Prescriptions for Network Strategy: Does evidence of network effects in cross-section support them?

Joel A.C. Baum*, Robin Cowan†, and Nicolas Jonard‡

*University of Toronto,
†University of Strasbourg and Maastricht University
‡University of Luxembourg

February 14, 2011

Abstract

While intuitively appealing (and common), drawing network strategy implications from empirical evidence of network performance effects in pooled cross-section is not necessarily warranted. This is because network positions can influence both the mean *and* variance of firm performance. Strategic prescriptions are warranted if empirically observed network effects reflect increases in *mean* firm performance. If network effects reflect increases in firm performance *variance*, however, such prescriptions are warranted only if the increase in the odds of achieving high performance are sufficient to compensate for the concomitant increase in the odds of realizing poor performance. Our simulation study, designed to examine network effects in both pooled cross-section and within-firm over time under a wide range of conditions, counsels caution in drawing implications for network strategies.

1 Introduction

At the intersection of network and strategy literatures, researchers are concerned with how patterns of interfirm alliances create network-based advantages for well-connected firms. Researchers have examined whether firms should occupy densely interconnected "closed " positions, which afford coordination and integration benefits by facilitating the ease of exchange and commonness of information among firms (Coleman, 1988, 1990), or in sparsely interconnected "open " positions, which confer information access and control benefits through conveyance of diverse information and resources and brokerage opportunities (Burt, 1992, 2000; Granovetter, 1973).

Although empirical evidence remains somewhat equivocal, the idea of "network effects" on firm performance are by now uncontroversial, and attention is focused increasingly on the identification of conditions under which open or closed network positions are more advantageous (e.g., Ahuja, 2000; Burt 1998, 2000; Rowley et al., 2000). What remains uncertain, however, is whether the evidence of "network effects" supports prescriptions for "network strategy," where the former is typically based on empirical estimates from panel or pooled time series and cross sectional data showing that firms occupying a particular type of network position at time t outperform firms that do not at time $t+1$, and the latter based on the inference that firms sustaining the type of network position shown beneficial in pooled cross-section, will outperform those that do not over the longer term.

Although such an inference is intuitively appealing, the implications of network effects in cross-section do not translate straightforwardly into implications for sustained network strategy. This is because network positions can affect both the mean and variance of a firm's performance distribution, which were shown by March (1991) to play different roles in competition for high performance. If a network position increases both the mean and variance of performance, the firm gains a performance advantage over its rivals. A network position advantage may also be gained if the increase in performance variance is sufficient to compensate for a decrease in mean performance, and vice versa. One of March's (1991) findings is that when competition is for primacy, increasing the variance of performance (rather than the mean) contributes increasingly to competitive advantage as the number of competitors increases.

But increased performance variance also exposes firms to a risk of very poor outcomes. The benefit of increasing firm performance variance thus depends importantly on the distribution of possible performance outcomes. The likelihood of achieving high performance, in particular, depends on the right-hand tail of the performance distribution; the left-hand tail is critical to experiencing poor performance. If the distribution is right-skewed, with only a small number of high performance outcomes available relative to poor ones, increasing performance variance improves a firm's chances of achieving one of the high outcomes, but the odds may be to too small to compensate for the reduction in the mean that results from the concomitantly increased odds of obtaining the poor ones.

The implication of these observations is that if the empirical evidence of network effects in pooled cross-section reflects the influence of firms' network positions at time t on the variance, rather than mean, of their performance outcomes at time $t + 1$, whether network positions associated with high performance in cross-section will also prove beneficial if sustained by firms over time depends on whether the increased odds of achieving high performance are nullified over time by an even larger increase in the risk of experiencing poor outcomes.

Three factors suggest both that empirical findings plausibly reflect performance variance (rather than mean) effects and that such nullification is likely. One is that empirical studies of network effects model the role of network position in achieving high performance in pooled cross-section, which is governed by performance variance, rather than avoiding low performance, which is governed by mean performance (March, 1991). A second is the commonly observed right-skewness of firm performance distributions, which suggest that increased performance variance may often raise firms' risk of low performance more than their chances of achieving high performance, and thus reduce their mean performance. The third is that analyses of performance in pooled cross-section are informed by performance variation across rather than within firms, and thus indifferent to which particular firms perform well in each cross section. As a result, while firms occupying a particular type of network position may tend to achieve high performance in each cross section, it may be different firms achieving the high performance in each cross section.

In this paper, we assess the extent to which network strategies can be inferred from network effects empirically observed in pooled cross-section. Our approach is simulation based, which allows us to examine performance in both pooled cross-section and within-firm over time over a wide range of performance distributions. We focus on firm performance effects of open and closed network positions, both because they are the main focus of contemporary empirical work, and because of their distinct performance improvement profiles. In particular, recent empirical research (e.g., Ahuja, 2000; Rothaermel and Deeds, 2004; Rowley et al., 2000) suggests that open network positions facilitate exploratory (nonlocal) search, which increases firms' performance variance, while closed network positions facilitate exploitive (local) search, which lowers firms' performance variance (Holland, 1975; March, 1991).

Our model is designed to characterize innovation networks and outcomes, for which the role of alliances for the acquisition of externally developed knowledge is well established (Powell et al., 1996; Rosenkopf and Almeida, 2003; Rosenkopf and Nerkar, 2001), exploitation and exploration activities are highly germane (Gilsing et al., 2008; Lavie and Rosenkopf, 2006; Rothaermel and Deeds, 2004; Rowley et al., 2000), and network effects have been an important focus of research (e.g., Ahuja, 2000; Powell et al., 2005; Rowley et al., 2000).

After developing our theoretical expectations regarding firms' network positions and performance variation more fully, we describe and validate our simulation model, demonstrating that, across a range of distributions of possible firm performance, the model replicates both the properties of 'real world' networks, as well as network effects on performance in cross-section obtained in empirical network studies. We then turn our attention to network strategies, and specifically, to the questions of whether 1) network positions associated with high performance in pooled cross-section are also beneficial if sustained in the long run, and 2) the skewness of firm performance distributions impacts the veracity of conclusions for network strategy drawn from network effects observed in pooled cross-section.

2 Network Positions and Performance Variance

Although network theorists agree that "better-connected " firms have a competitive advantage, there is disagreement regarding what "better-connected" means (Rowley et al., 2000). Coleman's (1988) closure argument implies that firms are better off occupying densely interconnected, closed network positions in which their partners are also partners, while in contrast Burt's (1992) structural hole argument, prescribes firms embed themselves in a sparsely connected, open network position comprised of disconnected partners. Rather than arguing the superiority of one network position over the other, Burt (1998: 45) suggests a contingency approach to reconcile the debate: "closure and hole arguments are not as contradictory as they might seem ... The ambiguity stems in large part from the different roles that social capital plays in the study populations with which each is justified." Open and closed network positions thus afford different benefits that are useful for different purposes, and understanding their effects requires consideration of the conditions under which firms are better off possessing the distinct benefits they afford.

Consistent with Burt's (1998) view, recent studies adopt contingency approaches in which the benefits of open and closed network positions depend on environmental conditions and task purposes (Ahuja, 2000; Rothaermel and Deeds, 2004; Rowley et al., 2000). These contingency approaches conceive the appropriate type of network position to depend on their differential value for exploitive and exploratory learning modes. That is, the degree to which firms are focused on exploiting existing technologies, skills, and information, or exploring the environment for emerging innovations and other significant changes.

