SMALL WORLDS IN CONTEXT: HOW GENERALIZABLE IS INTERORGANIZATIONAL NETWORK STRUCTURE?

Anindya Ghosh
Department of Organization and Strategy
School of Economics and Management
Tilburg University
5000 LE Tilburg
The Netherlands
ghoshmosai@gmail.com

Lori Rosenkopf
The Wharton School, University of Pennsylvania
3018 Steinberg Hall – Dietrich Hall
Philadelphia, PA 19104
(215) 898-6723
rosenkopf@wharton.upenn.edu

December 29, 2018

Keywords: small world, strategic networks, replication, comparative network structure, affiliation networks

We appreciate support from the Mack Institute for Technological Innovation at The Wharton School, comments from seminar participants at University of North Carolina-Chapel Hill and Bocconi University, and feedback from the DRUID18 conference.

SMALL WORLDS IN CONTEXT: HOW GENERALIZABLE IS INTERORGANIZATIONAL NETWORK STRUCTURE?

ABSTRACT

Our field's growing attention to interorganizational network structure frequently builds on the Watts and Strogatz (1998) small world model. Our literature has identified "small worlds" -- actual networks which simultaneously obtain relatively high clustering and short path length – in many contexts. Yet as we illustrate, these approaches use incommensurable methods and overgeneralized data, while yielding equivocal findings of structure on performance. To address these issues, we discuss underlying characteristics of primary network data and encourage researchers to detail their data and measures in order to facilitate comparative assessments. Further, we demonstrate a more rigorous approach, the "small world omega" statistic (Telesford *et al.*, 2011), which may serve as a basis for comparing small world networks more systematically. We close with suggestions for ongoing research.

Introduction

The long-held organization-theoretic belief that social structures shape the behaviors and outcomes of social actors has gained purchase in a host of strategic management studies of network effects on firm conduct (Baum, Calabrese, & Silverman, 2000; Gulati, Nohria, & Zaheer, 2000; Phelps, Heidl, & Wadhwa, 2012; Rosenkopf, Metiu, & George, 2001). Yet the wide contextual range of these studies, across not only varied industries and timeframes but also varied types of network ties and empirical approaches, yields findings that may not be commensurable or generalizable (Shi, Sorenson, & Waguespack, 2017).

As an example, consider how one popular network structure, the "small world," affects performance. In a study of alliance networks in 11 high-tech industries over 8 years, Schilling and Phelps (2007) demonstrated that the combination of network clustering and reach (the two essential characteristics of small worlds) was associated with higher firm innovative productivity, suggesting a *linear* relationship between "small-worldliness" and performance. In contrast, in a study of Broadway musical collaboration networks from 1945 to 1989, Uzzi and Spiro (2005) demonstrated that the small-worldliness of the artists' network has a *curvilinear* relationship with box-office receipts, which they attribute to the desirability of blending novelty and routines. Another contrast arises from a study of inventor collaboration networks across 337 metropolitan statistical areas over 21 years, where Fleming, King Iii, and Juda (2007) find *no clear effect* of small-worldliness on regional innovativeness, though they draw an interesting contrast between the network structures of Silicon Valley and Boston. Does this mean that the relationship between small-worldliness and performance differs for different types of networks? Or might it mean that the measurement of small worlds differs across network types, or even across researchers? How can we distinguish these issues?

To address these questions, we contend that further progress on the relation between network structure and performance requires some rethinking of our ideas about modeling, measuring and comparing network structures. We review the theory of small worlds and its application in strategic management, and then offer two divergent approaches that can help sort the variety of extant findings about small worlds. First, we argue that most network analyses in strategic management research are derived from what have been termed "affiliation networks" (Wasserman & Faust, 1994), and we differentiate affiliation networks by considering both the frequency of affiliative events as well as the number of participants in a typical affiliation. With this approach established, we then address the limitations of current small world measures and introduce a new network statistic ("omega") borrowed from brain science so that we can assess the variability in traditional small world network studies systematically. Extrapolating from these small world findings, we conclude with recommendations for more uniform approaches that would facilitate comparative work across, and generalization within, each network type.

The Watts-Strogatz Small World Model

While earlier work by social scientists studied the so-called "small world" problem, that is, how long are acquaintance chains between seemingly well-separated others (Milgram, 1967; Pool & Kochen, 1978), it was several decades later before network scientists formalized a model which might comport with the "six degrees of separation" concept. Watts and Strogatz (1998), heretofore "W-S", formalized the small world model as a regular (ring lattice) network with nominal rewiring to increase disorder, yielding networks characterized by both a short average path length and a high clustering coefficient simultaneously[1]. This was an important

[1] Mathematically, path length is the average distance between all pairs of vertices in the network (assuming a connected network where all nodes are reachable from each other), where distance between two vertices is the length of the shortest path between them. Watts and Strogatz define the clustering coefficient this way: "Suppose that a vertex

contribution to the graph-theoretic literature, as two key reference classes of graphs had each demonstrated one of these characteristics without the other. Specifically, Erdos-Renyi random graphs tended to exhibit short average path length, but did not have the clustering typical of social networks. In contrast, ring lattice graphs exhibited clustering, but had high path length. As a result, the W-S model represents a hybrid of these two types, and has received substantial attention across a wide variety of literatures.

As seen in Figure 1, in order to construct network candidates for small worlds, the Watts-Strogatz model starts with a lattice (characterized by n actors with k links to nearest neighbors) and then randomly rewires the links between nodes with a certain probability (p). When a link is selected for rewiring, one node of the dyad is held constant, while the other is randomly assigned to a new destination to preserve connectivity among all the nodes. When p=0, the network is a lattice, while for p=1, since all links are randomly rewired, the network generated is an Erdos-Renyi random graph[2]. Only for some moderately low values of p does the small world network emerge.

Insert Figure 1 about here

To understand these dynamics more fully, Figure 2 demonstrates the transitions between lattice, small world, and random networks for the example of n=1000 and k=20 by plotting a

---

v has k_v neighbors; then at most $k_v(k_v - 1)/2$ edges can exist between them (this occurs when every neighbor of *v* is connected to every other neighbor of *v*). Let C_v denote the fraction of these allowable edges that actually exist. Define *C* as the average of C_v over all *v*."

[2] The Erdos-Renyi graph is one particular type of many possible random graphs, where all edges are equally likely to be included, yielding a Poisson degree distribution. It is worth noting that different types of random graphs obtain if the random rewiring follows a different rule such as rewiring with a probability proportional to the degree of the node, which would result in a scale-free network.

normalized[3] clustering coefficient and average path length against log(p). When moving from left to right along the x-axis, we observe that there is a region in which the clustering coefficient remains high while the average path length decreases due to the shortcuts generated by the rewiring. As log p gets larger (that is, p approaches 1), the W-S networks approximate Erdos-Renyi random networks, and clustering drops significantly. Thus, small world networks occur in the range of p between ring lattices and random networks where the clustering coefficient remains high while the path length has dropped.

Insert Figure 2 about here

Ultimately, for a given n and k, varying p yields a family of W-S networks. The specific value of p determines whether this network is best classified as a lattice, a small world, or a random network. When n>>k (a key condition for all small world studies), then small worlds arise at relatively low values of p. In other words, only a small fraction of shortcuts are need to dramatically decrease path length while still preserving clustering. At the same time, it is important to note that while the lattice (p=0) has a uniform degree distribution (all nodes with degree k), as p grows and tends to 1, the rewired network will display a Poisson degree distribution (Barrat & Weigt, 2000). Said differently, at any value of p, the W-S network, by its construction, has little if any skewness in its degree distribution. Indeed, in a follow-up paper focused on social networks, Watts (1999) explicitly notes that a small world network is decentralized, as it would not be surprising for a centralized network to have short path length.

[3] Following Watts & Strogatz (1998) the average path length and clustering coefficients are normalized by the corresponding values for an equivalent ring lattice (i.e, for the network with p=0) since both path length and clustering are maximized at this value.

Application of the W-S Model in Management

Within the strategy and organization theory fields, the W-S model has captivated many researchers because this network structure seems better suited to efficient search and diffusion than other well-established networks such as a lattice or a randomly generated graph. The imagery of clusters and shortcuts also strikes a chord with our theorizing on social capital formation and exploratory search (Fleming *et al.*, 2007; Fleming & Marx, 2006; Rosenkopf & Padula, 2008; Schilling & Phelps, 2007; Verspagen & Duysters, 2004).

As a result, research examining small worlds has flourished over the past two decades in the strategy and organization domains, and can be grouped into three categories. First, scholars have identified small world structures from actual network data (e.g. Baum, Shipilov, & Rowley, 2003; Davis, Yoo, & Baker, 2003; Kogut & Walker, 2001; Verspagen & Duysters, 2004). Second, scholars have used a variety of measures to incorporate characteristics of network clusters and shortcuts to denote small world features in regression analyses, examining both the evolution of small worlds (e.g. Baum *et al.*, 2003; Gulati, Sytch, & Tatarynowicz, 2012; Rosenkopf & Padula, 2008) as well as the effect of "small-worldliness" on performance within and across networks (e.g. Fleming *et al.*, 2007; Schilling & Phelps, 2007; Uzzi & Spiro, 2005). Finally, scholars have used simulations to explore how processes of learning and diffusion occur on the W-S substrate across a range of parameters for n, k, and p (e.g. Fang, Lee, & Schilling, 2010; Rivkin & Siggelkow, 2007).

