
BIRDS OF A FEATHER, OR FRIEND OF A FRIEND? USING EXPONENTIAL RANDOM GRAPH MODELS TO INVESTIGATE ADOLESCENT SOCIAL NETWORKS*

STEVEN M. GOODREAU, JAMES A. KITTS, AND MARTINA MORRIS

In this article, we use newly developed statistical methods to examine the generative processes that give rise to widespread patterns in friendship networks. The methods incorporate both traditional demographic measures on individuals (age, sex, and race) and network measures for structural processes operating on individual, dyadic, and triadic levels. We apply the methods to adolescent friendship networks in 59 U.S. schools from the National Longitudinal Survey of Adolescent Health (Add Health). We model friendship formation as a selection process constrained by individuals' sociality (propensity to make friends), selective mixing in dyads (friendships within race, grade, or sex categories are differentially likely relative to cross-category friendships), and closure in triads (a friend's friends are more likely to become friends), given local population composition. Blacks are generally the most cohesive racial category, although when whites are in the minority, they display stronger selective mixing than do blacks when blacks are in the minority. Hispanics exhibit disassortative selective mixing under certain circumstances; in other cases, they exhibit assortative mixing but lack the higher-order cohesion common in other groups. Grade levels are always highly cohesive, while females form triangles more than males. We conclude with a discussion of how network analysis may contribute to our understanding of sociodemographic structure and the processes that create it.

The structure of social relations in a population emerges from both the distribution of personal attributes and the dynamics of interaction. For example, the frequency of interracial marriage reflects both the proportion of a population within each race and the norms influencing selection of marriage partners. When individuals can have multiple partners (e.g., friendships), the partnerships may be interdependent, generating complex social network structures. These population-level structures channel the spread of information, norms, disease, support, and other socially exchanged goods. The flow of such entities can in turn generate behavioral change in individuals and demographic change in populations.

A long line of research has examined the influence of social networks on individuals' demographic attributes and behaviors: fertility decline and contraception use (Kohler 1997, 2000; Kohler, Behrman, and Watkins 2001; Valente et al. 1997; Watkins and Danzi 1995), educational attainment (Mare 1991), migration (Massey 1988; Massey et al. 1987), health behavior (Alexander et al. 2001; Valente 2003), and infectious disease transmission

*Steven M. Goodreau, Department of Anthropology and Center for Studies in Demography and Ecology, Campus Box 353100, University of Washington, Seattle WA 98195; e-mail: goodreau@u.washington.edu. James A. Kitts, Graduate School of Business, Columbia University. Martina Morris, Department of Sociology and Center for Studies in Demography and Ecology, University of Washington. The authors thank Mark Handcock, David Hunter, Carter Butts, Marijtje van Duijn, Krista Gile, Deven Hamilton, and Pavel Krivitsky. Steven Goodreau and Martina Morris were supported by funding from the National Institutes of Health (R01-HD041877 and R01-DA012831). James Kitts and Martina Morris were supported by funding from the National Science Foundation (BCS-0433086), while James Kitts received additional NSF funding (IIS-0433637). This research uses data from Add Health, a program project designed by J. Richard Udry, Peter S. Bearman, and Kathleen Mullan Harris, and funded by Grant P01-HD31921 from the Eunice Kennedy Shriver National Institute of Child Health and Human Development, with cooperative funding from 17 other agencies. Special acknowledgment is due Ronald R. Rindfuss and Barbara Entwisle for assistance in the original design. Persons interested in obtaining data files from Add Health should contact Add Health, Carolina Population Center, 123 W. Franklin Street, Chapel Hill, NC 27516-2524 (addhealth@unc.edu). No direct support was received from Grant P01-HD31921 for this analysis.

(Morris and Dean 1994; Morris and Kretzschmar 1997). This literature provides evidence that social networks affect such outcomes and yields novel predictions about the nonlinear nature of these effects.

The reverse causal path is also evident: attributes and behavior may in turn affect selection of network partners. Thus, another branch of social network analysis examines the patterns of relations. Much of this work is descriptive, capturing structural regularities through summary parameters for network concepts such as mean partnership counts (May and Anderson 1988), core-periphery structure (Borgatti and Everett 1999), or clustering (Watts and Strogatz 1998). In other cases, the analysis is explicitly generative, positing a micro-level behavioral model that produces population-level network structure. Examples include “affect balance” in friendships (Davis 1967; Heider 1958) and “assortative mixing” in friendships (Hallinan and Williams 1989; Kandel 1978), marriage (Mare 1991), and sexual partners (Morris 1991) that produce population clustering.

Researchers have typically examined each generative process in isolation; for example, to evaluate triad closure (one’s friend’s friend is likely to become one’s friend), researchers measure the preponderance of relational triangles (three persons are all friends). While this correspondence of process to outcome may seem straightforward, we show in this article that the signatures of different generative processes are not independent. A single statistic (e.g., the triangle count) may reflect multiple processes, all of which need to be considered jointly for proper inference.

The main obstacle to joint inference for multiple network processes has been the lack of a general, tractable framework for statistical estimation. The potential dependence in tie formation leads to endogeneity, which complicates both estimation and intuition. While the formal representation of these processes has been well understood for decades (Frank and Strauss 1986; Wasserman and Pattison 1996), the estimation difficulties were only recently solved (Handcock 2003a; Snijders et al. 2006). Much network analysis has thus been descriptive in nature, relying on models for specific combinations of network processes with limited generalizability (Fershtman 1985; Snijders and Stokman 1987; Wasserman 1977) or on approximations (Frank and Strauss 1986; Wasserman and Pattison 1996) with poor statistical properties (Geyer and Thompson 1992; Handcock 2003a).

We review recent advances in statistical network analysis that surmount these problems and apply them to friendship data from the National Longitudinal Study of Adolescent Health (Add Health). Our goal is to identify the determinants of friendship formation that lead to pervasive regularities in friendship structure among adolescent students. We consider persons’ overall propensity to make friends (*sociality*) as well as their propensity to make friends based on their attributes (*selective mixing*) and based on their friends’ choice of friends (*triad closure*), given the constraints induced by local population composition. These mechanisms operate at different levels—individual, dyadic, and triadic, respectively—while population composition is a community-level constraint. Together, these form a natural set of predictors for the micro-level foundations of population-level relational structure. Analyzing dozens of schools allows us to explore how these effects operate within different population compositions and whether the friendship processes are general or idiosyncratic. Finally, we assess model fit, using both traditional measures and new network-specific diagnostics.

BACKGROUND

Population-level structure emerges as a result of interacting micro-level processes. Consider the observation that the number of triangles in friendship networks is typically greater than expected if ties were generated randomly. This pattern could be produced by a direct propensity for triad closure. However, by concentrating ties within attribute categories, assortative mixing can also create more triangles than in a random network. Variation in sociality can also produce this pattern, as a network with few active and

many less-active persons will have more triangles than one with the same ties distributed randomly. A single population-level signature can thus be produced by processes operating at individual, dyadic, or triadic levels. This illustrates a general challenge for inferring social dynamics from cross-sectional networks; there is rarely a neat correspondence of process and pattern, and statistical methods are needed to tease apart the micro-level foundations of structure.

Statistical methods for generalized network inference have evolved from simple models examining relational reciprocity (*mutuality*; see Holland and Leinhardt 1981) and assortative mixing (Fienberg and Wasserman 1981; Marsden 1981; Morris 1991) to a general flexible framework for estimation and inference in arbitrarily complex models. Known as exponential random graph (ERG) or p^* modeling, this approach models the probability that a relation exists (in logit form) as a linear function of predictors. This formulation looks similar to logistic regression, but in fact each tie is on both sides of the equation and often in multiple predictors, making tie probabilities recursively dependent. This dependence is not a weakness of ERG models but is an inherent feature of relational data, and we model it explicitly rather than treating it as a nuisance parameter. Recent advances in computationally intensive statistics provide practical methods for solving the resulting estimation problems.

What Leads to Sociodemographic Clustering in Networks?


Though there are more than 6 billion people on Earth, we all still live in relatively “small worlds.” There is a much better than average chance that our friends know each other, as relations tend to form discernible clusters, linked together by sparse “weak ties” (Granovetter 1973). This clustering impacts almost every aspect of our lives, from marriage choices and job opportunities to disease exposure. Sociologists have long noted that this clustering tends to connect sets of actors near one another in a space of sociodemographic attributes such as race or education (Blau 1977; McPherson, Smith-Lovin, and Cook 2001), a phenomenon we will call “sociodemographic clustering.” A variety of generative mechanisms may underlie local clustering in networks and its relationship to sociodemographic attributes. We will consider three mechanisms: sociality, selective mixing, and triad closure.

