An optimal test for strategic interaction in network formation games

Andrin Pelican and Bryan S. Graham *

November 12, 2024

*Pelican: e-mail: pelicanandrin@gmail.com. Graham: Department of Economics, University of California - Berkeley, 530 Evans Hall #3380, Berkeley, CA 94720-3880 and National Bureau of Economic Research, e-mail: bgraham@econ.berkeley.edu, web: http://bryangraham.github.io/econometrics/. We thank seminar participants at the University of California - Berkeley, the University of Cambridge, Harvard-MIT, ITAM, Jinan University, Yale University, Microsoft Research, Princeton University, Rice University, Boston University, the DC-MD-VA Econometrics Workshop, the World Congress of the Econometric Society, the North American Winter Meetings of the Econometric Society and the European Conference on Networks for feedback and suggestions. We also thank Daniele Ballinari, Vincent Boucher, Shuowen Chen, Enrico De Giorgi, Aureo de Paula, Péter Erdös, Bo Honoré, Hiroaki Kaido, Piotr Lukaszuk, Ulrich Mueller, Kaïla A. Munro, Gabriel Okasa, Seth Richards-Shubik, Xun Tang and Yi Zhang for their generous feedback and input on earlier drafts. A special thanks to Winfried Hochstättler for continuously pointing out the connection to the discrete mathematics literature and to Michael Jansson for help and insight into the nature of our testing problem. The paper has also benefited from discussions about the theory of strategic network formation with Sanjeev Goyal and Matt Jackson, multiple rounds of exceptionally detailed comments from the co-editor, as well as the feedback and suggestions provided by four anonymous referees. All the usual disclaimers apply. Financial support from NSF grant SES #1357499 is gratefully acknowledged by the second author. Portions of the research reported here were also undertaken during a visit to CEMFI by the second author with support from the Spanish State Research Agency under the María de Maeztu Unit of Excellence Programme (Project No: CEX2020-001104-M). This paper is a revised version of the co-authored chapter "Testing Strategic Interaction in Networks" which appeared in first author's 2019 doctoral dissertation at the University of St. Gallen. Earlier drafts of this paper circulated under the longer title "An optimal test for strategic interaction in social and economic network formation between heterogeneous agents".

Abstract

Consider a setting where N players, partitioned into K observable types, form a directed network. Agents' preferences over the form of the network consist of an arbitrary network benefit function (e.g., agents may have preferences over their network centrality) and a private, or dyadic, component which is additively separable in own links. This latter component allows for unobserved heterogeneity in the costs of sending and receiving links across agents (respectively out- and in- degree heterogeneity) as well as homophily/heterophily across the K types of agents. In contrast, the network benefit function allows agents' preferences over links to vary with the presence or absence of links elsewhere in the network (and hence with the link formation behavior of their peers). In the null model which excludes the network benefit function, links form independently across dyads in the manner described by Charbonneau (2017) among others. Under the alternative there is interdependence across linking decisions (i.e., strategic interaction). We show how to test the null with power optimized in specific directions. These alternative directions include many common models of strategic network formation (e.g., "connections" models, "structural hole" models etc.). Our random utility specification induces an exponential family structure under the null which we exploit to construct a similar test which exactly controls size (despite the the null being a composite one with many nuisance parameters). We further show how to construct locally best tests for specific alternatives without making any assumptions about equilibrium selection. To make our tests feasible we introduce a new MCMC algorithm for simulating the null distributions of our test statistics.

JEL Codes: C31, C57

Keywords: *Network formation, Locally Best Tests, Similar Tests, Exponential Family, Incomplete Models, Degree Heterogeneity, Homophily, Binary Matrix Simulation, Edge Switching Algorithms, Markov Chain Monte Carlo, Partitioned Graph*

In an economic model of *directed* network formation agents purposefully direct links to one another in order to maximize utility. A payoff function maps all possible network configurations into agent utilities. Agents use this payoff function to weigh the benefits of directing any particular link against the costs of doing so. A Nash Equilibrium (NE) network arises when all agents link choices are individually optimal given the choices made by other agents (e.g. Bala and Goyal, 2000).

Important examples of directed link formation in economics include firms choosing their suppliers (e.g., Atalay et al., 2011; Amelkin and Vohra, 2023), social media users choosing who to follow (e.g., Grandjean, 2016), adolescents selecting friends (e.g., Christakis et al., 2020), banks extending credit (or not) to firms (e.g., Marotta et al., 2015), and village households seeking assistance from peers in times of economic stress (e.g., De Weerdt, 2004). Jackson et al. (2017) present many other examples of networks in economics. Such data abound in the other social sciences as well (e.g., Apicella et al., 2012).

The utility an agent receives when she directs a link to another agent can be usefully divided into two components.[1] The first component is "private", or, more precisely dyadic. It is invariant to the presence or absence of other links in the network.[2] The second component is "social", or varying with the presence or absence of other links in the network.

An example of the first component is the payoff associated with a homophilous link (McPherson et al., 2001). This payoff component only depends on the attributes of the sending (ego) and receiving (alter) agents. Another example is associated with "degree heterogeneity": agents may vary systematically in their propensity to direct links, or in their attractiveness as link targets for others. Finally we might posit that the payoff from any particular link varies for idiosyncratic reasons, as in other random utility models (RUMs) of discrete choice (McFadden, 1974). Empirical models of network formation with these features were formally studied by Charbonneau (2017), Graham (2017), Dzemski (2018), Jochmans (2018) and Yan et al. (2018). These models are fundamentally *dyadic*: agents' network payoffs are a simple sum of link-specific payoffs and, crucially, invariant to the linking behavior of other agents.

