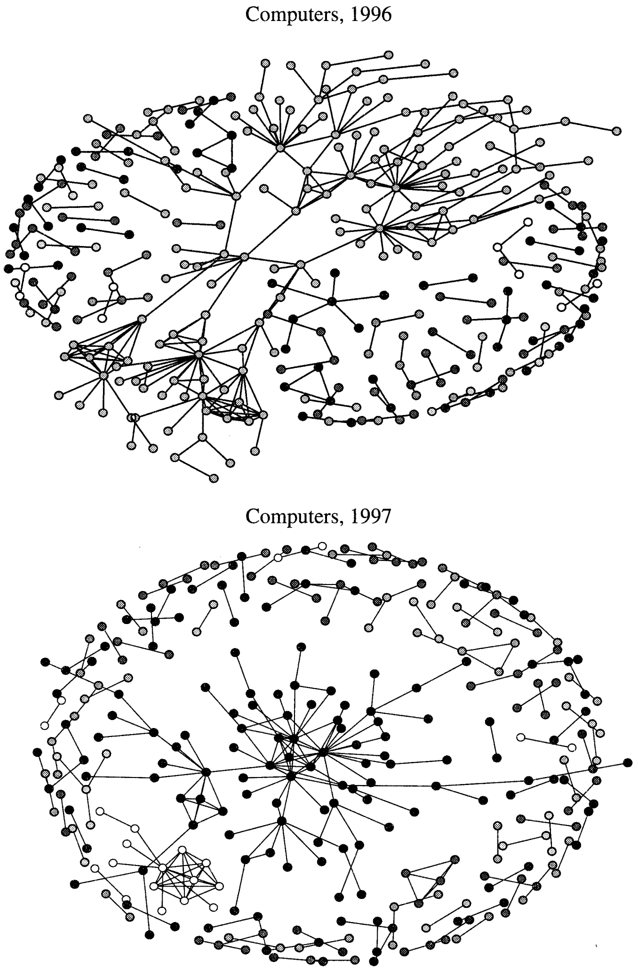

alliance activity in the industries, and overstated the amount of change in alliance activity over time.

Table 2 reports the negative binomial panel regression results for the three dependent variables ($Patents_{it+1}$; $Patents_{it+2}$; $Patents_{it+3}$). Because the random effects specification assumes that regressors and firm-specific effects are uncorrelated, we also provide results using firm fixed effects as a robustness check. Separate results are provided for three dependent variables. Models 1, 2, and 3 report the results using a one-year lag between the independent variables and firm patenting ($Patents_{it+1}$). Models 4, 5, and 6 report the results using a two-year lag ($Patents_{it+2}$), and Models 7, 8, and 9 report the results using a three-year lag ($Patents_{it+3}$). For each dependent variable, the first models (1, 4, and 7) include the control variables only, the second models add the direct effects of *Clustering* and *Reach* (Models 2, 5, and 8), and the third model adds the interaction term, *Clustering* × *Reach* (Models 3, 6, and 9). To conserve space, firm, industry, and time period effects, while estimated, are not reported.

Our sole hypothesis predicted a positive effect of the interaction of *Clustering* and *Reach* on firm patenting. The interaction term, *Clustering* × *Reach*, does not obtain statistical significance at conventional levels in the model specified with a one-year lag, using either fixed or random firm effects (Model 3). The coefficient for *Clustering* × *Reach* is positive and statistically significant in models using both two- and three-year lags (Models 6 and 9). This result holds for models using both fixed and random firm effects. Thus, our hypothesis received strong support in models using two- and three-year lags.[6]

To better understand the meaning of the interaction effect, the nature of the coefficients for *Clustering* and *Reach* in Models 6 and 9 in Table 2 must be understood. The estimated coefficients for *Clustering* and *Reach* in these models are *simple* effects rather than true main effects due to the significance of the interaction term (Jaccard and Turrisi 2003). Consequently, the effect of each on *Patents* is conditioned on the other variable taking on the value of zero. For example, the coefficient estimate of −0.022 for *Reach* in Model 6 (Random Effects) assumes that the value of *Clustering* is equal to zero (thus removing the impact of the interaction with *Reach*). Thus, the negative sign on the coefficient for *Reach* cannot be interpreted as a negative (main) effect of *Reach* on *Patents*. While the effect of *Reach* is indeed negative when *Clustering* is zero, the effect becomes positive when values of *Clustering* exceed 0.267 (the range of *Clustering* in the data is 0.0–0.8). Similarly, the effect of *Clustering* is negative (although not statistically significant) when *Reach* is equal to zero, but becomes positive for values of *Reach* greater than 1.224 (the range of *Reach* is 1.88–61.18).[7] The fact that the effects of both *Clustering* and *Reach* become positive when the other obtains a relatively small value, and increase in their positive effects with increases in the other, provides further support for our hypothesis. These mutually reinforcing effects are consistent with the symmetrical nature of multiplicative interaction effects (Jaccard and Turrisi 2003).