Sociality. There is heterogeneity among individuals in their propensity to establish friendship ties. We refer to this propensity as *sociality* and to individuals as more or less *social*. Sociality is not synonymous with *degree*, an individual’s number of ties, which is the outcome of a stochastic process shaped by both sociality and other factors. When sociality correlates with sociodemographic categories, highly social categories are disproportionately represented within partnerships, influencing population-level mixing patterns.

Selective mixing. Selective mixing is a dyad-level process by which pairs form (or break) relationships based on their combination of individual attributes. A common form is the greater propensity to partner with others having attributes similar to one’s own (*assortative mixing*). For example, romantic partners are often selected based on religious similarity. The opposite tendency (*disassortative mixing*) also exists, as in heterosexual mating. Other patterns of selective mixing are also found, such as asymmetric age matching among heterosexual partners. We distinguish process from outcome by the terms *selective mixing* and *mixing pattern*, respectively.

Assortative mixing is the most studied form of selective mixing. This process may occur because individuals are attracted to similar others (Byrne 1971) or because interaction with similar others (such as those who share a language) is easier (Carley 1991; Rogers and Bhowmik 1970). In the social network literature, the pattern resulting most directly from assortative mixing—a predominance of within-category ties—is often called *homophily* (Lazarsfeld and Merton 1954; McPherson et al. 2001). The term *homophily* is sometimes used to describe both process and outcome; here we restrict its use to outcome, since homophily may be affected by processes other than assortative mixing.

Figure 1. Triad Closure and Assortative Mixing

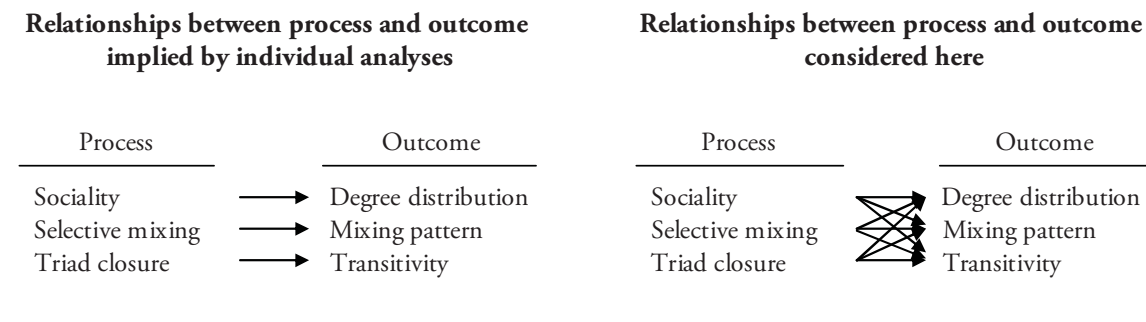


Triad closure. A *triad* is any set of three persons. The classic process of triad closure (Rapoport 1957) implies that triads containing two ties will tend to form the third, thus creating a *triangle* (a set of three persons who are all tied). This may be due to propinquity (two people encounter each other through their shared time with a third) or cognitive processes causing two people to value each other because of their agreement on a third (cf. structural balance theory; Cartwright and Harary 1966; Heider 1958). Again, we differentiate process and outcome, calling the former *triad closure* and the latter *transitivity*.

Assortative mixing can induce transitivity, since increasing the likelihood of within-category ties enhances the opportunity for completed triangles within categories, especially when groups are small. Conversely, triad closure may amplify homophily if there is already a tendency toward assortative mixing (Feld and Elmore 1982). If two ties A-B and B-C are both within an attribute category, the tie created by triad closure (A-C) will also be within-category (see Figure 1, panel a). Either assortative mixing or triad closure can contribute to this triad being closed (in contrast to the triad in Figure 1, panel b). Thus, an accurate estimate of either generative process requires controlling for the other, using information about the attribute composition of all relationships and the count of all triangles. Methods have been developed for generating null models for transitivity that condition specifically on individual variation in sociality (Fershtman 1985; Wasserman 1977) or actor attributes (Snijders and Stokman 1987), but they do not have the generality and flexibility of the ERG framework.

Researchers must also consider population composition when estimating the underlying network processes (Hallinan 1982), as the opportunity for partner selection is constrained by the available pool of partners. Among adolescents, for example, friendships generally form within schools, which are often relatively homogeneous for attributes like socioeconomic status when compared with the overall population. Even if students pick friends at random in each school, attribute homophily will result in the population as a whole. In our case, the Add Health sample represents U.S. adolescents attending 7th–12th grades, but the age, race, and sex composition in each school often differs from the sample as a whole. Since our analysis focuses on the network processes within schools, we condition on their composition.

In summary, we account for patterns in network structure by focusing on the generative dynamics of tie formation given the constraints of population composition. We examine three generative processes—sociality, selective mixing, and triad closure. The methods we use do not assume a simple correspondence between these micro-level processes and the

Figure 2. Process and Outcome in Social Network Models

Notes: Specific forms of selective mixing include assortative mixing and disassortative mixing. Corresponding specific forms of mixing pattern include homophily and heterophily.

population outcomes of degree, homophily, and transitivity, respectively (Figure 2). Rather, the ERG models allow us to tease these apart.

DATA

We use the friendship data from the first wave of Add Health, a sample of more than 90,000 U.S. students in grades 7 through 12, obtained in 1994–1995 through a stratified sample of schools. Subsequent waves include data on various outcomes of interest to demographers (marital and parental status; geographical, educational, and economic mobility; and health outcomes). Add Health offers an unprecedented opportunity to examine how adolescent social networks affect lifelong outcomes; understanding basic processes in these networks is an important first step.

The primary sampling units for Add Health were schools containing an 11th grade. If necessary, researchers randomly selected one or more feeder schools to obtain a school group comprising 7th–12th grades. For simplicity, we use the term *school* to describe a full 7th–12th grade configuration. All students were asked to complete the “in-school” questionnaire, giving an approximate census within school. The questionnaire provided a school roster and asked students to identify their five best male and five best female friends, in order of closeness. Students were allowed to nominate friends outside school or missing from the roster, or to stop before nominating five friends of either sex. Most students listed fewer friends than the maximum,¹ but for the remainder, there may be some truncation. The methodology is detailed in Resnick et al. (1997) and Udry and Bearman (1998) and at <http://www.cpc.unc.edu/projects/addhealth>. Some schools fell far short of the goal of a census of students; we focus on the 59 schools that came closest. This subset is (with one exception) the same as the subset used for the official public-access Add Health network files (for details, see National Longitudinal Study of Adolescent Health [2001] and our online supplement: <http://www.statnetproject.org/papers/published/GoodreauKittsMorris2008>).

We consider the network of mutual friendships, those dyads in which students nominate each other. Because they are reciprocated, mutual friendships are cross-validated and are likely to be stronger than one-way friendship nominations.² Focusing on mutual ties

1. Forty-nine percent nominated fewer than 5 female friends; 49% nominated fewer than 5 male friends; 73% nominated fewer than 10 total friends.

2. Other researchers (Adams and Moody 2007; South and Haynie 2004) have considered symmetrized or asymmetric relations in Add Health, which are interesting topics in their own right, but which represent

may thus reduce the truncation induced by the cap on five nominations of each sex. Given that the question asked for friends in order of closeness, any names that might have been nominated after the first five would presumably represent the weakest friendships and be less likely to represent the mutual friendships that we study here.

The Add Health survey collected substantial information on students' sociodemographic attributes and behaviors. To minimize the potential for endogeneity between attributes and friendship formation, we focus on the exogenous attributes of sex, race, and grade level. These are widely recognized bases of adolescent friendship clustering, but there is virtually no influence by friends on these attributes. We define "race" based on two questions about race and Hispanic origin: Hispanic (all races), black (non-Hispanic), white (non-Hispanic), Asian (non-Hispanic), Native American (non-Hispanic), and other (non-Hispanic). We abbreviate these to Hispanic, black, white, Asian, Native American, and other. The survey also gathered data on participation in clubs and activities. One might expect additional selective mixing on some of these variables, although determining the direction of causality is difficult. (For an analysis of these data that incorporates other person- and school-level variables, see Moody [2001].)

The survey contains multiple forms of missing data, including students who were not on the roster, did not fill out the survey, or left individual questions blank. We exclude the first two categories from our analysis; for the last, we allow an extra "NA" (not applicable) category for each attribute. More elaborate methods for handling missing data in networks are now in development (see Discussion). These methods leave us with a total of 63,408 students (71 to 2,209 students per school).

METHODS

ERG models originated in spatial statistics and network analysis (Besag 1974; Frank and Strauss 1986; Holland and Leinhardt 1981; Strauss and Ikeda 1990; Wasserman and Pattison 1996). Although the general theory is now established, applications have been limited for technical and theoretical reasons. The modeling class is little known in demography, so we provide some expository detail here. We then describe the problem of *model degeneracy* and detail the specific models we examine.