In some settings, however, agents may also value indirect links. For example, an arc from j to k may incidentally reduce the shortest path length from i to k, allowing agent i better access to k's information (e.g., Jackson and Wolinsky, 1996; Bala and Goyal, 2000). While arc jk is valued by i, this value is not incorporated into j's decision to direct the arc or not. Preferences of this type mean agents' decisions impose *externalities* on others. The detection

[1] In digraphs, or directed networks, it is customary to refer to edges as "arcs". Here we use the terms link, edge, arc, friendship, relationship etc. interchangeably.

[2] Note "other links" include those possibility directed by the sending agent to targets other than the one at hand.

of such externalities is the subject of this paper. Our setting is NE directed networks, as elegantly studied by Bala and Goyal (2000).[3]

Payoff functions with externalities feature prominently in formal theoretical models of network formation (cf., Jackson, 2008; Goyal, 2023). Equilibrium in network formation models with externalities may be analyzed using the tools of game theory. Indeed such models are typically called strategic network formation models. In what follows we say a network formation model is *strategic* if agents value indirect links or, equivalently, their optimal linking strategy varies with the linking behavior of others.

When links made by one agent alter the incentives for link formation faced by others, equilibrium network configurations may diverge from socially optimal ones (Goyal, 2023). This, in turn, suggests that well-designed interventions might make agents better off. In contrast, without a wedge between the private and social benefits of link formation, equilibrium and socially optimal networks will coincide. This paper introduces a test for whether agents' own incentives to form links vary with the choices of others. A rejection of our test, *under the maintained model*, indicates the presence of externalities, with their attendant implications for optimal policy design.

An overview of the test and its uses

Strategic network formation games are complicated. In a directed network with N agents, there are $n = N(N-1)$ strategic decisions to make, and hence a total of 2^n possible action profiles or network configurations; many of which may be Nash Equilibria (NE). In the seminal model of directed network formation introduced by Bala and Goyal (2000), for example, with $N = 5$ agents there are $1,069$ NE networks. Because of this combinatoric complexity, methods pioneered for the econometric analysis of discrete games with just a few players are not directly applicable – at least in practice – to network formation games.

In recent work, Christakis et al. (2020), Mele (2017), Miyauchi (2016), de Paula et al. (2018) and Sheng (2020) each proposed empirical models of strategic network formation.[4] Each of these models impose particular restrictions on the form of the network payoff function, the nature of any unobserved heterogeneity, and/or make assumptions about equilibrium selection. For example, Christakis et al. (2020) and Mele (2017) resolve incompleteness by assuming agents form links sequentially, allowing for the application of likelihood-based methods. Miyauchi (2016) requires a super-modular payoff function, de Paula et al. (2018) a payoff function which varies only with local network structure, while Sheng (2020) focuses

[3] Ideas presented below could be adapted to settings requiring other solution concepts, such as pairwise stable undirected networks with, or without, transfers (e.g., Bloch and Jackson, 2007; Jackson and Wolinsky, 1996).

[4] de Paula (2020) surveys work in this area and provides additional references.

on agents with a taste for transitivity. None of these papers incorporate agent-specific unobserved heterogeneity. Even with these restrictions, estimating the identified set for the parameters indexing the network payoff function in these models is challenging, as is conducting inference.[5]

Unlike this prior work, we do not consider set identification in this paper, focusing instead on the more modest goal of externality *detection*. Are preference externalities present or not? With this target question in mind, we introduce an econometric model of strategic network formation which simultaneously allows (i) for agents to value both direct and indirect links (i.e., a freely specified network benefit function), (ii) for the systematic returns to link formation to vary with observed dyad attributes, and (iii) for unobserved agent-specific correlated degree heterogeneity. While, relative to prior work, our focus is narrower, our model is richly featured.

Our setup maps neatly into the "costs versus benefits" payoff structures emphasized in theoretical models of strategic network formation (see, for example, Jackson (2008, Chapters 6 & 11) and Goyal (2023, Chapter 3)). Examples of models – (i) suitably enriched to include covariates, unobserved heterogeneity, and random link utility and (ii) adapted (if needed) to match our use of NE as a solution concept – encompassed by our framework include the "connections" model (e.g., Jackson and Wolinsky, 1996; Bala and Goyal, 2000), "structural hole" or "bridging" models (e.g., Goyal and Vega-Redondo, 2007; Kleinberg et al., 2008) and the favor exchange or "supported links" model of Jackson et al. (2012). We can also accommodate tastes for reciprocity, transitivity, network centrality and other forms of indirect link valuation.

We begin with a baseline dyadic logistic regression model for directed networks. Variants of this model have featured in applied social science research for decades (e.g., Bennett and Stam, 2000; De Weerdt, 2004). Early formal econometric analyses include those by Charbonneau (2017), Jochmans (2018) and Yan et al. (2018). Work that builds upon on large-N, large-T panel data research (e.g., Fernández-Val and Weidner, 2016) as well as conditioning arguments used in fixed-T panel settings (e.g., Chamberlain, 1980, 2010).