To assess and measure small-worldliness of networks, researchers in our field have utilized a family of measures which derive from the original words of Watts and Strogatz (1998: 440): "We find that these systems [small worlds] can be highly clustered, like regular lattices, yet have small characteristic path lengths, like random graphs." Indeed, they calculate clustering (C) and path length (L) measures for three actual networks (C. elegans, power grid, and actors)[4] and assert, "All three networks show the small-world phenomenon: $L \approx L_{random}$ but $C >> C_{random}$" (Watts & Strogatz, 1998 ; Table 1, pg. 441).

Building on this W-S approach, Kogut and Walker (2001) summarized these descriptive characteristics into a single statistic labeled the "Actual-to-Random Ratio for Length/Clustering"[5] (Table 2, pg. 325) with the intent of facilitating comparisons across networks. To clarify, this ratio is given by the formula $(C/C_r) / (L/L_r)$, where $C_r = k/n$ and $L_r = \ln(n)/\ln(k)$ are approximations of the average clustering coefficient and path length for an equivalent Erdos-Renyi graph with n nodes and average degree of k and n >> k (Watts & Strogatz, 1998)[6]. Thus, a network with high clustering (relative to the equivalent Erdos-Renyi random graph where clustering is low) and path length similar to that of a random network (which is low relative to that of the lattice) should manifest a ratio much greater than 1. After calculating that the three W-S networks obtained said ratios of 4.75 (C. elegans), 10.61 (power grid), and 2396.90 (film actors), they asserted that the German firms and German owners networks also represented small worlds due to their ratios of 22.46 and 100.48, respectively.

This ratio-based approach became commonplace, though many papers introduce variations in the labeling as well as the calculation and interpretation of the ratio. Davis et al. (2003) termed it the "SW quotient" in their analyses of director interlock and board interlock networks, identifying small worlds by seeking SW quotients "substantially greater than one" (p.313). Over a variety of temporal windows and network populations, they found SW quotients

[4] Formally, the statistics are computed for the main component (that is, the largest connected component) of the full network.
[5] This label is misleading, as the reported values are clearly the actual to random ratio for clustering/length.
[6] Rearranging terms, this statistic is given by the actual clustering/actual length ratio (C/L) normalized via multiplication by n(ln(n))/k(ln(k)).

between 4.55 and 11.87 for their moderately sized board interlock networks (n ranging from 177 to 581), and between 67.23 and 186.62 for their larger director interlock networks (n ranging from 1819 to 5853)[7]. Baum et al (2003) used the label "SW" for this value, noted a range of values obtained for this ratio over several studies including those discussed above, acknowledged that SW would vary with network size n, and suggested a threshold of SW>4 to identify small worlds in their networks of average size n=87 (p.708). Uzzi and Spiro (2005) utilized the Newman, Strogatz, and Watts (2001) bipartite correction to construct a similar statistic for a small world quotient, which they labeled "Q"[8]. Fleming et al (2007) constructed their small world measure with a different null model, normalizing clustering and path length from that expected in a ring lattice network, which yields lower values for their measure than what would obtain with the Erdos-Renyi null model. Kilduff *et al.* (2008) suggested that the appropriate small world threshold should be 4.75 (which we assume corresponds to the C. Elegans ratio).

Recent research has utilized still lower small world thresholds. For example, Baum, Cowan, and Jonard (2010) used a threshold of 2 in their simulation studies. Finally, Gulati *et al.* (2012: 467-468) reported in their literature review that networks are characterized as small-worlds for values ranging from 2 to 60. They analyzed computer industry networks, which varied in size from 374 to 6221 over the span of their data. Due to this wide variation, they normalized the small world quotient[9] by dividing the clustering coefficient ratio (C/C_r) by n and multiplying the path length ratio (L/L_r) by ln(n). The putative rationale for this approach is to remove size-

[7] Recall that the board interlock and director interlock networks are derived from the same affiliation data so the comparison of SW quotient values is less likely to be conflated by variations in n and k.

8 Note that the bipartite correction, developed for application to affiliation data, also controls for the actual degree distribution.of the network. In other words, the random graph used as the null model has the degree distribution fixed to that of the network being evaluated.

[9] While they term this quotient "Q", they do not use the fixed degree distribution as in Uzzi and Spiro (2005). Therefore in their case Q = (nC/k)/(ln(k)L/ln(n). The adjusted Q reduces to (C/k)/(ln(k)L), i.e., (C/L)/(ln(k)/k).

driven variation in the small world quotient, but no justification for this choice of normalization coefficients is provided.

Thus, as summarized in Table 1, small-world statistics and inferential methods in our field have varied with respect to null models, formulas, and inference without converging on a consensus approach. Without consensus, it is challenging to generalize findings, particularly when results conflict across studies. Of course, while this variation in approaches has prevented straightforward comparisons across different networks, examining the benefits and pitfalls of extant small world methods provides an opportunity to improve on them. Next we suggest two ways to reconsider studies of network structure, particularly in the small world scenario.

Insert Table 1 about here

Suggestion 1: Revisiting Network Data in Strategic Management

The vast majority of interorganizational networks analyzed in strategic management use archival data. Network ties in such data are usually derived from affiliations; for example, in alliance data, the actors, which are firms, are affiliated with events, which are alliances (Schilling, 2009). It is key to note that the primary network data, therefore, generates a rectangular matrix composed of actors and events; that is, a bipartite (2-mode) affiliation network (Wasserman & Faust, 1994). Yet most extant work in our field immediately transforms this bipartite structure into unipartite (1-mode) dyadic data representing ties between each pair of actors in each alliance. In this transformed data, the network tie between firms i and j is equal to the number of alliances in which both firms participate, and frequently binarized to indicate whether the two firms jointly participated or not[10].

[10] Recognize, however, that the transformation to dyads eliminates isolates from the analysis (such as firms not participating in alliances; firms with unique board members, and solo inventors on patents).

One virtue of the transformation of rectangular affiliation matrices to dyadic ties between actors for network analyses is that dyadic matrices of ties between actors lend themselves to small world studies and thus have the potential to enable generalizability across many types of social networks. As we will show, this secondary data has become the common currency in network studies; the underlying characteristics of the primary network data are not always fully reported, and therefore any putative generalizability may mask different underlying characteristics of the primary affiliation data. Therefore, in this section, we will examine how the size of the affiliation groups, the frequency of affiliation events, and the number of actors affiliating can vary across common network data contexts.

How big are the events? For a given event type, a key underlying property is the distribution of the number of actors participating in an event. We can begin this exploration by merely focusing on the mean of this distribution. Strategic alliance data typically reveal mostly dyadic relationships, with occasional larger ("multiparty") ones, making the average event size close to its natural lower bound of 2. For inventor collaboration networks, most patents have a small number of collaborators (though it should be noted that patents with only one inventor are removed from these datasets), and an average size of no more than 3 or 4. In contrast, Fortune 500 boards are more moderate in size, averaging about 10 members per board; movie projects and technical committee meetings are even larger.

What does event size mean for dyadic network structure? As event size grows, the number of imputed dyads created will grow with the square of event size. Clustering coefficients will be inflated due to the assumption that all actors affiliated with the event are dyadically connected. As a result, some researchers have wisely corrected for this inflation before assessing small world properties (Conyon & Muldoon, 2006; Newman *et al.*, 2001; Uzzi & Spiro, 2005), but many researchers have not.

How frequent are affiliation events? For any type of affiliation network data, another key underlying property is the number of affiliation events. Of course, this number will be a function of the researcher's design choices about the study timeframe as well as the network boundaries. Over the timeframe of the study data, how many events exist in the primary data? Consider a few extant studies that report primary event data. Fleming et al's patent collaboration network contains over two million patent events. (While part of this extreme size is a function of the data spanning more than twenty-five years, the yearly number of patents is still quite large even after subdividing the overall dataset by geographic region (MSA) and smaller time windows.) At the other extreme, Rosenkopf and Padula's cellular alliance study contained only 111 alliance deals over twelve years of data. Davis et al's director interlock networks, focused on the boards of Fortune 1000 firms which have public data available, include between 600 and 700 boards in a given year, each of which serves as one affiliative event.

Clearly, different researchers choose different timespans over which they amass their primary data, and longer timeframes necessarily include more events. Within these varying timespans, however, some conventions for the calculation of network parameters have emerged. Board and technical committee data are typically drawn only from the rosters of a given year, and this may well be because the larger event size creates more density in the derived dyadic networks. In contrast, alliance studies have converged on a five-year window of formations due to the challenges of estimating the duration and completion of most alliances. Similarly, Fleming uses a five-year window on collaboration in his patent data as well[11]. Thus, any comparison of event frequencies across network contexts will need to normalize the number of events by the conventional network window. All else equal, higher event frequency will yield greater network density.

How many actors participate overall? Watts (1999) suggested that networks with fewer than 1000 actors were not of interest with regard to the small world concept, because the point of small world theorizing was to suggest that network chains could be short even though networks were large; therefore small networks would have small chains by default. (Note that this has not stopped many researchers as industry-focused networks are frequently smaller than 1000.) Furthermore, Watts also suggested that for small worlds to be of interest, network density should be low, as short paths are common in dense networks.