The ERG Model Class

Social network data comprise a set of persons, their attributes, and their pairwise relationships (ties). We consider a binary relationship: presence or absence of a mutually nominated friendship. We define Y_{ij} as the random variable for the tie between persons i and j ($Y_{ij} = 1$ when i and j share a relationship, and 0 otherwise). These ties can be represented as an $n \times n$ matrix ($n =$ population size) called \mathbf{Y} ; diagonal entries Y_{ii} are undefined. For mutual friendship, a tie from i to j implies one from j to i ; thus, \mathbf{Y} is symmetric.

The ERG modeling class specifies the probability of a set of ties \mathbf{Y} given a set of persons and their attributes:

$$P(\mathbf{Y} = y | n \text{ actors}) = \frac{\exp\left(\sum_{k=1}^K \theta_k z_k(y)\right)}{c}. \quad (1)$$

The $z_k(y)$ terms represent model covariates, any set of K network statistics calculated on y and hypothesized to affect the probability of this network forming. Examples of possible z statistics include the number of ties, the number of ties between sixth graders, or the number of triangles. The θ coefficients determine the impact of these statistics for a specific set

substantively different phenomena than the mutual friendships we explore. The methods introduced in this article, however, can be used to analyze such cases as well.

of data. These are unknown parameters to be estimated. The denominator c represents the quantity from the numerator summed over all possible networks with n persons, constraining the probabilities to sum to 1.

Eq. (1) can be reexpressed as the conditional log-odds (logit) of individual ties:

$$\text{logit} \left(P \left(Y_{ij} = 1 \mid n \text{ actors, } Y_{ij}^c \right) \right) = \sum_{k=1}^K \theta_k \delta z_k(y), \quad (2)$$

where Y_{ij}^c denotes all dyads other than Y_{ij} , and $\delta z_k(y)$ is the amount by which $z_k(y)$ changes when Y_{ij} is toggled from 0 to 1. The presence of Y_{ij}^c in the conditional reflects the mutual dependence of ties. The logit formulation clarifies the interpretation of the θ vector: if forming a tie increases z_k by 1, then *ceteris paribus* the log-odds of that tie forming increase by θ_k . Note that a single tie may affect multiple z statistics.

The above formulations assume that both the number of actors and individual actor attributes are fixed. Although these assumptions can be relaxed, they are appropriate for these data.

The homogeneity assumption. In its most general form, this model could include $d - 1$ terms, where d is the number of dyads in the network. That would be a saturated model with every tie having its own probability. As in any statistical model, the goal is to find a parsimonious representation, and proposing homogeneous classes of dyads is a common form of parsimony. Any z statistic counting the aggregate number of some configuration (e.g., triangles) makes an implicit *homogeneity* assumption, that is, that all counted instances are equiprobable. This is similar to the assumption in linear regression that a covariate's effect is the same for all observations. In a demographically structured population, one may expect heterogeneities in social relations among persons to be associated with actor attributes. In this case, attribute-specific configurations may be included, such as the number of ties made by ninth graders or the number of ties between two ninth graders. Although the importance of age or other attributes in tie formation may seem obvious to demographers, the network literature has tended to focus on endogenous network effects like triad closure (for exceptions, see Burt 1990; Morris 1991).

Dyadic dependence and independence. When a model includes only terms that represent the composition of node attributes within ties, it is similar to traditional logistic or log-linear models for contingency tables (Koehly, Goodreau, and Morris 2004). Such models are said to exhibit *dyadic independence* because the probability of any tie does not depend on the value of other ties, only on the attributes of the two actors involved. As a result, Y_{ij}^c can be removed from the conditional in Eq. (2), yielding a set of independent observations amenable to standard statistical analysis.

A model positing an endogenous process of tie formation, by contrast, involves *dyadic dependence*. Examples include the triad closure models discussed in this article and the Markov model of Frank and Strauss (1986), which assumes that ties sharing an actor are dependent. These recursive dependencies require the conditional to remain in Eq. (2), rendering standard statistical analysis inappropriate. The full ramifications of a specific model for dependence may not be obvious; they can be derived, however, using methods described by Besag (1974), Frank and Strauss (1986), and Snijders et al. (2006).

Estimation and model degeneracy. Given a proposed model (set of z statistics), one would like to identify the θ vector maximizing model likelihood. However, the normalizing constant c in Eq. (1) is impossible to calculate for all but the smallest networks, preventing direct evaluation of the likelihood function. Approximations using logistic regression have been used (Besag 1974; Frank and Strauss 1986; Geyer and Thompson 1992; Strauss and Ikeda 1990; Wasserman and Pattison 1996), including an analysis of Add Health data (Moody 2001). In the special case of dyadic independence models, this approach (called *maximum pseudolikelihood estimation*, MPL) is equivalent to maximum likelihood estimation (MLE). For dyadic dependence models, however, MPL tends to perform poorly

(Geyer and Thompson 1992), and its theoretical properties are poorly understood (Handcock 2003a).

The likelihood can be approximated using Markov chain Monte Carlo simulation methods (MCMC; Geyer and Thompson 1992; Snijders 2002), which generate a sample from the space of possible networks to estimate the θ vector. This approach has additional benefits: the same algorithm may be used to simulate network realizations for a given θ vector with appropriate probability. A sample of such realizations provides a means for examining model fit (Hunter, Goodreau, and Handcock 2008). With some modifications, the algorithm may also be used to simulate a dynamic network, with partnerships forming and dissolving over time in ways consistent with the expected values of model statistics. It thus may be used to simulate the flow of information, disease, or other entities over networks described by the model.

Specifying a sensible model for dyadic dependent processes can be difficult; a seemingly reasonable model may in fact have very counterintuitive implications. For example, clustering in networks is often represented as a tendency to close triangles; common summary measures are the number of triangles, or the “clustering coefficient”—the number of triangles divided by number of triads with two or more ties (e.g., Watts and Strogatz 1998). The natural ERG model translation would be a z vector containing terms for the edge count (capturing overall sociality) and the triangle count or clustering coefficient (capturing the additional propensity for or against triad closure). However, the distribution of networks produced by this model, obtained through simulation, displays a bizarre pattern. For most combinations of θ values, the model produces networks in which either all ties exist or none do, and for θ values not producing such extreme networks, the distribution of networks is often bimodal; one mode has low density and high triad closure, while the other has high density and low triad closure, with the appropriate average density and triad closure. The observed network would almost never be produced by the model.

This behavior, termed *model degeneracy*, was explored in detail by Handcock (2003a, 2003b). The intuition behind degeneracy is relatively straightforward. If we specify a model that is unlikely to produce the observed network, then one of two things can happen when it is fit to data: either the MLEs do not exist and estimation does not converge (as in the extremal case above), or the MLEs exist but do not provide a good fit to the data (as in the bimodal case above). Degeneracy is not a shortcoming of the MCMC estimation procedure; nor does it necessarily imply that our intuition about the general process (e.g., that triangles tend to become closed) is wrong. It just implies that the process is not properly specified (e.g., that closure occurs homogeneously across all actor pairs). The solution is to specify a better-fitting model for the data, but this is less straightforward for networks than for other statistical contexts. In ERG modeling, a misspecified model can fail to converge, yielding no parameter estimates to guide model diagnosis or respecification.

Many of the potential solutions to degeneracy replace strong homogeneity assumptions with some form of heterogeneity that defines the scope or form of dependence. One plausible feature of the triad closure case may be a differential tendency for triad closure by attribute combinations—for example, that only triads comprising three actors with shared attributes tend to close. Another mechanism entails a decreasing marginal impact in the effect of triangles on tie formation. To capture this phenomenon, Hunter (2007) and Hunter and Handcock (2006) described a statistic called the *geometrically weighted edgewise shared partner distribution* (GWESP), a reparameterization of a statistic of Snijders et al. (2006). The shared partner distribution is an alternative approach to counting triangles. Two actors “share” a partner if both have a tie to the same partner, and each shared partner forms a triangle if the original pair are tied. In contrast to the census of triangles or the clustering coefficient (which produce a single measure for the whole network), the shared partner count is taken on each edge (hence the “edgewise”), producing a distribution of counts. The GWESP statistic defines a parametric form of this count distribution that gives

each additional shared partner a declining positive impact on the probability of two persons forming a tie. This approach has a clear interpretation and has been shown to work well in practice in overcoming model degeneracy and producing models that fit a wide range of data well, including the data considered here (Goodreau 2007; Hunter et al. 2008).

We use MPLE estimation for models with only dyadic independent terms and MCMC estimation for dyadic dependence models. Both are implemented in the *statnet* package for R (<http://www.statnetproject.org>), which introduces numerous features that make estimation fast and robust for the relatively large networks analyzed here (Handcock et al. 2008). ERG model fit can be assessed using traditional likelihood-based measures such as AIC, which we report here. However, such tests reveal little about the specific features of the data captured or not by each model. We thus turn to the new goodness-of-fit assessment methods described in Hunter et al. (2008). Consistent with our generative focus, we use the same MCMC algorithm and estimated model coefficients to simulate new networks. We then compare these simulated networks to the observed data on various structural properties to reveal which models are sufficient to produce which network properties.

