MODELS FOR NETWORK EVOLUTION

by

David L. Banks and Kathleen M. Carley

September 29, 1995

Abstract: This paper describes mathematical models for network evolution when ties (edges) are directed and the node set is fixed. Each of these models implies a specific type of departure from the standard null binomial model. We provide statistical tests that, in keeping with these models, are sensitive to particular types of departures from the null. Each model (and associated test) discussed follows directly from one or more socio-cognitive theories about how individuals alter the colleagues with whom they are likely to interact. The models include triad completion models, degree variance models, polarization and balkanization models, the Holland-Leinhardt models, metric models, and the constructural model. We find that many of these models, in their basic form, tend asymptotically towards an equilibrium distribution centered at the completely connected network (i.e., all individuals are equally likely to interact with all other individuals); a fact that can inhibit the development of satisfactory tests.

Keywords: triad completion, Holland-Leinhardt model, polarization, degree variance, network evolution, constructuralism

*David L. Banks is an Associate Professor in the Department of Statistics at Carnegie Mellon University, Pittsburgh, PA 15213.

Kathleen M. Carley is an Associate Professor in the Department of Social and Decision Sciences at Carnegie Mellon University, Pittsburgh, PA 15213.

The authors would like to thank Tom Snijders and David Krackhardt for their valuable comments.

1 Introduction

Much of the work in social network theory concerns the analysis of change in a single network. These analyses often entail attempts to explain change as a function of the network's structure (Sampson, 1968; Burt, 1980; Johnson, 1986; Coleman, Katz and Menzel, 1966). Although issues of change have been addressed in network terms (Granovetter 1973,1974), there are few formal models of network change (see Doreian, 1990). And there is even less understanding of how one can statistically examine different theories of network evolution. Thus we present a set of formal models of network change and demonstrate that relatively simple statistical techniques are available to determine whether or not the specific departure from the null binomial distribution, that is consistent with the model in question, is present in the data.

The set of models that we describe is not exhaustive; rather, these models derive from specific theories about socio-cognitive adaptation. By contrasting these theories in terms of reasonable mathematical representations of them we are able to provide a critique of the current state of theorizing about change in social networks. Furthermore, by formally representing a suite of theories, we are able to show how relatively simple techniques can be used to examine the extent to which these formal models fit any given data set.

Three prominent theoretical paradigms for network change derive from social comparison theory (Heider, 1958), exchange theory (Kapferer, 1972; Blau, 1967), and constructuralism (Carley, 1990, 1991; Kaufer and Carley, 1993). These social theories suggest general principles that may govern the addition or deletion of ties (or edges) between individuals (or nodes). The same paradigm may provoke several distinct mathematical models for network change, depending on how it is applied in a given social context. The paradigms are similar in that they all argue that tie construction is the result of a socio-cognitive process whereby individuals alter the set with whom they are likely to interact; however, these paradigms differ in the particular social agenda they posit as guiding this process.

In studying network change, it is important to employ statistical tests which differentiate among the paradigms. For example, Carley (1990) found that for the Kapferer (1972) dataset on change in a tailor shop, the constructural model offered a better fit to the data than alternative models of network change. However, this analysis left open the question of

whether there was a statistically significant difference between the models. More generally, progress in this area requires the development of formal models and a testing methodology.

This paper introduces a bestiary of models for network change. Most (but not all) of these models represent attempts to reify the three prominent paradigms into specific mathematical forms. Of course, the same paradigm sometimes leads to alternative models; for example, the exchange theory paradigm places primacy upon dominance relations, and thus provokes models for degree variance. The mechanism underlying these models depends on how competition for dominance is seen to develop; hence, a variety of mechanism have been postulated including charisma, power, popularity, etc. and in the extreme polarization, and even balkanization. Each of the models focuses on a single mechanism for change. Thus, a comparison of these models sets into stark contrast these various paradigms. In order to facilitate comparability, this paper treats only the case in which the number of potentially interactive individuals is fixed over the entire course of the study (in contrast, Doreian (1990) discusses issues that arise when the node set is permitted to vary). Similarly, it standardizes our exposition to restrict attention to models which describe directed edges.

The emphasis in this paper is on the specification of mathematical models, and the description of statistical tests which (under reasonable assumptions) enable the analyst to determine whether the bias in edge presence or absence, as suggested by the model, is present in the data. The discussion of statistical methods for model analysis is incomplete, since statisticians do not yet have methodology that enables goodness-of-fit tests for network-valued time series data. Nevertheless, as we demonstrate, it is possible to use relatively simple and available statistical techniques to begin to address these issues. Generally, the best we can do is falsify a particular model.

The problem of model selection is further complicated by the fact that different theoretical models often evolve towards the same equilibrium state. For example, simple forms of the triad completion model, the degree variance model, and the constructural model all lead, asymptotically, to a completely connected network or to observations which are in some kind of stochastic equilibrium around the completely connected network.

This stochastic equilibrium can be as simple as having each edge present independently with the same, typically large, probability p. Some models, such as the degree variance

model, allow these edge-probabilities to differ; for example, edges to less degree central nodes would be more fragile than those to highly degree central nodes. In triad completion models, edges may be deleted independently at a particular time point, but the inclination to achieve balance makes those missing edges more likely to reform at the next step than a constant p would suggest. A series of observations would show dependence across time; although, the asymptotic marginal measure for a single time point would have each edge present with probability p. The equilibrium result might be a random digraph with dependent edges.

All of the models which are treated in this paper emphasize the process of edge formation. In current social theories, edge removal is not usually discussed or modeled as carefully, and this paper reflects that limitation. At the same time, we note that this is an area that deserves more attention, especially since it is reasonable to guess that the breaking of existing ties is less casual than the formation of ties (at least for positive ties, such as friendship relations), and thus may conform more rigidly to predictive theories of behavior. But our purpose is to develop methods that enable direct statistical comparisons of competing models of network evolution, and so we stick close to the standard theoretical constructs. Nonetheless, it would not be hard to adapt some of the methods we propose to the analysis of models which incorporate processes of tie deletion.

Given the emphases of the current theories and the fact that most models ultimately lead to a densely connected network, a key conclusion of this survey is that it is generally important for analysts of network change to secure multiple observations early in the developmental trajectory rather than late. This is particularly true if, as these theories implicitly suggest, edge addition is a stronger force than deletion. Otherwise, it can become impossible to statistically distinguish the formal models or the underlying social theory.

2 A Bestiary of Models

Change is complex. Many researchers have proposed plausible forces which drive network change, and, in principle, all of these could act in combination. Major drivers include, but are not limited to, the following: individual learning (Carley, 1990, 1991), mimicry (Powell and DiMaggio, 1991), cognitive requirements for balance (Heider, 1958), social exchange processes (Blau, 1967), socio-demographic factors or personality factors (Homans,

1950; Blau, 1967), and cognitive limitations (March and Simon, 1958).

This section describes a range of models which are motivated by plausible social theories, and sketches statistical tests which may be used to discriminate among them. In reality, it is unlikely that any of the following models in and of itself is truly appropriate or sufficient for a given dataset given that each model embodies only a single mechanism for change. For example, the results from Carley's (1990) study of network change can be interpreted as suggesting that in any socio-cultural environment multiple mechanisms for tie addition and deletion are at work. When the social theory is broadly correct, it is reasonable to hope that the dominant behavior of the change will be reflected in the fitted model. However, the social theory may over-generalize or ignore contextual factors, or several theories may apply simultaneously with nearly equal force. In either case, the procedures we develop will be inadequate to the difficulties inherent in such applications. These procedures should therefore be viewed as a first step towards organizing the comparative analysis of different processes of change. Nonetheless, we hope that researchers can use these models, and the procedures for contrasting them, to identify which types of processes are strongly influential in the observed data.

In examining these models the reader should keep in mind that there is little agreement over what processes lead to edge addition (let alone deletion). Furthermore, there is disagreement over whether things such as reciprocity, isomorphism, and homomorphism are themselves processes affecting network change or are the result of other cognitive processes such as balance, homophily, and learning. Neither the Holland-Leinhardt models nor the metric models treated below make strong claims about whether reciprocity, isomorphism, or homomorphism are basic processes. The value of these models, in part, is that many socio-cognitive theories can be tested within the confines of these models. In contrast, the triad completion, degree variance, polarization/balkanization, and constructural models all make relatively strong claims about socio-cognitive processes. In these models reciprocity, isomorphism, and homomorphism are not basic processes.