The number of participants ("*n*") is of interest both in its overall magnitude but also in relation to the number of events. For example, in alliance networks, there seems to be a close correspondence between the number of firms and the number of alliance events. Likewise for inventor collaboration (between the number of inventors and the number of patents) and technical committees (between the number of firms and the number of committee meetings). In contrast, for boards of public Fortune 1000 companies, the number of directors participating is approximately 10 times the number of boards. This is a function of the uniqueness of a director interlock tie: since the individual director forms the interlock between firms by participating in both boards, the affiliator (director) is drawn from a larger set of individuals in contrast to the constrained set of organizations in a top company list.

[11] For an alternative approach see Uzzi & Spiro (2005), who eliminate actor ties after an actor has been absent from new musical productions for seven years.

The relationship between the event size, the frequency of events, and the number of actors defines network density. Alliance networks tend to be sparse because event size is small, so while the number of events may be of the same order of magnitude as the number of actors, the number of imputed dyads is far less than the number of possible dyads. The same is true of inventor collaboration networks. Boards are somewhat different: while event size is an order of magnitude larger, creating many more imputed dyads, the number of unique actors included across the boards yields an even greater number of possible dyads, so these networks are still sparse. Contrast this to technical committee networks, where event size is also larger: however, fewer events and actors conspire to yield high density.

Summary. Several representative networks[12] are plotted across the dimensions described above in Figures 3a and 3b, and it is obvious that event frequency and number of actors behave similarly[13]. The span of these networks over the multi-dimensional space suggests both similarities and differences in the application of small worlds across different tie types.

Insert Figure 3 about here

First, notice that the technical committee networks sit the most separately from other network contexts. Here, the data is usually bounded by some analog of an industry – all firms and their representatives participating in a forum such as a standards body or industry consortium. Network ties are generated by common participation by firms within committee meetings (events). As a result, the number of firms (actors) is small and the number of representatives

[12] The selected networks are drawn from the beginning, midpoint, and end of the longitudinal network datasets that we describe and explore in detail in the next section.

[13] For another perspective, we performed a principal components analysis over the three dimensions of interest, and found that event frequency and number of actors loaded heavily onto one principal component, while event size loaded heavily onto the other.

participating is moderate, while the technical committee meetings (events) that generate network ties are large in size and moderate in frequency (with the notable exception of the earliest data year which has much smaller committee sizes and is located near the board cluster). This leads to networks that are very dense and fully connected, so small worlds do not apply. In addition, the degree distributions do not exhibit much skew, being uniform in nature.

Board networks demonstrate both similarities to, and differences from, technical committee networks. Here, the data are bounded by utilizing a cross-industry list such as the Fortune 500 or 1000. Each firm (actor) has a board (event) of directors, and network ties are generated when firms share a common director ("board networks"), or when directors participate jointly in a board ("director networks"). While the number of events is low (one board per firm), the average event size is of the order of ten, thus the effect of dyadic transformation on clustering is very high. It is also important to note that not all firms are included in the interlock network; some firms are isolates because their board members do not participate on any other boards. Given this constriction, the main component is quite large, and includes 80-90% of the network participants. Both the board and director networks display some skewness in the degree distribution. The networks seem stable over time but new work by Chu and Davis (2016) suggests that demands for increasing diversity on boards are reducing clustering in these networks in recent years.

Alliance networks demonstrate different methodological choices and concerns. First, networks are usually identified by Standard Industry Classification (SIC) codes which typically constrain the number of actors and their partners to the hundreds or low thousands. The number of alliance dyads within a typical window of analysis (all formations within five years) is roughly comparable to the number of actors, but degree distributions are skewed. The size of the main component may vary from 10-70% (Rosenkopf & Schilling, 2007); within this main component, the presence of hubs shortens path length. At the same time, clustering is typically high, though this measure may be inflated in scenarios with many multi-party alliances by their artificial conversion to dyads (Newman *et al.*, 2001).

Finally, while patent collaboration networks vary in scale and scope from alliance networks, they nonetheless share underlying structural similarities. Patent collaboration networks are constructed from co-authored patents and the researcher determines the boundaries of the technology classes to be included as well as any other restriction such as geographic boundaries[14]. Here, the number of events (co-authored patents) and the number of inventors (actors) are larger, but their ratio remains low. Event size is larger than two and typically between 3 and 4. Thus the clustering effect due to the dyadic transformation is higher than the case of alliance data. Further, the main component in the inventor network is a much smaller proportion of the full network, ranging from 1% – 40%. The presence of multiple components of similar size therefore may pose challenges when looking at trends over time as the biggest component changes dynamically over time. The main component in this case also demonstrates a skewed degree distribution. At the same time, note that the largest alliance networks (computers) sit in the cluster of patent networks, underscoring key similarities in event size, frequency, and number of actors for this industry.

In short, it is clear that different types of affiliations, through their varied event sizes and frequencies relative to actors participating, can generate variety in network structures that is not

[14] While firms are not typically represented in collaboration networks because joint patents among firms are rare, firms have an important latent effect in these networks because most collaboration between inventors occurs when they are co-resident at firms. Longer paths of collaboration emerge due to inventor mobility and firm acquisitions or divestitures.

reducible to a traditional small-world calculation to be generalized across these contexts. These differences are summarized in Table 2. We next turn to a more rigorous approach to small world properties that might illuminate these differences more effectively.

Insert Table 2 about here

Suggestion 2: Reevaluating the Small World Statistics and Their Usage

Several characteristics of the Q-type small world statistics (that is, those that take a ratio of a normalized clustering coefficient to the normalized path length) highlight the opportunities to enhance them by utilizing rigorous measures and methodologies for inferring small worlds. First, since Q-statistics represent ratios of normalized clustering coefficient to normalized path length, they can conflate the effects of clustering and path length. Figure 4 shows that the value of SW presents an inverted-U curve as p ranges from 0 to 1 for varying n and k. This implies that identical SW-values may be obtained for WS networks of both low and high p, thereby ascribing the same characteristics to networks more similar to ring lattices as to networks more similar to random ones. Said differently, a SW-statistic for a network with a given number of nodes and edges cannot connote whether this network falls on the ring or random side of this figure without going back to the source data. In contrast, placing a real network unequivocally in the continuum between the two equivalent ideal end points of a ring graph (with high clustering and high path length) and a random graph (with low clustering and low path length) ideally requires a continuous, monotonically increasing statistic.

Insert Figure 4 about here

Second, as we have noted, the use of an Erdos-Renyi random network as a default null model does not correspond well to the skewed degree distributions observed in many networks. Relatedly, a ring lattice network provides a better null model for clustering coefficients; random clustering coefficients are so low that they inflate normalized clustering coefficients substantially, yielding Q-values with no clear upper bounds.

Finally, since traditional small world statistics are continuous and vary with the number of actors as well as the average degree of these actors, as we described earlier, our process for determining thresholds at which to distinguish small worlds from random networks has been arbitrary at best. Small world statistics are inherently probabilistic in that any null model utilized has an underlying distribution. At the least, we should report confidence intervals around these statistics to help us determine the variance around the mean value that is generally reported now. A tighter variance around the mean will give our inference greater validity. The specific thresholds or range of the statistic that corresponds to small world region needs to be determined based on both n and k rather than a single universal threshold. Confidence intervals coupled with a continuous statistic as described above will bolster the identification process of small world networks.

In sum, our field's approach to identifying small worlds can become more rigorous through the reporting of confidence intervals around a monotonically increasing statistic that uses distinct null models for the two components of the small world property while accounting for degree distribution. Fortunately, recent developments in network science arising from the study of the connectivity of neurons in the brain provide such an approach.

Introducing ω. Telesford *et al.* (2011) defined a new statistic, ω, to define the small world property in the spirit of the original WS model by comparing the clustering of the real network, C, to that of an equivalent lattice network, C_{latt}, and path length, L, to that of an equivalent random network, L_{rand}, as follows:

$$\omega = L_{rand} / L - C / C_{latt} \qquad (1)$$

Several important properties of ω which differ from the traditional small-world statistics must be noted. First, ω is composed from two distinct null models, the lattice which maximizes clustering and the random which minimizes path length. Further, each of these nulls is calculated via simulation, which allows them to incorporate degree distribution (unlike the W-S approximations used for SW).[15]

Second, as defined by the formula, the random null model appears in the numerator for the path length ratio, while the lattice null model appears in the denominator for the clustering ratio, and then these two ratios are differenced; let us examine each of these aspects in turn. For the path length ratio, since L_{rand} represents the shortest possible path length for a given degree distribution, this ratio will be closest to 1 for actual networks with path lengths close to those of

[15] An equivalent random network is created by assigning an edge to a node pair with uniform probability while maintaining the original degree distribution of the observed network. The observed network is used as a starting point and the edges are randomly rewired with probability one while ensuring that the network preserves the degree distribution. On average an edge is rewired at least 10 times. The resulting network is simulated 999 times and the average is used to calculate L_{rand}. Alternative, one can also use Newman et al. (2001) to calculate L_{rand} using the analytical solution they provide using polynomial generating functions derived from the degree distribution of the real network.