The dyadic logit model is useful for modeling homophily and degree heterogeneity. We augment this model with a network payoff term which additionally allows agents to value indirect links. The resulting model is quite complicated. Formally it is a very large complete information simultaneous move game. While we assume that the observed network is a NE,

[5] We wish to emphasize that these "critiques" reflect the inherent difficulty of the problem, not any deficiencies in the above cited papers. Indeed these researchers have shown considerable ingenuity in proposing ways to make methods designed for games with just a few players scale to the considerably more complicated many-player network setting. We also comment that these papers don't all use some same solution/equilibrium concept.

we make no auxiliary equilibrium selection assumptions.[6]

Let K be the number of support points in the distribution of observed agent attributes and N the number of agents in the network (That K is finite with $K \ll N$ is a strong assumption; one we return to below). Our model includes (i) K^2 "homophily" parameters, collected in the $K \times K$ matrix $\Lambda \overset{def}{\equiv} [\lambda_{kl}]$ for $k, l = 1 \dots K$, capturing how link returns vary systematically with ego and alter attributes (we define $\lambda \overset{def}{\equiv} \text{vec}(\Lambda'))$, (ii) two $N \times 1$ parameter vectors $\mathbf{A} \overset{def}{\equiv} [A_i]$ and $\mathbf{B} \overset{def}{\equiv} [B_i]$ for $i = 1 \dots N$, capturing, respectively, agent-specific out- and in-degree heterogeneity, and (iii) a scalar parameter, γ, measuring the extent to which agents value indirect links. Our model also includes (iv) an "equilibrium selection" function (This function assigns probabilities to all NE equilibria for every possible realization of the $n = N(N-1)$ link-specific random utility shocks, $\mathbf{U} \overset{def}{\equiv} [U_{ij}]$ for $i \neq j$ and $i, j = 1 \dots N$). Since we are agnostic about which NE is selected in the presence of multiple equilibria, this function is not specified by the analyst, but enters our analysis abstractly (see Theorem 1.1 below).

We treat $\delta = (\lambda', \mathbf{A}', \mathbf{B}')'$ as a (high dimensional) nuisance parameter and the equilibrium selection mechanism as a nuisance function. This focuses our attention solely on γ. While, in principle, an analysis of the identified set for γ might be possible, we instead focus on the one-sided hypothesis of $H_0 : \gamma = 0$ versus $H_1 : \gamma > 0$. Or, put differently, we identify the *sign* of γ.[7]

Our test involves comparing a statistic of the observed network (e.g., its transitivity index) with a critical value derived from a reference distribution. Natural questions are: (i) which reference distribution? (ii) how do I compute the critical value? (iii) which network statistic should I use? We provide answers to all three of these questions.

There is a long tradition in empirical work of using the Ërdos-Rényi model to generate the reference distribution. This invariably results in "straw man" tests since few real world networks are well-described by the Ërdos-Rényi model. To avoid spurious rejection of the no strategic interaction null, it is therefore important to use a richer baseline model; one that might actually describe a real world network. The dyadic logit regression model with agent-specific fixed effects is one such model. This model is commonplace in empirical network analysis (e.g., Bennett and Stam, 2000; De Weerdt, 2004; Hoff, 2005; Charbonneau, 2017;

[6] More precisely the observed network is either a pure strategy NE or in the support of a mixed strategy NE (in fact our results hold under an even weaker notion of equilibrium, as explained below).

[7] Our focus on one-sided hypotheses results in a particularly clean exposition and analysis, allows for the statement of some optimality results, and covers our main examples of interest. A researcher's specification of the network benefit function typically suggests whether the null-alternative pair $H_0 : \gamma = 0$ versus $H_1 : \gamma > 0$ or $H_0 : \gamma = 0$ versus $H_1 : \gamma < 0$ is most appropriate. Since a researcher can always replace a chosen network benefit function with its negative, we focus on the former case without loss of generality. Finally, as will become obvious, much of what follows extends naturally to two-sided hypotheses.

Bartolucci, 2024) (Perhaps because it represents a natural adaptation of familiar panel data logit models to the network setting). Like its panel counterpart, the dyadic logit model provides a simple, easy to interpret, and random utility based, model for describing edge formation. It is sufficiently flexible to match *any* observed in- and out-degree sequence as well as rich patterns of homophilous linking. These features are important since heavy-tailed degree distributions characterize many real world networks, as does homophily (e.g., Barabási, 2016; McPherson et al., 2001). While recent research explores more flexible models of dyadic link formation (e.g., Gao, 2020; Chen et al., 2021), the logit model remains a workhorse for practitioners.

Some models of *strategic* network formation tend to induce link clustering (e.g., Jackson et al., 2012), while others skewed degree distributions (e.g., Bala and Goyal, 2000). These same patterns of link formation can also arise for decidedly *non-strategic* reasons due to homophily and/or degree heterogeneity. The dyadic logit model with agent-specific ego (sender) and alter (receiver) effects therefore imposes a more stringent test for the no strategic interaction null. A further payoff is that our test can do double duty as an omnibus specification test for a widely-used method of network data modeling.

Because δ – the parameter indexing the dyadic logit null model – may range freely across its parameter space when $\gamma = 0$ our null hypothesis is a composite one. Test size equals the supremum of the rejection rate across all data generating processes (DGPs) with $\gamma = 0$. Because δ is high dimensional, the null model space is very large and constructing a test with uniformly good size and power properties is non-trivial (cf., Moreira, 2009). An additional non-standard feature of our testing problem is that the nuisance equilibrium selection function is only present under the alternative (cf., Andrews and Ploberger, 1994).

