Incorporating Structural Stigma into Network Analysis*

Francis Lee[†] and Carter T. Butts[‡§]

4/15/20

Abstract

A rich literature has explored the modeling of homophily and other forms of nonuniform mixing associated with individual-level covariates within the exponential family random graph (ERGM) framework. Such differential mixing does not fully explain phenomena such as stigma, however, which involve the active maintenance of social boundaries by ostracism of persons with out-group ties. Here, we introduce a new family of statistics that allows for such effects to be captured, making it possible to probe for the potential presence of boundary maintenance above and beyond simple differences in nomination rates. We demonstrate these statistics in the context of gender segregation in a school classroom, and introduce a framework for understanding the associated coefficients via network perturbation.

Keywords: social network analysis, homophily, stigma, social sanctions, xenophobia, ERGM

1 Introduction

Homophily (the tendency for individuals to be tied to others with similar characteristics) and segregation (the tendency for individuals with different characteristics not to be in contact, nor to be found within the same social or geographical settings) are widely studied structural phenomena across a range of contexts (see e.g. Blau, 1994; Schelling, 1971; Sakoda, 1971; Schaefer et al., 2017; McPherson et al., 2001; Goodreau et al., 2017; Butts, 2007). Within the specific context of cross-sectional social network models, these phenomena are generally treated within the rubric of *differential mixing*, i.e. as a net tendency for ties to occur at higher or lower rates within pairs of individuals with particular combinations of attributes, controlling for other factors. Mixing effects are easily incorporated into modeling frameworks such as the exponential family random graph models (ERGMs), and indeed such effects have been used to shed light on a variety of social phenomena including power in prison gangs (Schaefer et al., 2017), the structure of friendship ties (Goodreau et al., 2009), and HIV diffusion (Goodreau et al., 2017).

While differential mixing is an important dimension of social structure, it does not by itself account for all aspects of segregation within social networks. In particular, it does not capture the *active maintenance of group boundaries* observed in many settings. As Goffman (1963) famously noted, interacting with persons from other groups creates the risk that one's social identity will become "spoiled" via the association; that is, fellow in-group members may come to view ego as "contaminated" by associating with those having out-group membership, and their in-group status

*This research was supported in part by ARO award W911NF-14-1-055 and NSF award SES-1826589.

[†]Department of Sociology, University of California, Irvine

[‡]Departments of Sociology, Statistics, Computer Science, EECS, and Institute for Mathematical Behavioral Sciences, University of California, Irvine

[§]To whom correspondence should be addressed. SSPA 2145; UC Irvine; Irvine, CA 92697-5100; buttsc@uci.edu

may become suspect. This, in turn, may lead members of ego's in-group to withdraw contact from him or her. As Goffman observes, this mechanism is often sufficiently well-understood that ego will anticipate becoming ostracized for out-group contacts, and simply avoid creating them in the first place (especially where ego's in-group contacts are numerous, and the costs of ostracism correspondingly high). A related phenomenon was posited by Heider (1946) in the context of balance theory. Within an individual's mental model, group membership (a form of *unit relation* in Heider's terminology) acts as an implicitly positive tie; if a perceiver, X, views another individual A as having a positive-valance relation to some individual B belonging to group G, then this therefore creates a positive two-path (A→B, B→G) in X's mental model. Because closure of a positive two-path by a negative edge is unbalanced, X is predicted to perceive A as having a non-negative (and likely positive) association with G. Now, consider the case in which X has a strongly negative association with G and sees the A→G relation as positive. For X to form a positive tie to A is now unbalanced, and hence dissonant. Thus, the perception that someone is consorting with a member of a negatively perceived group can suppress willingness to interact with them positively, on purely affective (balance theoretic) grounds. Where group members perceive their own groups positively and others negatively, this mechanism leads to ostracism of boundary spanners.

While the social psychological mechanisms described by Goffman (stigma) and Heider (affective balance) are distinct, both have a common consequence for group structure: under the conditions where either is active, we should see that (ceteris paribus) (1) ego's propensity to form and maintain out-group ties is declining with their number of in-group ties; and (2) in-group members' propensities to form and main ties with ego are declining with ego's number of out-group ties. In this way, both mechanisms serve to *actively maintain group boundaries,* by discouraging boundary spanning and by marginalizing those who engage in it. This is in contrast with standard differential mixing effects, which (in the case of homophilous mixing) simply posit a constant reduced baseline probability of cross-group versus within-group ties. Homophilous mixing does not impose a conditional dependence between in-group and out-group ties, and does not per se lead to marginalization of boundary spanners. As such, general tendencies towards homophilous mixing do not capture active boundary maintenance of the sort described here.

In the remainder of this paper, we introduce a simple family of ERGM statistics that better captures active boundary maintenance, either in the general case (in-group versus out-group) or in the special case of stigmatized groups (where interactions with members of a particular group are subject to sanction by non-members). As we show, models with these statistics behave differently from models based on differential mixing alone, and their inclusion in an ERGM along with standard mixing terms can be used to test for the presence of boundary maintenance mechanisms. We also illustrate the use of these terms with a case study involving Parker and Asher's (1993) study of classroom friendships, demonstrating that such mechanisms do appear to be active in some social networks. Though our focus is on active boundary maintenance per se, the statistics considered here can also be used other purposes, some of which we also sketch. As ERGMs can be viewed as generative models for network structure, we conclude with the introduction of a framework for translating ERGM effects into ties gained or lost when the network is perturbed, which we demonstrate in the context of group boundary violations. This, too, can be generalized to other types of social mechanisms.

2 Measuring Boundary Maintenance with Inhomogenous Star Statistics

To develop a statistic for use in modeling boundary maintenance, we begin by assuming a set of N actors, each of whom is a member of some exogenously specified group (denoted A, B, C, etc.). We say that group B is *sigmatized* vis a vis the members of A if the boundary between A and B is actively maintained by the members of A. In the special case of homophilous interaction, each group is stigmatized vis a vis each other group (i.e., all boundaries are actively maintained). From our above discussion, we begin with the basic intuition that, where A actively maintains its boundary with B, the following should hold:

1. Let i be a member of group A and j a member of B. Then the conditional probability of an i, j edge declines with the number of ties from members of A to i.

2. Let i and j both be members of group A. Then the conditional probability of an i, j edge declines with the number of ties from j to members of group B.