The following discussion assumes that the data consist of successive observations upon a single social network (digraph) with a fixed set of m distinguished nodes. Edges are directed, and loops are disallowed (however, it is generally straightforward to extend the

models to cases with undirected edges and/or loops). There are a maximum of $r = m(m-1)$ edges in a directed, loopless network.

We let $X(t)$ denote the observed network at time t; here $X(t)$ is an $m \times m$ adjacency matrix with entries $x_{ij}(t)$. Specifically,

$$x_{ij}(t) = \begin{cases} 1 & \text{an edge runs from node } i \text{ to node } j \text{ at time } t \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

Since the edges are directed, the adjacency matrix will generally be asymmetric. The dataset $X(t_1), \ldots, X(t_n)$ is called a trajectory. In this paper we assume that one has observations only upon a single trajectory, but much of the discussion could be generalized to the analysis of multiple trajectories (cf. Sanil, Banks, and Carley (1994) for an examination of multiple trajectory data).

2.1 Holland-Leinhardt Models

Log-linear models for social networks were first proposed by Holland and Leinhardt (1981). We shall refer to the most general such exponential model as the Holland-Leinhardt model, and follow their lead in referring to a usefully simpler submodel as the p_1 model. The motivation behind these models is a nice combination of statistical tractability and social theory. Although the p_1 model was intended for use in the analysis of a single network at one point in time, Wasserman (1980) derives a generalization which defines a network-valued process.

The Holland-Leinhardt model employs parameters which flexibly capture such social properties as reciprocity and attractiveness, and it is one of the mainstays of network analysis. The model has been extensively developed by many authors, notably Wasserman and Galaskiewicz (1984) and Fienberg, Meyer, and Wasserman (1985). The latter find that the three chief drawbacks to its use are:

1. All dyads (pairs of nodes) form edges independently.

2. The model can be excessively flexible.

3. There is no principled way to perform a goodness-of-fit test.

Obviously, the last deficiency is common to all interestingly realistic models of network behavior, and the second provoked the development of the p_1 version of the Holland-Leinhardt model. The first concern is true of many mathematically tractable procedures and moving away from this assumption can reduce the chances of having a tractable model.

For a network observed at a single time-point, the Holland-Leinhardt probability model on $X = [x_{ij}]$ is

$$\mathbb{P}(X) = \exp[\sum_{i<j} \rho_{ij} x_{ij} x_{ji} \; + \; \sum_{i \neq j} \theta_{ij} x_{ij}\,] \prod_{i<j} n_{ij} \tag{2}$$

where $\rho_{ij} = \ln \frac{m_{ij} n_{ij}}{a_{ij} a_{ji}}$ for $i < j$, $\theta_{ij} = \ln \frac{a_{ij}}{n_{ij}}$ for $i \neq j$ with $n_{ij} = n_{ji}$ for $i > j$, and $m_{ij} + n_{ij} + a_{ij} + a_{ji} = 1$ for $m_{ij} = \mathbb{P}[x_{ij} = x_{ji} = 1]$, $n_{ij} = \mathbb{P}[x_{ij} = x_{ji} = 0]$, and $a_{ij} = \mathbb{P}[x_{ij} = 1, x_{ji} = 0]$. Essentially, this model describes a network in which each dyad forms edges according to its own probability distribution, and the behavior of one dyad confers no information about the behavior of any other dyad.

The p_1 model is obtained from (2) by adding the constraints that $\rho_{ij} = \rho$ for all $i < j$, $\theta_{ij} = \theta + \alpha_i + \beta_j$ for all $i \neq j$, and $\sum_i \alpha_i = \sum_j \beta_j = 0$. This restriction enables a satisfying interpretation of α_i as a productivity parameter which controls the outdegree at node i, β_j as an attractiveness parameter which controls the in-degree at node j, ρ as the average tendency towards reciprocation of edge formation in the network, and θ as a density parameter which controls the expected number of edges. Thus one can describe network behavior using theories that posit differential attractiveness of individuals and the predispositions to form reciprocal relationships, but the dyadic independence prevents the use of theories that involve cliquing, hierarchy, or transitivity. This limitation applies to both the p_1 model and the more general Holland-Leinhardt model.

Several strategies enable one to falsify the p_1 model upon narrow grounds. Fienberg et al. (1985) describe log-linear model tests that examine the tenability of the defining constraints. For example, the p_1 model is equivalent to fitting the "no three- or four-factor interaction" log-linear model, and the usual goodness-of-fit test applies. Similar tests check the constancy of ρ or θ, and enable precise diagnosis of any lack-of-fit that is discovered. However, all of these tests assume independent dyads.

To falsify the general Holland-Leinhardt model, one must weigh evidence of dyadic dependency. The most common test looks for an excess or deficit in the numbers of different

kinds of triadic relationships. Specifically, there are 16 patterns of edges possible for a given triplet of nodes. One can compare the observed counts of each type of pattern to their expected values conditional on the observed in- and out-degrees; if the observed counts differ significantly from expectations, this is evidence that the model fails to fit. Holland and Leinhardt (1975) detail the calculations for the expected values and covariance matrix of the U—MAN model; Holland and Leinhardt (1981) use this technique (together with conditioning and a multivariate normal approximation) to show that their model fits poorly to an example dataset, but is nevertheless a valuable first step. Of course, this procedure is only sensitive to model failures that manifest in the triad counts. However, Holland and Leinhardt's (1981) goal is to test the p_1 model, and it is not intended as a test of the more general Holland-Leinhardt model. In general, these tests that allow the researcher to examine triad completion at some level, are helpful but not very satisfactory. There is a dearth of theory on omnibus goodness-of-fit tests for these models and such work is necessary for a more satisfactory test.

The foregoing description sets the stage for the examination of network behavior through continuous time models. Holland and Leinhardt (1977) and Wasserman (1980) develop models in which dyadic behavior at a given time is independent, but across time, there is Markov dependency upon the past; also, for current work, see Leenders (forthcoming a and b). Let $\mathcal{S}$ be the set of all social networks on m distinct nodes, and for $s, s' \in \mathcal{S}$, define the transition probability function

$$\begin{aligned} \mathbb{P}_{s,s'}(t,h) &= \mathbb{P}[\, X(t+h) = s' \,|\, X(t) = s \,] \\ &= \prod_{i \neq j} \mathbb{P}[\, x_{ij}(t+h) = s'_{ij} \,|\, X(t) = s \,] + o(h) \quad \text{as} \quad h \downarrow 0 \end{aligned} \tag{3}$$

where s'_{ij} is one or zero according to whether an edge runs from node i to j in s'. This ensures that dyads change independently.

From this perspective, many dynamic models can be built through specification of the infinitesimal transition rates; i.e., one defines

$$\begin{aligned} \mathbb{P}[\, X_{ij}(t+h) = 1 \,|\, X(t) = s,\ s_{ij} = 0] &= h\lambda_{0ij}(s,t) + o(h) \\ \mathbb{P}[\, X_{ij}(t+h) = 0 \,|\, X(t) = s,\ s_{ij} = 1] &= h\lambda_{1ij}(s,t) + o(h). \end{aligned} \tag{4}$$

By choosing the functional form of the $\lambda_{.ij}$, analysts can often capture features predicted by social theory in a Markov model. For example, Leenders (forthcoming a and b) develops a dynamic model of friendship formation. Carley (1990) develops a dynamic model of interaction and group stability in which changes in network ties emerge from changes in knowledge and which, for a small number of individuals and ideas, can be represented as a Markov model. For example, Wasserman (1980) develops a dynamic model in which edge formation is governed by degree centrality, so that nodes with the largest in-degrees tend to attract more edges. A fuller account of this is given in subsection 2.4.

A drawback to the Markov approach is that it entails potentially delicate assumptions and may be impractical with a large number of states. For example, the large number of states in Carley s (1990) model led to the need to employ simulation in order to analyze the model under various conditions. For Holland-Leinhardt models it may, at times, be preferable to apply some simple techniques which directly examine their fit over time. We urge analysts to consider the following exploratory data analysis methods, if the goal is to examine data, before resorting to intensive modeling.

1. To check the adequacy of the assumption of dyadic independence, plot the significance probabilities of triad census tests against time. An approximate test based on the p_1 model is given in Holland and Leinhardt 1981.