Plots of the effect of the interaction on predicted values of $Patents_{t+2}$ and $Patents_{t+3}$ reinforce this interpretation. For ease of presentation and interpretation, we used the log-linear form of the negative binomial models in Table 2 (i.e., where the log of

[6] We also conducted a test of the hypothesis at the industry level rather than the firm level. In this test, we regressed the industry's average number of firm patents on the network- and industry-level variables. We obtained nearly identical results to those in Table 2. These results are available from the authors upon request.

[7] To calculate these effects, we used the log-linear form of the negative binomial models in Table 2 (i.e., where the log of the conditional mean function is linear in the estimated parameters). We followed the approach for calculating interaction effects described by Jaccard and Turrisi (2003, p. 23).

Table 2 Panel Negative Binomial Regression Models with Fixed and Random Effects ($N = 1{,}106$; $Obs = 3{,}444$)

	$Patents_{it+1}$			$Patents_{it+2}$			$Patents_{it+3}$		
	1	2	3	4	5	6	7	8	9
Fixed effects									
Constant	1.136** (0.354)	0.582 (0.359)	0.604 (0.360)	1.257** (0.327)	1.663** (0.333)	1.614** (0.324)	1.433** (0.337)	1.859** (0.369)	1.825** (0.368)
Presample Patents	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)
Density	−0.248 (1.154)	−0.624 (1.358)	−0.527 (1.468)	−0.411 (1.529)	−2.220 (1.808)	−2.637 (1.843)	−2.012 (1.861)	−1.598 (2.509)	−1.674 (2.134)
Centralization	−0.014 (0.008)	−0.014 (0.008)	−0.012 (0.008)	−0.018** (0.006)	−0.016* (0.007)	−0.035** (0.006)	0.019** (0.007)	0.019** (0.007)	0.019* (0.007)
Ind. R&D Intensity	2.739 (2.668)	2.867 (2.522)	2.877 (2.581)	0.741 (2.366)	−0.088 (2.373)	−0.246 (2.327)	−7.126** (2.478)	−6.754** (2.504)	−6.754** (2.504)
R&D Alliance %	−0.112 (0.275)	0.223 (0.275)	0.222 (0.289)	0.068 (0.217)	−0.131 (0.223)	−0.188 (0.191)	−0.040 (0.248)	−0.305 (0.264)	−0.312 (0.304)
Efficiency	−0.199** (0.068)	−0.189** (0.072)	−0.190** (0.073)	−0.303** (0.091)	−0.321** (0.095)	−0.327** (0.087)	−0.267** (0.097)	−0.272** (0.089)	−0.270** (0.088)
Betweenness	0.003 (0.006)	0.003 (0.005)	0.003 (0.005)	0.005 (0.006)	0.004 (0.007)	0.002 (0.006)	−0.001 (0.009)	−0.001 (0.010)	−0.001 (0.010)
Clustering		0.420** (0.136)	0.507* (0.235)		0.346** (0.127)	−0.141 (0.196)		0.234 (0.183)	−0.319 (0.279)
Reach		0.010** (0.003)	0.011** (0.003)		−0.012** (0.003)	−0.020** (0.004)		−0.007* (0.003)	−0.009* (0.004)
Clustering × Reach			−0.015 (0.030)			0.081** (0.023)			0.014* (0.007)
Log Likelihood	−4,646.65	−4,637.32	−4,637.12	−4,597.46	−4,586.78	−4,577.98	−4,468.75	−4,464.64	4,464.46
Random effects									
Constant	1.118** (0.309)	0.542 (0.339)	0.541 (0.339)	0.984** (0.307)	1.342** (0.303)	1.256** (0.290)	0.920** (0.296)	1.333** (0.331)	1.214** (0.321)
Presample Patents	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)
Density	1.444 (0.900)	0.250 (1.092)	0.243 (1.166)	0.527 (1.197)	−1.872 (1.394)	−2.451 (1.352)	−1.454 (1.434)	−1.286 (1.618)	−1.538 (1.654)
Centralization	−0.021** (0.006)	−0.020** (0.007)	−0.021** (0.007)	−0.021** (0.006)	−0.020** (0.006)	−0.027** (0.005)	0.016* (0.006)	0.017* (0.007)	0.013* (0.006)
Ind. R&D Intensity	0.887 (2.429)	1.030 (2.408)	1.027 (2.424)	−0.357 (2.231)	−0.818 (2.151)	−0.590 (2.135)	−8.029** (2.278)	−7.987** (2.343)	−8.101** (2.460)
R&D Alliance %	0.014 (0.230)	0.383 (0.214)	0.384 (0.222)	0.208 (0.215)	−0.017 (0.187)	−0.090 (0.158)	0.106 (0.220)	−0.139 (0.233)	−0.153 (0.274)
Efficiency	−0.342** (0.062)	−0.336** (0.069)	−0.336** (0.069)	−0.396** (0.079)	−0.436** (0.081)	−0.435** (0.073)	−0.297** (0.087)	−0.307** (0.080)	−0.312** (0.078)
Betweenness	0.008 (0.005)	0.007 (0.004)	0.007 (0.005)	0.003 (0.005)	0.004 (0.005)	0.001 (0.005)	−0.000 (0.008)	−0.001 (0.008)	−0.001 (0.008)
Clustering		0.554** (0.106)	0.548** (0.212)		0.485** (0.116)	−0.101 (0.186)		0.152 (0.159)	−0.422 (0.344)
Reach		0.008** (0.003)	0.008* (0.003)		−0.013** (0.003)	−0.022** (0.003)		−0.008* (0.003)	−0.011** (0.004)
Clustering × Reach			0.001 (0.028)			0.082** (0.019)			0.043* (0.020)
a	0.707** (0.047)	0.716** (0.047)	0.710** (0.048)	0.675** (0.047)	0.684** (0.048)	0.690** (0.480)	0.650** (0.046)	0.652** (0.046)	0.652** (0.046)
b	0.358** (0.021)	0.360** (0.022)	0.360** (0.022)	0.321** (0.019)	0.328** (0.020)	0.334** (0.02)	0.291** (0.018)	0.290** (0.018)	0.293** (0.018)
Log likelihood	−8,520.70	−8,509.78	−8,509.78	−8,425.33	−8,407.95	−8,392.95	−8,198.66	8,194.98	−8,193.03