ERG Model Specification

Our models include terms for sociality, selective mixing, and triad closure. The Add Health data provide an unusual opportunity for empirical replication: the models are fit to each school, so we may compare 59 sets of coefficients.³

Sociality. We infer sociality based on counts of ties observed: s represents the total number of ties, and k_i is the total number of ties for all persons with attribute value i . The s term acts as an intercept, and the coefficient for s represents the conditional log-odds of a tie for the reference category (in these models, reference categories are grade 7, white, and male). The k_i terms assume homogeneity within attribute class, allowing each race, sex, and grade to have different mean sociality.

Selective mixing. We consider two selective mixing dynamics. The first is a homogeneous propensity for assortative mixing across attribute categories (“uniform homophily”). The second is a propensity that is specific to individual categories (“differential homophily”). Statistics are as follows: first, h is the total number of ties between persons in the same attribute category, regardless of category. This *uniform homophily* is used for sex since there are only three tie types (MM, MF, FF); with main effects included, only one degree of freedom remains. Second, h_i is the total number of ties between persons both in attribute category i . There is one such statistic for each category of the attribute. This *differential homophily* is used for race and grade.⁴

Triad closure. For the reasons described above, we investigate triad closure using the GWESP statistic:

$$w = e^{\tau} \sum_{i=1}^{n-2} \left\{ 1 - \left(1 - e^{-\tau} \right)^i \right\} p_i,$$

where p_i equals the number of actor pairs who are friends and who have exactly i friends in common.

This statistic includes a parameter τ , in addition to the coefficient in θ associated with the statistic in the ERG formula (Eq. (1)). τ controls the geometric rate of decline in the

3. Note that some statistics do not appear 59 times because some schools do not include students in every category. In addition, for cases in which there were fewer than five students of a given category in a school, the coefficients are also excluded, since parameter estimates would be unstable and in extreme cases undefined.

4. Note that a reference category is necessary for sociality effects but not for selective mixing effects. This is because the k statistics across all categories for an attribute must sum to s , while there is no comparable constraint on the sum of the h statistics. Put another way, the number of ties summed across people of all races must equal the number of ties summed across people of all grades; but the total number of race-homophilous ties need not equal the total number of grade-homophilous ties.

effect of triad closure on tie probability for an increasing number of shared partners. The corresponding coefficient in θ measures the magnitude of triad closure for a given τ value. We adopt a value of 0.25 for τ , although results are robust to this choice.⁵

Because s represents the log-odds of a tie and thus has the number of non-ties in the denominator, the magnitude of s is directly affected by network density (the fraction of possible relationships that are realized). Friendship density declines as population size increases (more people does not mean proportionately more friends per person), so we expect the coefficient on s to decline with school size. The coefficients on the k and h statistics should be invariant to school size because they are defined relative to s . For instance, a coefficient for male-male ties captures the increase in the log-odds that two males become friends relative to the log-odds for other dyad types. Such a statistic will thus not vary systematically with school size or sex ratio—any significant variation can be interpreted as representing a process other than demographic constraint. This invariance does not necessarily hold for w , as will be seen in practice later in the article, nor for other dyadic dependence statistics generally.

To recapitulate, the s and k_i terms measure total edges and overall degree by attribute category, h and h_i measure uniform and differential homophily, and w measures transitivity. The θ coefficients associated with these statistics represent sociality, selective mixing, and triad closure, respectively.


We consider three models based upon this set of statistics: a demographic attribute (DA) model that considers only attribute-based, dyadic-independent processes (s , k_i , h_i); a triad closure (TC) model that considers only a homogeneous, dyadic-dependent process (s , w); and a full model combining the two (s , k_i , h_i , w). Our hypothesis is that selective mixing and triad closure terms (h_i , w) will be positive in all models. Because some ties are consistent with either process, the partial models (DA or TC) will overestimate the effect of the terms, while the full model will estimate the unique effects of each; thus, we predict that estimates will be higher in the partial models than in the full model. We also expect that both selective mixing and triad closure are important to friendship formation, so we expect the full model to fit better than either partial model. We do not expect any consistent differences in sociality terms between the DA and full model.

RESULTS

We begin by considering which of our models provides the best fit to our data by using the likelihood-based measures of AIC and likelihood-ratio tests. For all 59 schools, these measures indicate that the full model fits best by a large margin. The increase in log-likelihood between the partial and the full models ranges from 19.5 to 2,663.3, with one new term in the full model relative to the DA model, and up to 29 terms relative to the TC model.⁶ Because the full model fits best for all individual schools, we consider its coefficients to be our best estimates for the true magnitude of sociality, selective mixing, and triad closure. Figure 3 charts the sociality and selective mixing coefficients, using boxplots to show their range across the 59 schools. Sociality increases by grade until grade 12, though the apparent decline for 12th graders is due in part to nominations of more out-of-school friends. Grade-based selective mixing is consistently assortative (i.e., the selective mixing coefficient is positive), but is strongest among 7th graders and declines with senior-

5. This process of adopting a value for τ entailed randomly selecting a set of schools covering the range of school sizes and testing the model fit for τ values beginning at 0 and increasing by 0.05 until model fit no longer increased for the majority of schools. We began at 0 and moved upward because lower τ values are less likely to produce degenerate models (Goodreau 2007). In the range explored ($\tau = 0$ to 0.5), the specific value of τ had relatively little effect on coefficient estimates or on model fit.

6. Comparing the DA and TC models by AIC reveals that 22 schools are better fit by the former and 37 are better fit by the latter, with no clear relationship to school size or diversity.

Figure 3. Coefficients From the Full Model, Plotted Across All 59 Schools

Notes: Boxplots follow the Tukey method. Boxes represent quartiles; whiskers extend to the most extreme data point within 1.5 times the interquartile range from the edge of the box; and points represent outliers.

ity. Again, the trend reverses with 12th graders. Both grades 7 and 12 are more likely to be influenced by boundary effects.⁷

Sociality differences by race are much less pronounced. Blacks average lower sociality than other categories, but racial categories do not otherwise differ systematically. Selective mixing by race is, in general, weaker than selective mixing by grade and differs across categories. The strongest assortative mixing is found among blacks, followed by Asians, “others,” and whites. All four categories exhibit assortative mixing in more than 90% of schools. In contrast, Native Americans exhibit assortative mixing in only 72% of schools where they are present, usually to a relatively small degree, while Hispanics exhibit assortative mixing in only 63% of schools (see Table 1). “Other” race exhibits assortative mixing comparable to whites and Asians, even though one might have expected lower levels, given the category’s potentially high heterogeneity. This may indicate that within any given

7. One might choose to explore this by including another set of sociality terms that count in-school ties for all actors with a given number of nominations to friends outside of school. If these effects were increasingly negative for increasing numbers of out-of-school nominations, this would demonstrate the existence of the boundary effect and provide adjusted estimates for the sociality of each grade net of these effects. One would need to be careful in interpreting the coefficients in this case, given that we have no way of knowing if out-of-school names mentioned on surveys are in fact *mutual* close friendships, comparable to the relations that we study within the school.

Table 1. Selective Mixing Coefficients: Full Model

Variable	Number of Schools With Coefficient	Number of Schools With Coefficient > 0 (assortative mixing)	% of Coefficients > 0	Median Coefficient Value	Interquartile Range
Grade 7	57	57	100	4.820	4.404–5.530
Grade 8	58	58	100	4.032	3.414–5.366
Grade 9	59	59	100	2.856	2.300–3.516
Grade 10	59	59	100	1.860	1.503–2.323
Grade 11	59	58	98	1.552	1.258–1.934
Grade 12	58	58	100	2.249	2.033–2.593
White	58	53	91	0.777	0.357–1.288
Black	37	35	95	2.486	1.866–3.098
Hispanic	40	25	63	0.492	–0.245–0.966
Asian	23	22	96	1.722	1.055–2.212
Native American	36	26	72	0.403	–0.017–0.864
Other	16	15	94	1.478	0.462–2.106
Sex	59	59	100	0.720	0.669–0.811

school, the “other” category is relatively homogeneous, although in the absence of further information about these students’ identities, this conjecture is impossible to test.


Differences in sociality by sex are small, with females usually exhibiting slightly higher sociality than males. These values can be difficult to interpret since they are partly confounded by stratification in the survey question, which asked respondents to nominate up to five female friends and up to five male friends. The resulting effects are moderate relative to race and grade assortative mixing, but they are significant in every school.