According to March (1991: 85) "the essence of exploitation is the refinement and extension of existing competencies, technologies and paradigms. The essence of exploration is experimentation with new, uncertain alternatives." Exploitation involves using existing knowledge to improve organizational functioning by reducing variability in the quality or efficiency of current strategies, competencies, and procedures. In exploitation, the emphasis is on refining existing knowledge by gathering specific information that will provide deeper understanding in a particular area. The solution space is thus well defined, and search is local and highly specific. The classic exemplar is the well-known "experience curve" phenomenon in which firms reduce production cost and/or time by eliminating redundancies and inefficiencies through continuous tuning of internal practices and processes (Yelle, 1979; Argote, 1993). Exploration, in contrast, entails processes of concerted variation and experimentation to identify new ways of doing things and new things to do. In exploration, because the focus is on gathering information on identifying emerging innovations and alternative future options, the solution space is ill-defined, and search is wide, and a premium is placed on newer, more general information.

March (1991) contends that both processes are required, but that a trade-off must be made between how much to invest in refining existing technologies to stay competitive in their current markets in the short term, compared to developing new knowledge about novel technologies with which to compete in the long term when environmental demands change. The balance of resources firms allocate to exploitation and exploration thus tends to depend on environmental conditions (Lant, Milliken and Batra, 1992; Rowley et al., 2000). Environmental uncertainty, in particular, by affecting the predictability and frequency of change (Dess and Beard, 1984), influences the degree to which firms must emphasize refinement of existing knowledge and/or seek out new opportunities. In uncertain environments firms must allocate more resources to exploration than in more stable environments in which there is greater certainty about future directions and fewer environmental disturbances.

The distinct information requirements of exploitation and exploration suggest different prescriptions regarding the appropriate network position (Gilsing et al., 2008; Rowley et al., 2000). Open network positions, comprised of disconnected, non-redundant partners, are ideal for gaining access to diverse sources of information and knowledge, to facilitate identification of emerging opportunities and threats, and location of complementary knowledge (Powell et al., 1996; Mitsuhashi, 2003). Open network positions thus afford the firm unique information and perspectives from each of its partners that facilitates broad search for emerging innovations and alternative future options. In uncertain environments requiring large investments in exploration, sparsely connected, open network positions are thus advantageous.

Closed network positions, in contrast, inhibit firms' access to broader, divergent, distant, and less familiar approaches and critical to exploration (Uzzi, 1996, 1997). Although this limits their usefulness in meeting the demands of uncertain environments, the access to redundant and validating information they afford is essential to meeting the information requirements of exploitation (Dyer and Singh, 1998; Van de Ven, 1976; Walker et al., 1997). The ability to triangulate across multiple, redundant sources enhances evaluation of acquired information obtained from each source, aiding in refinement of existing knowledge. In certain environments requiring large investments in exploitation, closed network positions are thus advantageous.

The degree of uncertainty and rate of change in the environment thus influence appropriate network positions: Firms operating in a rapidly changing environment will benefit from open network positions facilitating exploration, while firms in a stable environment will benefit from closed network positions facilitating exploitation.

Importantly, the link between network positions and learning also suggests that closed and open network positions will impact firm performance variation distinctly. In particular, closed network positions, by facilitating exploitive learning, should reduce firm performance variation, while open network positions, by facilitating exploratory learning, should increase it. As a result, the relative performance effects of open and closed network positions ought to depend importantly on the distribution of possible performance outcomes. In particular, closed network positions should benefit firms when the distribution of possible performance outcomes favors low-variation firms, while open network positions should benefit firms when the distribution of possible performance outcomes favors high-variation firms.

When the distribution of possible performance outcomes is skewed highly to the right, a firm's odds of achieving high performance in cross-section increase with the variance of its performance outcomes (March, 1991). Consequently, in cross-section, the benefit of variance-increasing open network positions should increase with the right-skewness of the distribution of possible performance outcomes. In environments characterized by the possibility of disruptive, discontinuous innovations, and thus a right-skewed distribution of potential performance outcomes, open network positions should produce higher performance in cross-section. In contrast, in environments characterized primarily by incremental innovations, variance-reducing closed network positions should produce higher performance in cross-section. These predictions are consistent with arguments and empirical findings in studies of network effects (e.g., Ahuja, 2000; Gilsing et al., 2008; Rothaermel and Deeds, 2004; Rowley et al., 2000).

Over time, however, because variance-increasing open network positions increase firms' exposure to poor performance, mean performance for firms sustaining open network positions over time should decline with the right-skewness of the performance distribution. Moreover, because closed network positions provide protection against poor performance, the performance benefit of sustaining a closed network position over time should increase with the right-skewness of the performance distribution.

3 Model

Below, we design a simple model of partner selection in which firms ally for the purpose of innovating, and in so doing create an industry network. In the model, two forces shape the alliance network. The first is the process of alliance formation. The need to have complementary knowledge for successful innovation implies that partners will be neither too close together nor too far apart in knowledge space. Because the knowledge space is a metric space, this need for proximity will induce some amount of local correlation in the decisions to ally, yielding both repetition and cliquey-ness in firms' partnering decisions. The second force is innovation. When a firm produces a successful innovation, other firms are dislocated in the knowledge space. This process maintains a heterogeneous population of firms, scattered over the knowledge space, and also generates clique-spanning ties that lower network distances. Combined, the decisions of firms to ally with profitable partners, and their dislocation of in knowledge space in response to innovation, produce networks that exhibit features common to 'real world' networks.

More formally, consider a fixed, finite population of firms, located in a metric knowledge space. Firms form alliances with the goal of creating innovations, and for any firm, each of its alliances represents one R&D project.[1] An alliance is in essence an institution for knowledge sharing, so a firm forms an alliance with another firm in order to access the knowledge it needs to (potentially) make that project successful. The cost of an alliance is c, paid by both firms in any period in which the alliance exists. Firms are the locus of innovation. If an innovation takes place in a firm, it disrupts the activities of all other (non-innovating) firms. In response, firms rearrange their activities by deploying different knowledge than they did before the innovation. An innovation thus changes the value of the knowledge stocks of other firms; and it changes the value of partnerships between firms (represented by their specific knowledge endowments). We compress this into a single response, in which non-innovating firms are relocated in knowledge space, thereby changing the value of (potential) partnerships.

When a firm innovates, it receives a value equal to the total displacement of other firms in the knowledge space. The probability that a firm succeeds through a particular alliance is a positive, single-peaked function of the distance in knowledge space from the firm to its partner in that alliance. Knowing both of these facts, firms can calculate the expected benefit of any possible alliance as the product of the probability that the alliance will be successful and the total dislocation the firm can expect to impose on other firms in the industry, minus the alliance cost. Alliances form when expected profits are positive for both partners. In each period all actual and potential alliances are (re-)evaluated. If a potential alliance shows positive expected profits it forms, or is maintained. If an existing alliance shows negative expected profits it is terminated.

In the following sections we operationalize these assumptions, taking each part of the process individually.

[1] For simplicity we assume that firms undertake no R&D outside alliances.

3.1 Firms and innovation

A fixed, finite population of N firms is located in a K-dimensional metric knowledge space. Each firm is characterized by a knowledge endowment, $\mathbf{v}_i = (v_{i,1}, v_{i,2}, \ldots, v_{i,K}), i = 1, \ldots, N$. Firms do not move in knowledge space unless they are affected by innovations of other firms.[2]

We treat innovations as stochastic, and independent across firms. Each firm engages in knowledge sharing with its alliance partners, pooling the knowledge it gathers from them into a single, large research project. Each alliance formed by the focal firm contributes additively to the success rate of this research project. The marginal contribution of a particular alliance to the firm's project is its success rate, which in turns depends on the characteristics of the alliance.

Specifically, the success rate for any alliance depends of the "goodness of fit" of the alliance, which is assumed to be a single-peaked function of the Euclidean distance (with a maximum at finite distance) between the alliance partners in knowledge space.[3] Formally, we employ a bell-shaped (Gaussian) function to map knowledge-distance to one-period success rates, according to

$$\lambda_{ij} \equiv f(d_{ij}) = \bar{\lambda} \exp - \left[(d_{ij} - d^*) / \sigma \right]^2, \tag{1}$$

where d_{ij} is the Euclidean distance in knowledge space between i and j, and $\bar{\lambda} \ll 1$ is a scaling parameter we use to control the maximum success probability. Firm i's overall success rate is simply $\lambda_i = \sum_{ij \in g} \lambda_{ij}$, where g is the network of existing alliances during the period we

[2]their alliance partners (e.g., Baum et al., 2010), or possibly autarchically. Including this source of motion in the knowledge space has minor effects on some of the results we present below, but does not interact with the main effects we wish to demonstrate, and so for simplicity we omit it.