To implement this intuition within an ERGM context, we recall that the conditional probability of an i, j edge in a graph represented by adjacency matrix Y is given by

$$\Pr(Y_{ij} = 1 | Y^c_{ij} = y^c_{ij}, X, \theta) = \operatorname{logit}^{-1}\left(\theta^T \Delta_{i,j,t}(y, X)\right), \tag{1}$$

where $\operatorname{logit}^{-1}$ is the inverse logit function, Y^c_{ij} is the set of all edge variables other than Y_{ij}, θ is a parameter vector, t is a vector of sufficient statistics, and X is a set of covariates.[1] $\Delta_{i,j,t}$ here is the *changescore functional*, defined by

$$\Delta_{i,j,t}(y, X) = t(y^+_{ij}, X) - t(y^-_{ij}, X) \tag{2}$$

where y^+_{ij} and y^-_{ij} are matrices such that $y^+_{kl} = y^c_{kl}$ and $y^-_{kl} = y^c_{kl}$ for $(k, l) \neq (i, j)$ and otherwise $y^+_{ij} = 1$, $y^-_{ij} = 0$. Intuitively, $\Delta_{i,j,t}(y, X)$ simply returns the difference in the vector of sufficient statistics, t, obtained by setting y_{ij} to 1 versus 0, holding the rest of the network (i.e., y^c_{ij}) fixed. Since the conditional edge probability is monotone in the changescore, it follows that our goal is to select a statistic such that the difference in the statistic associated with adding or removing an edge satisfies our substantive criteria.

To define such a statistic, we begin by letting X be a vector of group memberships, such that X_i indicates the group to which vertex $i \in V$ belongs (with V being the vertex set). As before, we let A be a reference group, and B another group that is stigmatized vis a vis the members of A. We then define the following *AB-inhomgeneous 2-star statistic*:

$$t_{AB}(y, X) = \sum_{i \in V} \sum_{j \in \{V \setminus i\}} \sum_{k \in \{V \setminus i,j\}} y_{ij} y_{jk} I(X_i = A) I(X_j = A) I(X_k = B), \tag{3}$$

where I is an indicator for the truth value of its arguments. t_{AB} counts the number of 2-stars whose first two vertices belong to A, and whose third vertex belongs to B (i.e., that contain an edge from a member of A to a boundary spanner with an edge to a member of B). To see the effect of this statistic on edge probability, we compute the associated changescore function using Equation 2,

$$\Delta_{i,j,t_{AB}}(y, X) = t_{AB}(y^+_{ij}, X) - t_{AB}(y^-_{ij}, X) \tag{4}$$

$$= \begin{cases} \sum_{k \in \{V \setminus i,j\}} y_{jk} I(X_k = B) & (X_i = A) \vee (X_j = A) \\ \sum_{k \in \{V \setminus i,j\}} y_{ki} I(X_k = A) & (X_i = A) \vee (X_j = B) \\ 0 & \text{otherwise} \end{cases}. \tag{5}$$

[1] For brevity, we have tacitly taken the reference measure to be constant, since this will not affect our development.

Note that, when i and j are both in A, $\Delta_{i,j,t_{AB}}(y, X)$ counts the number of ties from j to B. Likewise, when i belongs to A and j to members of B, $\Delta_{i,j,t_{AB}}(y, X)$ counts the number of ties from members of A to i. Thus, when the associated parameter θ_{AB} is negative, t_{AB} reduces the conditional log odds of a tie from one member of A to another by $|\theta_{AB}|$ for each tie that the second member of A has to the stigmatized group. Likewise, this same condition reduces the conditional log odds of a tie from a member of A to the stigmatized group by $|\theta_{AB}|$ for each tie that the would-be sender receives from other members of A. It follows, then, that t_{AB} satisfies our desiderata for a statistic implementing active maintenance by group A vis a vis a focal stigmatized group.

Having treated the case of one group stigmatized vis a vis another, we can further generalize to other cases. Arguably the most important case is that of homophilous interaction (here better thought of as *xenophobia* in the structural sense used e.g. by Butts (2007)), where each group is stigmatized vis a vis each other group. To implement this effect, we propose a generalized *inhomogeneous 2-star* statistic,

$$t_{I2P}(y, X) = \sum_{i \in V} \sum_{j \in \{V \setminus i\}} \sum_{k \in \{V \setminus i,j\}} y_{ij} y_{jk} I(X_i = X_j) I(X_j \neq X_k). \tag{6}$$

As the name implies, t_{I2S} counts the number of 2-stars with the central actor breaching their group boundary and forming a tie to another group. To understand the action of t_{I2S} on the graph, we consider its changescore function,

$$\Delta_{i,j,t_{I2S}}(y, X) = t_{I2S}(y^+_{ij}, X) - t_{I2S}(y^-_{ij}, X) \tag{7}$$

$$= \begin{cases} \sum_{k \in \{V \setminus i,j\}} y_{jk} I(X_j \neq X_k) & X_i = X_j \\ \sum_{k \in \{V \setminus i,j\}} y_{ki} I(X_k = X_i) & X_i \neq X_j \\ 0 & \text{otherwise} \end{cases}, \tag{8}$$

which is clearly the number of outgoing ties from j to out-group members (when i and j are in the same group), or the number of incoming ties from in-group members to i (when i and j are in different groups). As before, these have a suppressive effect on tie probability when the associated parameter, θ_{I2S}, is negative, implementing the notion that in-group ties reduce the propensity to bridge to out-group members, and (likewise) that out-group ties reduce the propensity to be selected by in-group members. Inclusion of t_{I2S} hence provides a way of modeling active boundary maintenance in the context of homophilous interaction. A minimal example of an inhomogeneous 2-star is shown graphically in Figure 1.

2.1 Other Types of Inhomogeneous Two-stars

The statistics considered above constitute particular classes of inhomogeneous 2-stars, chosen on substantive theoretical grounds. Beyond these, we note that other classes of inhomogeneous 2-stars also exist. In the directed case, there are indeed 15 distinct classes of inhomogeneous 2-stars involving two vertices from one group (A) and a third from another (B), as enumerated in Figure 2. These differ both in terms of whether the B-group member is at the center of the star, and in terms of the dyadic relationships between the central vertex and the peripheral vertices. For the boundary maintenance phenomena considered here, only 2-stars with a peripheral B-group member are relevant: those with a bridging B-group member are can be viewed as parameterizing the tendency of (respectively) A-group members to have shared B-group partners, or for B-group members to have concurrent relationships with multiple A-group members. These other forms may have utility in other settings, as may forms in which all parties are members of distinct groups (not shown).

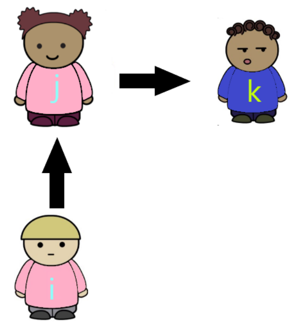

Figure 1: The above depicts an inhomogeneous 2-star, where the relevant group identity is determined based on the color of one's clothing. An actor j wearing pink sends a cross-group tie to actor k wearing blue. Actor j also receives a tie from another actor i wearing pink.