The triad closure (GWESP) coefficients for the full model are plotted against school size in Figure 4. This term is consistently positive and tends to increase with school size, at least until a threshold of approximately 750 students. Interpreting this value requires some care. Consider a GWESP coefficient of 1.8, the approximate value at which it asymptotes in larger schools. If a tie will close one triangle, and all actor pairs in that triangle currently have no shared partners, the log-odds of the tie are increased by 5.4 relative to an otherwise similar tie that would not close a triangle. If it would close two triangles, the log-odds are increased an additional 4.0 relative to closing one. These diminishing returns continue for each additional triangle.⁸

The relationship between network size and the GWESP coefficient results from the nonlinear nature of triadic effects. Friends in smaller schools actually average more friends in common than those in larger schools (1.1 among the 10 smallest vs. 0.6 among the 10 largest), but this is more likely to happen by chance in smaller schools: the probability of closing triangles by chance declines with the square of network size. We intuitively

8. To get a sense of the interpretation of the τ parameter, consider the following: for the same coefficient 1.8 but for $\tau = 0.0$, the two changes in log-odds would be 5.4 and 3.6 (instead of 5.4 and 4.0 for $\tau = 0.25$). For $\tau = 1.0$, they would be 5.4 and 4.7. That is, the closer τ is to 0, the more quickly the influence of triangles on tie probability tapers off for triangles beyond the first one that a node is in. In the other direction, as τ approaches infinity, the values do not taper off at all, but remain at 5.4 for each additional triangle.

Figure 4. Triad Closure (GWESP) Coefficient: Full Model



tively understand this “small town” phenomenon, where everyone knows everyone else. The relationship between magnitude of effect and school size is not an artifact; it reflects the fact that higher-order relational processes operate differently in differently sized communities. If persons generally prefer some level of social closure, they must exert more effort in larger populations to create it.

Comparing coefficients across models. We can compare these estimates to those obtained from the partial models to see whether simultaneous estimation attenuates the coefficients, as hypothesized. Table 2 presents an analysis of a single school, detailing changes in all coefficients over nested models. Figure 5 then shows coefficients for three statistics over all schools. Tables for the remaining schools and comparable plots for the remaining coefficients appear in the online supplement.

Analysis of nested models on all 59 schools shows that the selective mixing coefficients from the full model are less than the corresponding coefficients in the DA model for grade and sex in all schools (e.g., see Grade 7 in Figure 5, plot a), and for white, black, and Asian races in most schools. The apparent decline in selective mixing averages between 5% and 15% for all grades, sex, and the above three race categories, suggesting that the homophily we observe is partly a function of triad closure.

For the Hispanic, Native American, and other race categories, our hypothesis of attenuated effects is not supported; selective mixing *increases* on average when a triad closure term is included in the model (e.g., see Hispanics in Figure 5, plot b). Interpretation of this pattern is somewhat nuanced, but one of its features is that homophilous ties in this category are less likely than others to be found in triangles. Thus, once the triad closure term for all students is included, the magnitude of selective mixing necessary to explain observed homophily becomes greater. As an example, consider the probabilities implied by our model coefficients that two seventh-grade Hispanic boys in School 11 will form a relationship. In the DA model, this value is 1.4%. In the full model, it is 1.1% if the tie will not close a triangle, and 47.4% if it will. Clearly, closing a triangle dramatically increases its probability, as it does for all racial categories. But the DA value (1.4%) lies closer to the nontriangle value (1.1%) and further from the triangle value (47.4%) than is true for other categories. The DA value provides a weighted average of the other two, capturing

Table 2. Model Coefficients: "School 18"

	DA Model	TC Model	Full Model	% Change, Partial to Full Model
Sociality				
Intercept	-11.720	-6.997	-11.188	
Grade 8	0.334	—	0.325	-3
Grade 9	1.227	—	1.042	-15
Grade 10	1.882	—	1.636	-13
Grade 11	1.699	—	1.456	-14
Grade 12	1.306	—	1.149	-12
Black	-0.345	—	-0.233	-33
Hispanic	0.836	—	0.743	-11
Asian	-0.097	—	-0.050	-49
Native American	0.752	—	0.602	-20
Other	0.573	—	0.448	-22
Female	0.346	—	0.168	-51
Selective Mixing				
Grade 7	4.825	—	4.142	-14
Grade 8	4.367	—	3.648	-16
Grade 9	2.577	—	2.165	-16
Grade 10	1.466	—	1.094	-25
Grade 11	1.831	—	1.552	-15
Grade 12	3.046	—	2.226	-27
White	1.041	—	0.865	-17
Black	3.874	—	3.166	-18
Hispanic	-0.470	—	-0.642	37
Asian	2.320	—	2.423	4
Native American	0.721	—	0.831	15
Other	NA	—	NA	
Sex	0.785	—	0.677	-14
Triad Closure				
GWESP	—	2.142	1.726	-19
Likelihood	-6,073.584	-5,725.892	-5,043.718	-5,043.718
AIC	12,201.170	11,455.780	10,143.440	10,143.440

Note: Coefficients for "NA" category are not shown.

the proportion of homophilous ties that are in triangles. Since it tends to fall closer to the nontriangle value for Hispanic-Hispanic ties than for other races, such ties are less likely to appear in triangles.

Figure 5, plot c shows the apparent change in the GWESP coefficient between the TC and full models. The average attenuation here is larger than for any of the attribute terms (25.9%), suggesting that selective mixing has a stronger impact on triangle counts than triad closure has on homophily.

Figure 5. Example Plots of Selective Mixing Coefficients: Full Model Versus Partial (DA and TC) Models


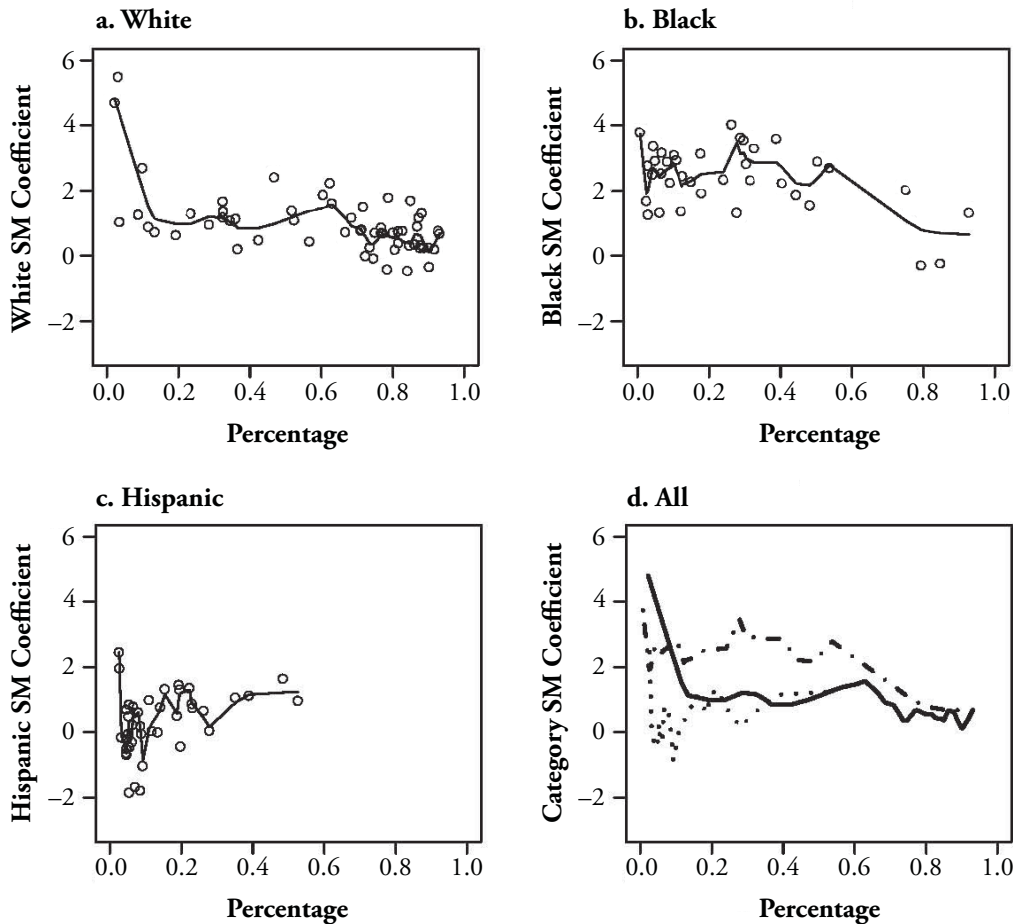


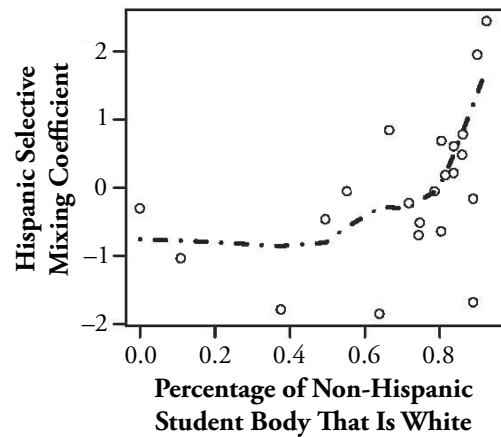
Figure 6. Selective Mixing (SM) Coefficients, by Proportion of School in Category: Full Model



Note: Lines represent lowess curves.