[3]Several empirical studies have shown, in a variety of contexts, that alliance or merger success is driven by partners' relative knowledge endowments. This is commonly formalized using "distance in knowledge space", which is measured in a variety of ways. Additionally, recent work has shown that the probability that a pair of firms *forms* an alliance is an concave in their distance in knowledge space. See for example (Ahuja and Katila, 2001; Mowery et al., 1996, 1998; Mueller et al., 2009; Rothaermel and Boeker, 2008; Schoenmakers and Duysters, 2006; Stuart, 1998).

consider. The population-wide arrival rate of innovations is $\lambda = \sum_{i=1}^{N} \lambda_i$. Provided success rates are small enough ($\lambda_{ij} < \bar{\lambda} \ll 1$), any firm's overall success rate λ_i is also small, and so is the population-wide rate λ, so that one firm at most will succeed in any given period.

3.2 Dislocation following an innovation

Following an innovation, the knowledge landscape changes, altering the value of different types of knowledge and different knowledge combinations. We assume that an innovation by firm i disrupts the status quo for all other firms. Firms respond by deploying new knowledge in their activities, and so an innovation changes both where firms are located in space and which combinations of knowledge are valuable. The extent to which a non-innovating firm is affected by an innovation is determined both by how close it is to the innovating firm, and by the disruptiveness of the innovation. We formalize this by assuming that following an innovation, firms are dislocated in the knowledge space. More precisely, following an innovation by firm i, firm i itself is not moved but all other firms are dislocated as a function of their Euclidean distance to i in knowledge space. Any firm k is relocated, uniformly at random within a disk centered on k's pre-innovation location. The disk has radius

$$r_k = \exp\left(-d_{ik}/\theta\right), \tag{2}$$

where θ is a parameter governing the innovation "regime "of the industry; that is, whether the industry tends towards disruptive or incremental innovations. The direction of movement of firm k is determined by an angle drawn uniformly at random in $[0, 2\pi)$. Thus the expected magnitude of the dislocation of k is $r_k/2$. Total expected dislocation, which we employ as a measure of the "value" of an innovation by firm i, is thus simply

$$V_i = \sum_{k \neq i} r_k = \sum_{k \neq i} \exp\left(-d_{ik}/\theta\right)/2. \tag{3}$$

Increasing θ amounts to increasing the value of r_k, the radius of the destination disk of any dislocated firm, and thus the total amount of dislocations firm i can expect to impose on the rest of the industry, as well as the right-skewness of its performance distribution.

3.3 Expected profits and the decision to ally

The value to firm i of a particular alliance ij is the marginal contribution of that particular alliance to the firm's expected profit, net of alliance cost c. The marginal contribution of alliance ij to the success rate of firm i is simply λ_{ij}. (We assume that innovations do not compete in any after-market, so firms can treat each alliance decision independently from all others, both their own and other firms'.) If i innovates, the value of that innovation is the sum of the dislocations it imposes on other firms, as in Equation 3. We can thus write the expected value of alliance ij to firm i as

$$\pi_i^{ij} = \lambda_{ij} V_i - c, \tag{4}$$

the product of marginal contribution of alliance ij to the success rate of firm i by the dislocation i expects to impose on other firms, net of alliance cost c.

An alliance will form if expected profits are positive to both partnering firms. Rewriting, the alliance (formation or continuation) condition is thus simply

$$\lambda_{ij} V_i \geq c \text{ and } \lambda_{ij} V_j \geq c. \tag{5}$$

In each period the industry network thus consists of the alliances formed by all firm pairs satisfying Constraint 5. However, in a dynamic industry, we observe that knowledge portfolios, alliance partnerships and the industry network all change over time. At the aggregate level, the industry network structure is determined by (current) knowledge stocks of the firms in it. But the network structure is instrumental in the innovation within the industry. Innovations force firms to respond, and to redeploy their knowledge, possibly using different knowledge in reaction to the new market conditions created by the innovation. Thus innovations change firms' position in the knowledge space, by forcing them to use different knowledge than they had done previously. So at the aggregate level network structure and the industry knowledge profile, coevolve. In the sections that follow, we explore this coevolution through numerical simulation of the model.

4 Numerical implementation of the model

Being a model of complex coevolution, the model does not lend itself to analytical solution. Thus we examine its behavior numerically, with the following settings.

We consider a population of $N = 100$ firms, and use as the knowledge space the real plane. Thus each firm's knowledge endowment is a pair of positive real numbers, $\mathbf{v}_i = (v_{i,1}, v_{i,2})$, for all $i = 1, \ldots, 100$. At the outset, all firms hold knowledge endowments distributed uniformly over the unit square. We set the optimal distance to $\delta^* = 0.5$, and σ, the parameter that governs how fast success probabilities fall as firm-pairs deviate from the optimal distance, at 0.2. The upper bound to the success rate of an alliance (which is achieved when participants are at the optimal distance) is $\bar{\lambda} = 0.005$, as visible in Equation 1. The cost of forming and maintaining an alliance is set to $c = 0.002$.

The final parameter, θ, controls the innovation regime of the industry. The idea is that in some industries innovations are mostly incremental, and will not have large effects on the relevance of different types of knowledge or knowledge combinations within the industry. Here, θ will be small, and innovations will have little disruptive effect on non-innovating firms. By contrast in other industries, where there is scope for more dramatic changes (young, technology-based industries, for example), innovations can be large, and potentially disrupt many firms and alliances. Here θ will be large.[4]

Our main interest lies in whether the network and firms in it perform differently under different innovation regimes. As the disruptiveness of the innovation regime increases, so too does the total dislocation a firm can expect to impose on the rest of the industry, and concomitantly the right-skewness of the firm performance distribution. Consequently, we expect that, in cross-section, the performance benefit of open (closed) network positions will increase (decrease) with the disruptiveness of the innovation regime, but that, in the long run, the relationships will be reversed. To examine these predictions, we consider 20 different values of θ, evenly spaced between 0.08 and 1.0 on a log-scale.

The experiment is run by initializing the firms' knowledge endowments uniformly at random over the unit square. Alliances form each period according to Constraint 5. Innovations arrive randomly, determined by the independent Poisson processes of the set of alliances as described in Subsection 3.1. We discard the initial 100 periods, to avoid any possible spurious effects arising from the initialization, and run the alliance formation/innovation process for 2500 periods. At each value of θ we replicate the experiment 48 times, and in the figures below we display plots showing the median, first, and third quartiles of the 48 observations.

[4] On disruptive and incremental innovations, see for instance Anderson and Tushman (1986) and Abernathy and Clark (1985).

5 Results

5.1 Three snapshots

As a first step in the analysis, we display in Figure 1 three illustrative networks. These are networks for period 1000 for each of 3 values of θ. From the figure, it is apparent that more disruptive innovation regimes are associated with sparser networks.

With incremental innovation ($\theta = 0.08$) we observe one connected component and many isolated firms. Firms become isolated when they are, through an unlucky movement following an innovation, relocated out of reach of other firms and so can form no partnerships. It is difficult for them to reintegrate because they do not innovate themselves, and since the regime is incremental no other firms are relocated near to them. In the intermediate regime ($\theta = .25$) the pattern is very different: there are very few isolated firms, and the industry consists of one single giant connected component. In the disruptive regime ($\theta = 1.0$), a distribution of connected components emerges and the network is far less dense than in the other regimes. Innovations are potentially disruptive here, so firms may be moved a long distance in knowledge space following an innovation. They can be temporarily isolated, but following a highly disruptive innovations, there is a reasonably good possibility that other firms will be relocated to their neighborhoods, and they can again form partnerships.[5] In all three networks, the small world ratio (re-scaled clustering over re-scaled path length) is significantly greater than one, except for the incremental regime.