Under the logistic assumption on the random component of link utility, using a classic exponential family conditioning argument, we introduce a family of similar tests. Similarity means that the size of our test equals α *regardless of where we are in the null model space*. Put differently, our test is correctly-sized under null models with Λ values that result in "very little" homophily as well as under null models with Λ values that result in "lots of" homophily. Similarity also means that the size of our test does not vary with the configuration of the in- and out-degree heterogeneity vectors, $\mathbf{A}$ and $\mathbf{B}$.

We provide an *exact* characterization of the null distributions of the test statistics in this family and, crucially, a feasible Markov Chain Monte Carlo (MCMC) algorithm for simulating from them. Simulating the null distribution requires drawing a binary adjacency matrix uniformly at random from the set of all adjacency matrices satisfying certain constraints. Constrained binary matrix simulation has numerous applications in biology, psychology, ecology and other fields (cf., Sinclair, 1993; Blitzstein and Diaconis, 2011). Unfortunately, extant simulation algorithms cannot be used to simulate the null distribution needed here; our algorithm is therefore novel and of independent interest.

We also derive the form of the locally best test under the alternative $H_1 : \gamma > 0$. Remarkably we are able to do this while remaining agnostic about equilibrium selection. Finally, because our test is exact, we also side-step difficult issues that arise when undertaking asymptotic analysis in the single network context (see Graham (2020) for references and discussion).

Possible applications of the methods introduced in this paper include:

1. **Assessing model adequacy or specification testing:** The researcher believes the baseline null model is adequate for the setting at hand, but wishes to report an omnibus goodness-of-fit test (similar to the practice of reporting the Sargan-Hansen J-Statistic in the context of GMM estimation). While a rejection in this setting is interpreted as evidence against the baseline null model, it not interpreted as evidence in favor of any particular alternative. Our use of "classic" sufficiency arguments separates the the information in the data relevant for estimation of δ – the model parameter under the null – from that relevant for assessing model adequacy (cf., Barndorff-Nielsen and Cox, 1994, p. 29). As is well-known, it is not possible to construct a test with good power in all directions of mis-specification (Lehmann and Romano, 2005, Theorem 14.6.1). The researcher's choice of test statistic should therefore, at least heuristically, reflect those directions of mis-specification of most concern. A failure to reject in this setting suggests that the dyadic logistic regression model is correctly specified.[8]

2. **Detecting strategic interaction of a specific form:** The researcher's primary interest is in the specified model and she wishes to sign identify γ. In this example the analysts undertakes empirical work under the maintained assumption that *the true model is either in the null model space or in the specified alternative model space.* The data are used to determine which case prevails. This knowledge is actionable. For example, knowledge that $\gamma > 0$ may be sufficient to justify policies which subsidize link formation.

3. **Cataloging "unusual" network features:** The researcher wishes to assess whether certain features of the network in hand are "unusual". In contrast to the first case, here the researcher suspects that the network in hand is not well-described by the baseline one, but, in contrast to the second case, she remains somewhat agnostic about the form of the true model. The null model defines a set of reference networks with certain

[8]Dyadic regression analysis has a long history in economics going back, at least, to the work of Tinbergen (1962). See Graham (2020) for a survey and references. We note that this use case has the potential to introduce pre-testing bias if researchers only report their results conditional on accepting the null.

properties identical to those in the network of interest (e.g., the in- and out- degree sequences, numbers of links between agents with different covariate configurations). The researcher can compare features of interest in their network (e.g., diameter, reciprocity, support) with their distributions across the null reference set to assess whether their network is, indeed, "unusual". There is a long history, as noted earlier, of comparing network statistics to their expected value under an Ërdos and Rényi null. Here we provide a more realistic reference null distribution. See Holland and Leinhardt (1976) for a discussion of this type of analysis in sociology, Section 5 of Jackson et al. (2012) for an example from economics; Milo et al. (2002) for an example from computational biology, and Gotelli (2000) for a discussion of applications to species co-occurrence analysis in ecology. Researchers undertaking this last type of analysis might be best described as doing structured data exploration.

While our focus is on strategic network formation, it seems likely that the ideas developed below could be adapted to design tests appropriate for other incomplete econometric models. In recent work Chen et al. (2018) and Kaido and Zhang (2019) introduced likelihood ratio type tests applicable to incomplete models. Our test, in contrast, is a conditional score test. Conditioning, while requiring exponential family structure, is helpful in settings with a high dimensional nuisance parameter (cf., Moreira, 2009). Our score-based approach may also have computational advantages in settings where likelihood evaluation under the alternative is difficult (e.g., when enumeration of all NE is impossible).

Outline of the paper

Section 1 introduces our model of strategic network formation. We begin by defining agent preferences and characterizing equilibrium networks. Section 2 outlines our approach to testing. We first characterize the exact distribution of any statistic of the adjacency matrix under the null. Next we derive the form of the locally best test statistic for specific alternatives. Although we characterize the exact null distribution of our test statistics, for reasons of practically, we approximate this distribution by simulation. Section 3 outlines our new Markov Chain Monte Carlo (MCMC) algorithm for generating random draws from the required null distribution. Section 4 illustrates our methods in the context of the Nyakatoke risk-sharing network studied by De Weerdt (2004) and others. Section 5 finishes with a short discussion of limitations of our methods as well as a few thoughts on possible areas for additional research.