It is noteworthy that the inhomogeneous 2-stars are closely related to the *brokerage roles* enumerated by Gould and Fernandez (1989), though brokerage differs consequentially from star-configurations in requiring an absence of closure, giving brokerage roles very different structural and substantive properties. Brokerage roles (and extensions thereof) have been considered in an ERGM context by e.g. Yaveroğlu et al. (2015). We also note that the B-central inhomogeneous 2-stars are closely related to some of the multilevel network motifs described by Wang et al. (2016), especially the 2A and 2B stars which describe the propensity for a member of one group to form many (respectively few) ties to another group. Although our focus in the remainder of this paper is on the specific statistics introduced earlier in this section, we note that there is considerable room for theoretical and empirical examination of the broader space of inhomogeneous star terms.

3 Simulation Study

To demonstrate the capacity of the inhomogeneous two-star statistic to implement group boundary maintenance, we simulate group structure under three scenarios. One is a simple model without active boundary maintenance, including only edge and node mixing terms. The second adds a boundary maintenance by incorporating an additional inhomogeneous 2-star effect. The third increases the tendencies to mix within group and penalizes a tendency to mix between groups as well as the inhomogeneous 2-star effect. We simulated networks using the *ergm* package in R 3.5.0 (Hunter et al., 2008), with a custom user extension from *ergm.userterms* (Hunter et al., 2013). Networks consisted of 2 group identities, each assigned to 20 individuals. Networks were simulated using the Tie/No Tie proposal, and with a thinning interval of 1000. 1000 networks were simulated for each parameter set. The simulation coefficients are presented below in Table 3.

	Edges	Nodematch	Nodemix (Between Groups)	Inhomogeneous 2-Star
Model 1	-2	.75	-2.5	0
Model 2	-2	.75	-2.5	(-1,-0.95,...,0)
Model 3	-2	3	-1	(-1,-0.95,...,0)

Table 1: Parameters for the simulation experiment. Models 2 and 3 incorporate an inhomogeneous 2-star term, whose coefficient varies from -1 to 0 in 0.05 increments.

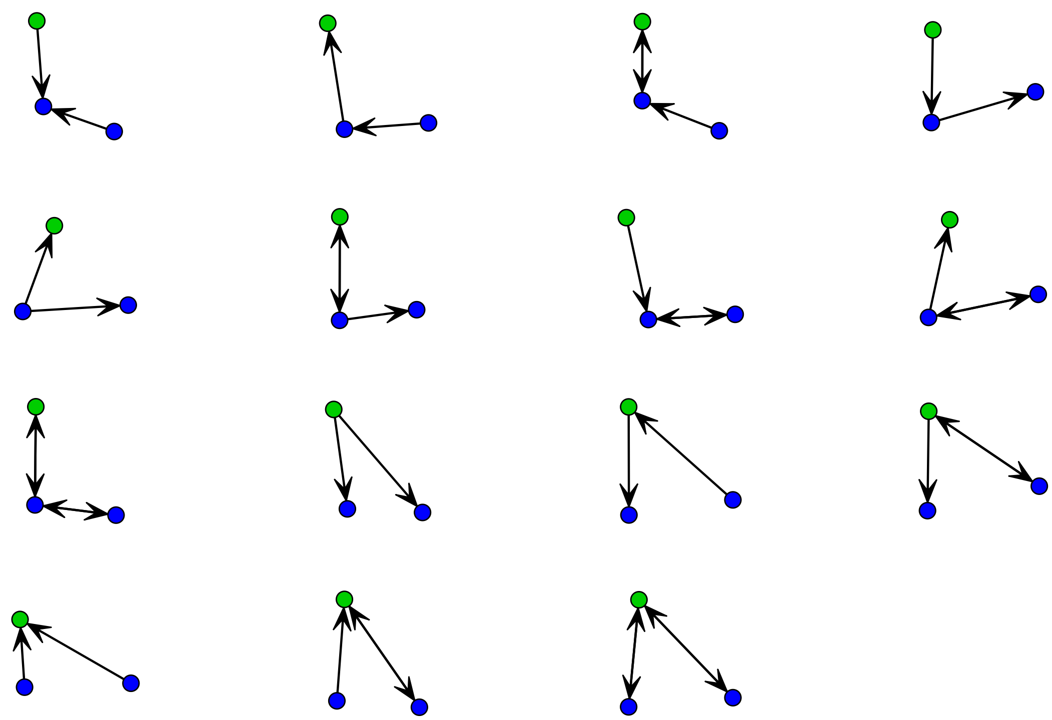

Figure 2: The above describes the 15 different isomorphic classes of the inhomogeneous 2-star for directed case. The remainder of this paper will deal with the second isomorphic class as representative of the second order stigmatic process.

From our table of simulation parameters, we can see that in Model 1, the log odds of a tie existing between two individuals who are part of different groups will be -2.5. In Model 2, the log odds of a tie existing between two individuals who are part of different groups will be -2.5 plus the inhomogeneous 2-star coefficient times the number of within-group ties Individual 1 has, plus the inhomogeneous 2-star coefficient times the number of within-group ties Individual 2 has. As the inhomogeneous 2-star coefficient varies from -1 to 0, this will penalize the log-odds by the number of within-group ties that they have. We examine the degree of group separation in each case by means of the E-I index, which measures the tendency for ties to be made between versus within groups (Krackhardt and Stern, 1988). An E-I index of 1 indicates all ties are between groups, -1 indicates that all ties within group, and 0 indicates exact equality of between versus within group ties. We simulate 1000 networks in each condition, and calculate the mean E-I index. The resulting E-I index values (means and 95% simulation intervals) are shown in Figure 3.

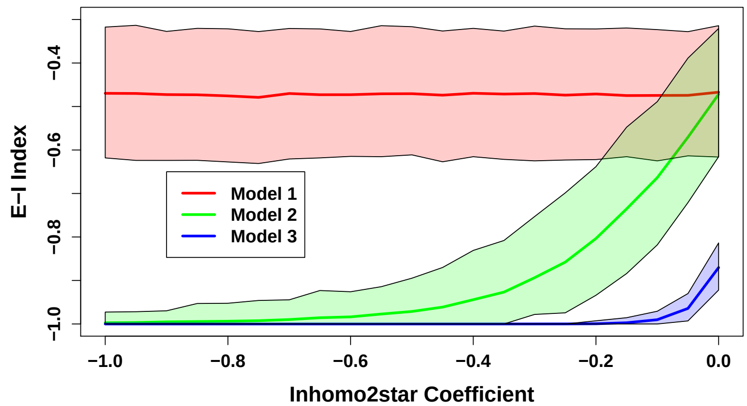

Figure 3: The mean E-I Index of the three simulations (solid lines) and 95% simulation intervals. Model 1 (red) includes edges and node mixing terms only. Model 2 (green) incorporates the inhomogeneous 2-star alongside the parameters of the first model. Model 3 (blue) intensifies this tendency of mixing towards one's own group and disassociation with the other group.