An equivalent ring lattice network is generated using a Markov-chain algorithm that maintains node degree and interchanges edges with uniform probability; although, these interchanges are only materialized if the resulting adjacency matrix for the network has entries that are closer to the main diagonal. The clustering coefficient is maximized over several user defined iterations. Thus the new matrix is retained only if the resulting network has a higher clustering coefficient after each step. This ensures that clustering is maximized and a lattice-like structure that has high values of clustering coefficient and long path length is obtained. The resulting lattice-like network is simulated 999 times and the average is used to calculate C_{latt}.

their random equivalents, and it will decrease geometrically with increases in the actual network's path length. In contrast, for the clustering ratio, since C_{latt} represents the highest possible clustering coefficient, this ratio will approach 1 for highly clustered networks and decrease toward zero for unclustered networks. Since differencing these two ratios yields omega, values of ω close to zero connote "small-worldliness" as $L \approx L_{rand}$ and $C \approx C_{latt}$. Positive values indicate a graph with more random characteristics because $L \approx L_{rand}$, and $C \ll C_{latt}$, while negative values indicate a graph with more ring-like characteristics as $L \gg L_{rand}$, and $C \approx C_{latt}$. Therefore, unlike the traditional small world statistics, omega generates a monotonically increasing statistic where its possible values are restricted to the interval -1 to 1 regardless of network size (Telesford *et al.*, 2011).

Figure 5 illustrates the monotonically increasing behavior of ω for the Watts-Strogatz model as the rewiring probability is increased from 0 to 1 (lattice to random) for networks of different sizes. Several other features of this graph are notable, particularly in comparison to the behavior of SW in Figure 4. First, ideal small worlds (ω close to 0) are observed at low values of p comparable to the peak values of SW. Second, the variation of ω with n is far less pronounced, in contrast to the dramatic variation in the SW value with n. Third, while Telesford et al (2011) suggest that values of ω between -0.5 and +0.5 may be considered "small-worldly" for larger networks, it is notable that for lower values of n (100 and 250 in Figure 5), the actual lower bound of ω does not reach the theoretical minimum of -1.0. In these cases, a smaller range around 0 is indicated for a small world classification.

Insert Figure 5 about here

However, ω also has important limitations. First, while traditional null model statistics are briskly calculated through formulaic approximations, the calculation of null model values for ω requires the actual network data, and generating ω requires us to run simulations. The availability of code to run such simulations does mitigate to some extent this limitation. However, for large networks these simulations have to manipulate very large matrices and hence are computationally intensive, often taking days to generate the statistic. Second, ω does not eliminate the need for discretion in determining appropriate small world thresholds; however, the confidence intervals generated by the method help assess significant changes in the measure. Finally, ω should only be applied to networks that are susceptible to the small world classification; it should not be calculated for high-density networks.

Application to actual networks

For four extant studies with networks susceptible to small world behavior, we set out to replicate the authors' original small world findings, and calculate ω for comparison. In each case, we calculate ω by simulating 999 instances of the equivalent lattice and random null models defined above using the actual network as the starting point. For each instance, the value of the ω statistic is computed and then used to generate its distribution. We use the code available at the Brain Connectivity Toolbox website in our simulations.

Alliance networks. Rosenkopf & Padula (2008: 675, Table 1) study wireless telecom alliance networks from 1991 to 2002. They demonstrate that the main component of the wireless telecom network starts with only nine nodes in 1991 but grows rapidly to 25 by 1993 and 94 by the end of their study period[16]. We graph their SW coefficients from 1993 to 2002 in Figure 5a[17], which illustrates how they infer the emergence of small worlds from the dramatic increase in the small-world coefficient from 3.9 to 12.8 between 1993 and 1995, and SW remains high for the rest of the time period under observation with a peak around 1999. Simple correlations over the 1993-2002 window demonstrate that SW is driven by the ratio of actual to random clustering (correlation = .99) rather than the ratio of actual to random path length (correlation = -.11).

To calculate ω for the telecom network from 1993-2002, the authors provided the Rosenkopf and Padula (1998) network data, and the results are displayed in Figure 6b. Several insights emerge from the comparison of the omega and SW approaches. First, observe that omega rises from an initial value of approximately -0.4 in 1993 until it approaches the ideal small world omega value of 0.0 during the last three years of the study period; the distribution of the yearly statistics also tighten along with their rise. There is no downturn in omega, in contrast to the 1999 peak observed with the SW statistic. Recalling that the interval of omega within which a small world might be identified for small networks (n<100) should be substantially tighter than the [-0.5, +0.5] range proposed by Telesford et al (2011), we might conclude that while small worlds emerge in the telecom alliance data, they arise more gradually and somewhat later than the SW measure might have led us to propose. Further, the small worlds are driven by clustering, and they arose from what was originally a more lattice-like structure, as the 1993 omega value of -0.4 is close to the lower bound of omega for networks of size 100.

[16] This number of firms in the main component represents the following percentage of the active firms: 56% in 1991; 63% in 1993; 78% in 2002.

[17] We do not include the statistics for 1991 and 1992 as the number of nodes in the network is so small as to negate the relevance of these statistics.

Insert Figure 6 about here

Thus, for these smaller-sized alliance networks, the use of omega largely reinforces the extant findings, merely adding nuance about timing and suggesting that any peak observed in SW may be an artifact of the statistic rather than an important characteristic of the network. As a contrast to the smaller-sized alliance networks, we attempt to replicate Gulati et al.'s (2012) study of computer industry alliance from 1996 to 2005, which suggests that small worldliness peaks at some point in time and then subsides (see Figure 2b, page 458 where they show what they label "small world Q" following this pattern). Recall from our earlier discussion that their "small world Q" is an adjusted SW statistic. We follow their stated data collection strategy and reconstruct their computer industry networks (SIC 3571–3579 & SIC 3612–3699) using data from SDC Platinum. We are able to collect data on 6842 alliances. The total number of firms in the network increases from 808 to 5688 over the study period, and the equivalent increase in the main component is from 395 to 3656 firms (49% to 64% of active firms). These counts differ from those provided by Gulati et al (2012), where their reported yearly counts ranged from 374 to 6221 firms, and further communication with the authors revealed that their main component size ranged from 43 to 1822 firms (11% to 29% of active firms).[18]

Figure 7a displays the (non-adjusted) SW statistic for our data. Similar to those observed for cellular alliance networks, SW grows dramatically from a low level at the beginning of the

[18] We discussed with the database provider (Thomsen Reuters SDC) how the database from which the alliances were drawn may have continued to evolve as backdated alliances are added to the database, growing the number of alliances and firms. At the same time, backdated mergers or acquisitions would decrease the number of firms. We also discussed these data discrepancies with the authors and learned that they had included a variety of other alliances from business services or with missing codes. Such a tacit, labor-intensive approach may increase the completeness of the network data yet also precludes replicability.

observation period. However, no peak and subsequent recession of the SW value occurs. We calculated the adjusted SW per the authors' method (as discussed earlier) and the result, in Figure 7b, appears monotonically decreasing with one slight upward tick in 1999. There is no peak in 2000, so differences in the primary data may be yielding substantial differences in both the SW and adjusted SW statistics.

Insert Figure 7 about here

While the lack of replicability is a matter of concern, it is still valuable to examine how the use of omega adjusts the inferences drawn by using SW on the computer alliance networks we were able to create, as depicted in Figure 7c. Here we observe that while each network obtains values of omega that reside in the small world interval, after a modest rise in value, the network seems to stabilize for the remainder of the study period (while values oscillate slightly, confidence intervals indicate that this variance is insignificant). Despite the fact that these computer alliance networks are substantially larger than those from the cellular industry, the use of omega reveals small-worldliness in both settings while also indicating that this property emerges from clustering.[19]

[19] As a further test of generalizability across industries, we examined SDC alliance data for the chemical (SIC 28XX) and semiconductor (SIC 3674) industries, shared by the authors of Ghosh, Ranganathan, and Rosenkopf (2016). Similar trends obtain: the total number of firms, the number of firms in the main component, and the small world statistic SW monotonically increase over time. This SW trend is mainly driven by the normalized clustering coefficient which grows concomitantly with n. Adjusted SW, which multiplies the SW statistics by ln(n)/n, again exhibits an opposite and decreasing trend as we observed with our computer data, underlining our concerns about the validity of the measure rather than merely differences in our primary data as compared to the original authors'.

At the same time, ω demonstrates different trends across the different industries. While all converge into the small world neighborhood (-0.5 to +0.5) over time, the semiconductor small worlds emerge from more random networks, while the computer small worlds emerge from more clustered networks. Interestingly, the chemical networks are close to ω =0 throughout the study period, which may reflect the relative maturity of that industry.

Inventor collaboration networks. We next examine the Silicon Valley and Boston regional inventor networks previously analyzed by Fleming et al. (2007, their Figure 4, page 943). Here the authors assert "…that inventor collaboration structure has become more small-world for many MSAs [Metropolitan Statistical Areas], Silicon Valley in particular, since the early 1990s." They do not consider particular thresholds for small worlds; rather, they note a dramatic upswing in their small world coefficient (as noted previously, normalizing the ratio by the ring lattice rather than the Erdos-Renyi random network), which rises to over 100 for Silicon Valley and over 20 for Boston. We replicated these findings by obtaining data as published in the Harvard Dataverse (Lai, D'Armour, & Fleming, 2009) and following their stated procedures. Our results are displayed in Figure 8a, and appear to track their data quite closely.

Before examining omega, it is important to note how inventor collaboration networks compare to the alliance networks we have examined above. These networks grow much larger over time; in the data we collected, the number of inventors collaborating grow from 2,017 to 19,661 in Silicon Valley and 2474 to 11,081 in Boston over the 25-year study period. However, the main components of these networks are much smaller; ranging from 3% to 57% of inventors in Silicon Valley and 3% to 31% of inventors in Boston. In other words, as patenting intensity grew over the study period, the number of inventors in the collaboration networks grew dramatically over the study period, yielding networks larger in size like the computer networks.