Proofs as well as some Monte Carlo simulation results are collected in a Supplemental Web Appendix. This appendix also includes a discussion of some additional applications of our MCMC simulation algorithm.

Readers interested primarily in applications can read Section 1, the first part of Section 2, and the empirical illustration of Section 4. The balance of the paper can be read later (perhaps after viewing the Python Jupyter Notebook available in the supplemental materials).

1 An family of empirical models of strategic network formation

1.1 Notation and setup

A directed graph $G(\mathcal{V},\mathcal{A})$ consists of a set of vertices (agents) $\mathcal{V}=\{1,\ldots,N\}$ and a set of ordered pairs of nodes, respectively called *egos* and *alters*, $\mathcal{A}=\{(i,j),(k,l),\ldots\}$ for $i\neq j$, $k\neq l$, and $i,j,k,l\in\mathcal{N}$. The elements of $\mathcal{A}$ correspond to those arcs, or directed links, present in $G(\mathcal{V},\mathcal{A})$.

In what follows we typically work with the adjacency matrix $\mathbf{D}=[D_{ij}]$ where

$$D_{ij}=\begin{cases}1 & \text{if } ij\in\mathcal{A}\\ 0 & \text{otherwise}\end{cases}. \tag{1}$$

Since we rule out self-links, the diagonal of $\mathbf{D}$ consists of structural zeros.

Let $G-ij$ denote the network obtained by deleting link ij from G (if present), and $G+ij$ the network one gets after adding this link (if absent). Let $\mathbf{D}\pm ij$ denote the adjacency matrix associated with the network obtained by adding/deleting link ij from G.

The set of all $2^{N(N-1)}$ possible adjacency matrices on N labeled vertices is denoted by $\mathbb{D}_N$. Hence $\mathbf{d}\in\mathbb{D}_N$ is a feasible network wiring or, equivalently, a game outcome. Let $\mathbf{d}_i$ be the i^{th} row of $\mathbf{d}$, or a pure strategy selection for agent i (i.e., a binary vector indicating which edges she chooses to direct). A pure strategy profile for all players other than i is denoted by $\mathbf{d}_{-i}$. We will sometimes refer to "players other than i" as i's *peers*.

For each agent there are $M\stackrel{def}{\equiv}2^{N-1}$ possible actions, corresponding to all possible configurations of links she may direct towards her peers. A mixed strategy for agent i, $\sigma_i=(\pi_{1i},\pi_{2i},\ldots,\pi_{Mi})'$, is probability distribution on these M possible pure strategy selections; $\sigma=(\sigma_1,\sigma_2,\ldots,\sigma_N)'$ is a mixed strategy profile for all N agents, while σ_{-i} is the strategy profile of agent i's peers.

1.2 Payoff function

The utility or payoff agent i gets from network $\mathbf{d}$ is

$$\nu_i\left(\mathbf{d}_i, \mathbf{d}_{-i}; \theta, \mathbf{U}_i\right) = \underbrace{\gamma_0 g_i\left(\mathbf{d}\right)}_{\text{Network Benefit}} - \underbrace{\sum_j d_{ij} c_{ij}\left(X_i, X_j; \delta, U_{ij}\right)}_{\text{Link "Costs" (i.e., private utility)}} \tag{2}$$

with $g_i\left(\mathbf{d}\right)$ a known, but not necessarily closed-form, function of the network adjacency matrix, normalized such that $g_i\left(\mathbf{0}\right) = 0$, $\theta = \left(\gamma, \delta'\right)'$, and the link "costs" function taking the form

$$c_{ij}\left(X_i, X_j; \delta, U_{ij}\right) = -\left[A_i + B_j + X_i' \Lambda_0 X_j - U_{ij}\right] \tag{3}$$

where X_i is a $K \times 1$ vector of mutually exclusive group membership indicators that *is observed* by the econometrician and $\mathbf{U}_i = \left(U_{i1}, \ldots, U_{ii-1}, U_{ii+1}, \ldots, U_{iN}\right)'$ is agent i's vector of idiosyncratic logistic preference shocks over the $N-1$ possible links she can direct (and $\mathbf{U} = \left(\mathbf{U}_1', \ldots, \mathbf{U}_N'\right)'$).[9] All agents observe their own, as well as their peers', preference shock vectors. As is standard in game theory (e.g., Fudenberg and Tirole, 1998), we use, in a small abuse of notation, $\nu_i\left(\sigma_i, \sigma_{-i}; \theta, \mathbf{U}_i\right)$ to denote agent i's *expected* utility under the mixed strategy profile $\sigma = \left(\sigma_i, \sigma_{-i}\right)$.

The first term in (2) captures how agent i's utility varies with the entire structure of the network; this may include benefits from direct, as well as indirect connections. The second term in (2) captures the purely private benefits to i associated with directing an arc to j. That is, the component of utility associated with arc ij that *does not vary* with the presence or absence of links elsewhere in the network.