First, we note that the inclusion of active boundary maintenance effects has a substantial impact on group segregation, above and beyond what is produced by mixing effects alone (red interval). We also observe that as we vary the inhomogeneous 2-star's negative effect, we see a relatively smooth segregation response (with segregation increasing rapidly as inhomogeneous two-stars are suppressed). Moreover, in the strongly xenophobic model (blue), we observe that the convergence to full segregation occurs very quickly as the parameter falls. We also see effects on the degree of variation in the E-I index. The simulation intervals show some variation in the E-I index even with a large penalty to inhomogeneous 2-stars in the moderately segregated mixing regime, whereas in the strongly segregated regime, the variation in realized segregation is sharply suppressed. This suggests that active boundary maintenance can *amplify and reinforce* mixing effects, making its impact on segregation larger when there is already some degree of mixing inhibition on which it can act.

4 Empirical Example

To show the empirical utility of the inhomogeneous two-star term for measuring stigmatization, we examine the Parker and Asher (1993) friendship dataset, wherein a group of fifth grade students were asked to list the friends they had within their class. Within the network, there were 22 children, with 13 boys and 9 girls. Coding by gender, we see that girls and boys both have a tendency to homophilously associate, but the presence of one boy who is only friends with girls and receives no ties from boys suggests that a stigmatic process may be at play: such a pattern would be exceedingly unlikely to occur from simple mixing effects alone. To more formally investigate this possibility, we proceed by fitting an ERGM with inhomogeneous 2-star effects.

We model friendship nominations as follows. Friendship networks are known to be driven by strong norms of reciprocity (mutual tie nomination) and social phenomena such as transitive closure also feature heavily in these networks. We can model these phenomena with statistics that calculate the number of reciprocal ties and the edgewise shared partner distribution. Qualitatively,

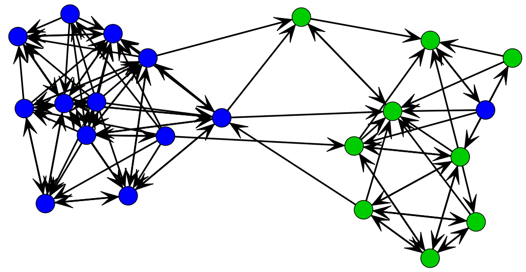

Figure 4: Friendship nominations among fifth-graders from Parker and Asher; blue nodes are boys, and green nodes are girls. The strong pattern of segregation, together with the observation that pupils with more cross-group ties tend to have fewer within-group ties, suggests active boundary maintenance by gender.

we observe from Figure 4 that friendship nominations seem to be very obviously segregated within gender, except for a lone boy who has no male friends and multiple female friends. If the second-order process of stigma by association were absent, it would be extremely unlikely that we would simultaneously see so few cross-group ties not involving this student and so many cross-group ties involving this student, nor the coincidence of this anomaly with the student's being ostracized from his own group. By contrast, this is precisely the type of anomaly that would be expected under active boundary maintenance.

To test for this possibility, we include an inhomogeneous two-star term by gender with a baseline model including within and between group mixing, reciprocity, and an effect for triadic closure (GWESP). As AIC comparison of this model versus a model with distinct within-group mixing rates by gender favored homogeneity, we employ a single base rate for within-group ties. (Note that we do not require an explicit edge term, as the baseline tie probability across groups is set by nodemix terms, and the (homogeneous) within-group probability is set by a nodematch term.) In our model exploration process, we used best subset selection to exhaustively search what subset of model terms would generate models with the lowest AIC. The model terms we assessed were nodemix, globalized inhomogeneous 2-star, edges, mutuality, geometrically weighted in-degree, geometrically weighted out-degree, geometrically weighted edgewise shared partnership distribution, and geometrically weighted dyadwise shared partnership distribution. We then assessed the model adequacy of the 50 models with the lowest AIC. The selected model has the lowest AIC and passes ERGM model adequacy checks as shown in Figure 5.[2] We present the model coefficients in Table 4. Models were estimated using the *statnet* package in R 3.5.0 (Handcock et al., 2008). The MCMC sample size was 10,000 and the MCMC interval was also 10,000.

The fitted model may be interpreted as follows. The base log-odds of same-gender nominations is approximately -1.9, with boy-girl nominations having a nominally (but not significantly) higher log-odds of -1.6 and girl-boy nominations having a substantially lower base log-odds of -3.5. We thus see an overall gender asymmetry in nominations, with boys tending to treat girls similarly to other boys, and girls tending to treat boys as substantially less attractive nomination targets than other girls (ceteris paribus). A strong and significant reciprocity effect is present, although the

[2]Note in particular that adding geometrically weighted degree effects to the final model does not materially alter the inhomogeneous 2-star effect, and the degree terms are not significant (in addition to being unfavored by AIC). The inhomogeneous 2-star effect is not hence a side effect of degree suppression.

	Estimate	Std. Error	Pr(> $\|z\|$)
Inhomo2star.Gender	-0.386	0.117	0.00095
Nodemix.Boys.Girls	-1.648	0.435	$< 1e-5$
Nodemix.Girls.Boys	-3.495	0.726	$< 1e-5$
GWESP	0.475	0.225	0.035
GWESP (α-decay)	0.803	0.331	0.0153
Mutuality	1.102	0.418	0.00846
Nodematch.Gender	-1.948	0.389	$< 1e-5$

Table 2: Estimated coefficients for friendship nominations. The null deviance is 640.5 with 462 degrees of freedom and the residual deviance of this model is 346.5 with 455 nominal degrees of freedom.

estimated effect is not large enough to result in even odds of reciprocation holding out other factors. A positive GWESP term indicates a tendency towards transitive closure, with the fairly large decay parameter (apx 0.8) indicating that the impact of multiple shared partners on closure probability is substantial. All of the above are fairly typical of homophilous friendship networks. With respect to the inhomogeneous two-star, we see a significant negative effect (apx -0.4, $p < 0.001$), indicating the presence of active boundary maintenance. Intuitively, this implies that (ceteris paribus) *each* cross-gender tie sent by a student, j, decreases the conditional log-odds of an (i, j) nomination from a student i of the same gender by roughly 0.4; likewise, each incoming (i, j) nomination decreases the conditional log-odds of a cross-gender tie sent by j by the same margin. Thus, even a small degree of group embeddedness is predicted to inhibit cross-group nominations, and even fairly minimal extension of cross-group ties is predicted to lead to substantial ostracism by one's own group. We explore this phenomenon further in the next section, by means of conditional simulation.