Insert Figure 8 about here

The calculation of ω for the Silicon Valley and Boston inventor networks, shown in Figure 8b, presents contrasting interpretations. While the traditional small world statistics suggest increasing small-worldliness of both networks, the ω statistic demonstrates a trend away from small worlds into the lattice region during the latter years of the study period. Said differently, clustering increases dramatically, becoming the dominant characteristic of these networks, while path length remains relatively consistent. Of course, differences between the two regions are still observed. Here the Silicon Valley network ω starts already in the small world region (-0.2), peaking around near 0 in 1981 before descending into the lattice-like region by the mid-nineties. In contrast, the Boston network starts in the lattice-like region, passing into the small world region with a maximum value of 0.1 in 1985 and then similarly descending into the lattice-like region again in the mid-nineties.

In sum, the rising number of inventors in each network generates local clustering, which is consistent with the follow-on analyses conducted by Fleming et al. in their work. It is also worth noting that the omega values generated by the inventor network, despite their tight confidence intervals, are more erratic year-to-year in comparison to those of the alliance network. We attribute this difference to the larger flux of actors in the inventor network; while firm-focused networks drawn from ongoing alliance collaborations demonstrate reasonably similar membership and structure on a year-to-year basis, individually-focused networks drawn from patent collaborations derive from inventors who come and go far more frequently and copiously over the study period, yielding more obvious changes in structure.

Board interlocks. Davis *et al.* (2003) show that the networks of firms through board interlocks as well as the networks of directors through common participation demonstrate small world structure from 1982 to 2001. We replicated their approach by collecting data on boards of public Fortune 1000 companies from RiskMetrics for the period between 1999 and 2011. The board networks contain large main components that encompass 80-90% of the full network, in contrast to the smaller and more variable main component memberships for patent and alliance networks. The number of firms in board networks is largely constant as about 55-60% of the Fortune 1000 firms have at least one director interlock but the number of unique directors has increased from 1999 with a slight downward trend in 2011. Figure 9a displays the SW findings which corroborate their work. The SW statistic decreases from 1999 to 2011, although one would still infer the networks as small world using this statistic. Omega on the other hand tells a different story, as seen in Figures 9b and 9c. While the director network would still qualify as a small world network using this statistic, the board interlock network can no longer be inferred as a small world network. Rather it is closer to a random network as has also been shown by other papers such as Conyon and Muldoon (2006) and Newman *et al.* (2001).

Insert Figure 9 about here

Discussion

Our work has demonstrated how findings on network structure have been overgeneralized due to less-than-systematic measurement approaches. By illustrating the differing characteristics of various affiliation networks, offering a more rigorous small world network measure, and attempting to replicate prior work to explore these issues, several implications for future research emerge.

Clearer descriptions of primary affiliation network data. As most network analysis using archival data uses affiliative network data, in our view, it is imperative to provide the details of the primary data that leads to the processed one mode network. At a minimum, researchers should provide the number of actors, the number of events, the distribution of event size, and the timeframe of network analyses. Such detail allows both the researcher and the reader to place the network of interest in the context of the variety of networks we have described. As one example, IMDb data is available for researchers to explore movie project affiliation networks. Here, movie events can be large in size, and the number of movies in a given year may be high, but the number of unique individuals participating in movies would be high as well. If one were to analyze the main component of movies tied by common actors over a specific time period, would we expect its structure be most like that of board networks, or technical committee networks? And in turn, should the small world property even be applicable?

To enable replication, researchers should also describe the methods they used to procure their primary data. Our attempts to reproduce the Gulati et al (2012) alliance dataset illustrated the complications of generating a comparable dataset from an established primary source at a later point in time. Several points are relevant here. Since source data may be continually updated with backdated events, researchers should recognize that the most recent years in their datasets may still be incomplete. In a post-hoc attempt to recreate the Rosenkopf and Padula network data from the original source, we also found that the final two years of network data indicated substantial changes in the number of firms included, but these changes did not affect the general identification of the small world presence despite slight differences in the time trends at the end of the window. We also note that patent data, drawn from successful patent applications, can be subject to similar concerns. This is not to say that researchers should only analyze older data; rather, that researchers would be wise to assess whether results are robust to inclusion or removal of the most recent segments of their data (Shi *et al.*, 2017) as well as whether results are robust across varied databases (Schilling, 2009).

There are other important things that researchers can do to improve replicability of their datasets. While it is typical for researchers to note some characteristics of data selection such as SIC codes and treatment of M&A activity, a better solution than general description of these methods is the inclusion of the actual queries submitted to the databases. As our field moves more toward automated selection of data, this approach is straightforward, but the challenge of understanding more tacit approaches to data collection remain. Here, interrater reliability approaches are critical.

Of course, publishing the actual dataset for other researchers' use, as Fleming has done, alleviates the need for researchers to try to replicate the dataset. That said, we understand that the common approach in our field has been to develop a proprietary dataset and harvest it for some time before allowing others to access it, if ever. Field-level efforts to create incentives for greater sharing can be helpful here.

Identification and evolution of small worlds. By applying omega to the three types of networks originally classified as small worlds, it is clear that ω is a more conservative statistic than the family of SW and Q measures. Said differently, using a ratio with only the random network as the null leads us to classify many networks with high clustering as small world even though the path length might not be as low as a random network. Since the more rigorous omega approach addresses both the null model and ratio issues, we recommend that future research reexamine other networks identified as small worlds in our literature using this approach.

Another issue of concern for network replications is their frequent focus on the main component of interconnected actors rather than the fuller set of partially disconnected actors. Small world analyses, with their particular emphasis on path length, require connected actors.

Yet in many networks, the main component may only represent a small part of the overall network activity. In particular, during the early stages of network evolution, there may be many different network components of similar sizes, and which is formally determined as "main" by slightly larger size may lead to significant flux in actors and structure over time periods unless some of them consolidate into a relatively large and stable giant component. Accordingly, researchers should analyze and report relative component sizes as well as actor flux in the main components while assessing whether these trends yield discontinuities in small world analyses.

Additionally, examining how networks evolve endogenously, and in particular how their small world nature changes over time, has received special attention ever since the seminal Kogut and Walker (2001) paper extolling their robustness in the face of change. A majority of the papers corroborate such stability as well as increasing small-worldliness with time (Fleming *et al.*, 2007; Rosenkopf & Padula, 2008). On the other hand, Gulati *et al.* (2012) emphasized a rise and fall of small-worldliness as their computer industry network exhibits an inverted V shape using their adjusted Q measure, yet our computer industry network did not obtain such a curve using the adjusted Q measure and our ω statistic exhibited more or less stable values in the latter years. Future research should examine whether such patterns are indeed common and the mechanisms that drive them, while being sensitive to short-term versus long-term trends.

Performance implications of small world properties. Finally, the reported effects of small worlds on performance vary widely across studies (Uzzi, Amaral, & Reed-Tsochas, 2007). As we discussed in our introduction, Uzzi and Spiro (2005) examined collaboration networks for Broadway musicals, demonstrating an inverted U-shaped relationship between Q and revenue. In contrast, Fleming, King III, & Juda (2007) and Fleming and Marx (2006) found no relationship between the small world characteristics of inventor collaboration networks and patent productivity. In the alliance literature, Verspagen and Duysters (2004) suggested that alliance networks have small world properties while theorizing that small world networks would be more efficient by generating social capital (via clustering) as well as transmitting knowledge and information (via short path length), but did not directly test the effect on performance. Building on these ideas, however, Schilling and Phelps (2007) argued that the more a firm was embedded in an industry-wide alliance network with high clustering and short average path lengths (reach), the more likely it was to gain access to knowledge important for innovation. Rather than using a ratio from the SW family, they regressed firm patent productivity on the interaction of clustering x reach in industry alliance networks, yielding a positive association[20]. These discrepancies in findings across studies are largely unresolved, and may be related to the fundamental issue of identifying small-worldliness. Future research on the performance effects of small world structure should examine whether using omega-like statistics clarifies outcomes.

Relatedly, since the (fully decentralized) Watts-Strogatz small world model has become the default substrate on which simulation studies in our field have been performed (e.g. Fang *et al.*, 2010; Rivkin & Siggelkow, 2007; Shore, E., & D., 2015), our findings suggest the value of reexamining simulation work using alternative substrates which adjust for degree distribution. Just as Schilling and Fang (2014) demonstrate how that “hubbiness” matters for learning using the Xulvi-Brunet model (Xulvi-Brunet & Sokolov, 2004) as their substrate, it is critical for us to identify whether system- and node-level outcomes vary with the network model chosen. Given

[20] It is worth noting that these authors did not calculate small world statistics because they included firms beyond the main component, requiring harmonic adjustments to calculate path lengths between otherwise disconnected firms. This again raises the question of when it is wise for authors to focus their attention on merely the main component.

the skewness inherent in alliance or citation networks, for example, it is an important question whether the use of Watts-Strogatz substrates to examine learning or diffusion biases findings due to their lack of skew.