   If significance probabilities are small, or if consistent trends appear, then the fundamental Holland-Leinhardt assumption is suspect. Bear in mind that consistency of trend is difficult to spot in correlated significance probabilities, and that the strength of the correlation will probably diminish as the time steps increase. The examination of trend is important since (1) If there are many time points, a few low significance probabilities will occur by chance, but these will form no pattern; (2) As the network moves to stochastic equilibrium (if it does), the significance probabilities should become larger with time, so a young network may seem not to fit the Holland-Leinhardt model, but an older network might; (3) If one sees a pattern, this can do much to suggest the kind of model which may do better than the Holland-Leinhardt, so the examination of trend acts as a diagnostic for the kind of model failure observed.

2. If dyadic independence holds, one should check for drift in the model parameters. To decide whether the same p_1 model holds at each t, one can estimate the parameters ρ, θ, α_i, and β_j for each observation. If these show no trend, then information may be pooled to obtain sharper estimates. In contrast, if there is dyadic dependence, then even without a discernible parameter drift more elaborate models are needed. To check for parameter drift plot the parameter estimates and their standard errors against time. A reasonably important trend should be visually apparent.

   If one is not using a p_1 model, but rather the more general model of (2), then one can still look for trends in the switching times of each edge. Of course, this will require observations at a substantial number of well-spaced time points.

Should the dyads appear independent at each observation, and if the model parameters show discernible drift, then it becomes worthwhile to consider the use of some more elaborate, possibly Markovian, model.

2.2 Metric Models

Metric models for social networks were proposed by Banks and Carley (1994). The motivation derives from statistical practice rather than social theory, and thus, in analogy with the mean and variance of conventional statistics, they build a model which is naturally parameterized in terms of a central network and a dispersion about that network. Sanil et al. (1994) extended the metric model to the analysis of social networks evolving over time. The general approach they describe does not assume dyadic independence and so can be used with theories that posit differential attractiveness of individuals and the predisposition to form symmetric relationships, as well as theories that involve cliquing, hierarchy, or transitivity. We summarize this work, and describe a further extension which may increase the range of applicability.

The cornerstone of a metric model is a probability measure over the set of $\mathcal{S}$ of directed networks on m distinct nodes which has the following form:

$$\mathbb{P}_{(s^*,\sigma)}[\, s \,] = c(s^*, \sigma) e^{-\sigma\, d(s,s^*)} \qquad \forall\, s \in \mathcal{S} \tag{5}$$

where $d : \mathcal{S} \times \mathcal{S} \to \mathbb{R}^+$ is a metric, s^* is the unknown central network, and σ is a measure

of concentration ($1/\sigma$ is the corresponding dispersion parameter). Thus the probability of observing a particular social network s diminishes as its distance from s^* increases. Also, for σ=0, one has uniform measure on $\mathcal{S}$, whereas $\sigma = \infty$ implies one observes s^* with probability one.

This formulation is related to Mallows' method (1957) of setting probabilities on the set of permutations. The choice of the metric d is crucial, and should reflect a sense of nearness appropriate to the context of the study. For example, if one wants to model reciprocity, then the distance metric should declare two networks which differ by a reciprocating edge change to be closer than two networks which differ by a non-reciprocating edge change. Similarly, one can model the degree centralization effect by selecting a metric that makes two networks which differ by an edge change that increases the in-degree of a popular node closer than two networks which differ by an edge change that increases the in-degree of an unpopular node.

If one has an i.i.d. sample $s_1, \ldots, s_n$ from (5), then the maximum likelihood estimates of s^* and σ must satisfy:

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^{n} d(s_i, \hat{s}^*) &= \frac{\sum_{s \in \mathcal{S}} d(s, \hat{s}^*) e^{-\hat{\sigma} d(s, \hat{s}^*)}}{\sum_{s \in \mathcal{S}} e^{-\hat{\sigma} d(s, \hat{s}^*)}} \\ \hat{s}^* &= \text{argmin}_{s \in \mathcal{S}} \; \hat{\sigma} \sum_{i=1}^{n} d(s_i, \hat{s}^*) \; + \; n \log \sum_{s \in \mathcal{S}} e^{-\hat{\sigma} d(s, \hat{s}^*)}. \end{aligned}$$

Solving these equations requires enumeration of the elements of $\mathcal{S}$. For even moderate m it quickly becomes impossible to maximize the likelihood through complete enumeration; unless the metric is unusually tractable, one must resort to numerical search.

A useful characterization of (5) is possible in a common simple case. Frank and Strauss (1986) used the Hammersley-Clifford theorem to show that all probability measures on $\mathcal{S}^+$, the set of undirected networks, can be written in the form

$$\mathbb{P}_D[\, s \,] = c \exp[\sum_{A \subseteq C} \alpha_A] \quad \forall \; s \in \mathcal{S}^+ \tag{6}$$

where c is a normalizing constant and α_A is a nonzero constant if and only if A is a clique of the nonrandom dependence graph D. In our context, the nodes of D are all possible undirected edges on m vertices; a clique in D is a subset of the vertex set of D that is either a singleton set or has the property that all pairs of elements are connected by edges in D.

The graph D determines the dependence structure between random edges; if D connects a specific pair of edges, then those edges in s are conditionally dependent given the other edges in s.

For undirected networks, the dependence graph D is edgeless if and only if one has a Holland-Leinhardt model. To show that the Holland-Leinhardt model is a special case of the metric model, identify the network s with its adjacency matrix; thus $s_{ij} = 1$ if and only if an undirected edge links nodes i and j. The Frank and Strauss representation enables the following result.

Theorem 1: A probability measure on $\mathcal{S}^+$ is Holland- Leinhardt (i.e., has independent, but possibly differently distributed, dyads) if and only if it can be written as (5) for a semimetric of the form

$$d_{HL}(s_1, s_2) = \sum_{i<j} a_{ij}|s_{1ij} - s_{2ij}| \tag{7}$$

where the $a_{ij} \geq 0$. (A semimetric allows zero distance between distinct objects; it corresponds to taking $a_{ij} = 0$ for some (i, j).)

Proof: See the appendix.

The semimetric in Theorem 1 may be viewed as a weighted version of the Hamming metric on $\mathcal{S}^+$. As the proof shows, the a_{ij} are the log odds ratios of the edge probabilities, where the numerator of the ratio contains the least probable outcome. The theorem states that the assumption of independent dyads which characterizes the Holland-Leinhardt model is equivalent to a concept of distance based upon a weighted sum of edge discrepancies. This contrasts with alternative metrics, entailing edge dependencies, which are sometimes more appropriate. An analogous result holds for directed networks.

The dynamic version of the metric model for directed edges implements the intuition that network evolution is more likely to occur by short steps than large leaps. To formalize this, Sanil et al. (1994) proposed the discrete-time Markov model

$$\begin{aligned} \mathbb{P}[d(X(t), X(t-1)) = j | X(t-1), \ldots, X(1)] &= \mathbb{P}[d(X(t), X(t-1)) = j | X(t-1)] \\ &= \binom{N}{j} f(\gamma, t)^j (1 - f(\gamma, t)^{r-j} \end{aligned} \tag{8}$$

where $f(\gamma, t)$ such that $0 \leq f(\gamma, t) \leq 1$ is a parametric function controlling the expected

rate of change, and the metric on $\mathcal{S}$ is

$$d(s_1, s_2) = \mathrm{tr}[(s_1 - s_2)^T (s_1 - s_2)] = \sum_{i \neq j} |s_{1ij} - s_{2ij}| \tag{9}$$

for s_1, s_2 the adjacency matrices identified with networks s_1, s_2. Also, N is $m(m-1)$ if directed edges are allowed and there are no loops. Whereas, N is $m(m-1)/2$ if undirected edges occur and there are no loops. This metric implies that whether or not edge changes are independent can change over time; the starting point can be determined nonrandomly or by some dependent random mechanism. If the rate of change is slow, then this dependence will persist. Moreover, this metric implies that the conditional distribution of X_t given $X_(t-1)$ is such that all edges change independently with the same probability f(gamma,t). Therefor, the number of edge changes is a for $f(\gamma, t)$ at time t. If $f(\gamma, t)$ changes smoothly over time, then the sufficient statistic consists of the number of edge changes at each time point, since there is information in the entire trajectory. In typical applications, $f(\gamma, t)$ is monotonic decreasing in t. For example, if $f(\gamma, t) = \exp(-\gamma t)$, then as $t \uparrow \infty$, the expected rate of change slows and probability mass would concentrate upon a single network.