5 Investigating Boundary Maintenance via Network Perturbation

To provide another way of looking at how boundary maintenance operates, we use the generative properties of the ERGM family to perform a *perturbation analysis* of the Parker and Asher model, in effect asking how particular changes to the network - such as the creation of a cross-group tie - would be predicted to affect the conditional probability of observing other network features. Specifically, we focus on two types of hypothetical scenarios which we find fruitful for understanding what the social forces found above mean for friendship networks. Each represents, in a basic sense, a violation of group boundary norms; the question is then how such a perturbation would be predicted to propagate to others' relationships with those involved in the violation, all other things being held constant. The scenarios we consider are as follows:

1. The tie variable y_{AB} is toggled "on" (In this case, A initiates a cross-group tie to non-neighbor B)
2. An exogenous covariate x_A for actor A changes value.

We note that one value of computational modeling is that we can examine unusual scenarios in terms of their impact on an identified set of social mechanisms; thus, while gender change (scenario two) would be unusual in this population, and would be accompanied in an actual event by other social consequences, here we can use a hypothetical gender change that involves *only* the specified mechanisms as a tool to probe the effect of active boundary maintenance on friendship. Both scenarios should be viewed as thought experiments rather than actionable predictions, though they shed some light on what would be expected to happen were no other mechanisms operational.

Scenario 1: y_{AB} is toggled on

Although one could imagine wanting to examine the change in any network statistic associated with an exogenous tie change, an intuitive quantity that one may wish to examine (especially given the setting of stigma) is the expected change in the in-degree of A. That is, given that we add a tie from A to out-group member B, what are the social consequences for person A (ceteris paribus) in terms of incoming friendship nominations? Specifically, we examine the difference

$$\mathbf{E}[\sum_{j}^{J} Y_{jA}|\theta, X, Y_{.A}^{C} = (y_{AB}^{+})_{.A}^{C}] - \mathbf{E}[\sum_{j}^{J} Y_{jA}|\theta, X, Y_{.A}^{C} = (y_{AB}^{-})_{.A}^{C}], \tag{9}$$

i.e. the difference in the conditional expectation of A's in-degree when the new tie is present, versus when it is absent (holding the rest of the network constant), under our friendship model. The coefficients θ are derived from our above-fit friendship model, and the exogenous covariates X and the observed network y come from the dataset. y_{AB} refers to the tie variable being sent by actor A to actor B. $(y_{AB}^{+})^{C}$ and $(y_{AB}^{-})^{C}$ refer to every single tie variable being fixed except for y_{AB} which is either toggled on or off, respectively, and $y_{.A}^{C}$ refers to all elements of y except the incoming ties to A. We represent this graphically in Figure 6. Our interest is on the predicted states of the green edge variables (here represented schematically as "missing") given the toggled tie (blue), holding the rest of the network (black) fixed. We summarize the difference in the conditional behavior of the randomized edge variables when the focal edge is toggled in terms of A's expected in-degree change, averaging over all possible realizations of the vector of incoming ties in each case.

$$\begin{array}{c c} & \begin{array}{cccc} A & B & C & D \end{array} \\ \begin{array}{c} A \\ B \\ C \\ D \end{array} & \begin{pmatrix} 0 & 1 & 1 & 1 \\ NA & 0 & 0 & 0 \\ NA & 1 & 0 & 1 \\ NA & 0 & 1 & 0 \end{pmatrix} \end{array}$$

Figure 6: Sociomatrix representation of the assessment of the expected change in in-degree as a result of ego A's sending a cross-group tie to B (blue element) to whom they were not previously tied. Green edge variables are free to vary in the conditional simulation, while black elements are fixed.

Within the context of the Parker and Asher dataset, we show the distribution of expected degree change when a cross-group tie is made in Figure 7.

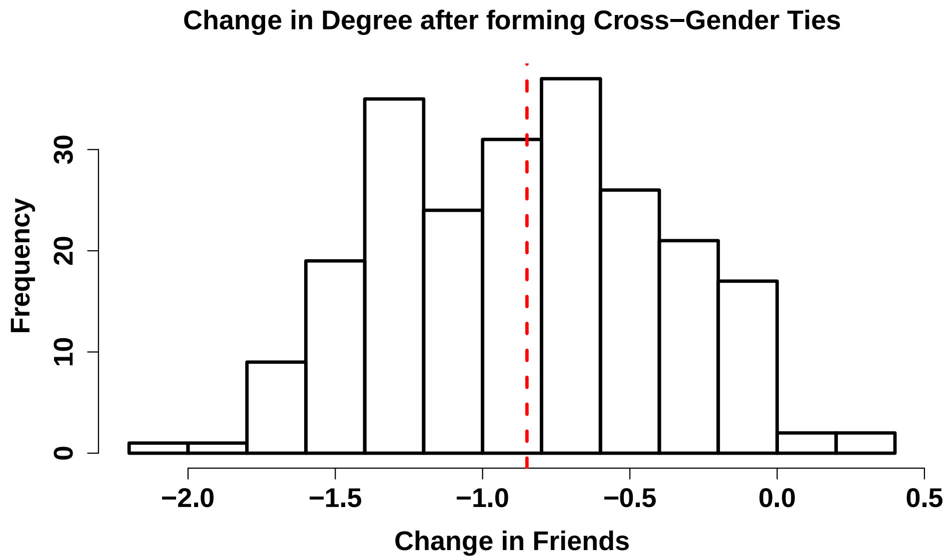

Figure 7: Distribution of in-degree change for a random actor A, conditional on the rest of the network being at its observed state and a cross-gender tie added from A to random individual B. The dotted red line depicts the mean change in degree. On average, forming a cross-group tie is predicted to result in loss of incoming friendship nominations.

As we can see in this scenario, sending a new cross-gender tie to a random member of the other group results in an average loss of 0.81 friendship nominations. Given that the average number of friendship nominations is 4.682, in expectation, an individual who forms a cross-gender tie can expect to lose roughly 17% of their friendship nominations. This outcome results despite the effects of reciprocity and GWESP, which tend to add edges in response to additional edges; here, these are outweighed by the boundary maintenance mechanism (implemented via the inhomogeneous two-star effect), which makes A less attractive to their in-group when they send more ties to outsiders.

Scenario 2: x_A for an actor A changes value

Similar to the above, we may observe the predicted consequence for A's incoming friendship nominations when his or her membership status is changed, holding all else fixed. Though this can be viewed as an abstract thought experiment, it also provides a simplified scenario for what might happen when an individual previously thought to have a "normal" identity (Goffman, 1963) is revealed to have an identity that is stigmatized with respect to his or her previous in-group (e.g. HIV stigma, sexual orientation, gender orientation, criminal background). In the context of the student network, gender change has an additional complication: it involves A's switching to a stigmatized group from the perspective of A's former in-group members, but at the same time switching away from a stigmatized group from the perspective of their former out-group members. We could hence expect that such a change will have little impact on A's indegree, since ties lost from one group could be made up by the other. On the other hand, unilateral change in A's group membership also converts his or her outgoing ingroup ties to outgroup ties, which may provoke ostracism.