[5]Observing the network evolve over time, we do see that firms break off from the large component, form smaller components, sometimes later rejoining the large component. Membership in the large connected component and the smaller components is relatively fluid in a disruptive regime.

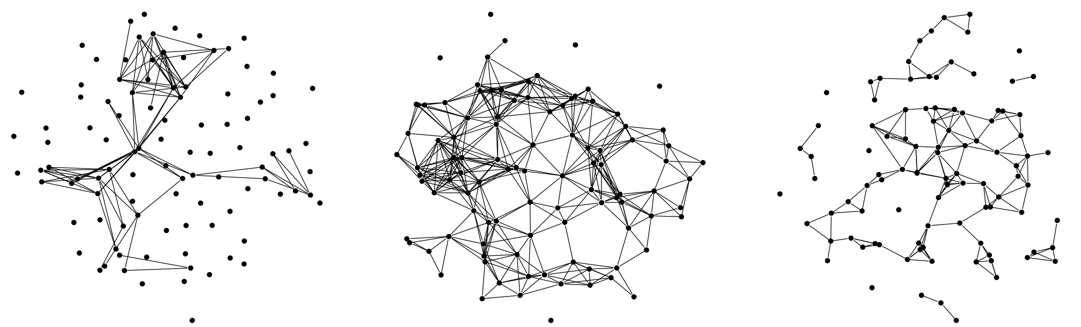

	$\theta = 0.08$	$\theta = 0.25$	$\theta = 1.0$
Av.Deg.	0.845	3.96	1.72
Trans.	0.16	0.54	0.55
APL	3.61	3.76	2.89
Reach	0.06	0.32	0.13
SWR	32.0	4.34	13.7
AKD	1.21	2.21	3.92

Figure 1: Three representative networks for different values of disruptiveness, each in period 1000. Firms are plotted in knowledge space. We show statistics for Average Degree (Av.Deg.), Transitivity (Trans.), Average path length (APL), Reach, Small World Ratio (SWR), and Average distance in knowledge space between pairs of firms (AKD)

5.2 Network structures

Next we examine in some detail the networks that emerge in the model, to assess the extent to which the model replicates the properties of 'real world' networks. Because there is considerable randomness in the dynamics, we mimic a standard empirical strategy, aggregating a set of one-period networks to build a more representative, weighted network. The weight of an edge in this cumulative network is the number of times this edge (which represents an alliance) was active over the set of periods we consider. This weighted network captures the structural difference between an alliance that lasts a single period, and one that is sustained (or repeated) over many periods. In the figures that follow we display several standard descriptive network statistics computed from the weighted network based on the final 500

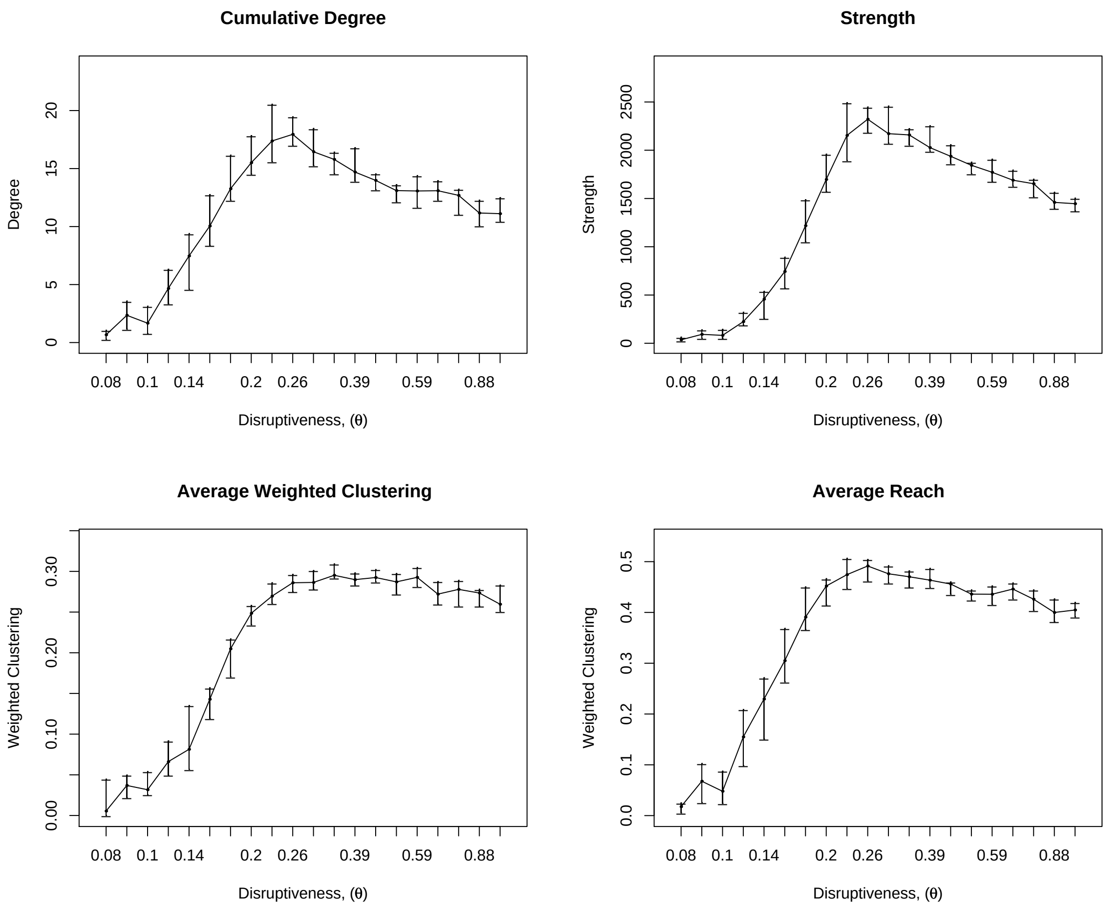

Figure 2: Network statistics for the final 500 periods, aggregated into one network

periods of the simulations.

Figure 2 presents whisker plots (showing median, first, and third quartiles) of degree (i.e., number of distinct partners over history), strength (i.e., number of alliances formed by the average firm over its history), clustering (or transitivity; i.e., proportion of alliance partners who are also partners), and average reach (sum of inverse network distances to all other firms) in the weighted network, as functions of the disruptiveness of the innovation regime θ.[6]

[6]Reach is computed only relative to firms within the same component as the focal firm.

Disruptiveness is non-monotonically related with each of the network statistics. The number of partnerships formed by the average firm (Strength), and the number of different partners it has over its history (Degree) both rise and then fall with θ. This is driven by two antagonistic effects of disruptiveness. As we move from less to more disruptive regimes (increasing θ), any innovation has the potential of causing larger relocations on non-innovating firms, all else equal (the destination disk of a dislocated firm has a radius that increases with θ). This implies that any alliance, whatever the probability of success, has a higher expected value when θ is higher (see Equation 4). This static effect will increase networking. At the same time, because dislocations can be larger, over time firms tend to spread out more in the knowledge space. As this spread takes place, more and more firm pairs are too far apart to have reasonable probabilities of success, and firms are also in general too far apart for dislocations to be large, resulting in lower expected value for any alliance. This dynamic effect will tend to reduce networking over time, and more rapidly so in more disruptive innovation regimes.