Notes. All models include firm, time period, and industry effects. Standard errors are in parentheses.
$^{*}p < 0.05$, $^{**}p < 0.01$ (two-tailed tests for all variables).

the conditional mean function is linear in the estimated parameters) to calculate these effects. Figure 2 presents the interaction plot of *Clustering* and *Reach* to illustrate the magnitude of the interaction effect. The "Low Clustering" line shows the slope of the effect of *Reach* on *Patents* when the value of *Clustering* is set to one standard deviation below its mean. The end points of the line are calculated at one standard deviation below and above the mean of *Reach*. The "High Clustering" line represents the effect of *Reach* on *Patents* when the value of *Clustering* is set to one standard deviation above its mean. Consistent with the results in Models 6 and 9 of Table 2, increases in *Reach* increase the positive effect of *Clustering* on *Patents*. The symmetrical case of plotting low and high *Reach* lines for low and high values of *Clustering* (not shown) provides similar results.

To assess the magnitude of the interaction effect we employed the estimated marginal effects ($e^{\beta X}\beta$). The magnitude of the interaction effect when both component variables increase one standard deviation above their means for the model employing a two-year lag and random effects is 1.00 patents (or 2.3%). For the model specified with a two-year lag and employing firm fixed effects, this yielded an increase of 0.98 patents (for the average firm), or 2.3%. The magnitude of the interaction effect is smaller in the models using a three-year lag. Thus, the size of the interaction effect in absolute terms is fairly small in our data and appears to realize its peak within two years. Based on these results, we speculate that the effect of network structure as a medium of knowledge diffusion decays over time. While a particular structure may persist over time, the knowledge that diffuses through it has limited benefit as actors absorb and apply these knowledge flows to productive ends.

The results related to the control variables also merit discussion. The effect of betweenness centrality on subsequent firm patenting failed to achieve statistical significance in any of the estimated models. In contrast, efficiency had a significant negative effect on firm patenting in all models. This result suggests that the presence of structural holes in a firm's ego network of alliance relationships has a deleterious effect on its inventive output. This is consistent with results obtained by Ahuja (2000) and Soh (2003). To our knowledge, our study represents the largest panel data investigation of this relationship.