Under (8), the reduced log-likelihood equation is

$$\ell(\gamma | X(1), \ldots, X(n)) = \sum_{t=1}^{n-1} [d_t \ln f(\gamma, t) + (r - d_t) \ln(1 - f(\gamma, t))] + \ln P(X(1)) \tag{10}$$

where $d_t = d(X(t+1), X(t))$. Then, assuming that $P(X(1))$ is independent of γ, the maximum likelihood estimate $\hat{\gamma}$ satisfies

$$\sum_{t=1}^{n} \left(\frac{\partial f}{\partial \hat{\gamma}} \right) \frac{f(\hat{\gamma}, t) - (d_t / r)}{f(\hat{\gamma}, t)(1 - f(\hat{\gamma}, t))} \quad = \quad 0. \tag{11}$$

There is an appealing interpretation of (11). If we consider the binomial distribution at each time point, then one has a weighted sum of the difference between the estimated and the observed proportion of relations that change at each time step, divided by the variance of the estimate. In other words, the equation seeks the $\hat{\gamma}$ which minimizes a weighted sum of standardized discrepancies. The weights on each term are the rate of change of $f(\gamma, t)$ with respect to γ at time t. For general functions $f(\gamma, t)$, one solves (11) numerically (multiple solutions are possible).

To test the fit of a model of the form (8), one can use the fact that under the null hypothesis that the model is correct, one knows the distribution of the number of edges

that change at each step. Define

$$p_t = \mathbb{P}[X \leq d_t] = \sum_{j=1}^{d_t} \binom{r}{j} f(\hat{\gamma}, t)^j (1 - f(\hat{\gamma}, t)^{r-j} \qquad t = 1, 2, \ldots, n-1. \tag{12}$$

Were γ known, this would be the probability integral transformation, and the $p_1, \ldots, p_n - 1$ would be an i.i.d. sample from a discrete distribution that is nearly the continuous distribution $\mathcal{U}[0, 1]$. But the effect of estimating γ should be small, so the Kolmogorov-Smirnov test enables an approximate assessment of whether the p_t are (nearly) uniformly distributed. If they are not, then the model does not hold.

The principal advantage of the model in (8) is its tractability, which derives from its close connection to the binomial distribution. But it possesses the undesirable property that for a slowly decreasing function f, its asymptotic distribution concentrates on a network that is uniformly distributed on $\mathcal{S}$.

Before leaving the metric perspective, we propose an analog of the linear model decomposition which may have value in describing social networks. In particular, it offers a plausible model for measurement error.

Consider measurements taken upon handshaking relationships in an office. It is arguable that the measurements consist of three components: a fixed component, corresponding to relationships enforced by the organizational structure; an evolving component, corresponding to the formation and dissolution of friendships; and a random component, corresponding to reporting error or chance encounters in the elevator. This suggests the model

$$X(t) = (M \oplus Y(t)) \otimes E(t), \tag{13}$$

where $X(t)$ is the observed network at time t, M is the constant relationship network derived from the organizational structure, $Y(t)$ is the set of evolving relationships (whose stochastic structure could be governed by whatever social theory seems most apt), and $E(t)$ represents measurement error. Here $\oplus$ denotes the operation of edge union, and $\otimes$ denotes the operation of addition modulo 2 on the components of the adjacency matrices. The error networks $E(1), \ldots, E(n)$ are i.i.d. according to (5), where the central network s^* is the edgeless network.

It is unclear whether this linear models approach will prove useful, but we are optimistic.

For certain metrics it is easy to obtain the maximum likelihood estimates of the components in (13), but for others one must resort to numerical techniques.

2.3 Triad Completion Models

Triad completion models are derived from the work by Heider (1958) on balance. Heider's theory pertains to the perceptual level. According to this theory, individuals seek to minimize perceived imbalance. Thus, if Isaac considers himself a friend of Johanna, and Isaac perceives that Johanna is friends with Kirsten, then if Isaac is not friends with Kirsten there will be imbalance. There are several points here. First, analyses of balance, to be true to the original Heiderian formulation, should use socio-cognitive data on each individual's perception of all other pairs (see Carley and Krackhardt, forthcoming). Second, balance theories are all concerned with the directedness of relations; however, the theories differ in whether and how (at the sociometric level) they expect the tie to be reciprocated. According to Heider, two individuals are prone to interact at time t_{q+1} if, at time t_q, they are both interacting with the same third partner (mutual friend). Heider's theory accounts for both positive and negative relations, and predicts that both kinds evolve towards balance. Our present focus, in this paper, is exclusively upon positive relations. This restriction has little practical impact, since researchers rarely collect data on negative ties (for exceptions see Sampson, 1968, and Blau, 1967). The role and existence of negative ties is an important study for future investigation.

Balance theory, and the related triad completion models, have been extensively studied in the social sciences. This has led to a series of interrelated models; these include Heider's "balance model" (1946, see also Cartwright and Harary, 1956), Davis' "clustering" model (1967, 1970), Holland and Leinhardt's "transitivity" model (1971), Davis and Leinhardt's "ranked cluster's of cliques" model (1972), Newcomb's "positive balance" model (1968), and Johnsen and McCann's "direct positive influence" model (1982). Johnson (1986b) gives a discussion of the relationships among these models. We shall not recapitulate the details of these variations, but rather focus upon procedures for testing standard features common across triad models.

To clarify the exegesis, we shall again briefly violate our convention and discuss networks

with undirected edges. This deviation reflects the fact that there are 16 kinds of triadic relationships for directed edges, but only four kinds for undirected edges, and thereby we can avoid conceptual clutter entailed by keeping account of direction. It also means that issues of reciprocity will not arise. (For directed edges, we note that while a balancing process does not guarantee reciprocity, at least in the short run, it does tend towards reciprocity in the long run. We also note that, theoretically, balance is an alternative to reciprocity as a process for creating network change.) But all of the points that we make in the undirected case extend in a straightforward way to the more complicated directed case.

For networks with undirected edges, triad completion models assert that changes between $X(t_q)$ and $X(t_{q+1})$ are likely to be edge additions that link nodes which both have an edge to a common node. Formally, if $x_{ij}(t_q) = 0$, then

$$\begin{aligned} &\mathbb{P}[x_{ij}(t_{q+1}) = 1 \mid \exists\ k \text{ such that } x_{ik}(t_q) = 1 \text{ and } x_{kj}(t_q) = 1] > \\ &\mathbb{P}[x_{ij}(t_{q+1}) = 1 \mid \not\exists\ k \text{ such that } x_{ik}(t_q) = 1 \text{ and } x_{kj}(t_q) = 1]. \end{aligned} \tag{14}$$

A subnetwork of three nodes that are completely linked is called a completed triad.

Variations of the triad completion model may add some or all of the following conditions:

A.1 Edges that are part of a completed triad are less likely to disappear in the next time step than those that are not.

A.2 For a triad (i, j, k), the occurrence of the first edge in the subnetwork has the same probability as the occurrence of the second edge, provided there is no completed triad formed with a fourth node at either addition.

A.3 For a triad (i, j, k) with exactly two edges between them, the disappearance of the first edge has the same probability as the disappearance of the second edge, provided that neither edge forms a completed triad with a fourth node.

The addition of A.1 implies that completed triads are relatively stable. The use of all four conditions is the strongest form of the triad completion model, and greatly restricts evolutionary behavior. Naturally, more complicated sets of elaborating conditions can arise when edges are directed.

One can use simple tests to corroborate the triad completion model. For example, at time t_q on an undirected network, there are $r/2$ possible changes. Let r_1 be the number of potential edge additions that would complete a triad, and r_2 be the number of potential edge additions that would not complete a triad, Also, let Y_1 count the number of observed edge additions that completed a triad in $X(t_{q+1})$, and Y_2 count the number of observed edge additions that did not complete a triad. Then to verify the defining condition (14), one can use the elementary test of binomial proportions to decide whether the sample proportion Y_1/r_1 is significantly greater than Y_2/r_2. Each of the three subsidiary conditions A.1-1.3 can be examined in a similar fashion.

However, these simple tests disguise several deficiencies:

1. If the triad completion model holds, then asymptotically, the observations will be in stochastic equilibrium around the completely connected network (the variance will depend upon the propensity to delete edges). At this point, it is extremely difficult to distinguish between the triad completion model and any model that implies high edge density since r_1 and r_2 will be small, lowering the power of the test of binomial proportions. Indeed in a very dense network the power of any test will be low.