The equation below describes the expected in-degree change of A conditional on A's changing exogenous covariate and all of the other ties remaining fixed:

$$E[\sum_{j}^{J} Y_{jA}|\theta, X', Y_{.A}^{C} = y_{.A}^{C}] - E[\sum_{j}^{J} Y_{jA}|\theta, X, Y_{.A}^{C} = y_{.A}^{C}], \tag{10}$$

with $X'_i = X_i$ for all $i \neq A$, and X'_A being "flipped" to the alternate state. We represent this graphically in Figure 8. Every color has the same meaning as in scenario 1, but in this case the external covariate on A is set by the researcher to its new value.

$$\begin{array}{cc} & \begin{array}{cccc} \color{blue}{A} & B & C & D \end{array} \\ \begin{array}{c} \color{blue}{A} \\ B \\ C \\ D \end{array} & \begin{pmatrix} 0 & 0 & 1 & 1 \\ \color{green}{NA} & 0 & 0 & 0 \\ \color{green}{NA} & 1 & 0 & 1 \\ \color{green}{NA} & 0 & 1 & 0 \end{pmatrix} \end{array}$$

Figure 8: Sociomatrix representation of the assessment of the expected change in in-degree as a result of ego A's covariate changing (blue). Green edge variables are free to vary, while those in black are held fixed.

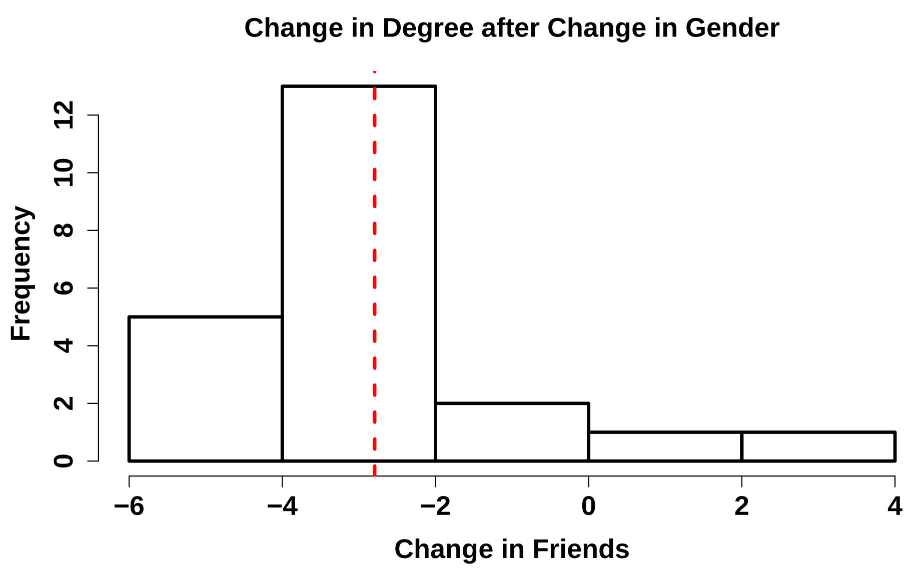

Figure 9: Distribution of in-degree change conditional on the rest of the network being constant when a random individual's gender is changed. The dotted red line depicts the mean change in in-degree. On average, swapping group membership leads to tie loss.

We can see from Figure 9 that, on average, switching a random student's gender leads to a loss of 2.79 friendship nominations (approximately 59%). The core driving mechanism here is the sudden reconsideration of ego's formerly within-gender nominations as cross-gender ties, making ego an unattractive nomination target for members of his or her new in-group; although reciprocity and closure effects involving ties to and among ego's former in-group will, on average, help preserve some nominations from that group, this is not sufficient to overcome the lack of nominations from the new in-group. While actual gender transitions are far more complex than treated here, this simple thought experiment highlights a general challenge for actors crossing actively maintained boundaries: *the same ties that once bound them to their in-group become liabilities to acceptance once a transition is made, while their incoming ties from former in-group members (now out-group members) are at high risk of being lost.* Overcoming this liability hence requires far more extensive efforts by ego to reorganize his or her personal network, including both outreach to the new in-group and elimination of ties to the old in-group. For strong ties, this may impose substantial emotional or other costs, making the assumption of a new group identity substantially more difficult than it might otherwise be.

6 Discussion

As we have shown, inhomogeneous two-star terms are effective tools for representing stigmatic or other active group boundary maintenance processes within network models, capturing the tendency of group members to ostracize one of their own when he or she associates with members of a stigmatized out-group. As they amplify the impact of differential mixing, it is reasonable to suspect that these two types of group maintenance effects will often co-occur in practice. Indeed, we see exactly this in the context of the Parker and Asher study, which shows a combination of simple mixing inhibition and active maintenance. It is plausible that this effect has been missed in prior studies of homophily, and testing for possible inhomogeneous two-star effects is something that should be considered where it is plausible that stigmatic or related processes are at work. In such contexts, we would note that a *lack* of inhomogeneous two-star effects may be just as interesting as their presence: their absence may suggest that homophily is driven either by "neutral" mechanisms like differential contact, or by strictly positive selection for alters with similar attributes.

As the effects studied here are extensions of the conventional 2-star statistic, it is reasonable to consider whether they pose risks of degeneracy. Degeneracy in ERGMs refers to a phenomenon that arises from unrealistic model specifications, in which probability becomes heavily concentrated on a small number of (usually implausible) structures as graph size increases. Degenerate models are informative, in that they indicate that the theory embodied by the proposed model leads to emergent network structures that are not realistic; given that one is attempting to model an observed system, however, avoiding degenerate models is a natural goal.[3] Previous explorations of standard k-star terms have shown that (when not balanced as in alternating k-star sequences Hunter (2007); Snijders et al. (2006)) they typically lead to degenerate distributions when paired with positive parameter values Handcock (2003). When $\theta > 0$, the terms used here can pose a similar risk of degenerate behavior. Intuitively, this can be understood as follows: when $\theta > 0$, the more ties that an individual has to out-group members, the more attractive he or she is to in-group members; and, likewise, the more strongly an individual is tied to his or her in-group, the greater his or her tendency to have out-group ties. This putative mechanism (if not checked by other effects) can lead to a positive feedback loop where boundary spanning encourages more in-group attention, which in turn feeds more boundary spanning, until a density explosion Butts (2011) occurs. This is implausible for a typical social network.

When $\theta < 0$, however, our mechanism is very different: out-group ties tend to inhibit in-group

[3] That said, if a seemingly plausible model for a given network turns out to be degenerate, one should view this as theoretically meaningful, just as in other cases in which a proposed model fails to explain one's data. It should also be noted that degenerate models can in some cases be realistic, e.g. Grazioli et al. (2019).

ties (and vice-versa), leading to self-limiting behavior. In this regime, inhomogeneous two-star terms cannot fuel a density explosion (because they do not enhance tie formation), and even their potential inhibitory effects are limited by bounds on mean degree (since penalties to tie formation scale with incoming and/or outbound ties). Thus, in the regime of interest here, model degeneracy is not a major concern.