Among the other variables in the models, most were not consistent in terms of sign and significance. This might be due, in part, to the moderate-to-large correlations among the network measures (i.e., *Centralization, Density, Reach, Clustering*, and *Clustering × Reach*). This multicollinearity might influence the robustness of our main finding because parameter estimates are unstable to very small changes in the data when substantial collinearity is present, sometimes resulting in the signs on estimated coefficients

Figure 2 **Graph of Interactions for Random Effects Models,** $Patents_{t+2}$ **and** $Patents_{t+3}$

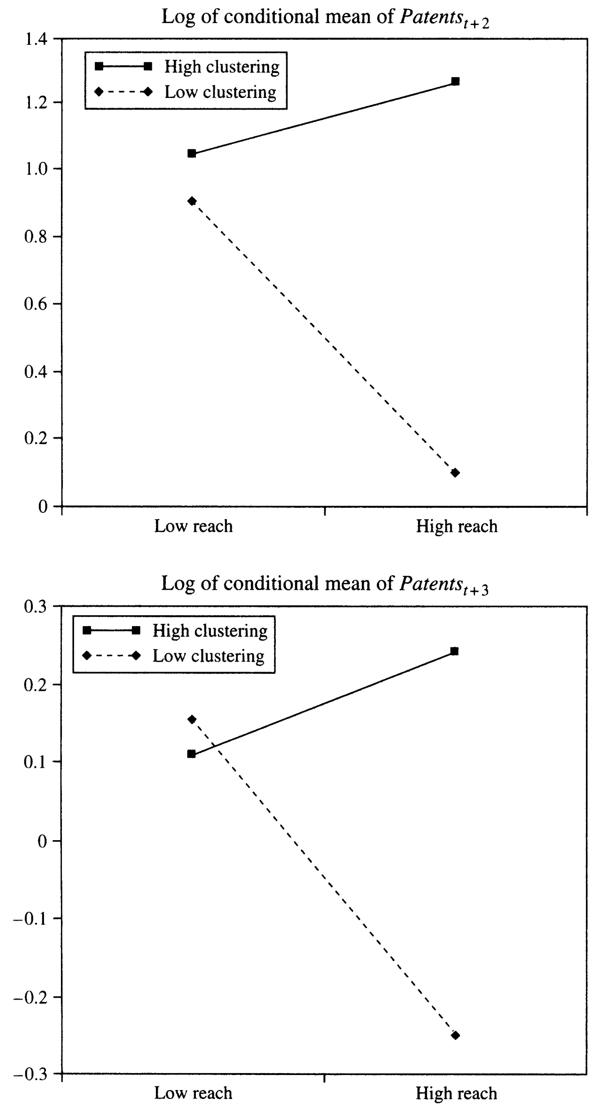

to flip (known as the "wrong sign" problem) (Gujarati 1995). To examine the influence of multicollinearity on our main result, we reran each of the models in Table 2 with *Centralization* removed and, alternatively, with *Density* removed (not reported here). The results for *Reach*, *Clustering*, and *Clustering* × *Reach* remained substantively unchanged across all models.

Finally, the *Presample Patents* variable was positive and significant in all models, indicating its importance as a control for firm-level unobserved heterogeneity. Furthermore, several time period and industry dummies (not reported) were consistently significant in all models.

Robustness of Results

One concern regarding our results is that we were not able to control for differences in firm R&D because nearly 42% of our sample firms were privately owned during some portion of the sample. Prior research has found that patent stock measures and annual R&D expenditures are highly correlated (e.g., Trajtenberg 1990). We measured a firm's patent stock as the total number of patents obtained by firm i in the four years prior to and including year t. Due to the extremely high correlation between this variable and *Presample Patents* ($r = 0.94$), we reestimated all of our models using the time-varying patent stock variable in place of *Presample Patents*. As might be expected (due to the substantial correlation between the two variables) our results (not reported) did not substantively change from those reported in Table 2.

For our second robustness check, we analyzed the data using a Poisson fixed-effects estimation procedure. We did so to address the concern identified in Footnote 5. This approach controls for all unobserved time-invariant sources of heterogeneity. In this analysis, we *excluded* all time-invariant variables and obtained qualitatively similar results (not reported) to those presented in Table 2.