In theoretical work $g_i\left(\mathbf{d}\right)$ is often called the *network benefit* function, while $c_{ij}\left(X_i, X_j; \delta, U_{ij}\right)$ would be typically associated with the cost of forming edge ij (e.g., Jackson, 2008; Goyal, 2023). These costs are generally assumed constant in theory research, while – as is appropriate given the empirical context – they are heterogeneous across agents and links here.[10] The placement of a negative sign in front of the second term in (2) is in keeping with the "costs" nomenclature of the theory literature, but is without loss of generality: $-c_{ij}\left(X_i, X_j; \delta, U_{ij}\right)$ is simply the portion of the payoff i gets from directing an edge to j that is invariant to all other linking decisions.

While the benefit-cost nomenclature is useful for developing intuitions about the form of NE in this setting, to reiterate, what is essential here is that the first term may vary

[9] More generally X_i enumerates the support points of a collection of (observed) discrete agent-specific regressors (or a partition of this support into K regions).

[10] Johnson and Gilles (2000) study the implications of cost heterogeneity on equilibrium network structure in the "connections" model.

arbitrarily with $\mathbf{d}$, and hence with peer actions, while the second term is invariant to peers' actions and, furthermore, additively separable in own actions. In what follows we call the (negative of the) j^{th} summand in the second part of (2) the *baseline utility* that i gets from directing edge ij.

Baseline utility

Considering baseline utility first, we see it is increasing in the heterogeneity terms, assumed unobserved by the econometrician, A_i and B_j. Agents with high values of *out-degree heterogeneity* A_i get a large amount of baseline utility from any link they send. In a social network context high A_i agents are "extroverts". Agents with high *in-degree heterogeneity* B_j, in contrast, are especially attractive targets, or alters, for links sent by others. In a social network high B_j agents are "prestigious" or "popular".[11]

The $X_i'\Lambda_0 X_j \overset{def}{\equiv} W_{ij}'\lambda_0$ term allows baseline utility to depend on whether agents assortatively match on their attributes (we define $W_{ij} \overset{def}{\equiv} (X_i \otimes X_j)$ and recall that $\lambda \overset{def}{\equiv} \mathrm{vec}\,(\Lambda')$). The elements of the $K \times K$ matrix $\Lambda = [\lambda_{kl}]$ parameterize the systematic utility generated by links, say, from group k to group l. For example, in a social network girls might, all things equal, prefer other girls as friends. The Λ_0 matrix parameterizes homophily (or heterophily) of this type.

In our fixed-N setting $\{A_i\}_{i=1}^N$ and $\{B_i\}_{i=1}^N$ are fixed-dimensional parameters. We further treat $\{X_i\}_{i=1}^N$ as non-stochastic in what follows. Note that, as in fixed-effect panel data analysis, $\{(A_i, B_i)'\}_{i=1}^N$ and $\{X_i\}_{i=1}^N$ may freely correlate.[12]

The final component of baseline utility is idiosyncratic; we assume that the $\{U_{ij}\}_{i\neq j}$ are independent and identically distributed (iid) logistic random variables. The logistic assumption generates exponential family structure which we exploit when forming our test.

Equation (2) with $\gamma = 0$ gives agent preferences under our baseline or null model (essentially the dyadic link formation model studied by Charbonneau (2017)). This model, when fitted by maximum likelihood, can successfully match many features of real world networks. Specifically arbitrary in- and out-degree sequences and assortative linking patterns on discrete agent attributes (cf., Graham, 2020). The model cannot accommodate homophilous sorting on latent attributes (a limitation which may affect the interpretation of a test rejection). It also maintains a logistic assumption on U_{ij}, a restriction relaxed by Gao (2020).

While (2) with $\gamma = 0$ underpins a growing empirical literature on networks, we are

[11] Alternatively we can think of high A_i agents as being able to direct links at low cost, and high B_j agents as being low cost alters.

[12] An implication is that $\{(A_i, B_i, X_i')'\}_{i=1}^N$ need not be i.i.d. There is no requirement, for example, that the agents in the network are a random sample from some population.

Figure 1: Network benefit function examples

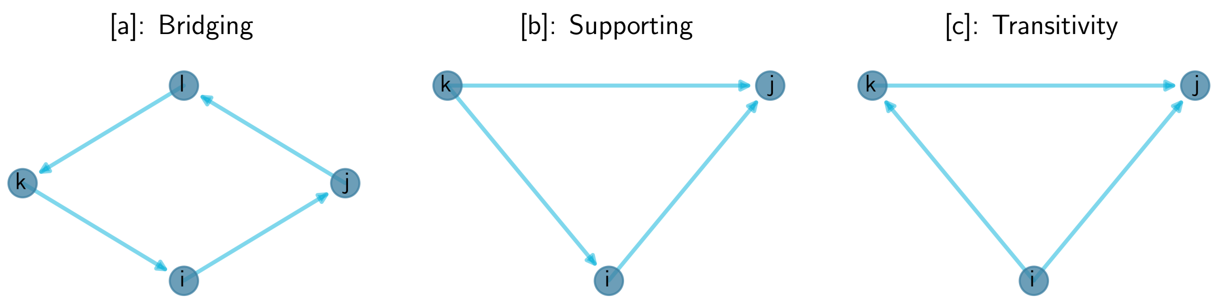

Source: Authors' calculations.
Notes: Panel [a]: agent i is a bridge from k *to* j and agent l is a bridge from j to k. Panel [b]: edge ij is supported by agent k. Panel [c]: adding edge ij generates a transitive triad.

especially interested in settings where this model does not provide a good description of the network in hand.