What we observe in Figure 2 is that for small θ (incremental innovation regimes) the first effect dominates, and for larger θ the second dominates. There is a narrow interval, roughly for $0.2 \leq \theta \leq 0.3$, where the two effects cancel out and both the number of partnerships formed by and number of different partners of the average firm peak. The same effects are at work on average weighted reach (or closeness), which correlates strongly with degree in networks in general, and thus follows a pattern very similar to strength as well.

Also of interest with respect to strength and degree is the extent to which firms revisit old partners, a common finding in empirical studies (e.g., Gulati, 1995; Gulati and Gargiulo, 1999). We can gauge the prevalence of such behavior by considering the ratio of strength to degree. Although this ratio varies somewhat with the disruptiveness of the innovation regime, over the final 500 periods of the simulations used to compute the statistics, the ratio of strength to degree is roughly 100 to 1. In other words, the average firm allies with each of its partners roughly 100 times over the final 500 periods. This results as the need for proximity in the knowledge space induces some amount of local correlation in firms' decisions to ally, yielding repetition in firms' partnering decisions.

Average weighted clustering (or transitivity) follows the same pattern of rising and then falling with θ. The explanation is again in terms of static and dynamic effects of θ. When θ increases the total dislocation that a given firm can expect to cause also increases, and so it will be willing to form more partnerships, at distances that can deviate from the optimal distance by larger amounts. But this is also true for any of that firm's potential partners. So each firm forms more alliances, and so will its partners, including alliances with other partners of the focal firm. This will increase the proportion of a firm's partners who are also partners. However, as noted, a larger θ implies that over time firms spread out more in knowledge space. Pairwise distances on average increase faster with larger θ, lowering the expected value of any partnership, and in turn lowering clustering. The interaction of these two effects explains the observed non-monotonic relationship with θ.

Empirical studies also indicate that R&D alliance networks typically display small world characteristics (e.g., Powell et al., 2005). Figure 3 shows the weighted small world ratio, which is the ratio of rescaled weighted clustering to rescaled weighted path length, again aggregating all the alliance activity of the final 500 periods into one network.[7] What we

[7]In an unweighted network, the shortest path length between any two nodes is the number of steps on the shortest path connecting them. In a weighted network, weights can be treated as costs, and the shortest path between two nodes the least-cost path. In our case, edge weights are frequencies of interaction, which we transform into costs by taking reciprocals. For the network in aggregate, the average least-cost (shortest) path length is the mean of the least cost (shortest) paths over all node-pairs.
Because path length and clustering both respond mechanically to changes in network density, to examine small world properties, network statistics must be "re-scaled". That is, average path length and clustering are normalized by the same statistics in an equivalent random network. "Equivalent" here refers to a network

observe in the figure is that small world networks emerge in our model: the small world ratio is always larger than one, most strongly so when the innovation regime is highly disruptive.

Such excess clustering is to be expected from the model because alliance formation is dictated by distance considerations in a metric space. If A is at the right distance from both B and C, then the likelihood that B is at the right distance from C is higher than the likelihood that two random points in that space are at the right distance. Thus the networks we observe should be more clustered (i.e., contain more transitive triples) than equivalent (in the sense of having the same number of firms and alliances) random networks. In the model, small world networks emerge from the conjunction of randomness in the innovative process and the quest for profitable partners: Alliances bring proximate firms together, creating cliques; innovation bounces firms apart and can create clique-spanning ties. Neither force dominates, so small worlds emerge.

Overall then, the emergent networks produced by our simulation model display the conduct (repeated ties and transitivity) and properties (clustering, reach, and short distances) characteristic of observed, 'real world' alliance networks.

5.3 Network positions and performance

There are two ways to think about the relationship between network position and performance. In the first, the locus is the innovation: what determines the magnitude of an innovation? This is the pooled cross-section approach. In the second, the locus of analysis is the firm: what determines the magnitude of a firm's innovative performance? This is the

having the same number of nodes and density (in our case an Erdös-Renyi random network is appropriate). Specifically, for each network in our data we generate 100 E-R random networks of the same size and density, and distribution of edge weights. We average the average least-cost path length and clustering of those 100 graphs to create "equivalent" statistics which are used to normalize average path length and clustering in the networks in our data.

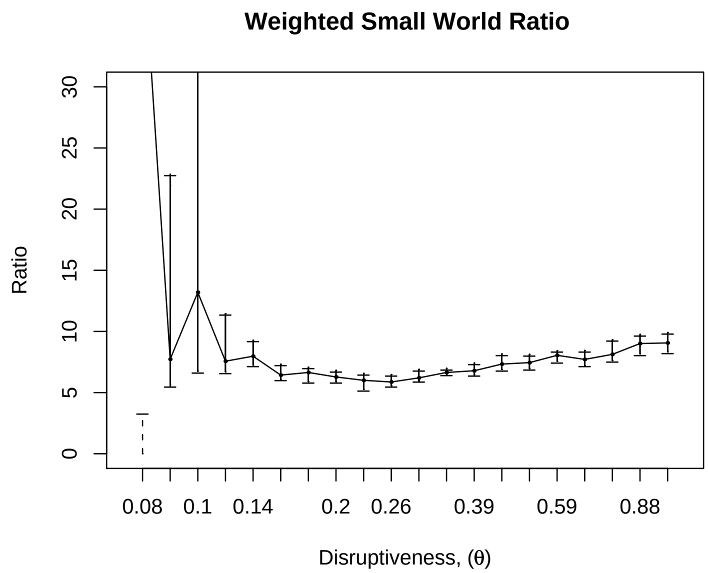

Figure 3: Small world ratio versus disruptiveness. (Median and inner quartiles, over 48 observations for each value of θ.)

within-firm, over time approach, which is relevant for strategic prescriptions. We ask both questions and examine whether or not the answers are consistent, and thus whether or not evidence of network effects in pooled cross-section data supports prescriptions for network strategy over time.

5.3.1 Pooled cross-section

What is the relationship between the size of an innovation at time $t + 1$ and the network position of the innovating firm at time t? The size of innovation is measured by the dislocation costs that a successful alliance imposes on the other firms in the industry. To measure the closure of a firm's network position, we use clustering; to measure its openness, we use betweenness centrality. [8]

Each time an innovation occurs, we collect information on the successful firm's position in the innovation network at the start that period (i.e., time t) and the total dislocation imposed by its innovation on other firms during that period (i.e., $t+1$). For completeness and comparison purposes, we also collect information on degree and reach. Pooling these observations, we compute rank correlation coefficients to assess the relationship between firm's network position and innovative performance in cross-section. These are shown in Figure 4. [9]

Innovations that impose a large relocation on other firms are always associated with high degree and high reach — the correlations between innovation size and both the degree of the innovating firm and its closeness to other firms are positive across all values of disruptiveness. The correlations increase quickly with initial increases in θ and then flatten out or slowly decrease with further increases. The two correlations respond to changes in θ in almost the same way. This is because closeness is a monotonic (inverse) transform of the distances of

[8] The neighborhood *clustering* of firm i is the proportion of neighbors of i who are neighbors of each other. It is written

$$c_i^g = \frac{\#\{jk \in g : j, k \in N_i^g\}}{n_i^g(n_i^g - 1)/2}.$$

The *clustering coefficient* $c^g = \sum_i c_i^g/n$ is the average taken over all the firms.
Betweenness centrality of firm i, b_i^g, is the sum, over all possible pairs $k, l \in N - \{i\}$, of the proportion $p_{k,i,l}^g$ of shortest paths between k and l that run through i, i.e.,

$$b_i^g = \sum_{k,l \neq i} p_{k,i,l}^g.$$

[9] We thus focus on network effects among innovating firms. In doing so, we over-sample on success, as only observations in which some strictly positive amount of dislocation is imposed on the rest of the industry are considered. Including non-innovating firms in the sample adds a large number of zero observations. Qualitatively the patterns remain, but the addition of the zeros attenuates the correlations significantly. Consequently we restrict attention to innovating firms. Although this removes from consideration the innovate-or-not aspect of alliances, it maintains relationships between size of innovations, or profits, and network structures, which are in themselves interesting, and the relationship between "size of success" and disruptiveness is qualitatively the same whether or not non-innovators are included.