There is one striking change that we did not anticipate. From the DA to the full model, the intercept s increases moderately, while the sociality term for females decreases by at least 23% in all schools. The interpretation of this is straightforward: the more females in the friendship dyad (MM vs. MF vs. FF), the more likely that dyad is to be in a triangle. Once triad closure is included, female sociality declines, since more of their ties are generated by triad closure than are male ties; the intercept (which partly models male ties, as male is the reference category for sex) declines. These effects are consistent with previous work identifying differences in friendship nomination patterns by sex (Dunphy 1963) and are especially interesting in the light of work demonstrating the Durkheimian implications of social cohesion on suicidality in these data (Bearman and Moody 2004).

Racial selective mixing and the prevalence of sociodemographic categories. Although assortative mixing among blacks appears the highest in simple boxplots (Figure 3), this compounds two effects: minority status and the impact of minority status on assortative friendship selection. The variation in racial composition across the 59 schools gives us some leverage to tease these two effects apart. Figure 6 examines the racial selective mixing coefficients from the full model against the fraction of student body in the given racial category, allowing us to compare how different categories cluster when in the minority or majority.

Figure 7. Hispanic Selective Mixing, by Proportion White: Full Model

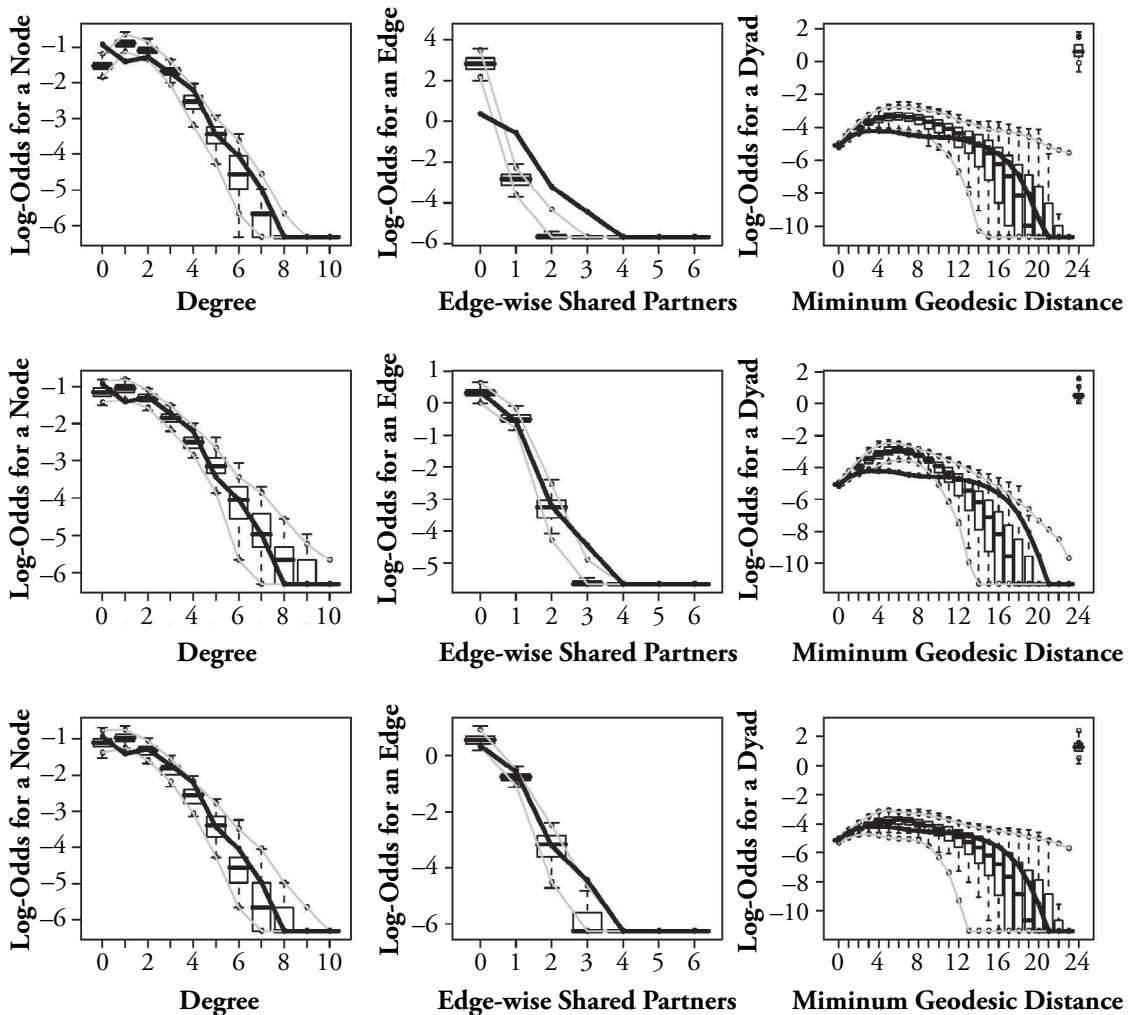
Note: The line represents a loess curve.

Although blacks have higher average selective mixing coefficients than whites across all schools, here we see that the effect depends on the relative prevalence of the racial category. When either whites or blacks compose more than about 50% of the student body, they exhibit roughly equal levels of assortativity. In the 10%–50% range, blacks appear more assortative than whites. However, when each category represents a small minority in their school, the pattern switches dramatically; black assortativity stays roughly the same or declines slightly, while at small numbers, whites become extremely assortative—far more than blacks at the same concentration. By far the highest racial selective mixing coefficients in all of our models are among whites in the two schools where whites are a small minority.

For Hispanics, negative selective mixing is concentrated in those schools where Hispanics represent a small minority. However, it is not always the case that a small minority of Hispanics have disassortative mixing. A plausible fulcrum for this inconsistency may be whether the Hispanic population is itself homogeneous (dominated by a single origin: e.g., Mexican) or heterogeneous. Since the Add Health questionnaire asked Hispanic students to select their origin from a list of six options, we test this conjecture by constructing a homogeneity index representing the probability that any two randomly selected Hispanic students share their origin. Examining the relationships between this measure and Hispanic homophily, either across all schools or limited to those less than 10% Hispanic (where the inconsistency seems to occur) fails to show a significant relationship (results not shown). An alternative hypothesis is that Hispanics might be assortative when the remainder of the population is homogeneous—that is, Hispanic identity is most salient when Hispanics are the only minority subpopulation in the school. The reverse is that when multiple racial minorities are prominent, individual Hispanic students might identify and socialize with students of some other racial minority, creating Hispanic disassortative mixing. Figure 7 plots the magnitude of the Hispanic selective mixing coefficient against the percentage of non-Hispanic white students for all schools that are less than 10% Hispanic.⁹ Although the plot is only suggestive, for a small subpopulation of

9. A plot replacing “white” with “largest non-Hispanic racial category” (not shown) demonstrates a similar but less clear pattern.

Figure 8. Goodness-of-Fit Plots: "School 18"



Notes: The first row depicts the demographic attribute (DA) model. The second row represents the triad closure (TC) model. The third row represents the full model.

Hispanic students to be assortative, it appears necessary but not sufficient that the student body is otherwise largely non-Hispanic white.

Goodness of fit. We present an example of goodness-of-fit plots for one school in Figure 8. Each plot compares the observed data to 100 simulated networks (generated from a given ERG model and its estimated coefficients) for a given network statistic. The three rows represent the DA, TC, and full models, respectively; the three columns represent the three statistics that we compare. In each case, the dark solid line represents a given statistic from School 18, while the boxplots represent the same statistic for the 100 simulated networks.

The first plot in each row shows the degree distribution, tabulated across all actors in the network. By including sociality terms in two of our models, we were in effect modeling mean sociality by category; but we did not explicitly consider the full degree distribution, one of the most heavily studied features of network data. The second plot in each row represents the distribution of shared partners (number of friends in common) tabulated across all friendship pairs; this provides a sense of the level and scale of clustering. Our GWESP

term is a parametric version of this distribution, guaranteed to capture its mean; like degree, however, there is no guarantee that our models will fit the full distribution. The third plot considers a higher-order network statistic that is not directly related to any of the model terms: the distribution of geodesic distances. Geodesic distances are the pairwise path distances between students (friends are distance 1, friends of friends who are not friends directly are distance 2, and so on). This distribution can vary substantially across networks and is often used to gauge the diffusion potential of a network.