A third concern regarding our results is that they may be influenced by the presence of persistent serial correlation in the residuals. This could result from temporally stable unobserved firm effects (Greene 1997), or from reverse causality running from firm invention to industry-level network structure (e.g., clustering or reach), manifesting in the lagged network variables. We explicitly address the first potential source of serial correlation by including firm fixed effects. Unreliable estimates may also result from unobservables that vary systematically over time. Serial correlation in the errors would persist even after controlling for stable firm effects. We examined this possibility in two ways. First, we regressed our measures of clustering and reach, and their interaction on annual firm patent counts using a linear panel data model. We did so using contemporaneously measured firm patents and one-, two-, and three-year lags of firm patents. We found no significant relationship between firm patents and clustering, reach, or their interaction in any of these models. Next, we aggregated firm patents to the industry level using the average annual patent count across firms in the industry. The idea here is that as industry inventiveness increases, so does the likelihood that firms in such industries form alliance networks with high clustering and reach. We ran the same specifications as those using firm patents and found no significant effects.

Discussion

We argued that two key structural properties of large-scale networks, clustering and reach, play important roles in network diffusion and search. Clustering enables even a globally sparse network to achieve high information transmission capacity through locally dense pockets of closely connected firms. Reach increases the quantity and diversity of information available to firms in the network by bringing the information resources of more firms within relatively close range. We thus argued that networks that have both the high information transmission capacity enabled by clustering, and the high quantity and diversity of information provided by reach, should facilitate greater innovation by firms that are members of the network. We tested this argument using longitudinal data on the innovative performance of a large panel of firms operating in 11 industry-level alliance networks. The results indicated support for our argument: the combination of clustering and reach was associated with significantly higher firm patenting. The results were stronger for models employing a two- and three-year lag versus a one-year lag, suggesting that firms do not quickly realize the innovation benefits of collaboration (Stuart 2000). These results were robust to a number of controls and model specifications.

Our results support much of the theory developed in recent work on small-world networks (Cowan and Jonard 2003, Uzzi and Spiro 2005). Our results are consistent with Uzzi and Spiro's argument that the cohesion and connectivity of a small-world network enable the circulation of creative material that can be recombined into new creative products. Our argument that the heterogeneity of knowledge distributed across clusters enhances innovation is similar to Uzzi and Spiro's argument that the different conventions and styles used in different clusters are a valuable source of diversity in the network.[8]

This research has a number of contributions. First, whereas previous alliance network research has examined the impact of a firm's network position or the structure of its immediate network neighborhood on firm innovation, our study is the first that we know of to examine the influence of the structure of industry-level alliance networks on firm innovation. The results of this study also inform the debate over whether innovation is enhanced by network density or efficiency (see Ahuja 2000): Both *local* density and *global* efficiency can exist simultaneously, and it is this combination that enhances innovation.

Finally, our results speak to the literature on knowledge spillovers. Knowledge spillovers represent an externality in which the knowledge produced by one firm can be appropriated, at little cost, by other firms (Jaffe 1986). Empirical evidence indicates that spillovers are important in explaining innovation and productivity growth (Griliches 1992). However, spillovers are not equally accessible to or appropriable by all firms. Prior research has shown that spillovers tend to be spatially bounded: Their effect is more pronounced for firms conducting research in similar technological domains (Jaffe 1986) and geographic locations (Feldman 1999). Our results add to this literature by suggesting that interfirm networks may be an important mechanism of knowledge spillovers, and that the specific pattern these relationships exhibit can have important consequences for the innovativeness of networked firms.

We do not wish to overstate our results—this study has a number of limitations. Our findings may be influenced by our assumption of average alliance duration. If alliances endure, on average, for more than three years, then the connectivity of our observed networks will be biased downward. This bias may influence our results. Unfortunately, due to data limitations, we were unable to explore this possibility.[9] A limitation of our theoretical focus is that we ignore the influence of network characteristics other than structure. We do not address the properties of the alliances themselves (e.g., strength, governance structure, scope). Different types of relationships may be better or worse for searching for, versus transferring, knowledge (Hansen 1999). In addition, different types of relationships will be more or less costly to maintain, and thus affect the efficiency of network structure for knowledge creation. We do not examine how