2. The test does not combine information across time periods; one can test for triad completion tendencies between t_1 and t_2, but one would like to combine this with information on the change from t_2 to t_3, and so forth, in order to make a statement about the adequacy of the model over the entire trajectory.

3. It is unclear how one should handle cases such as the following:

   - At time t_q, edges link i and j to k, but no edge links i to j; at time t_{q+1}, the edge linking i to j is added, but the edge linking i to k is removed.
   - If at time t_q, an edge between i and j would complete triads with nodes k and l, does that edge have a greater probability of occurring than an edge that completes only a single triad?

   Many possible conventions could remove these ambiguities, but the best solution probably depends upon the social context of the problem. Again, this is a place where our

statistical models easily outrun our social theory.

4. In each test, the null model is that the counts are binomial with parameter values embodying a different relationship than that implied by the particular condition under examination. This null model offers a decent qualititative approximation to more complex models of behavior, but interpretation requires caution. Thus the tests are chiefly corroborative—rejecting the null supports the triad completion model, but does not exclude similar models; failure to reject the null invalidates the triad completion model, but does not prove that a simple binomial model is correct.

We describe a natural methodology for addressing some of these concerns, which may be useful when one wants a more stringent examination of the data than simple binomial tests enable.

The first test we proposed above (based on triad counts) is usually too simplistic, but it pertains to the model which leads to the triad census tests invented by Holland and Leinhardt (1978). Their approach is a conditional test; for the case of undirected links, one fixes the number of edges, and then decides whether the observed number of completed triads is greater than would be expected from i.i.d. uniform placement of the edges. (In usual applications this is slightly more complex; one has directed links, and takes weighted counts of the different kinds of partial triads.) In contrast, our recasting of the test for this situation depends upon the edge census, not the triad census.

We now propose a second test, that is related to the first, but adds important refinements. This test is tailored to reject the null model against more refined specifications of the alternative. These refinements embody sensible extensions of the broad social theory, such as might arise in plausible situations. For example, this test repairs the second deficiency by examining multiple networks over time. It begins to address the third deficiency by suggesting a mechanism for thinking clearly about the different types of edge additions and deletions that may affect behavior. However, the first and fourth deficiency still apply. This second test is based on the recognition that an edge addition to $X(t_q)$ may complete a triad, or it may provide the second link in some triad but does not complete any triads, or it may provide at most the first link in all triads. Count the number of such possible additions; denote them by $r_1(t_q)$, $r_2(t_q)$, and $r_3(t_q)$, respectively. Similarly, an edge dele-

tion in $X(t_q)$ may break one or more triads, or it may break one or more pairs of edges on a common node, or it may remove an isolated edge. Count the number of such possible deletions; denote them by $r_4(t_q), r_5(t_q)$ and $r_6(t_q)$, respectively. Then $\sum\ r_i(t_q) = r/2$.

Similarly, let $Y_i(t_q)$ denote the actual number of edge changes of each of the six types that occurred in $X(t_{q+1})$. Then $(Y_1(t_q)/r_1(t_q), \ldots, Y_6(t_q)/r_6(t_q))$ is a vector of sample proportions; it is a straightforward exercise to calculate its covariance matrix. Then one can perform a simultaneous test (cf. Miller, 1981) to see whether the orderings imposed by the conditions chosen to capture the sense of the social context are actually rejected by the data. If they are violated, then the triad completion model fails.

To pool the results of such a test over time, one can use Fisher's procedure (Fisher, 1934, p. 101). For each t_q (except the last), one obtains a significance probability p_q for the hypothesis that the conditions hold at time t_q. It is reasonable to assume these significance probabilities are independent, and under the null hypothesis that the triad completion model holds, each p_q is uniformly distributed on $[0, 1]$. Then $-2 \sum\ \ln p_q$ has the chi-squared distribution with $2(q-1)$ degrees of freedom under the null, and tends to be larger when the model fails. This permits a test of the model which combines information across the entire span of observation.

Two final notes on this topic. First, one could use the continuous-time Markov process model in (4) to describe triad completion behavior. To do this, one would allow the $\lambda_{0ij}(s,t)$ and $\lambda_{1ij}(s,t)$ to increase or decrease depending upon whether the formation of the edge completes a triad or removes an isolated edge, respectively. A similar approach is suggested by Leenders (forthcoming b). But testing the fit of this model would not be easy. Secondly, since triad completion models are close transitivity models one could use a transitivity statistic (Snijders, 1990).

2.4 Degree Variance

Degree variance models derive in part from the work by Blau (1967) and others on exchange theory. These models assume that nodes differ in their intrinsic probability of attracting an edge. Nodes with high degree centrality tend to attract many edges, while those with low degree centrality tend to attract few. You might think of this process as involving individual

or structural attributes like charisma, popularity, or power. In its purest form, the theory makes no predictions about reciprocity.

Kapferer (1972) used a degree variance model to argue that, over time, there would be an increase in interaction between senior and junior workers, and an increase in interaction between supervisors and those who directly reported to them. Mathematical formulations of exchange theory have been less closely examined than the triad completion model, although some studies have been undertaken by Cook, Emerson, Gilmore, and Yamagishi (1983), Friedkin (1993), Bonacich (1987), Markovsky, Willer, and Parron (1988), and Skvoretz and Willer (1993) which look at this with an over time perspective. Much of the current theoretical debate in this area concerns the basis for attributing degree variance, which in turn determines the propensity for attracting and maintaining ties and whether or not ties will be reciprocated.

For clarity, and since one expects evolving degree variance to be observable in networks at fixed times, we temporarily ignore evolution to concentrate on measuring degree variance in single networks. Indeed, much of the extant work in this area takes exactly this approach; e.g., the p_1 model, and variations of the U—MAN model such as the approximation proposed by Holland and Leinhardt (1981) in which theU—MAN model is conditioned on the variances of the in and the out degrees and on the correlation between in and out degrees. This particular literature contains tests of the models relative to the null hypotheses that there are no node differences and edges are i.i.d.. Additionally, Snijders (1991) provides a test of these models as null hypotheses.

A standard approach to assessing degree variance is due to Snijders (1981), who proposed a test based on the comparison of the observed variance in the in-degrees to the variance expected under a model of completely random edge connection (conditioning on the observed numbers of edges). That procedure seems more sensitive to the null model than a method based on mixtures of binomials; additionally, the mixture of binomial approach obtains substantially more information about any degree variant structure which may exist (such as the number of components, and the proportions of nodes at each of the different levels of degree centrality).

It is also possible to test a degree variance model as a submodel of the p_1 model, in

which all of the attractiveness parameters β_j are zero. But other versions of the degree variance model are possible, and these need not incorporate the very strong independence assumptions of the p_1 model. For example, the mechanism underlying degree variance may be that individuals look around to see who is more popular or powerful and then throwing their vote in with the rest of the crowd. Thus an analyst should be cautious in interpreting log-linear model tests about the β_j coefficients.

Unlike these models we focus on identifying components in a mixture distribution. We see this as an important step to understanding changes in degree variance over time as much of the theoretical literature in this area focuses on how classes of nodes and not just individual nodes change (see for example Kapferer 1972). To examine degree variance at a fixed time t_q, let Z_i be the number of in-edges to node i in $X(t_q)$. If one assumes that in-edges occur independently and that all nodes are equally degree central, then the $Z_1, \ldots, Z_m$ are a random sample from a binomial distribution $Bin(m, p)$, where p is the unknown probability of an in-edge. Alternatively, if the degree variance model holds and nodes have differential levels of attractiveness, then the $Z_1, \ldots, Z_m$ are a random sample from a mixture of binomials; formally, the probability mass function for Z_i is given by

$$\mathbb{P}[Z_i = z] = \sum_{j=1}^{c} \pi_j \binom{m}{z} p_j^z (1 - p_j)^{m-z} \tag{15}$$

where c is the unknown number of components in the mixture, and the π_j are the mixing proportions. Perhaps $c = m$, in which case each node has a unique degree, or perhaps c is much smaller. The latter case might arise if there were a small number of different kinds of nodes; for example, if there were one supervisor and many subordinates, then the number of components might be two. The supervisor would have a high degree centrality score, and the others would have lower but equal degree centrality scores.