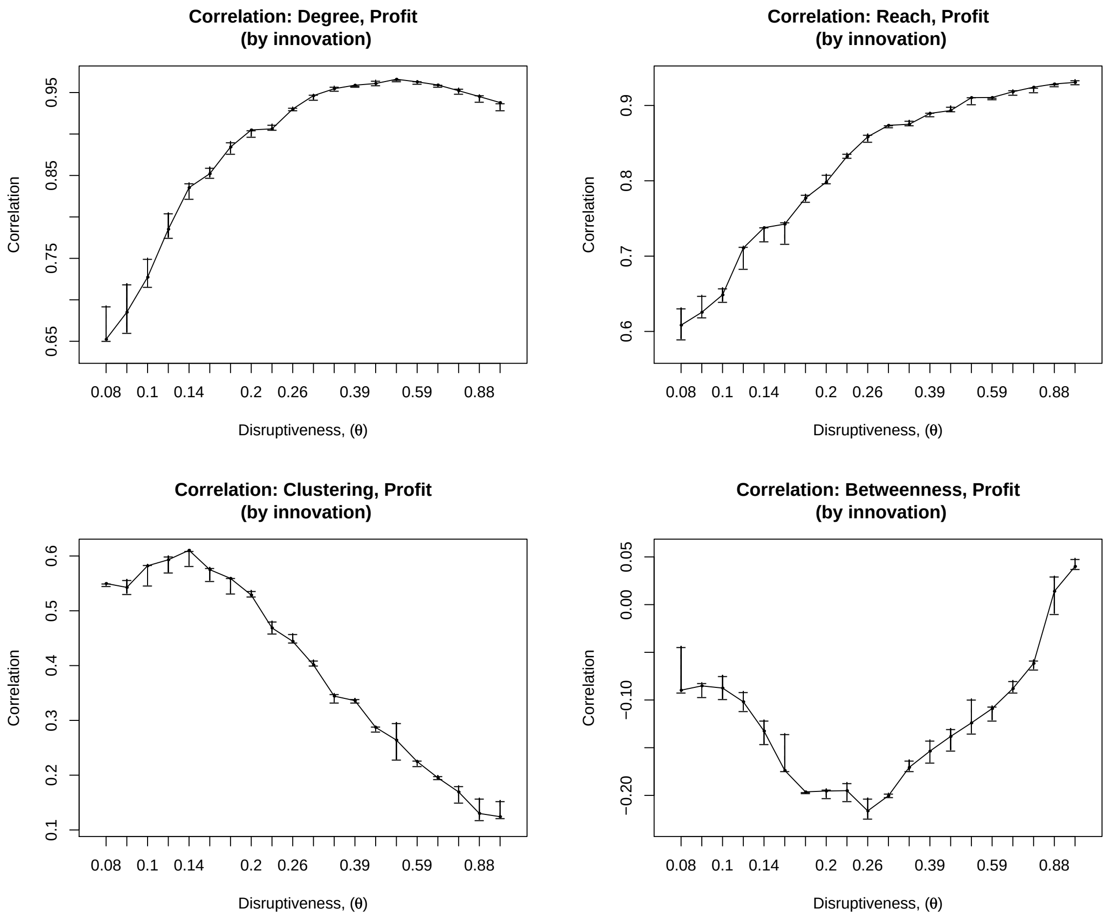

Figure 4: Correlations between one-period profits and ego-network characteristics of innovating firms. (One-period profits are defined as the total dislocation caused by the innovation minus the cost of all alliances the innovating firm participates in in that period.)

the focal firm to all other firms in network space, and as such depends strongly on degree. The correlation coefficient of degree and closeness exceeds .80 in all our experiments; thus, the degree of a firm explains to a very large extent its proximity to other firms in the network (the firm's reach).

Like reach, firm profit (total dislocation minus cost) is a monotonic (inverse) transform of the distances of the focal firm to all other firms, but this time distances are in knowledge space. The observed positive correlations thus reflect a positive association between proximity in knowledge space and proximity in network space. As alliances form based on considerations of distances in knowledge space (both the success probability of an alliance, and the dislocation that firms participating in the alliance can expect to impose on other firms depend on distances in knowledge space), the observed positive associations are expected.[10] The strength of the association, of course, depends on θ. When θ is small, the dislocation a firm can expect to impose on other firms is always small, and the success probability plays an essential role. Success probability, though, is maximized when the partners are at a distance of 0.5 of each other, a distance at which there is in effect no more dislocation(when θ is very small, $\exp(-0.5/\theta) \approx 0$). Thus the association between degree (or reach) and profit will be weak for small θ. As θ increases, a lower probability of success is compensated by a larger expected dislocation. So firms will increasingly form alliances at distances that deviate significantly from the optimal one. In addition, dislocations extend further, and thus a firm's innovation can also afect its own alliance partners. This results in a stronger positive correlation between degree (or reach) and profits. When θ becomes large, firms spread out widely, the success probability of any alliance declines, and the association between degree (or reach) and profits weakens somewhat.

Clustering is strongly positively correlated with one-period profits when θ is small, but falls rapidly as θ grows. To understand this, consider an extreme: a very low θ, which implies that dislocations will be small, demands that for an alliance to have positive expected profit it must be at the optimal distance in knowledge space. Here a firm's partners will be located almost strictly on a circle centered on the focal firm, with radius equal to the optimal distance. If a

[10] More specifically, the dislocation that a firm participating to an alliance can expect to impose on other firms falls as the distance to each other firm increases. As per the success probability, it is only about being at the "right" distance.

firm has few partners, then the probability they are at the optimal distance from each other is very small. Thus low degree will be correlated with low clustering. To have a reasonable probability that (some of) a firm's partners are at the optimal distance from each other, that firm must have many partners. Thus if a firm has high clustering, it has many nearby firms. These are the firms that can be dislocated if it innovates. High clustering firms will have high dislocations, and so high profits. As disruptiveness increases, firms become less stringent about optimal distance, there is more tolerance with respect to the optimal distance in terms of alliance profitability, and the "circle" becomes a fatter and fatter doughnut of feasible distances. As this happens, the effect just described is attenuated, and the correlation between clustering and profit falls.

Betweenness shows a non-monotonic pattern. Correlations are negative and falling, then rising and becoming slightly positive. Low betweenness is associated with essentially two types of firms: those having many triangles in their neighborhoods (thus many shortest paths between other pairs of firms do not run through them) and loose-ends. In an incremental regime (small θ), based on the above discussion of clustering, having many triangles is associated with having large degree and thus large dislocations imposed on others, so betweenness is not creating an advantage. It is even less so when θ initially increases. However, after some point the correlation of clustering and degree has fallen enough, and the negative association of betweenness and performance weakens, eventually reaching a point where the correlation is actually positive, pointing at the value of open positions.

Thus, consistent with recent theory and empirical evidence of the contingent benefits of network positions in pooled cross-section (e.g., Gilsing et al., 2008; Rowley et al., 2000), firms in the model tend to benefit more from closed network positions in incremental innovation regimes, and more from open network positions in disruptive innovation regimes.

5.3.2 Within firm, over time

Rather than focusing on the effects of firms' network positions at time t on their performance at time $t+1$, consider instead the effects of firms' network strategies and performance over time. Firms may, for example maintain positions within dense cliques, or deliberately seek brokering positions between otherwise disconnected firms, or they may simply seek to form many alliances. In our model, of course, firms do none of these things: they simply seek knowledge complementarities. As a result of this pursuit, combined with their locations in the knowledge space, however, firms vary in their network positions, and these variations can persist over time.