[8] Uzzi and Spiro's data and analysis are different from ours in some important ways. First, as they point out (2005, p. 470, Footnote 8), in a mature small-world network such as theirs, the path length changes little over time, behaving like a fixed effect with a constant value near one. This means that that their principle finding is driven primarily by temporal variation in clustering. Our networks, by contrast, exhibit significant cross-sectional and temporal variation in path length and network size, leading to great variation in our measure of reach. Second, our networks are far less dense than their network. Their network becomes sufficiently dense and clustered that it leads to excessive cohesion and homogenization of material, and a decline in creative performance. In essence, such a globally dense network has the advantages and disadvantages we argued would exist within each cluster. To investigate this effect in our data, we reestimated each of our models, replacing our interaction term with the quadratic version of clustering (i.e., clustering^2). This variable was not significant in any model; thus, we have no evidence of a parabolic effect of clustering in our data. We speculate that our networks never reach a high level of density, and thus are at less risk of excessive cohesion. Finally, and perhaps most importantly, Uzzi and Spiro's network is composed of individuals, whereas our networks are composed of firms. Some of the dynamics that lead to deleterious effects of cohesion (for example, strong feelings of obligation between friends leading to an "assistance club" for ineffectual members of the network) are far more likely in the relationships between individuals than between firms.

[9] We did not collect alliance formation data prior to 1990 because SDC data prior to that time is inconsistent. We chose to end observations of patents in 2000 (implying our alliance observation ended in 1997) because the lag between patent application and grant date is two to four years (which was toward the end of our data collection).

the attributes of the firms shape the flow of knowledge (see Owen-Smith and Powell 2004). We have also not addressed the potential impact of the nature of knowledge that is being accessed, transferred, and recombined in the network. Different characteristics of knowledge (e.g., tacit versus explicit, complex versus simple, etc.) can influence the knowledge creation and innovation process (Zander and Kogut 1995). Network structure may also differentially interact with different dimensions of knowledge. For example, the high density of clusters may facilitate the search and transfer of tacit, complex knowledge, but the relatively few connections to other clusters may make such search and transfer problematic. These aspects of relationships and knowledge will likely be important in fully understanding the relationship between interfirm knowledge networks and knowledge creation, but are beyond the scope of our paper.

Another limitation of our work is that the generalizability of our main result is likely to be limited to industries that make frequent use of alliances. Networks characterized by extreme sparsity may not have a sufficient degree of connectedness to observe clustering or meaningful reach. However, the implications of our results are not necessarily limited to alliance relationships. Because firms are connected via other relationships, the global structure of such relationships may influence firm innovativeness. For example, firms are often connected by interpersonal collaborative relationships among individual inventors. The extent to which the global structure of these relationships is characterized by clustering and reach may have implications for the inventiveness of individual inventors and their firms (Fleming et al. 2004). Furthermore, because knowledge can flow between firms through other mechanisms such as individual mobility, geographic clustering, participation in technical committees, or learning from information made public through patenting, it is possible that some of the knowledge creation advantages of a particular alliance network structure might spill over to other industry (or nonindustry) participants. Each of these limitations represents an exciting area for future research.

Acknowledgments

This research was supported by the National Science Foundation under Grant SES-0234075. The authors are grateful for the suggestions of Juan Alcácer, Laszlo Barabasi, Joel Baum, Bill Greene, Anne Marie Knott, Dan Levinthal, Bill McKelvey, Mark Newman, Joe Porac, Lori Rosenkopf, Rob Salomon, Kevin Steensma, Kate Stovel, and Duncan Watts.

References

Ahuja, G. 2000. Collaboration networks, structural holes, and innovation: A longitudinal study. *Admin. Sci. Quart.* **45** 425–455.

Allison, P. D., R. Waterman. 2002. Fixed-effects negative binomial regression models. *Soc. Methodology* **32** 247–265.

Anand, B. N., T. Khanna. 2000. Do firms learn to create value? The case of alliances. *Strategic Management J.* **21** 295–315.

Basberg, B. L. 1987. Patents and the measurement of technological change: A survey of the literature. *Res. Policy* **16** 131–141.

Baum, J. A. C., T. Calabrese, B. S. Silverman. 2000. Don't go it alone: Alliance network composition and startups' performance in Canadian biotechnology. *Strategic Management J.* **21** 267–294.

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

Blundell, R. R., R. Griffith, J. Van Reenen. 1995. Dynamic count data models of technological innovation. *Econom. J.* **105** 333–344.

Borgatti, S. P., M. G. Everett, L. C. Freeman. 2002. Ucinet for Windows: Software for social network analysis. Harvard, Analytic Technologies.

Brown, J., P. Duguid. 1991. Organizational learning and communities of practice: Towards a unified view of working, learning and innovation. *Organ. Sci.* **2** 40–57.