The problem of identifying the components in a mixture distribution is classically difficult, especially when the number of components is unknown. Recently, there has been considerable progress on this front. Crawford et al. (1992) describe a Bayesian approach to this problem, and Roeder (1992) gives a frequentist approach. Either avenue offers interesting possibilities for advancing social network data analysis. Both offer methods for assessing the accuracy of their inferences.

Symmetrically, one could consider a corresponding study of the tendency of nodes to send out-edges to other nodes. If one counts the number of out-edges for each node at time t_q, then one can perform an exactly analogous analysis to distinguish the model of independent equiprobable out-edges from a model in which the nodes are unequally sheepish. As before, one assesses the adequacy of a simple binomial model against a mixture of binomials. The extent to which reciprocity holds again depends on the exact model.

To assess degree variance over time, one wants to see whether some nodes maintain consistently high probabilities of receiving in-edges, and/or consistently lower probabilities of losing edges. A variety of techniques can be used. For example, if there are only two time points then a test for the null hypothesis of equal binomial probabilities can be used, such as a chi square test. Alternatively, Snijders (1990) proposed a goodness of fit test for a model of network change in which the nodes varied in their tendency to receive and send edges. Additionally, the researcher can look for over time trends in tie addition and deletion. Whether nodes vary consistently in their tendency to receive or send ties can be assessed by inference on binomial proportions, using procedures similar to the procedures described in subsection 2.3 One has to count the number of nodes that have not sent edges to i at time period t_q, and record the proportion of those that send edges by time t_{q+1}. By comparing the vector of such proportions for node i to a similar vector for node j, one can make a simultaneous test of whether the tendency of node i to receive or send ties is greater than that of node j.

As described in subsection 2.1, Wasserman (1980) developed a Markov model for network evolution that adjusts the form of the infinitesimal transition rates to build in certain types of evolutionary behavior. One example which is detailed in that paper is the "popularity model," which implements the basic degree variance strategy. The structure of the rates is

$$\begin{aligned} \lambda_{0ij}(s,t) &= \lambda_0 + \pi_0 s_{+j} \\ \lambda_{1ij}(s,t) &= \lambda_1 + \pi_1 s_{+j} \end{aligned}$$

where s_{+j} is the in-degree of node j, λ_0 and λ_1 reflect the number of edges overall, and π_0 and π_1 measure the influence of the current number of in-edges and non-edges on node j's popularity. It turns out that this model is tractable, and enables Wasserman to calculate

the asymptotic equilibrium distribution of the network in terms of the model parameters.

Another over time test of degree variance would examine whether the variance in node centrality increases or decreases over time. Assuming limited time to interact the variance in node centrality should increase over time if popularity/attractiveness/charisma thesis holds. This captures a different aspect of node variance than the mixture of binomials test does. The chief difficulty in practice is that one typically has only two or three time points, so a regression test of increasing variance with only two or three x values will be unpersuasive.

At the component level, networks at two points in time can be compared by first locating the components, and then determining whether the components have become increasingly differentiated in average degree centrality and the extent to which the members of the component (the individual nodes) are the same over time. Differentiation of density can be handled as previously discussed. Similarity in membership can be calculated by a chi square test over group members being the same or different at the two time periods, or by doing a QAP correlation on the membership matrices.

Generally speaking, in equilibrium, it becomes relatively difficult to distinguish nodes, or to distinguish degree variance models from certain types of Holland-Leinhardt models. This occurs because the in-degree of nodes tends to increase, so that all nodes ultimately appear highly and equally degree central; similarly, it is difficult to distinguish the later stages of a degree variance model from a p_1 model when density is high.

2.5 Polarization and Balkanization

Polarization and Balkanization also follow from theoretical conceptions of power, and may be viewed as variations on the degree variance model. Polarization occurs when the society splits into two groups each centered around a specific node, as happens when there are opposing cliques, each with a powerful leader. Balkanization occurs when the society splits into a large number of groups, each centered around a specific node. Theories of polarization and balkanization have traditionally been more conspicuous in the political than the social sciences. Such theories rarely argue either for or against reciprocity. Rather, they are more concerned with the overall structure of the network than particular dyadic relations.

Polarization has not been well-defined mathematically. Broadly, for a single time point, it is said to occur if there is a set $\mathcal{S}_2$ consisting of two nodes such that nearly every other node is linked to exactly one of the two. In contrast, Balkanization has not been defined at all. For our development, it is a natural generalization of polarization, and occurs if there is a set $\mathcal{S}_k$ consisting of more than two nodes such that nearly every other node is linked to exactly one member in that set.

Let $s, s' \in \mathcal{S}_k$ (where k may be 2). Then variations on the polarization and balkanization models can add some or all of the following conditions:

B1 If nodes i and j both link to s, then the probability of a link between i and j is greater than if i linked to s and j linked to s'.

B2 The probability that an edge between nodes i and j is removed is greater if i has an edge to s and j has an edge to s', than if i and j are both linked to s.

B3 For node i the probability that a new in-edge links to s is an increasing function of the in-degree of s. Alternatively, one might argue that the probability for a node i to have a new edge to s is also a function of the involvement of node i with other nodes already linked to nodes in $\mathcal{S}_k$; for example, it may be dependent on the number of paths of length 2 from i to these nodes.

B4 If node i has no links to any node in $\mathcal{S}$, then the probability that it forms such a link is a decreasing function of the cardinality of $\mathcal{S}$.

Notice that extensions B1 and B2 of the definitions add features that make the network behavior resemble networks evolving under triad completion. They also result in s and s being two poles when $k = 2$. Also, B3 says that the rate polarization or balkanization speeds up as the trajectory progresses. Finally, B4 is based on the principle that in individuals or states have limited capabilities to interact and so are less likely to form involvements to any particular node as the number of nodes increases.

As before, one can test models of this kind through tests on simultaneous binomial proportions. The methodology is very similar to that of subsections 2.3 and 2.4, and thus it will not be detailed here. If $\mathcal{S}_k$ is not known *a priori* and must be estimated from the data,

then the involvement of this discrete parameter makes the problem significantly harder and invites the use of Bayesian models. In many respects, polarization and balkanization models represent hybrids of the degree variance model that has already been treated.

Additionally one can construct a test around the poles. The polarization for known poles is characterized by a strong negative association between them. Given the null hypothesis of a random network a test of independence can be used for an kxk table where the cell values are the number of links between the nodes in $\mathcal{S}_k$.

2.6 Constructural Model

Constructural models were initially proposed by Carley in 1986, and more precise mathematical formulations appeared in 1990 and 1991 (see Kaufer and Carley, 1993, for an extended treatment). Similar arguments centering around cohesion and structural similarity were advanced by Friedkin (1993).

According to the constructural paradigm, social network change is the result of the ongoing learning process carried out by all individuals in the community. Specifically, the strength of the tie from individual i to individual j is an increasing function of the ratio of the amount of information i and j share to the amount of information i shares with all other individuals. Reciprocity results in the long run as individuals come to share more information with each other and similar others. However, and importantly for the model, reciprocity need not occur in the short run. As individuals interact and communicate information, the quantity of shared information between the interactants increases; however, the strength of the tie need not increase, if others in the community acquire common information even more rapidly. Ultimately, as long as there are no barriers to communication or forgetfulness, all individuals in the society end up with equal propensity to form ties. In this model, unlike the others discussed, the actual network (who actually interacts with whom) and the expected network (the probability that individual i will choose to interact with individual j in the next time step) are distinguished. The expected network is a weighted network and can be converted to a simple digraph by transforming all probabilities greater than some cutoff to 1 and all others to 0. Asymptotically, the probability of an actual interaction in the observed network comes to equal the value in the expected network. Thus,

asymptotically, the observed network is in stochastic equilibrium around the completely connected network. Moreover, the expected network becomes at equilibrium a completely connected network. In equilibrium, reciprocity is assured.

The constructural theory is based on several well explored social tenets, including homophilia and social relativity. Homophilia was the interaction process described by Homans (1950) as the basis for the production of friendship, and by Durkheim (1912) as the basis for the production of moral solidarity. Homophilia has been empirically demonstrated on a number of dimensions, including age, sex, race, education, social class, and occupation (Lazarsfeld and Merton, 1954; McPherson and Smith-Lovin, 1987; Coleman, 1957, and Laumann 1966). Similarly, social relativity has been viewed as a normative tendency under which the individual chooses actions and attitudes according to a process that involves the evaluation of self relative to others (cf. Burt, 1982, p.1-16). This tendency has also received strong empirical support (e.g., Merton, 1957; Sherif, 1935; Festinger, 1954).