The school illustrated in Figure 8 typifies the qualitative trends seen in these plots across the schools for our three models. For all schools, the DA model sharply underestimates the observed numbers of shared partners; selective mixing is never sufficient to capture friendship clustering. On the other hand, the TC model rarely captures geodesic distances well, generally overestimating the number of people reachable by short (length 2–5) paths and underestimating those unreachable or reachable only by long paths. The full model comes close to matching all higher-order statistics in some smaller schools. In the larger schools, it does not perfectly match the data, but it comes closer to the full set of goodness-of-fit metrics than any partial models. We return to model fit in the Discussion section.

DISCUSSION

Our analysis of adolescent friendship networks reveals evidence that both selective mixing and triad closure operate across a wide range of sociodemographic settings to structure the process of mutual friendship formation. These processes interact, generating a complex set of effects. The attributes of grade and sex always generate both strong assortative mixing and within-category triad closure. The effects of race vary: white, black, and Asian students typically exhibit assortative mixing and within-category triad closure, representing structural cohesion within their subpopulations, but other categories are more complicated. Hispanic students (by far the most numerous of remaining categories) can display random (i.e., unbiased) or even disassortative mixing. Even when Hispanics exhibit assortative mixing, it is often coupled with a relative lack of within-category triad closure, reducing their higher-order cohesion.

We also find that the higher-order processes governing friendship formation can depend on the school's overall demographic profile. For example, the extent of Hispanic cohesion appears to be partly shaped by the homogeneity of the school's non-Hispanic population: the more homogeneous the non-Hispanic population, the more cohesive the Hispanic friendships. Similarly, the strength of assortative mixing for whites is inversely related to their relative share of the student body: the more they are in the minority, the more segregated their friendships become. This relation does not hold for blacks: assortative mixing is strongest for blacks when they compose an intermediate proportion of the population. Perhaps blacks are less assortative than whites when they are a small minority because this status is familiar to them.

Although the models we examined reproduce structural features in the observed networks relatively well, the lack of fit for some features suggests directions for future investigation. Chief among these directions is consideration of attribute-specific triad closure. Our method of comparing coefficients from two models provided indirect evidence that triad closure differs by category, but this could be examined directly with a model that includes attribute-specific GWESP statistics. Another direction is modeling the prevalence of students having no reciprocated friendships. Our models reproduce the observed distribution of mutual friendship counts per person surprisingly well, given that they include only terms for category-specific mean sociality. But they consistently underestimate the proportion of actors with no reciprocated friendships. One could capture this tendency for social isolation with a single term for number of actors of degree 0. This would identify the cross-school variation in the preponderance of loners beyond that expected and may improve other aspects of model fit, but it would provide little insight into the factors

that produce social isolation. Additional models and data could help clarify these factors, as well as determine whether some of this reported social isolation actually stems from item nonresponse.

The full model also does not capture the social distance structure of the friendship networks for larger schools. These networks are more “stretched out” than the models predicted, with more actors reachable only by very long paths (as seen in the last panel of Figure 8). This may be because we modeled cross-grade ties homogeneously, while the data suggest that the greater the difference in grade, the less likely students are to become friends. An additional term capturing absolute difference in grade would capture this effect and likely improve the fit to the social distance distribution. Including an effect for mixing by school (i.e., across- vs. within-school) for the multischool school groups would also likely help.

Another area for future work is to also consider unreciprocated friendship nominations. It is hard to predict how the structure of such relations would differ from those we observe here, but comparing the two in transitivity, for instance, would give us a richer overall picture of friendship formation. It would also allow for investigation of mutuality and hierarchy in networks simultaneously with the processes we have considered here. Moody (2001) considered some of these effects for the Add Health data, albeit in the pseudolikelihood framework; reexamining those processes using the likelihood methods now available for ERG models, while also taking account of missing data, would be an important next step.

One further avenue is to formalize the analysis of the effects of macro-level variables, such as network size or composition on ERG model coefficients across schools. To do so, one could in theory employ hierarchical linear modeling techniques; developing formal methods for applying this framework to ERG coefficients is an area for future research.

Our modeling goal here was to move beyond network description to uncover the generative processes that underlie network formation. We believe that the ERG modeling framework offers a general set of tools to aid in this goal, and the Add Health survey (with 59 replicate schools) has a unique structure for exploiting these new capabilities. By focusing on relatively exogenous personal attributes (race, sex, and grade), we were able to identify some of the demographic correlates of friendship formation and to rule out the kind of endogenous feedback processes between friendship choices and personal attributes found in other contexts. By estimating all models within the relatively small units of schools, we were able to get closer to ensuring that all pairs of students are equally able to form friendships, so their choices are less confounded by gross heterogeneities in opportunity. By simultaneously estimating effects at the individual, dyadic, and triadic levels, we were able to disentangle the effects of each process from the others. This analysis moves us closer to quantitatively comparing the strength of effects within and across the levels on which student friendship preferences operate.

What this particular analysis cannot do is provide estimates of actual friendship preferences among students. Part of this is due to the simplicity of the models we examine, but the real limitation is data. For example, the strong grade homophily we observe probably reflects the opportunity structure of grade-segregated classrooms as much as, if not more than, student preferences for friends in the same grade. The patterns of mixing by race also reflect a combination of opportunity and preference, since schools rarely offer a full range of potential friends of every race. In general, it is very difficult to distinguish preference from opportunity in a cross-sectional observational network study (Feld 1981, 1982; McPherson and Smith-Lovin 1987), though it is worth noting that this limitation is not unique to network data. Experimental data are required to observe the preferences that guide friendship choice; there, preferences may be revealed by design. The ERG modeling tools can be used for analyzing such experimental data, and the coefficients would then have a simpler preference interpretation.

The ERG modeling framework places the statistical analysis of networks on a solid and accessible footing, giving researchers a powerful new tool for understanding the micro-level

processes governing the relational structure of populations. The models can incorporate demographic determinants of social relations, along with a wide range of endogenous network processes. The issue of model degeneracy has been identified and analyzed, removing the obstacle to full likelihood estimation that hobbled these models for over a decade. And the computationally intensive algorithm used for parameter estimation has the benefit of also providing a method for generating simulated networks, both cross-sectional and dynamic. Additional issues remain to be solved, such as the handling of missing and sampled network data, but progress is being made on this and other methodological fronts (see Gile and Handcock 2006). The result is that population scientists now have a set of general tools for inferring from the relational structure of a population the social and demographic processes that generated it.

REFERENCES

- Adams, J. and J. Moody. 2007. "To Tell the Truth: Measuring Concordance in Multiply Reported Network Data." *Social Networks* 29:44–58.
- Alexander, C., M. Piazza, D. Mekos, and T. Valente. 2001. "Peers, Schools, and Adolescent Cigarette Smoking." *Journal of Adolescent Health* 29:22–30.
- Bearman, P.S. and J. Moody. 2004. "Suicide and Friendships Among American Adolescents." *American Journal of Public Health* 94:89–95.
- Besag, J. 1974. "Spatial Interaction and the Statistical Analysis of Lattice Systems." *Journal of the Royal Statistical Society Series B* 36:192–236.
- Blau, P.M. 1977. *Inequality and Heterogeneity: A Primitive Theory of Social Structure*. New York: Free Press.
- Borgatti, S.P. and M.G. Everett. 1999. "Models of Core/Periphery Structures." *Social Networks* 21:375–95.
- Burt, R.S. 1990. "Kinds of Relations in American Discussion Networks." Pp. 411–52 in *Structures of Power and Constraint*, edited by C. Calhoun, M.W. Meyer, and W.R. Scott. Cambridge: Cambridge University Press.
- Byrne, D.E. 1971. *The Attraction Paradigm*. New York: Academic Press.
- Carley, K. 1991. "A Theory of Group Stability." *American Sociological Review* 56:331–54.
- Cartwright, D. and F. Harary. 1966. "Structural Balance: A Generalization of Heider's Theory." *Psychological Review* 63:277–93.
- Davis, J.A. 1967. "Clustering and Structural Balance in Graphs." *Human Relations* 20:181–87.
- Dunphy, D.C. 1963. "The Social Structure of Urban Adolescent Peer Groups." *Sociometry* 26: 230–46.
- Feld, S.L. 1981. "The Focused Organization of Social Ties." *American Journal of Sociology* 86:1015–35.
- . 1982. "Social Structural Determinants of Similarity Among Associates." *American Sociological Review* 47:797–801.
- Feld, S.L. and R. Elmore. 1982. "Patterns of Sociometric Choices: Transitivity Reconsidered." *Social Psychology Quarterly* 45(2):77–85.
- Fershtman, M. 1985. "Transitivity and the Path Census in Sociometry." *Journal of Mathematical Sociology* 11:159–89.
- Fienberg, S.E. and S. Wasserman. 1981. "Categorical Data Analysis of Single Sociometric Relations." *Sociological Methodology* 12:156–92.
- Frank, O. and D. Strauss. 1986. "Markov Graphs." *Journal of the American Statistical Association* 81:832–42.
- Geyer, C.J. and E.A. Thompson. 1992. "Constrained Monte Carlo Maximum Likelihood for Dependent Data." *Journal of the Royal Statistical Society Series B* 54:657–99.
- Gile, K. and M.S. Handcock. 2006. "Model-Based Assessment of the Impact of Missing Data on Inference for Networks." Working Paper No. 66. Center for Statistics and the Social Sciences, University of Washington. Available online at <http://www.csss.washington.edu/Papers/wp66.pdf>.