Here we are therefore interested in what might be conceived as firms' "network strategies". From a strategic standpoint, a firm following a "closure strategy" , for example, would be expected to maintain a level of clustering over time that is, on average, higher than other firms in the network. Evidence of a closure strategy would thus be higher than average clustering over an extended period of time. To assess the performance of such a strategy, the relevant indicator would be average firm performance over the same period. Within this general framework, we can examine the long run correlations between firms' "network strategies" and performance by examining firms' average network positions and performance over an extended period of time. For each firm, we therefore calculated average degree, reach, clustering, and betweenness, and total profits over the firm's history (i.e., 2500 periods, with the initial 100 periods discarded as before). We then computed rank correlations between each firm's average performance and network position for the population of firms. These long-run, within-firm correlations are presented in Figure 5, across the range of disruptiveness, and superimposed on the cross-sectional correlations from Figure 4 to facilitate comparison.

As the figure shows, "network strategies", as we have defined them, and long run firm performance are indeed correlated, and for some ranges of disruptiveness the correlations are strong – in both positive and negative directions. Importantly, the long-run correlations contrast sharply with the cross-section correlations.

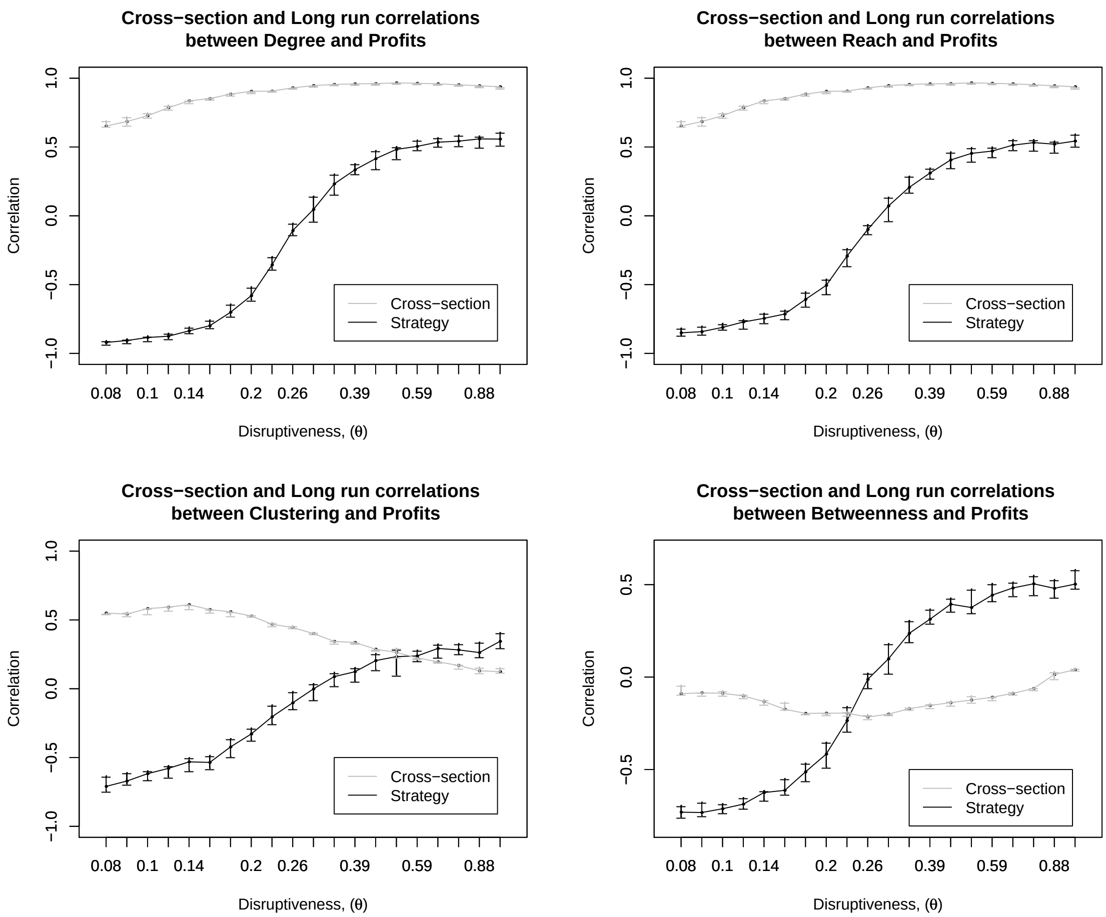

Figure 5: Comparison of cross-section (grey) and long-run (black) correlations between innovation performance and network position.

All four network attributes share a similar pattern: long-run, within-firm correlations are markedly negative for small θ (incremental innovation regimes) and markedly positive for large θ (disruptive innovation regimes). When θ is large pooled cross-section and long-run correlations tend to agree, but they strongly disagree when θ is small.

To understand this divide, note first that, for any level of disruptiveness (i.e., any value of θ), firms are making mistakes in their alliances. The positive-expected-profit principle firms follow when deciding on which alliances to form is applied in a non-stationary environment. Movement of firms in the knowledge space creates (possibly large) differences between expected and realized profits. The dynamic effect of θ is to spread firms out in the knowledge space faster for larger θ; the scope of firms' mistakes thus increases with θ since the difference between tomorrow's environment and today's grows with θ. At the same time, the static effect of θ increases the expected dislocation — as well as the likelihood of larger realized dislocations — resulting from any alliance.

As an alliance costs c in any innovation regime, for small θ, alliances will thus be marginal in the sense that their expected profit (expected dislocation minus cost c) is positive but small, as will be realized profits. For larger θ, alliances yield greater expected profits, as expected dislocation increases while cost remains constant, and so larger realized profits will also tend to result. Consequently, as θ increases, a given value of degree yields higher expected profits, and more frequently realizes greater profits.

Reach and betweenness correlate strongly with degree, and the correlations involving them follow patterns which are very similar to the pattern exhibited by degree. Clustering is less-strongly correlated with degree, and shows a weaker pattern. The pattern is ultimately driven by degree, though through a different mechanism. As discussed above, when θ is small, a firm's partners will be located on a circle centered on the focal firm, with a radius equal to the optimal distance. Consequently, only high degree firms will have high clustering. But in such an environment, high degree (i.e., high alliance cost) firms perform poorly over the long run. As θ increases, the connection between clustering and degree weakens, and so the relationship between clustering and profit follows that of degree less strongly.

Returning to the predictions in Section 2, the results support the prediction regarding closed network positions, with the correlations for clustering in pooled cross-section and within-firm over time, increasing and decreasing, respectively, with disruptiveness. The prediction for open network positions is not supported, however, as correlations for betweenness in both pooled cross-section and over time increase with disruptiveness above moderate levels.

6 Discussion and Conclusion

We have observed substantial discrepancies in correlations between network position and performance in pooled cross-section and within-firm, over time, and as a result, prescriptions drawn from the two types of correlations would be very different in some environments. As our discussion of the simulation results suggests, part of the explanation for these discrepancies may lie in problems associated with the use of expected profits as a partner selection criterion in a dynamic environment. But we have also advanced a more profound explanation that resides in the nature of performance benefits firms derive from their network positions. And, in particular, whether network positions alter firm performance through its mean or variance.

Open network positions facilitate broad search for emerging innovations and future options vital to exploratory search (Powell et al., 1996). Closed network positions, yield access to the redundant and validating information essential to exploitive search (Dyer and Singh, 1998). Appropriate network positions thus depend on the degree of uncertainty and rate of change in the environment (Gilsing et al., 2008; Rowley et al., 2000): Firms operating in rapidly changing environments will benefit from open network positions, which aid in exploration, while firms in stable environments will benefit from closed network positions, which aid in exploitation.

This link between network positions and learning suggests, as we observed, that open and closed network positions affect firm performance distinctly, with closed positions reducing firm performance variation, and open positions increasing it. Closed network positions should therefore benefit firms when the distribution of possible outcomes favors low performance variation, and open positions when the distribution favors high performance variation.