Network benefit function

When $\gamma > 0$, the first term in (2) – the *network benefit* function $g_i(\mathbf{d})$ – enriches the baseline model to allow agent preferences over links to vary with the presence or absence of links elsewhere in the network. The researcher is free to specify the network benefit function as desired. A few selected examples, drawn from recent theoretical work on strategic network formation, gives a sense of the range of possibilities.

Example 1.1. (Connections) In a seminal paper, Jackson and Wolinsky (1996), introduced the *connections model*. In a directed variant, Bala and Goyal (2000) set $g_i(\mathbf{d}) = \sum_{i \neq j} \phi\left(\ell_{ij}\left(\tilde{\mathbf{d}}\right)\right)$ where $\tilde{\mathbf{d}}$ is the undirected network obtained from $\mathbf{d}$ (i.e., $\tilde{\mathbf{d}} = \left[\tilde{d}_{ij}\right]$ with $\tilde{d}_{ij} = 1 - (1 - d_{ij})(1 - d_{ji})$), $\phi : \{1, 2, \ldots, N-1\} \to \mathbb{R}$ is a known function with $\phi(k) > \phi(k+1) > 0$ for any $k = 1, 2, \ldots, N-1$, and $\ell_{ij}\left(\tilde{\mathbf{d}}\right)$ the shortest path length between agents i and j in $\tilde{\mathbf{d}}$. Agents prefer to be close to other agents in the network in order to easily access their information, but also wish to maintain as few links as possible, since links are costly to direct. Strong externalities arise in this model: edge ij may incidentally reduce the shortest path length between agents k and l, but such benefits are not internalized by agent i. Also, since information flows bidirectionally, both agents i and j benefit from edge ij, while the cost is shouldered by i alone.

Example 1.2. (Structural Hole / Bridging) Kleinberg et al. (2008) introduce a model of network formation inspired by Burt's (1995) theory of "structural holes". Burt (1995) argued that individuals that connect disparate groups within a network gain "bridging", "middle-person" or inter-mediation benefits. Such benefits arise from lying on a (shortest) path connecting two agents not directly connected themselves. Citing empirical evidence, Kleinberg et al. (2008) emphasize the special benefits of lying on length two paths between disconnected agents.[13] If $d_{ki}d_{ij}(1-d_{kj})=1$, then i serves as a "bridge" between k and j (see Panel [a] of Figure 1.2). The summation $\sum_l d_{kl}d_{lj}(1-d_{kj})$ yields a count of the total number of bridging agents between k and j. While agents benefit from serving as a bridge between two agents, these benefits decline in the number of other agents also serving as bridges for the same (directed) dyad. This yields a network payoff function of the form $g_i(\mathbf{d}) = \sum_j \sum_{k\neq j} \phi\left(d_{ki}d_{ij}(1-d_{kj}), \sum_l d_{kl}d_{lj}(1-d_{kj})\right)$ with $\phi(0,k) \equiv 0$ and $\phi(1,k) > \phi(1,k+1) > 0$ for $k = 1,\ldots,N-2$. See Goyal and Vega-Redondo (2007) for a related model.[14]

Example 1.3. (Supported Links, Transitivity, Reciprocity) Jackson et al. (2012) introduce a model where agents value supported links. Edge ij is supported by agent k if $d_{ij}d_{ki}d_{kj}=1$ (see Panel [b] of Figure 1.2). This configuration allows agent k to monitor, or referee, relationship ij, making it more valuable. This suggests a network benefits function of $g_i(\mathbf{d}) = \sum_j d_{ij}\left(\sum_k d_{ki}d_{kj}\right)$. If, instead, agents value reciprocity we would set $g_i(\mathbf{d}) = \sum_j d_{ij}d_{ji}$; while if they value transitivity in links (see Panel [c] of Figure 1.2) we would set $g_i(\mathbf{d}) = \sum_j d_{ij}\left(\sum_k d_{ik}d_{kj}\right)$.

Marginal utility

Let, in an abuse of notation, $\nu_i(\mathbf{d}) \equiv \nu_i(\mathbf{d}_i, \mathbf{d}_{-i}; \theta, \mathbf{U}_i)$; the marginal utility of arc ij for agent i equals

$$MU_{ij}(\mathbf{d}) = \begin{cases} \nu_i(\mathbf{d}) - \nu_i(\mathbf{d}-ij) & \text{if } d_{ij}=1 \\ \nu_i(\mathbf{d}+ij) - \nu_i(\mathbf{d}) & \text{if } d_{ij}=0 \end{cases} \tag{4}$$

Marginal utility measures the utility gain (loss) to agent i from adding (subtracting) link ij holding the structure of all other links in the network constant (including any other links agent i directs). The component of marginal utility associated with the network benefit function $g_i(\mathbf{d})$ plays an important role in our analysis. Define the marginal network payoff

[13] "[T]here appears to be much less measurable benefit to u if it is the internal node on a path between two nodes at graph distance greater than two" (Kleinberg et al., 2008, p. 285).