Burt, R. S. 1992. *Structural Holes*. Harvard University Press, Cambridge, MA.

Burt, R. S. 2001. Bandwidth and echo: Trust, information, and gossip in social networks. A. Casella, J. E. Rauch, eds. *Networks and Markets: Contributions from Economics and Sociology*. Russel Sage Foundation, New York, 30–74.

Cameron, A. C., P. K. Trivedi. 1986. Econometric models based on count data: Comparisons and applications of some estimators and tests. *J. Appl. Econometrics* **1** 29–53.

Cameron, A. C., P. K. Trivedi. 1998. *Regression Analysis of Count Data*. Cambridge University Press, Cambridge, UK.

Coleman, J. S. 1988. Social capital in the creation of human capital. *Amer. J. Sociol.* **94** S95–S120.

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

Dee, T. S., D. C. Grabowski, M. A. Morrisey. 2005. Graduated driver licensing and teen traffic fatalities. *J. Health Econom.* **24** 571–589.

Deeds, D. L., C. W. L. Hill. 1996. Strategic alliances and the rate of new product development: An empirical study of entrepreneurial firms. *J. Bus. Venturing* **11** 41–55.

Dosi, G. 1988. Sources, procedures, and microeconomic effects of innovation. *J. Econom. Literature* **26** 1120–1171.

Dyer, J. H., H. Singh. 1998. The relational view: Cooperative strategy and sources of interorganizational competitive advantage. *Acad. Management Rev.* **23** 660–679.

Feldman, M. P. 1999. The new economics of innovation, spillovers, and agglomeration: A review of empirical studies. *Econom. Innovation New Tech.* **8** 5–25.

Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Sci.* **47** 117–132.

Fleming, L., A. Juda, C. King, III. 2004. Small worlds and regional innovative advantage. Harvard Business School Working Paper Series, 04-008, Harvard University, Cambridge, MA.

Freeman, C. 1991. Networks of innovators: A synthesis of research issues. *Res. Policy* **20** 499–514.

Freeman, L. 1979. A set of measures of centrality: I. Conceptual clarification. *Soc. Networks* **1** 215–239.

Gomes-Casseres, B., J. Hagedoorn, A. Jaffe. 2006. Do alliances promote knowledge flows? *J. Financial Econom.* **80** 5–33.

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

Granovetter, M. S. 1992. Problems of explanation in economic sociology. N. Nohria, R. Eccles, eds. *Networks and Organizations: Structure, Form, and Action*. Harvard Business School Press, Boston, MA, 25–56.

Greene, W. H. 1997. *Econometric Analysis*, 3rd ed. Prentice Hall, Upper Saddle River, NJ.

Griliches, Z. 1990. Patent statistics as economic indicators: A survey. *J. Econom. Literature* **28** 1661–1707.

Griliches, Z. 1992. The search for R&D spillovers. *Scandinavian J. Econom.* **94**(Suppl.) 29–47.

Gujarati, D. N. 1995. *Basic Econometrics*, 3rd ed. McGraw-Hill, New York.

Gulati, R. 1998. Alliances and networks. *Strategic Management J.* **19** 293–317.

Gulati, R., M. Gargiulo. 1999. Where do interorganizational networks come from? *Amer. J. Sociol.* **104** 1439–1493.

Hansen, M. T. 1999. The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Admin. Sci. Quart.* **44** 82–111.

Hansen, M. T. 2002. Knowledge networks: Explaining effective knowledge sharing in multiunit companies. *Organ. Sci.* **13** 232–250.

Hargadon, A. B. 1998. Firms as knowledge brokers: Lessons in pursuing continuous innovation. *California Management Rev.* **40**(3) 209–227.

Hargadon, A. B., A. Fanelli. 2002. Action and possibility: Reconciling dual perspectives of knowledge in organizations. *Organ. Sci.* **13** 290–302.

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.

Hecker, D. 1999. High technology employment: A broader view. *Monthly Labor Rev.* **122**(6) 18–28.

Heckman, J. J. 1979. Sample selection bias as a specification error. *Econometrica* **47**(1) 153–161.

Henderson, R., K. Clark. 1990. Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Admin. Sci. Quart.* **35** 9–30.

Jaccard, J., R. Turrisi. 2003. *Interaction Effects in Multiple Regression*. Sage, Thousand Oaks, CA.

Jaffe, A. 1986. Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits and market value. *Amer. Econom. Rev.* **76** 984–1001.