Conclusion

Our field's research on network structure is at a crossroads. To develop mid-range theory, it is critical that we assess generalizability by facilitating replication and comparability across studies. For the particular example of small world networks, we demonstrated how this construct has been measured in different ways in the organization theory and strategy literatures which ultimately have obscured important differences in network contexts. We illustrated these differences in network contexts by focusing attention back on the characteristics of the primary affiliation data that researchers typically transform to dyadic networks in order to abstract to a small-world generalization. We also discussed the inconsistencies in the measurement of small-worldliness and suggested a more rigorous approach toward assessing small-world properties. Statistics such as omega begin to address some of the weaknesses of traditional approaches, but measuring these characteristics of networks is still a blend of art and science; therefore our findings indicate the importance not only of consistent statistics, but also of consistent data. Much work remains to be done to understand both the antecedents and consequences of network structure across the wide variety of affiliation contexts where it emerges.

Table 1. Small world examples from the management literature

Authors	*Network*	*Null Model*	*Period*	*n*	*k*
Kogut and Walker (2001)	German firms	Erdos renyi	1993-1997	291	2.02
Baum et al. (2003)	Canadian I-banks	Erdos renyi	1952-1957 1969-1974 1985-1990	53 41 142	1.36 2.22 3.83
Davis et al. (2003)	US Co. interlocks	Erdos renyi	1982 1999	195 195	6.80 7.20
Verspagen and Duysters (2004)	Strategic alliances	Erdos renyi	1980-1996	5504	5.29
Uzzi and Spiro (2005)	Actors in Broadway Musicals	Fixed degree distribution random graph	1945-1989	2029	Not reported
Fleming et al. (2007)	US patenting inventors	Ring Lattice	1986-1990	7069	4.73
Schilling and Phelps (2007)	US alliances in 11	Do not use Small World Explicitly	1992-2000	171	3.11
Rosenkopf and Padula (2008)	Telecom Industry Strategic alliances	Erdos renyi	1991-2002	94	3.45
Gulati et al. (2012)	Computer Industry Strategic alliances	Erdos renyi	1996-2005	900	3.5

Table 2. Affiliation network characteristics across four sample contexts

	Inventor Collaboration	***Board Interlocks***	***Technical Committee***	***Alliance Networks***
Network Actors (nodes)	Inventors	Firms	Firms	Firms
Event (where actors affiliate)	Patent	Board of Directors	Technical Committee Meeting	Alliance Contract
Affiliators	Inventors	Directors	Firm representatives	Firms
Event size	Low	Medium	High	Lowest
Typical sampling frame	5 years	1 year	1 year	5 years
Event Frequency	Highest	High	Medium	Low
# of actors	High	Medium	Low	Low
# of actors/event size	High	Medium	Low	Medium
# of affiliators	High	High	Medium	Low
Density	Low	Medium	High	Low
Main component: % nodes	Low	High	High	Medium
Degree distribution skewness	High	Medium	Low	High

Figure 1. Small world formalization from Watts and Strogatz (1998)

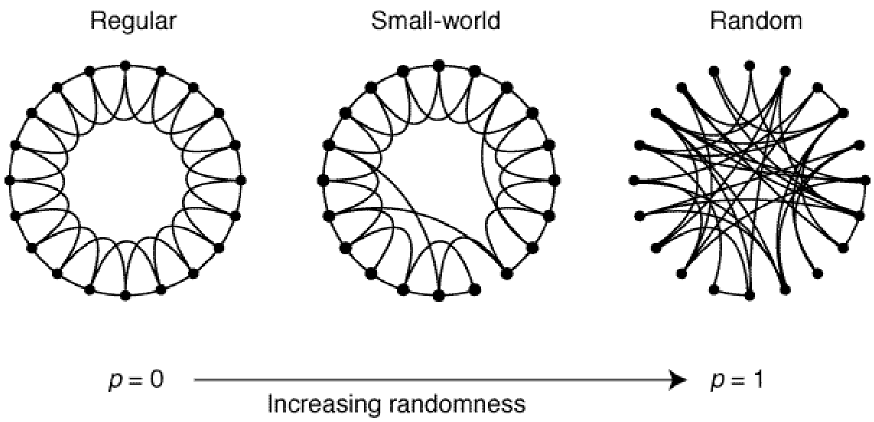

n=20, k=4

Figure 2. Normalized Clustering Coefficient and Path Length versus Prob (Rewiring)

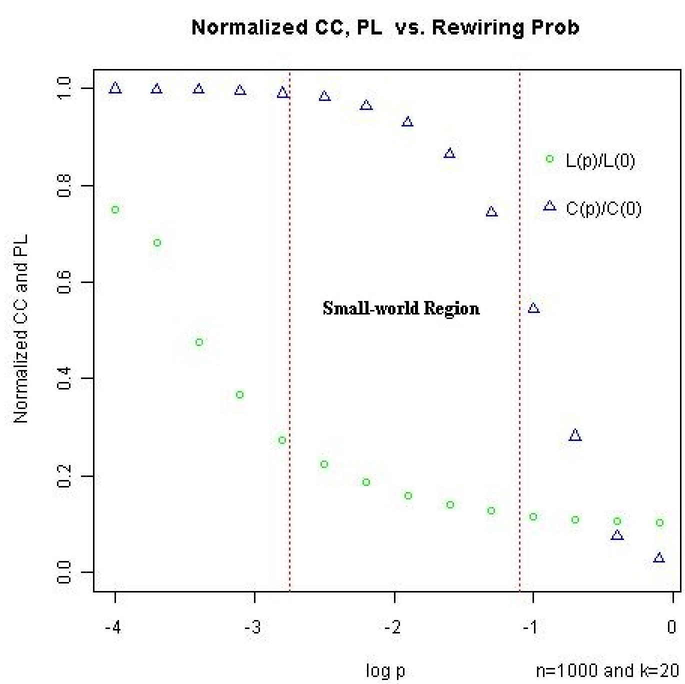

Figure 3. Networks clustered by primary affiliation data characteristics

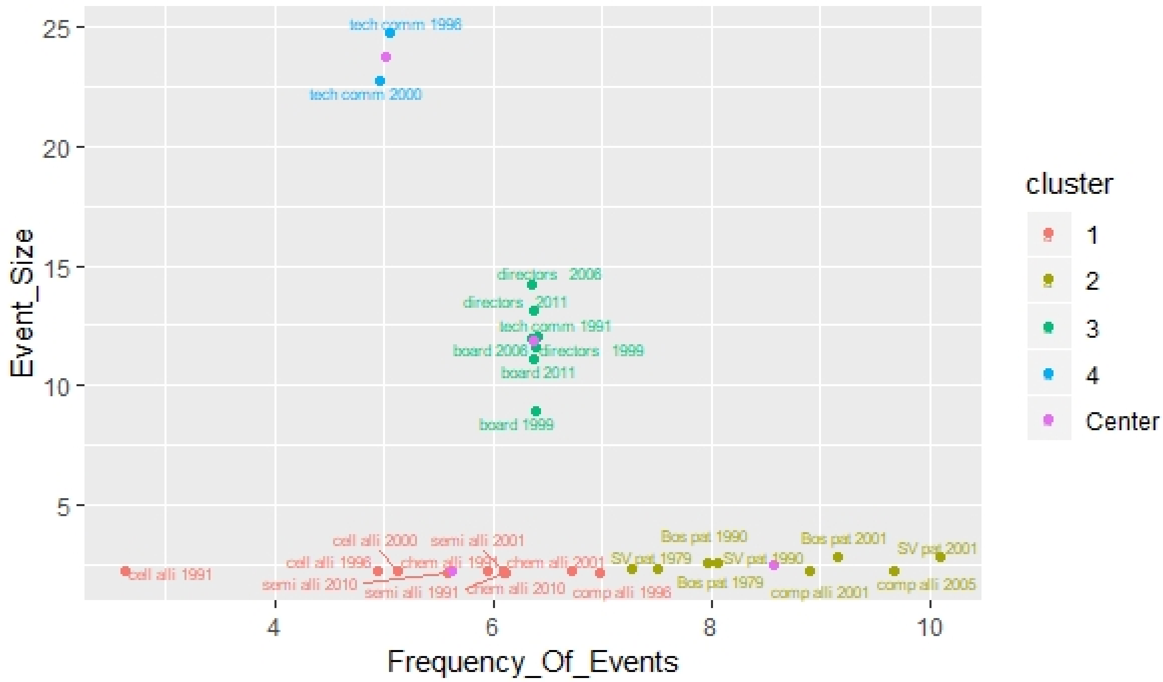

Figure 3a. Plot Event Size vs. Frequency of Events (log)