Following March (1991), we proposed that right-skewed distributions of possible outcomes favor firms occupying variance-increasing open network positions in the competition for high performance in pooled cross-section. Over time, however, because high variance also increases exposure to poor outcomes, mean performance for firms sustaining open network positions declines with the right-skewness of the performance distribution. Moreover, because variance-reducing closed network positions provide protection against poor outcomes, the performance benefit of sustaining a closed network position over time increases with the right-skewness of the performance distribution.

To examine these predictions, we designed our model to permit variation in the disruptiveness of the innovation regime — from incremental to radical — in order to alter the right-skewness of the distribution of possible performance outcomes. The possibility of disruptive innovations generates a right-skewed distribution of potential performance outcomes in which firms occupying open network positions should achieve higher performance in pooled cross-section. In contrast, if only incremental innovations are possible, variance-reducing closed network positions should produce higher performance in cross-section. These predictions are consistent with arguments and pooled cross-section findings in empirical studies of network effects (e.g., Gilsing et al., 2008; Rowley et al., 2000), and our results support them. Our analysis thus identifies performance variability as a general mechanism that can potentially account for the observed empirical findings in pooled cross-section.

The results also support the within-firm, over time prediction for closed network positions.

Thus, while the correlations for clustering in pooled cross-section decrease with disruptiveness, the within-firm over time correlations for clustering increase with disruptiveness (see Figure 5). Prescriptions regarding closed network positions drawn from the pooled cross-section correlations would thus counsel firms to adopt closed network strategies in incremental innovation regimes, while correlations based on firms' long-run strategies indicate that these prescriptions are exactly opposite to the behavior of firms successful in the long run.

The prediction for open network positions is not supported, however, as correlations for betweenness centrality in both pooled cross-section and over time increase with disruptiveness above moderate levels.

Although we have focused here on network position effects the phenomenon we articulate is more general, applying to any firm characteristic (R&D intensity, for instance) that affects firm performance in pooled cross-section by altering its variance.

We began with the observation that, while intuitively appealing (and common), drawing network strategy implications from empirical evidence of network performance effects in pooled cross-section is not necessarily warranted. As we explained, this is because network positions may influence both the mean *and* variance of firm performance. Although strategic prescriptions are warranted if network effects observed empirically in pooled cross-section reflect increases in *mean* firm performance, if network effects instead reflect increases in firm performance *variance*, such prescriptions are warranted only if the increase in the odds of achieving high performance are sufficient to compensate for the concomitant increase in the odds of realizing poor performance. Our analysis suggests that network effects may indeed reflect changes in firm performance *variance*, and moreover, that the increased odds of achieving high performance may be *insufficient* to compensate for the concomitant increase in the odds of realizing poor performance. This suggests the exercise of caution in drawing implications for network strategies.

References

[1] Ahuja, G. 2000. Collaboration networks, structural holes, and innovation: A longitudinal study. Administrative Science Quarterly, 45: 425-455.

[2] Argote, L. 1999. Organizational Learning: Creating, Retaining and Transferring Knowledge. Boston MA: Kluwer.

[3] Baum J.A.C., Cowan, R. and Jonard, N. 2010. Network-Independent Partner Selection and the Evolution of Innovation Networks, *Management Science*, 56, 2094-2110.

[4] Baum J.A.C., and Ingram P. 2002. Interorganizational learning and network organization: Toward a behavioral theory of the interfirm. In M. Augier, J.G. March (Eds.). The Economics of Choice, Change, and Organization: Essays in Memory of Richard M. Cyert: 191-218. Cheltenham, UK: Edward Elgar.

[5] Burt, R. S. 1992. Structural Holes: The Social Structure of Competition. Cambridge: MA: Harvard University Press.

[6] Burt, R. S. 2000. The Network Structure of Social Capital, Research in Organizational Behavior, Vol. 22: 345-423. New York: JAI-Elsevier Science Inc.

[7] Coleman, J. S. 1988. Social Capital in the Creation of Human-Capital. American Journal of Sociology, 94: S95-S120.

[8] Coleman, J. 1990. Foundations of the Social Theory. Cambridge, MA: Harvard University Press.

[9] Dess, G.G., and Beard, D.W. 1984. Dimensions of organizational task environments. Administrative Science Quarterly, 29: 5273.

[10] Dyer, J.H., and Singh, H. 1998. The Relational View: Cooperative Strategy and Sources of Interorganizational Competitive Advantage. Academy of Management Review, 23(4): 660-679.

[11] Gilsing,V., Nooteboom, B., Vanhaverbeke, W., Duysters, G., and van den Oord, A. 2008. Network embeddedness and the exploration of novel technologies: Technological distance, betweenness centrality and density. Research Policy, 37: 1717-1731.

[12] Granovetter, M. 1973. The strength of weak ties. American Journal of Sociology, 78: 1360-1380.

[13] Gulati, R. 1995. Social structure and alliance formation patterns: A longitudinal analysis. Administrative Science Quarterly, 40: 619-652.

[14] Gulati, R., and Gargiulo, M. 1999. Where do inter-organizational networks come from? American Journal of Sociology, 104: 1439-1493.

[15] Holland, J.H. 1975. Adaptation in Natural and Artificial Systems. Ann Arbor MI: University of Michigan Press.

[16] Lant, T.K., Milliken, F.J., and Batra, B. 1992. The role of managerial learning and interpretation in strategic persistence and reorientation: An empirical exploration. Strategic Management Journal, 13: 585608.

[17] Lavie, D., and Rosenkopf, L. 2006. Balancing exploration and exploitation in alliance formation. Academy of Management Journal, 49: 797-818.

[18] Levinthal, D.A., and March, J.G. 1993. The myopia of learning. Strategic Management Journal, 14 (Winter Special Issue): 95-112.

[19] March, J. G. 1991. Exploration and exploitation in organizational learning. Organization Science, 2: 71-87.

[20] McEvily, B., and Zaheer, A. 1999. Bridging ties: A source of firm heterogeneity in competitive capabilities. Strategic Management Journal, 20: 1133-1156.

[21] Mitsuhashi, H. 2003. Effects of the social origins of alliances on alliance performance. Organization Studies, 24: 321–339.

[22] Powell, W. W., White, D., Koput, K., and Owen-Smith, J. 2005. Network dynamics and field evolution: The growth of interorganizational collaboration in the life sciences. American Journal of Sociology, 110: 1132-1205.

[23] Rosenkopf, L. and Nerkar, A. 2001. Beyond local search: Boundary-spanning, exploration, and impact in the optical disk industry. Strategic Management Journal, 22: 287-306.

[24] Rosenkopf, L., and Almeida, P. 2003. Overcoming local search through alliances and mobility. Management Science, 49: 751-766.

[25] Rothaermel, F., and Deeds, D.L. 2004. Exploration and exploitation alliances in biotechnology: A system of new product development. Strategic Management Journal, 25: 201-221.

[26] Rowley, T., Behrens, D., and Krackhardt, D. 2000. Redundant governance structures: An analysis of structural and relational embeddedness in the steel and semiconductor industry. Strategic Management Journal, 21: 369-386.

[27] Uzzi, B. 1996. The sources and consequences of embeddedness for the economic performance of organizations: The network effect. American Sociological Review, 61: 674-698.

[28] Uzzi, B. 1997. Social structure and competition in interfirm networks: The paradox of embeddedness. Administrative Science Quarterly, 42: 35-67.

[29] Van de Ven, A. H. 1976. On the nature, formation and maintenance of relations among organizations. Academy of Management Review, 1: 24-36.

[30] Walker, G., Kogut, B., and Shan, W. 1997. Social capital, structural holes and the formation of an industry network. Organization Science, 8: 109-125.

[31] Yelle, L. E. 1979. The learning curve: Historical review and comprehensive survey. Decision Sciences, 10: 302-328.