As with the general case of the 2-star statistic, which is a special case of the more general k-star family, it is possible to construct higher-order versions of the statistics described here. A collection of such terms could subsequently be employed in curved statistics, analogous to the geometrically weighted k-stars often employed to model degree distributions Hunter (2007). Such statistics could facilitate the modeling of diminishing marginal effects of boundary spanning on tie formation, rather than the constant incremental effect of the inhomogeneous two-star. However, because inhomogeneous stars of order greater than 2 can also vary in their composition, determining how stars of varying inhomogeneity should be weighted is non-trivial. This would seem to be a potentially fruitful direction for subsequent work.

The network perturbation framework we introduced is not limited to calculating expected degree change and can be generalized to any network statistic of interest. Given the cross-sectional nature of the data and the assumptions involved, interpreting the differences in expected degree as dynamic predictions must be viewed as heuristic (at best); however, such "what if" experiments provide important insights into the behavior of structural mechanisms that would be difficult to obtain otherwise. In particular, here they allow us to see that, for the Parker and Asher classroom, tie-forming dyadic and triadic closure mechanisms are outweighed by the effects of ostracism, and, ceteris paribus, perturbations leading ego to span group boundaries (whether by extending nominations or by changing their own status) would proximately be predicted to be associated with loss of friendship nominations. Comparative analysis of such effects, combined with ethnographic or perceptual data, could be useful in assessing the extent to which individuals' perceived pressure to conform to group norms or refrain from cross-group interaction in fact correlate with the plausible risks of deviation.

Finally, we note that, while the active boundary maintenance mechanisms described here provide one route to the suppression of inhomogeneous two-stars, it is possible that other social mechanisms may also produce such an effect: as always, interpretations of estimated ERGM parameters in terms of putative generative mechanisms should be made in the context of what else is known regarding the network in question. For instance - as with other forms of dependence in network data - the distribution of inhomogeneous two-stars may be affected by mixing effects on unobserved covariates. For instance, in the Parker and Asher case shown above, it is not impossible that the apparently ostracized boy in fact has some specific personal characteristic that simultaneously leads him to form numerous ties with girls and be subject to exclusion by the other boys due to mixing alone, without any additional social process. Such possibilities can never be entirely ruled out on the basis of cross-sectional data. However, such an explanation would require one to posit not only the presence of one or more unobserved covariates with strong mixing effects, but also to propose that the values of such covariates happen to be distributed in such a way as to mimic what would be expected from ostracism. Such accounts are difficult to falsify. While our view is thus that active boundary maintenance often provides a plausible and parsimonious account of inhomogeneous two-star suppression, it is important to be open to other possibilities. Investigations of the mechanisms of two-star suppression using dynamic data will undoubtedly refine our understanding on this issue.

7 Conclusion

The ERGM framework provides a powerful framework to study the effects of multiple social processes in shaping a network structure. The inhomogeneous two-star term provides a simple and elegant means of incorporating active maintenance of group boundaries implied e.g. by some sociological theories of stigma easily into model-based network analysis. Our simulation analyses demonstrate that this family of mechanisms operates differently from (and interacts with) simple differences in mixing rates, and our illustrative empirical analysis shows that the mechanisms are operative in at least some settings. It is hoped that terms like that introduced here will serve to broaden the reach of parametric network modeling, and clarify the connections between structural biases and the social mechanisms that generate them.

It should be noted, that the strength of this approach for evaluating stigma in a population does not evaluate the respondent's level of perceived stigmatization. Rather, it evaluates the level of *structural* stigma at the network level. Considering this can be evaluated at both the ego-level and the complete network level, this provides a considerable utility, both in understanding an ego's propensity to stigmatize and in yielding a global metric of tendency towards stigmatization. The ability to quantify stigma without appealing to attitudinal measures provides additional benefits. First, from a data collection perspective, indirectly inferring stigmatization from network structure avoids the need to ask sensitive questions of respondents regarding negative experiences with other network members, potentially increasing response rates and allowing the phenomenon to be studied in settings for which direct queries would be problematic (e.g., workplaces). Second, attitudinal items involving stigma may be subject to stronger self-presentation biases than more neutral network items, and errors with respect to the latter can be reduced by approaches such as multiple informant designs (Butts, 2003). Finally, if the sociological phenomenon of interest is some behavioral change dependent on one's social ties (i.e. an influence process or contagion process), then it is irrelevant whether or not the ego perceives themselves as stigmatizing others, the relevant mechanism is dependent on the propensity for them to form ties across group and how that propensity is affected by the ties to their own group. Thus, an explicitly structural measure provides a more direct way of understanding how a stigmatic process affecting network structure can subsequently impede diffusion across group.

We conclude that the inohomogeneous 2-paths provide a simple but powerful family of statistics for explaining how cohesive subgroups based on some ascriptive characteristic (i.e. gender, race, etc.) arise from tendencies towards active group boundary maintenance. It is to be hoped that incorporating such terms into models of group structure will provide a useful tool for assessing when and how social groups actively combat boundary spanning, versus simply being kept apart by mixing alone.

8 References

Blau, P. M. (1994). *Structural Contexts of Opportunities.* University of Chicago Press.

Butts, C. T. (2003). Network Inference, Error, and Informant (In)Accuracy: a Bayesian Approach. *Social Networks*, 25(2):103–140.

Butts, C. T. (2007). Models for Generalized Location Systems. *Sociological Methodology*, 37(1):283–348.

Butts, C. T. (2011). Bernoulli Graph Bounds for General Random Graphs. *Sociological Methodology*, 41:299–345.

Goffman, E. (1963). *Stigma: Notes on the Management of Spoiled Identity.* Simon and Schuster, New York.

Goodreau, S. M., Kitts, J. A., and Morris, M. (2009). Birds of a Feather, Or Friend of a Friend?: Using Exponential Random Graph Models to Investigate Adolescent Social Networks. *Demography*, 46(1):103–125.

Goodreau, S. M., Rosenberg, E. S., Jenness, S. M., Luisi, N., Stansfield, S. E., Millett, G. A., and Sullivan, P. S. (2017). Sources of Racial Disparities in HIV Prevalence in Men who have Sex with Men in Atlanta, GA, USA: A Modelling Study. *The Lancet HIV*, 4(7):311–320.

Gould, R. V. and Fernandez, R. M. (1989). Structures of Mediation: A Formal Approach to Brokerage in Transaction Networks. *Sociological Methodology*, pages 89–126.

Grazioli, G., Yu, Y., Unhelkar, M. H., Martin, R. W., and Butts, C. T. (2019). Network-based Classification and Modeling of Amyloid Fibrils. *Journal of Physical Chemistry, B*, 123(26):5452–5462.