According to the constructural model, the strength of the tie between individuals will oscillate as they learn new information. However, in a connected society, ultimately everyone will know everything that anyone knows. A consequence is that at equilibrium all individuals are tied to all other individuals with exactly the same strength of tie. The constructural model differs from the other models of network change in that it differentiates between the tie among individuals as they perceive it and the tie between individuals as it is observed by others.

The formal analysis of this model for even moderate numbers of nodes or pieces of information goes beyond any conveniently simple theory. Given the current state of social network assessment, probably the best that one can do is to simulate the evolutionary behavior of such networks under a constructural model, and then judge whether the results are consistent with observed data.

The general strategy for such simulation tests is to select a constructural model, start many simulations at the first observed sample point, and then count the proportion of trajectories that are close to the observed trajectory. If the proportion is high, then the model is corroborated. If the proportion is small, then the model fails. There are very delicate questions regarding the extent to which one can use the sample data to estimate

the parameters in the simulations; Snijders (forthcoming) describes a method that addresses such problems.

3 Conclusions

We have presented a bestiary of models for network evolution. The theoretical bases for these models vary. Often the rationales are sufficiently general that an entire class of models follows from the same theoretical paradigm; e.g., one can motivate degree variance and its extreme forms(polarization, and balkanization) from the exchange perspective. Conversely, in their pure form the diverse theoretical perspectives often result in models with indistinguishable long-term behavior. Basic versions of balance theory, exchange theory and constructuralism can all lead to completely connected networks, or networks that oscillate randomly in the neighborhood of the completely connected network.

High levels of asymptotic connectivity are not necessarily a problem in the short run. Many of the models derive from theories that predict some decrease in ties; however, they simply do not predict as much of a decrease as is often observed. Carley (1990), in her re-analysis of Kapferer's (1972) data on change in a tailor shop, found that there were a substantial number of cases where the strength of the tie between two individuals weakened over time, or the tie ceased to exist. However, she found that exchange theory, Heiderian balance theory, and constructuralism were all limited in their ability to account for such decreases. Even though formal representations of these theories could account for some of the observed decrease, they were inadequate to the magnitude of decrease that was observed. This weakness is shared by many of the theories treated in this paper, and we urge the pursuit of theory and modeling that emphasizes mechanisms for the removal of edges.

With respect to the addition and deletion of edges, as a field we run the risk of our methodology outstripping our theory. In the theoretical literature more attention has been focused on the addition of edges than on their removal, although there are allusions to "deletion" processes, such as limited cognition, limited time, limited resources, moving, firing, and death. Each of these processes would result in a different prediction about which edges were to be deleted. Clearly, any statistical elaboration of a theory could arbitrarily

add one or more of these mechanisms. Our point, in part, is that such elaboration should not be done arbitrarily; rather, theoreticians and empirical researchers need to work to develop a clearer understanding of the deletion processes. Once these processes are understood then we can generate the appropriately elaborated statistical model. Ideally, the theoretical and methodological understanding of edge removal and deletion processes would co-evolve. Finally, whether or not edge additions and deletions can be treated symmetrically in a model depends on the underlying theory chosen. We expect that in most models edge addition and deletion will not be treated in a symmetrical fashion.

A second, more methodological, point is that in order to test theories of change (and the associated models), the researcher must generally capture data well before the group reaches equilibrium. From a practical standpoint, such an equilibrium may be unreachable; experience suggests that external factors often intervene well before the simplicity of asymptopia is attained. Nonetheless, our analysis suggests that the closer the group is to this final state, the more indistinguishable key models become. Specifically, all models that emphasize mechanisms for the formation of edges over mechanisms for their removal tend, with time, to become more dense; thus predicted networks are qualitatively similar, and there is little power for tests of alternative models. A caveat here. If the model contains edge deletion mechanisms then the equilibrium may not be the completely connected network. Such models would entail theoretical extensions beyond those discussed here. Nevertheless, given such models it should be possible to distinguish between competing theories even when the group is in equilibrium. This suggests the need for longitudinal data with large enough time steps to observe change but small enough that external factors will not have intervened and disrupted the equilibrium. See also the discussion by Leenders (1995, p. 175).

A third suggestion is that we must recognize the power of the computer; it can obviate the need to seek tractable models which are too delicate for realistic applications. Simulation analysis has been widely adopted in many sciences, and our formal models have not been so successful that we should be slow to follow suit. There is every reason to hope that they will be able to discriminate more finely than the currently popular tools.

Issues of ultimate prediction aside, a central concern is which model (and hence which

underlying theory) best fits a given dataset. Currently, there are few omnibus goodness-of-fit tests available for social network applications. Thus, a possible strategy, and one we have taken in this paper, is to use extant statistical techniques tailored to be sensitive to the specific kinds of departures from the null binomial model specified by the theoretical model. This specificity enables researchers to identify the kinds of social theories that seem to be most influential in directing the changes in the network. The techniques for doing this are fairly straightforward. However, in many cases, it may be the case that several models provide acceptable descriptions of one's data. In this case, there is no magic solution. Further progress in this domain requires a better understanding of goodness-of-fit tests.

4 Appendix

Proof: Assume the model in (5) holds. Since the semimetric in (7) is a sum, using it as the metric in the model enables one to factor the probability mass function into terms that depend only upon each edge. Thus edges are independent and the model is Holland-Leinhardt.

Going the other way, assume the Holland-Leinhardt model holds. If p_{ij} is the probability of an edge between nodes i and j, then

$$\mathbb{P}[s] = \prod_{i<j} p_{ij}^{s_{ij}} (1 - p_{ij})^{1 - s_{ij}}. \tag{16}$$

For the model in (5), the case when $s = s^*$ shows that the normalizing constant is

$$c(s^*, \sigma) = \prod_{i<j} \max\{p_{ij},\, 1 - p_{ij}\}. \tag{17}$$

Now define

$$a_{ij} = \begin{cases} \log \frac{p_{ij}}{q_{ij}} & \text{if } p_{ij} \leq .5 \\ \log \frac{q_{ij}}{p_{ij}} & \text{if } p_{ij} > .5 \end{cases} \tag{18}$$

so that the expression in model (5) can be rewritten as

$$\mathbb{P}[\, s \,] = \prod_{i<j} \max\{p_{ij},\, 1 - p_{ij}\} \;\times \prod_{e_{ij} \text{ in } s, \text{ not in } s^*} \frac{p_{ij}}{q_{ij}} \;\times \prod_{e_{ij} \text{ in } s^*, \text{ not in } s} \frac{q_{ij}}{p_{ij}}. \tag{19}$$

Now suppose that e_{ij} is an edge in s, so that the corresponding term in the Holland-Leinhardt product is p_{ij}. If $p_{ij} > .5$ then e_{ij} is also an edge in s^*, and thus the only term

that appears in (19) is p_{ij}, from the first product. However, if $p_{ij} \leq .5$, then e_{ij} is not in s^*, and the q_{ij} in the first product is multiplied by p_{ij}/q_{ij} in the second, yielding p_{ij} as required. Similar argument holds when e_{ij} is not an edge in s, and thus the Holland-Leinhardt model can be written in the indicated form. $\Box$.

5 References

Banks, D.L. and Carley, K.M. (1994). "Metric Inference for Social Networks," *Journal of Classification*, 11, 121-149.

Blau, P.M. (1967). *Exchange and Power in Social Life*, Wiley, New York.

Bonacich, P. (1987). "Power and Centrality: A Family of Measures," *American Journal of Sociology*, 92: 1170-82.

Burt, R.S. (1980) "Innovation as a Structural Interest: Rethinking the Impact of Network Position on Innovation Adoption," *Social Networks*, 4: 337-355.

Burt, R.S. (1982). *Toward a Structural Theory of Action*, Academic Press, New York.

Carley, K.M. (1986). "An Approach for Relating Social Structure to Cognitive Structure," *Journal of Mathematical Sociology*, 12: 137-189.

Carley, K.M. (1990). "Group Stability: The Socio-Cognitive Approach," in Lawler, E., Markovsky, B., Ridgeway, C. and Walker, H. (eds.), *Advances in Group Processes*, JAI Press, Greenwich, CN.

Carley, K.M., (1991). "A Theory of Group Stability," *American Sociological Review*, 56: 331-354.