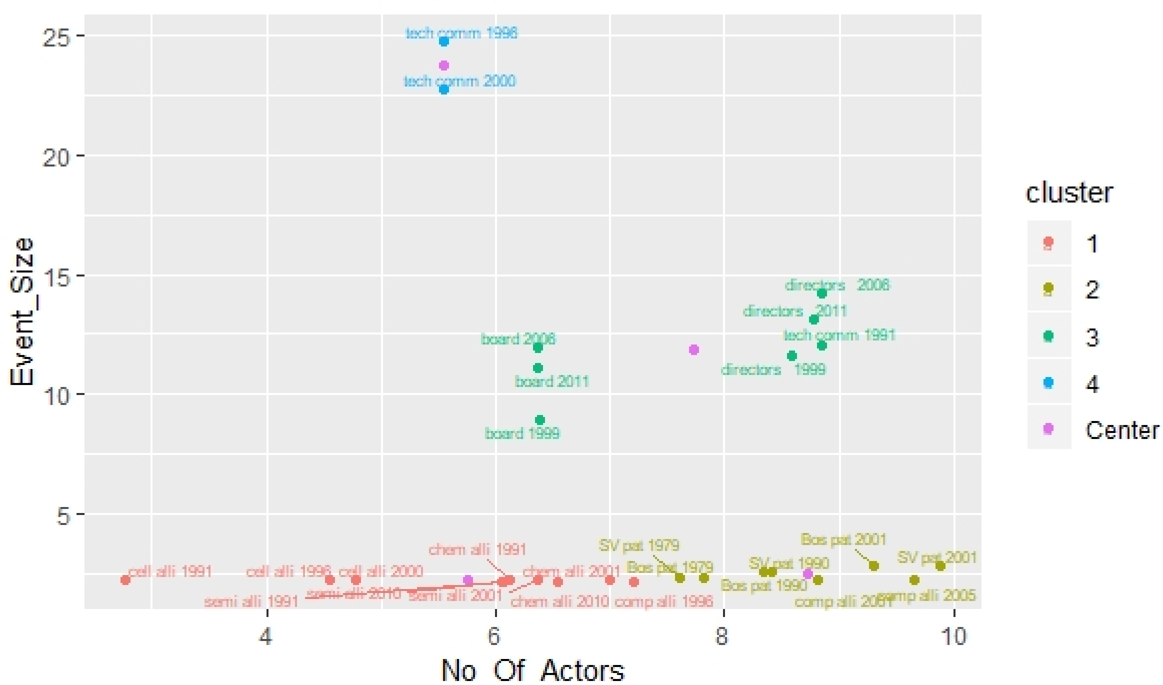

Figure 3b. Plot Event Size vs. Number of Actors (log)

Figure 4. SW versus Prob (Rewiring) for different number of network nodes n

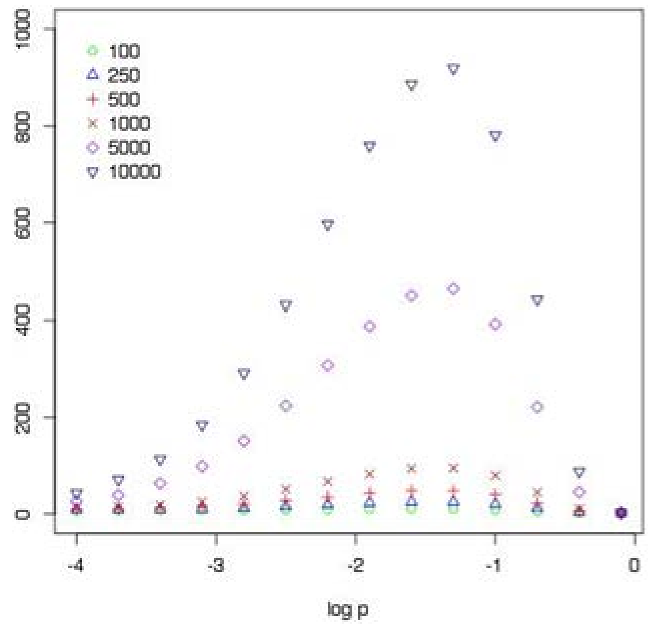

Figure 5. ω versus Prob (Rewiring) for different number of network nodes n

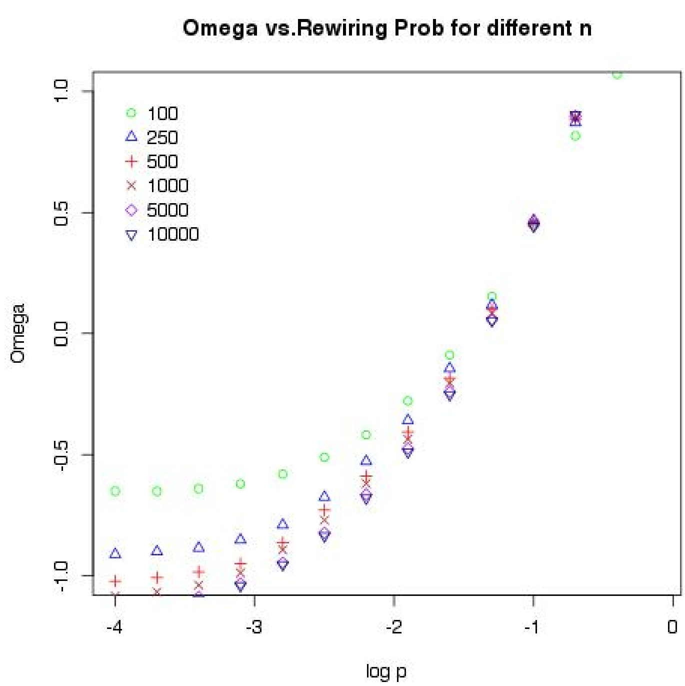

Figure 6a &b. Wireless Telecom Network SW & ω evolution

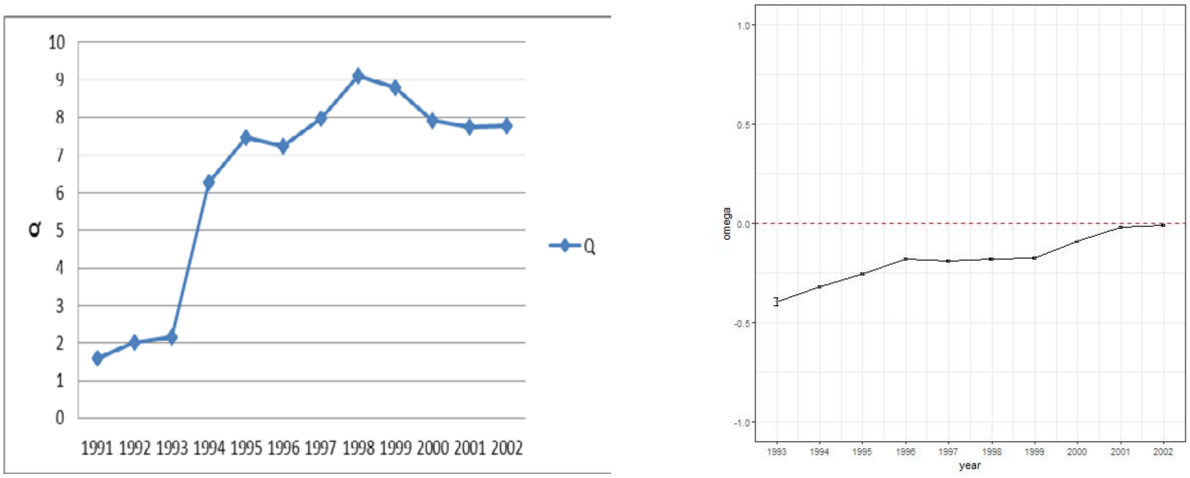

Figure 7a-c. Computer Industry Network evolution 1996-2005

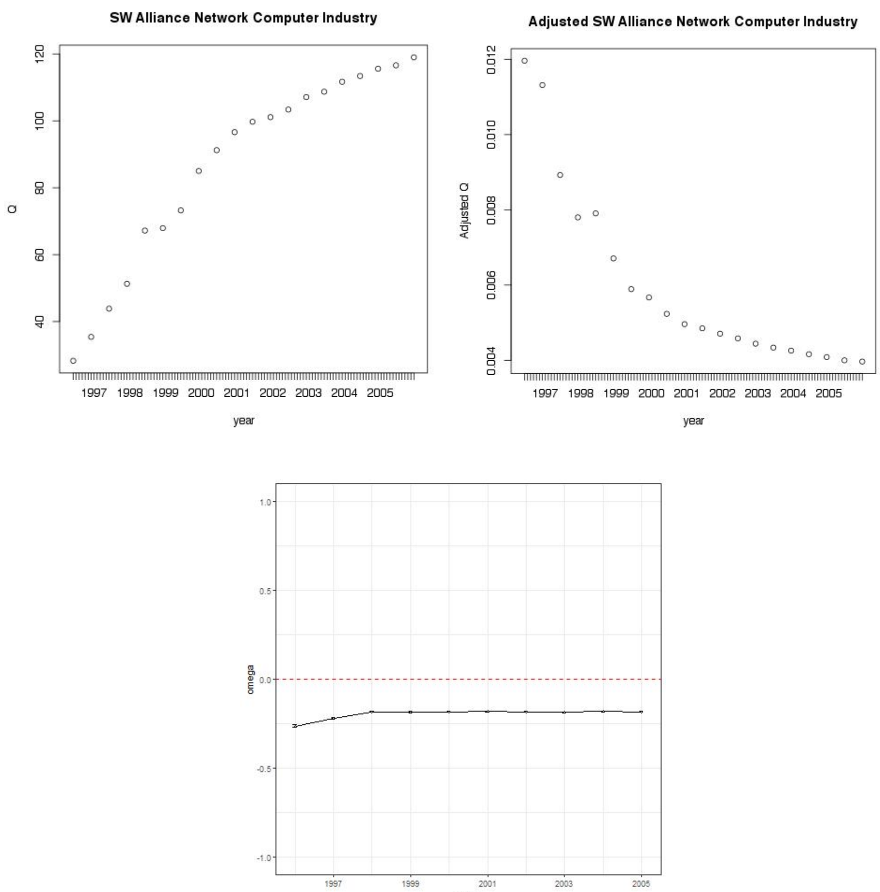

Figure 8. Replication of Figure 4 Fleming et al. (2007) – Silicon Valley (blue) & Boston (red)