Handcock, M. S. (2003). Statistical Models for Social Networks: Inference and Degeneracy. In Breiger, R., Carley, K. M., and Pattison, P., editors, *Dynamic Social Network Modeling and Analysis*, pages 229–240. National Academies Press, Washington, DC.

Handcock, M. S., Hunter, D. R., Butts, C. T., Goodreau, S. M., and Morris, M. (2008). statnet: Software tools for the representation, visualization, analysis and simulation of network data. *Journal of Statistical Software*, 24(1):1–11.

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

Hunter, D. R. (2007). Curved Exponential Family Models for Social Networks. *Social Networks*, 29:216–230.

Hunter, D. R., Goodreau, S. M., and Handcock, M. S. (2013). ergm.userterms: A Template Package for Extending statnet. *Journal of Statistical Software*, 52(2):1–25.

Hunter, D. R., Handcock, M. S., Butts, C. T., Goodreau, S. M., and Morris, M. (2008). ergm: A Package to Fit, Simulate and Diagnose Exponential-Family Models for Networks. *Journal of Statistical Software*, 24(3).

Krackhardt, D. and Stern, R. N. (1988). Informal Networks and Organizational Crises: An Experimental Simulation. *Social Psychology Quarterly*, 51:123–140.

McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*, 27(1):415–444.

Parker, J. G. and Asher, S. R. (1993). Friendship and Friendship Quality in Middle Childhood: Links with Peer Group Acceptance and Feelings of Loneliness and Social Dissatisfaction. *Developmental Psychology*, 29(4):611.

Sakoda, J. M. (1971). The Checkerboard Model of Social Interaction. *Journal of Mathematical Sociology*, 1(1):119–132.

Schaefer, D. R., Bouchard, M., Young, J. T., and Kreager, D. A. (2017). Friends in Locked Places: An Investigation of Prison Inmate Network Structure. *Social Networks*, 51:88–103.

Schelling, T. C. (1971). Dynamic Models of Segregation. *Journal of Mathematical Sociology*, 1(2):143–186.

Snijders, T. A. B., Pattison, P. E., Robins, G. L., and Handcock, M. S. (2006). New Specifications for Exponential Random Graph Models. *Sociological Methodology*, 36.

Wang, P., Robins, G., and Matous, P. (2016). Multilevel Network Analysis Using ERGM and its Extension. In *Multilevel Network Analysis for the Social Sciences*, pages 125–143. Springer.

Yaveroğlu, O. N., Fitzhugh, S. M., Kurant, M., Markopoulou, A., Butts, C. T., and Pržulj, N. (2015). `ergm.graphlets`: A package for ERG modeling based on graphlet statistics. *Journal of Statistical Software*, 65(12).

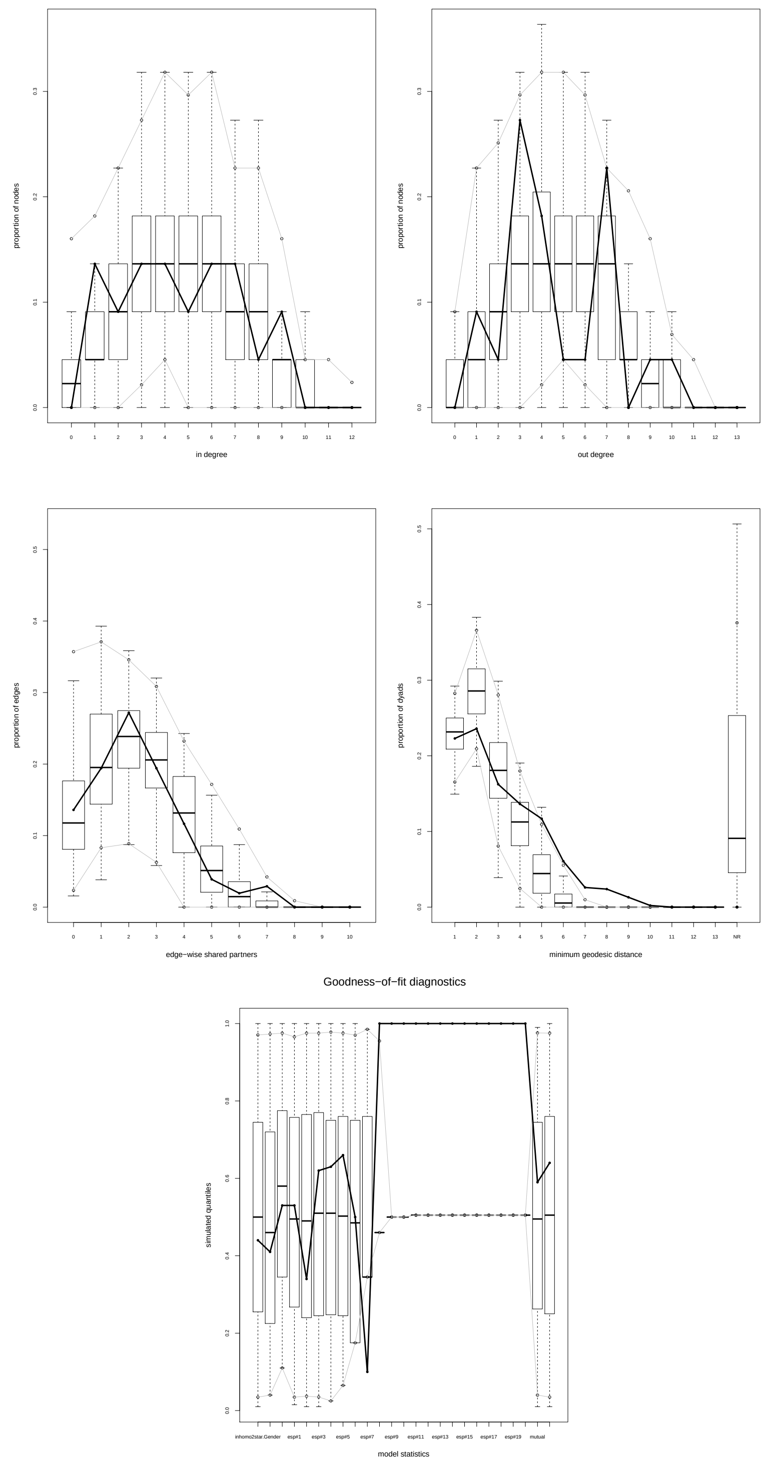

Figure 5: Adequacy checks for the Parker and Asher model. By simulating networks from the estimated model, we test whether various aspects of the simulated networks (e.g. the in/out-degree distribution, the geodesic distribution, the shared partner distribution, and the mean statistics) could potentially reproduce those same aspects observed in the dataset. The black lines depict the observed network statistics from the dataset, and the gray lines depict 95% simulation intervals under the estimated model.