Carley, K.M. and D. Krackhardt (forthcoming). "Cognitive Inconsistencies and Non-symmetric Friendship," *Social Networks*.

Cartwright, D. and Harary, F. (1956). "Structural Balance: A Generalization of Heider's Theory," *Psychological Review*, 63: 277-293.

Coleman, J. (1957). *Community Conflict*, Free Press, New York.

Coleman, J.S., Katz E., and H. Menzel (1966). *Medical Innovation: A Diffusion Study*, Bobbs-Merrill Company, New York.

Cook, K.S., Emerson, R.M., Gilmore, M.R., and Yamagishi, T. (1983). "The Distribution of Power in Exchange Networks: Theory and Experimental Results," *American Journal of Sociology*, 89: 275-305.

Crawford, S., DeGroot, M., Kadane, J., Small, M. (1992). "Modeling Lake Chemistry Distributions: Approximate Bayesian Methods for Estimating a Finite Mixture Model," *Technometrics*, 34: 441-453.

Davis, J.A. (1967). "Clustering and Structural Balance in Graphs," *Human Relations*, 20: 181-187.

Davis, J.A. (1970). "Clustering and Hierarchy in Interpersonal Relations: Testing Two Graph Theoretical Models on 742 Sociomatrices," *American Sociological Review*, 35: 843-851.

Davis, J.A. and Leinhardt, S. (1972). "The Structure of Positive Interpersonal Relations in Small Groups," in J. Berger et al. (eds.), *Sociological Theories in Progress*, Vol 2, 218-251, Houghton Mifflin, Boston, MA.

Doreian, P. (1990). "Mapping Networks Through Time," in Weesie, J. and Flap, H. (Eds.), *Social Networks Through Time*, 245-264, ISOR/University of Utrecht, Belgium.

Durkheim E. (1912/1954). *The Elementary Forms of Religious Life*, Free Press, New York.

Festinger L. (1954). "A Theory of Social Comparison Processes," *Human Relations*, 7: 117-140.

Feinberg, S., Meyer, M., and Wasserman, S. (1985). "Statistical Analysis of Multiple Sociometric Relations," *Journal of the American Statistical Association*, 80: 51-67.

Fisher, R.A. (1934). *Statistical Methods for Research Workers*, 5th edition, Oliver & Boyd, Ltd., London.

Frank, O. and Strauss, D. (1986). "Markov Graphs," *Journal of the American Statistical Association*, 81: 832-842.

Friedkin, N. (1993). "Structural Bases of Interpersonal Influence in Groups," *American Sociological Review*, 58: 861-872

Granovetter, M.S. (1973). "The Strength of Weak Ties," *American Journal of Sociology*, 68: 1360-1380.

Granovetter M.S. (1974). *Getting a Job: A Study of Contacts and Careers*, Harvard University Press, Cambridge, MA.

Heider, F. (1946). "Attitudes and Cognitive Organization," *Journal of Psychology*, 21: 107-112.

Heider F. (1958). *The Psychology of Interpersonal Relations*, Wiley, New York.

Holland, P.W. and Leinhardt, S. (1971). "Transitivity in Structural Models of Small Groups," *Comparative Group Studies*, 2: 107-124.

Holland, P.W. and Leinhardt, S. (1975). "Local Structure in Social Networks," *Sociological Methodology, 1976*, ed. D. Heise, Jossey-Bass, San Francisco.

Holland, P.W. and Leinhardt, S. (1977). "A Dynamic Model for Social Networks," *Journal of Mathematical Sociology*, 3: 85-111.

Holland, P.W. and Leinhardt, S. (1978). "An Omnibus Test for Social Structure Using Triads," *Sociological Methods and Research*, 7: 227-256.

Holland, P.W. and Leinhardt, S. (1981). "An Exponential Family of Probability Distributions for Directed Graphs," *Journal of the American Statistical Association*, 76: 33-65.

Homans G.C. (1950). *The Human Group*, Harcourt, Brace and Co., New York.

Johnson, E.C. and McCann, H.G. (1982). "Acyclic Triplets and Social Structure in Complete Signed Digraphs," *Social Networks*, 3: 251-272.

Johnson, E.C. (1986). "Social Networks and Innovation Adoption: A Look at Burt's Use of Structural Equivalence," *Social Networks*, 8: 343-364.

Johnson, E.C. (1986). "Structure and Process: Agreement Models for Friendship Formation," *Social Networks*, 8: 257-306.

Kapferer, B. (1972). *Strategy and Transaction in an African Factory: African Workers and Indian Management in a Zambian Town*, University of Manchester Press, Manchester, England.

Kaufer, D.S. and Carley, K.M. (1993). *Communication at a Distance: The Effect of Print on Socio-Cultural Organization and Change*, Lawrence Erlbaum Assoc., Hillsdale, NJ.

Laumann, E. O. (1966). *Prestige and Association in an Urban Community: An Analysis of an Urban Stratification System*, Bobbs-Merrill, New York.

Lazarsfeld, P. and Merton, R. (1954). "Friendship as Social Process: A Substantive and Methodological Analysis," in Berger, M., Abel, T., and Page, C. (Eds.), *Freedom and Control in Modern Society*, D. Van Nostrand, Princeton.

Leenders, R. (1995). *Structure and Influence. Statistical Models for Dynamics of Actor Attributes, Network Structure and their Interdependence.* Amsterdam: Thesis Publishers.

Leenders, R. (forthcoming a). " Models for network dynamics: A Markovian framework, " *Journal of Mathematical Sociology*

Leenders, R. (forthcoming b). " Dynamics of friendship choices and best-friendship choices, " *Journal of Mathematical Sociology*

March, J.G. and Simon, H.A. (1958). *Organizations*, Wiley, New York, NY.

Mallows, C. (1957). "Non-Null Ranking Models I," *Biometrika*, 44: 114-130.

Markovsky, B., Willer, D., and Parron, T. (1988). "Power Relations in Exchange Networks," *American Sociological Review*, 53, 220-236.

McPherson, J.M. and Smith-Lovin, L. (1987). "Homophilly in Voluntary Organizations: Status Distance and the Composition of Face-to-Face Groups," *American Sociological Review*, 52: 370-379.

Merton, R.K. (1957, 1968). "Continuities in the Theory of Reference Groups and Social Structure," in Merton, R.K. (Ed.), *Social Theory and Social Structure*, Free Press, New York.

Miller, R.G. (1981). *Simultaneous Statistical Inference*, Springer-Verlag, New York.

Newcomb, T.M. (1968). "Interpersonal Balance," in R.P. Abelson et al. (eds.), *Theories of Cognitive Consistency: A Sourcebook*, 28-51, Rand-McNally, Chicago.

Powell, W.W. and DiMaggio, P.J. (1991). *The New Institutionalism in Organizational Analysis*, The University of Chicago Press, Chicago.

Roeder, K. (1992). "Semiparametric Estimation of Normal Mixture Densities," *Annals of Statistics*, 20: 929-943.

Sampson, S.F. (1968). *A Novitiate in a Period of Change: An Experimental and Case Study of Relationships*, doctoral dissertation, Cornell University.

Sanil, A., Banks, D.L., and Carley, K. (1994). "Models for Evolving Fixed-Node Networks: Model Fitting and Model Testing," to appear in *Social Networks*.

Sherif, M. (1935). "A Study of Some Factors in Perception," *Archives in Psychology*, 187.

Skvoretz, J. and Willer, D. (1993). "Exclusion and Power in Exchange Networks," *American Sociological Review*, 58: 801-818.

Snijders, T. (1981). "The Degree Variance: An Index of Graph Heterogeneity." *Social Networks*, 3: 163-174.

Snijders, T. (1990). "Testing for Change in a Digraph at Two Time Points." *Social Networks*, 12:359-373.

Snijders, T. (1991). "Enumeration and Monte Carlo methods for 0-1 matrices," *Psychometrika*, 56(3): 392-417.

Snijders, T. (forthcoming). "Stochastic actor-oriented models for network change," *Journal of Mathematical Sociology*

Wasserman, S. (1980). "Analyzing Social Networks as a Stochastic Process," *Journal of the American Statistical Association*, 75: 280-294.

Wasserman, S. and Galaskiewicz, J. (1984). "Some Generalizations of p_1: External Constraints, Interactions, and Non-Binary Relations," *Social Networks*, 6: 177-192.

Wasserman, S. and Iacobucci, D. (1988). "Sequential social network data," *Psychometrika*, 53: 261-282.