Kogut, B. 2000. The network as knowledge. Generative rules and the emergence of structure. *Strategic Management J.* **21** 405–425.

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

Laumann, E. O, P. V. Marsden, D. Prensky. 1983. The boundary specification problem in network analysis. R. S. Burt, M. J. Minor, eds. *Applied Network Analysis*. Sage, Beverly Hills, CA, 18–34.

Levin, R., A. Klevorick, R. Nelson, S. Winter. 1987. Appropriating the returns from industrial research and development. *Brookings Papers Econom. Activity, Microeconomics* **3** 783–820.

Mowery, D. C., J. E. Oxley, B. S. Silverman. 1996. Strategic alliances and interfirm knowledge transfer. *Strategic Management J.* **17** 77–91.

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

Newman, M. E. J., S. H. Strogatz, D. J. Watts. 2002. Random graph models of social networks. *Proc. Natl. Acad. Sci. USA* **99** 2566–2572.

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** 5–21.

Phelps, C. 2003. Technological exploration: A longitudinal study of the role of recombinatory search and social capital in alliance networks. Unpublished dissertation, New York University, New York.

Powell, W. W., L. Smith-Doerr. 1994. Networks and economic life. N. J. Smelser, R. Swedberg, eds. *The Handbook of Economic Sociology*. Princeton University Press, Princeton, NJ, 368–402.

Powell, W. W., K. W. Koput, L. Smith-Doerr. 1996. Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology. *Admin. Sci. Quart.* **41** 116–145.

Robinson, D. T., T. E. Stuart. 2007. Network effects in the governance of strategic alliances. *J. Law, Econom., Organ.* **23**(1) 242–273.

Rosenkopf, L., P. Almeida. 2003. Overcoming local search through alliances and mobility. *Management Sci.* **49** 751–766.

Rosenkopf, L., A. Metiu, V. P. George. 2001. From the bottom up? Technical committee activity and alliance formation. *Admin. Sci. Quart.* **46** 748–772.

Sampson, R. 2004. The cost of misaligned governance in R&D alliances. *J. Law, Econom., Organ.* **20** 484–526.

Schmookler, J. 1966. *Invention and Economic Growth*. Harvard University Press, Cambridge, MA.

Shan, W., G. Walker, B. Kogut. 1994. Interfirm cooperation and startup innovation in the biotechnology industry. *Strategic Management J.* **15** 387–394.

Smith-Doerr, L., J. Owen-Smith, K. W. Koput, W. W. Powell. 1999. Networks and knowledge production: Collaboration and patenting in biotechnology. R. Leenders, S. Gabbay, eds. *Corporate Social Capital*. Kluwer Academic Publishers, Norwell, MA, 331–350.

Soh, P.-H. 2003. The role of networking alliances in information acquisition and its implication for new product performance. *J. Bus. Venturing* **18** 727–744.

Stuart, T. E. 2000. Interorganizational alliances and the performance of firms: A study of growth and innovation rates in a high technology industry. *Strategic Management J.* **21** 791–812.

Trajtenberg, M. 1987. Patents, citations, and innovations: Tracing the links. Working Paper 2457, National Bureau of Economic Research, Cambridge, MA.

Trajtenberg, M. 1990. A penny for your quotes: Patent citations and the value of innovations. *RAND J. Econom.* **21** 172–187.

Uzzi, B. 1997. Social structure and competition in interfirm networks: The paradox of embeddedness. *Admin. Sci. Quart.* **42** 35–67.

Uzzi, B., J. Spiro. 2005. Collaboration and creativity: The small world problem. *Amer. J. Sociol.* **111** 447–504.

Vonortas, N. S. 1997. Research joint ventures in the US. *Res. Policy* **26** 577–595.

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

Watts, D. J. 1999. Networks, dynamics, and the small-world phenomenon. *Amer. J. Sociol.* **105** 493–528.

Wellman, B. 1988. Structural analysis: From method and metaphor to theory and substance. B. Wellman, S. D. Berkowitz, eds. *Social Structures: A Network Approach*. Cambridge University Press, Cambridge, UK, 19–61.

Yamaguchi, K. 1994. The flow of information through social networks: Diagonal-free measures of inefficiency and the structural determinants of inefficiency. *Soc. Networks* **16** 57–86.

Zander, U., B. Kogut. 1995. Knowledge and the speed of the transfer and imitation of organizational capabilities: An empirical test. *Organ. Sci.* **6** 76–92.