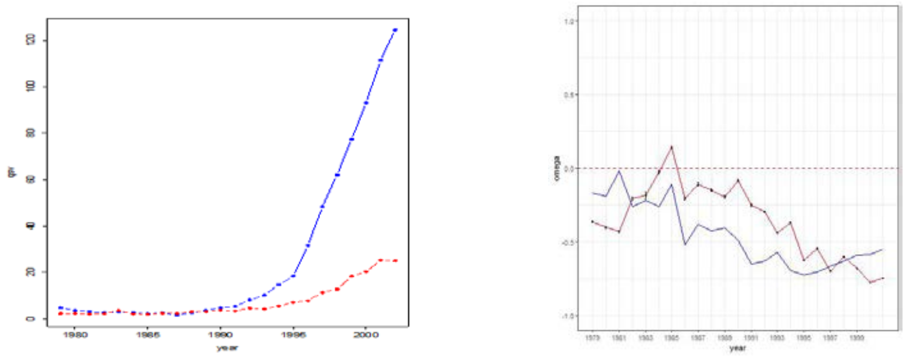

Figure 9a-c. Small world Statistics across time in Fortune 500 board networks

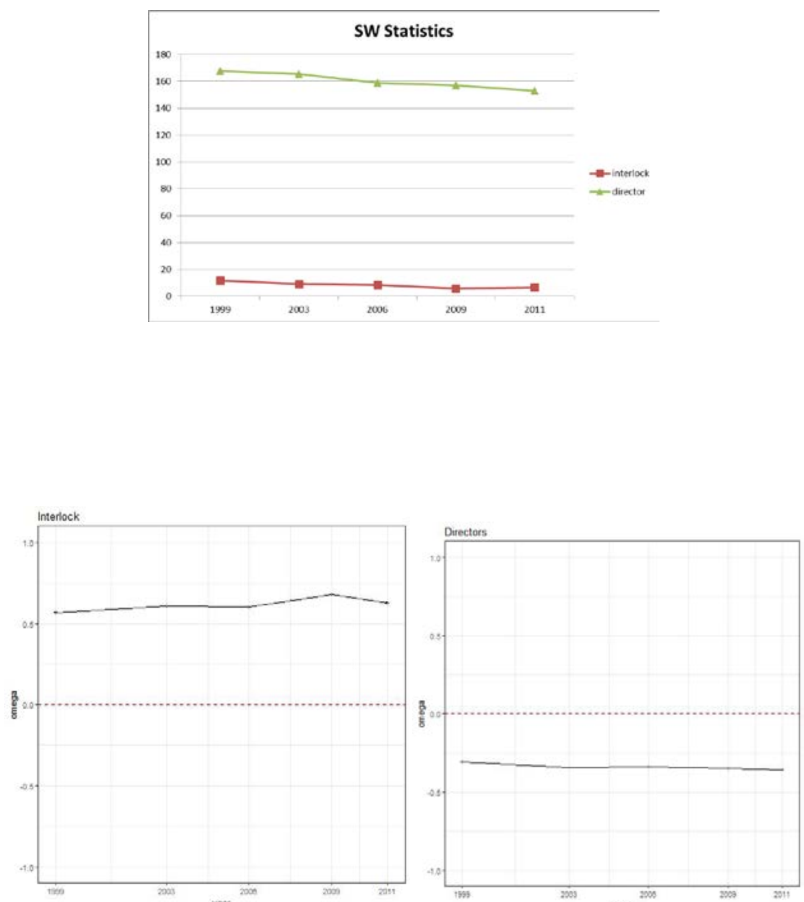

References

Barrat A, Weigt M. 2000. On the properties of small-world network models. *The European Physical Journal B-Condensed Matter and Complex Systems* **13**(3): 547-560.
Baum J, Calabrese T, Silverman B. 2000. Don't go it alone: Alliance network composition and startups' performance in Canadian biotechnology. *Strategic Management Journal* **21**(3): 267-294.
Baum J, Cowan R, Jonard N. 2010. Network-Independent Partner Selection and the Evolution of Innovation Networks. *Management Science*: mnsc.1100.1229.
Baum JAC, Shipilov AV, Rowley TJ. 2003. Where do small worlds come from? *Industrial and Corporate Change* **12**(4): 697-725.
Chu JSG, Davis GF. 2016. Who Killed the Inner Circle? The Decline of the American Corporate Interlock Network. *American Journal of Sociology* **122**(3): 714-754.
Conyon MJ, Muldoon MR. 2006. The small world of corporate boards. *Journal of Business Finance & Accounting* **33**(9-10): 1303-1321
Davis GF, Yoo M, Baker WE. 2003. The Small World of the American Corporate Elite, 1982-2001. *Strategic Organization* **1**(3): 301.
Fang C, Lee J, Schilling MA. 2010. Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science* **21**(3): 625-642.
Fleming L, King Iii C, Juda AI. 2007. Small Worlds and Regional Innovation. *Organization Science* **18**(6): 938.
Fleming L, Marx M. 2006. Managing Creativity in Small Worlds. *California Management Review* **48**(4): 6.
Ghosh A, Ranganathan R, Rosenkopf L. 2016. The impact of context and model choice on the determinants of strategic alliance formation: Evidence from a staged replication study. *Strategic Management Journal* **37**(11): 2204-2221.
Gulati R, Nohria N, Zaheer A. 2000. Strategic networks. *Strategic management journal* **21**(3): 203-215.
Gulati R, Sytch M, Tatarynowicz A. 2012. The rise and fall of small worlds: Exploring the dynamics of social structure. *Organization Science* **23**(2): 449-471.
Kilduff M, Crossland C, Tsai W, Krackhardt D. 2008. Organizational network perceptions versus reality: A small world after all? *Organizational Behavior and Human Decision Processes* **107**(1): 15-28.
Kogut B, Walker G. 2001. The small world of germany and the durability of national networks. *American sociological review* **66**(3): 317-335.
Lai R, D'Armour A, Fleming L. 2009. The careers and co-authorship networks of US patent-holders, since 1975 (version 3). Data set, Patent Network Dataverse. *Harvard Institute for Quantitative Social Science (distributor), Harvard Business School, Boston. Accessed September* **3**: 2014.
Milgram S. 1967. The small world problem. *Psychology Today* **2**(1): 60-67.
Newman MEJ, Strogatz SH, Watts DJ. 2001. Random graphs with arbitrary degree distributions and their applications. *Physical Review E* **64**(2): 026118.

Phelps C, Heidl R, Wadhwa A. 2012. Knowledge, Networks, and Knowledge Networks A Review and Research Agenda. *Journal of Management* **38**(4): 1115-1166.
Pool IS, Kochen M. 1978. Contacts and influence. *Social Networks* **1**(1): 5-51.
Rivkin JW, Siggelkow N. 2007. Patterned interactions in complex systems: Implications for exploration. *Management Science* **53**(7): 1068-1085.
Rosenkopf L, Metiu A, George VP. 2001. From the bottom up? Technical committee activity and alliance formation. *Administrative Science Quarterly* **46**(4): 748-772.
Rosenkopf L, Padula G. 2008. Investigating the Microstructure of Network Evolution: Alliance Formation in the Mobile Communications Industry. *Organization Science* **19**(5): 669-687.
Rosenkopf L, Schilling M. 2007. Comparing alliance network structure across industries: observations and explanations. *Strategic Entrepreneurship Journal* **1**(3-4): 191-209.
Schilling MA. 2009. Understanding the alliance data. *Strategic Management Journal* **30**(3): 233-260.
Schilling MA, Fang C. 2014. When hubs forget, lie, and play favorites: Interpersonal network structure, information distortion, and organizational learning. *Strategic Management Journal* **35**(7): 974-994.
Schilling MA, Phelps CC. 2007. Interfirm Collaboration Networks: The Impact of Large-Scale Network Structure on Firm Innovation. *Management Science* **53**(7): 1113.
Shi Y, Sorenson O, Waguespack DM. 2017. Temporal Issues in Replication: The Stability of Centrality-Based Advantage. *Sociological Science* **4**: 107-122.
Shore J, E. B, D. L. 2015. Facts and Figuring: An Experimental Investigation of Network Structure and Performance in Information and Solution Spaces. *Organization Science* **26**(5): 1432-1446.
Telesford QK, Joyce KE, Hayasaka S, Burdette JH, Laurienti PJ. 2011. The ubiquity of small-world networks. *Brain Connectivity* **1**(5): 367-375.
Uzzi B, Amaral LAN, Reed-Tsochas F. 2007. Small-world networks and management science research: a review. *European Management Review* **4**(2): 77-91.
Uzzi B, Spiro J. 2005. Collaboration and Creativity: The Small World Problem. *American Journal of Sociology* **111**(2): 447-504.
Verspagen B, Duysters G. 2004. The small worlds of strategic technology alliances. *Technovation* **24**(7): 563-571.
Wasserman S, Faust K. 1994. *Social network analysis: Methods and applications*. Cambridge Univ Pr.
Watts DJ, Strogatz SH. 1998. Collective dynamics of 'small-world' networks. *Nature* **393**(6684): 440-442.
Xulvi-Brunet R, Sokolov I. 2004. Reshuffling scale-free networks: From random to assortative. *Physical Review E* **70**(6): 066102.