- Goodreau, S.M. 2007. "Advances in Exponential Random Graph (p^*) Models Applied to a Large Social Network." *Social Networks* 29:231–48.
- Granovetter, M. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78:1360–80.
- Hallinan, M.T. 1982. "Classroom Racial Composition and Children's Friendships." *Social Forces* 61:56–72.
- Hallinan, M.T. and R.A. Williams. 1989. "Interracial Friendship Choices in Secondary Schools." *American Sociological Review* 54:67–78.
- Handcock, M.S. 2003a. "Assessing Degeneracy in Statistical Models of Social Networks." Working Paper No. 39. Center for Statistics and the Social Sciences, University of Washington. Available online at <http://www.csss.washington.edu/Papers/wp39.pdf>.
- . 2003b. "Statistical Models for Social Networks: Degeneracy and Inference." Pp. 229–40 in *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*, edited by R.L. Breiger, K.M. Carley, and P. Pattison. Washington, DC: National Academies Press.
- Handcock, M.S., D.R. Hunter, C.T. Butts, S.M. Goodreau, and M. Morris. 2008. "ergm: A Package to Fit, Simulate and Diagnose Exponential-Family Models for Networks." *Journal of Statistical Software* 24(3):1–29.
- Heider, F. 1958. *The Psychology of Interpersonal Relations*. New York: Wiley.
- Holland, P.W. and S. Leinhardt. 1981. "An Exponential Family of Probability Distributions for Directed Graphs." *Journal of the American Statistical Association* 76(373):33–50.
- Hunter, D.R. 2007. "Curved Exponential Family Models for Social Networks." *Social Networks* 29:216–30.
- Hunter, D.R., S.M. Goodreau, and M.S. Handcock. 2008. "Goodness of Fit of Social Network Models." *Journal of the American Statistical Association* 103(481):248–58.
- Hunter, D.R. and M.S. Handcock. 2006. "Inference in Curved Exponential Family Models for Networks." *Journal of Computational and Graphical Statistics* 15:565–83.
- Kandel, D.B. 1978. "Homophily, Selection, and Socialization in Adolescent Friendships." *American Journal of Sociology* 84:427–36.
- Koehly, L.M., S.M. Goodreau, and M. Morris. 2004. "Exponential Family Models for Sampled and Census Network Data." *Sociological Methodology* 34:241–70.
- Kohler, H.P. 1997. "Learning in Social Networks and Contraceptive Choice." *Demography* 34:369–83.
- . 2000. "Fertility Decline as a Coordination Problem." *Journal of Development Economics* 63:231–63.
- Kohler, H.P., J.R. Behrman, and S.C. Watkins. 2001. "The Density of Social Networks and Fertility Decisions: Evidence From South Nyanza District, Kenya." *Demography* 38:43–58.
- Lazarsfeld, P.F. and R.K. Merton. 1954. "Friendship as a Social Process: A Substantive and Methodological Analysis." Pp. 8–66 in *Freedom and Control in Modern Society*, edited by M. Berger, T. Abel, and C.H. Page. New York: Van Nostrand.
- Mare, R.D. 1991. "Five Decades of Educational Assortative Mating." *American Sociological Review* 56:15–32.
- Marsden, P.V. 1981. "Models and Methods for Characterizing the Structural Parameters of Groups." *Social Networks* 3(1):1–27.
- Massey, D.S. 1988. "Economic Development and International Migration in Comparative Perspective." *Population and Development Review* 14:383–413.
- Massey, D.S., R. Alarcon, J. Durand, and H. Gonzalez. 1987. *Return to Aztlan: The Social Process of International Migration From Western Mexico*. Berkeley, CA: University of California Press.
- May, R.M. and R.M. Anderson. 1988. "The Transmission Dynamics of Human Immunodeficiency Virus (HIV)." *Philosophical Transactions of the Royal Society of London Series B—Biological Sciences* 321(1207):565–607.
- McPherson, J.M. and L. Smith-Lovin. 1987. "Homophily in Voluntary Organizations: Status Distance and the Composition of Face-to-Face Groups." *American Sociological Review* 52:370–79.

- McPherson, M., L. Smith-Lovin, and J.M. Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology* 27:415–44.
- Moody, J. 2001. "Race, School Integration, and Friendship Segregation in America." *American Journal of Sociology* 107:679–716.
- Morris, M. 1991. "A Log-Linear Modeling Framework for Selective Mixing." *Mathematical Biosciences* 107:349–77.
- Morris, M. and L. Dean. 1994. "Effect of Sexual-Behavior Change on Long-Term Human-Immunodeficiency-Virus Prevalence Among Homosexual Men." *American Journal of Epidemiology* 140:217–32.
- Morris, M. and M. Kretzschmar. 1997. "Concurrent Partnerships and the Spread of HIV." *AIDS* 11:641–48.
- National Longitudinal Study of Adolescent Health. 2001. *Network Variables Codebook*. Chapel Hill, NC: University of North Carolina.
- Rapoport, A. 1957. "A Contribution to the Theory of Random and Biased Nets." *Bulletin of Mathematical Biophysics* 19:257–71.
- Resnick, M.D., P.S. Bearman, R.W. Blum, K.E. Bauman, K.M. Harris, J. Jones, J. Tabor, T. Beuhring, R.E. Sieving, M. Shew, M. Ireland, L.H. Bearinger, and J.R. Udry. 1997. "Protecting Adolescents From Harm. Findings From the National Longitudinal Study on Adolescent Health." *Journal of the American Medical Association* 278:823–32.
- Rogers, E.M. and D.K. Bhowmik. 1970. "Homophily-Heterophily: Relational Concepts for Communication Research." *Public Opinion Quarterly* 34:523–38.
- Snijders, T.A.B. 2002. "Markov Chain Monte Carlo Estimation of Exponential Random Graph Models." *Journal of Social Structure* 3(2):1–40.
- Snijders, T.A.B., P.E. Pattison, G.L. Robins, and M.S. Handcock. 2006. "New Specifications for Exponential Random Graph Models." *Sociological Methodology* 36:99–153.
- Snijders, T.A.B. and F.N. Stokman. 1987. "Extensions of Triad Counts to Networks With Different Subsets of Points and Testing Underlying Random Graph Distributions." *Social Networks* 9:249–75.
- South, S.J. and D.L. Haynie. 2004. "Friendship Networks of Mobile Adolescents." *Social Forces* 83:315–50.
- Strauss, D. and M. Ikeda. 1990. "Pseudolikelihood Estimation for Social Networks." *Journal of the American Statistical Society* 85:204–12.
- Udry, J.R. and P.S. Bearman. 1998. "New Methods for New Research on Adolescent Sexual Behavior." Pp. 241–69 in *New Perspectives on Adolescent Risk Behavior*, edited by R. Jessor. Cambridge, England: Cambridge University Press.
- Valente, T.W. 2003. "Social Network Influences on Adolescent Substance Use." *Connections* 25:1–4.
- Valente, T.W., S.C. Watkins, M.N. Jato, A. VanderStraten, and L.P.M. Tsitsol. 1997. "Social Network Associations With Contraceptive Use Among Cameroonian Women in Voluntary Associations." *Social Science & Medicine* 45:677–87.
- Wasserman, S.S. 1977. "Random Directed Graph Distributions and Triad Census in Social Networks." *Journal of Mathematical Sociology* 5(1):61–86.
- Wasserman, S. and P. Pattison. 1996. "Logit Models and Logistic Regressions for Social Networks: I. An Introduction to Markov Graphs and p^* ." *Psychometrika* 60:401–25.
- Watkins, S.C. and A.D. Danzi. 1995. "Women's Gossip and Social Change: Childbirth and Fertility Control Among Italian and Jewish Women in the United States, 1920–1940." *Gender & Society* 9:469–90.
- Watts, D.J. and S.H. Strogatz. 1998. "Collective Dynamics of 'Small-World' Networks." *Nature* 393:440–42.