[14] We could, inspired by Freeman (1977), also consider the model where agents directly value their network betweenness centrality such that $g_i(\mathbf{d}) = \frac{1}{(N-1)(N-2)} \sum_{j,k\in\mathcal{N}\setminus\{i\}} \frac{\text{\# of shortest paths from agents } j \text{ to } k \text{ which pass through } i}{\text{\# of shortest paths from agents } j \text{ to } k}$.

associated with agent i directing a link to j as

$$s_{ij}(\mathbf{d}) = \begin{cases} g_i(\mathbf{d}) - g_i(\mathbf{d} - ij) & \text{if } d_{ij} = 1 \\ g_i(\mathbf{d} + ij) - g_i(\mathbf{d}) & \text{if } d_{ij} = 0 \end{cases} \tag{5}$$

Using (2) and definition (5) yields a marginal utility for arc ij of

$$MU_{ij}(\mathbf{d}) = A_i + B_j + W_{ij}'\lambda_0 + \gamma_0 s_{ij}(\mathbf{d}) - U_{ij}. \tag{6}$$

As it features in the computation of the optimal test statistic introduced below, it is helpful to derive the form of $s_{ij}(\mathbf{d})$ for the example network benefit functions introduced earlier.

Example 1.1. (Connections) In the connections model, when i directs a link to j she weakly reduces her shortest path length to all other agents in the network. In this model $s_{ij}(\mathbf{d}) \geq 0$ for all $\mathbf{d} \in \mathbb{D}_N$. While there is no closed form expression for $s_{ij}(\mathbf{d})$ in the connections model, it is straightforward to compute shortest path lengths between agents numerically (many network manipulation software libraries include routines to do this). If removing (adding) arc ij increases (decreases) i's distance to many other agents in the network, then $s_{ij}(\mathbf{d})$ will be large.

Example 1.2. (Structural Hole / Bridging) For the bridging network benefit function $s_{ij}(\mathbf{d})$ equals

$$s_{ij}(\mathbf{d}) = \sum_{k \neq j} \phi\left(d_{ki}(1 - d_{kj}), 1 + \sum_{l \neq i} d_{kl} d_{lj}(1 - d_{kj})\right).$$

The marginal utility of edge ij is therefore increasing in the number of agents k which direct edges to i, but not to j. It is decreasing in the number of agents l and k in which edges kl and lj are present (but edge kj is not).

Example 1.3. (Supported Links, Transitivity, Reciprocity) In the support model $s_{ij}(\mathbf{d}) = \sum_k d_{ki} d_{kj}$, which is simply a count of how many agents would support edge ij if it were formed. When agents have a taste for transitivity we have instead

$$s_{ij}(\mathbf{d}) = \sum_k d_{ik} d_{kj} + \sum_{k \neq j} d_{ik} d_{jk}$$

which is a count of how many transitive triads (involving agent i) would be created if edge ij is added. Finally if agents have a taste for reciprocity we have $s_{ij}(\mathbf{d}) = d_{ji}$; indicating that the marginal utility of edge ij varies with the presence or absence of the reciprocal edge ji.

1.3 Equilibrium networks

We assume that the observed network $\mathbf{D}$ coincides with the equilibrium outcome of an N-player complete information game. Each agent (i) observes $\{(A_i, B_i, X_i')\}_{i=1}^N$ and $\{U_{ij}\}_{i\neq j}$ and then (ii) decides which, out of the $N-1$ other agents, to send links to. Agents may play mixed strategies.

A mixed strategy profile σ^* is a NE when $\theta = \theta_0$ and $\mathbf{U} = \mathbf{u}$, if for all $i = 1, \ldots, N$,

$$\nu_i\left(\sigma_i^*, \sigma_{-i}^*; \theta_0, \mathbf{u}_i\right) \geq \nu_i\left(\mathbf{d}_i, \sigma_{-i}^*; \theta_0, \mathbf{u}_i\right) \tag{7}$$

for all possible pure strategy selections $\mathbf{d}_i$. We assume that the *observed* network $\mathbf{D}$ is either a pure strategy NE or in the support of a mixed strategy NE.[15]

Assumption 1.1. (DATA GENERATING PROCESS)

1. (NON-STOCHASTIC $\mathbf{X}$) Let $\mathcal{V} = \{1, \ldots, N\}$ be the N agents in the network in hand, each with a *fixed* (i.e., non-stochastic) group membership of X_i.

2. (LOGISTIC PREFERENCE SHOCKS) Let $\mathbf{U} = [U_{ij}]_{i\neq j}$ be an $N(N-1)$ vector of i.i.d. logistic link preference shocks observed by all agents.

3. (NASH EQUILIBRIUM) Let $\theta_0 \in \Theta$ be the parameter indexing the payoff function (2). The observed network $\mathbf{D}$ is either a pure strategy NE or contained in the support of a mixed strategy NE of the strategic form game $\left(\mathcal{V}, \mathbb{D}_N, \{\nu_i(\cdot, \cdot; \theta_0, \mathbf{U}_i)\}_{i\in\mathcal{V}}\right)$.

Treating $\mathbf{X}$ as non-stochastic simplifies both notation (allowing us to suppress, for example, the dependence of payoffs on $\mathbf{X}$) and analysis (see the proof to Theorem 1.1 below). It is also without loss of generality. In our setting there is no asymptotic thought experiment and the fixed N sampling distribution we consider conditions on $\mathbf{X}$ throughout. Randomness across hypothetical replications of the network formation game are to due soley to variation in $\mathbf{U}$ and any randomness in NE selection (see below).

[15] Observe that agent i must consider 2^{N-1} different pure strategy deviations in order to verify that their chosen strategy is optimal. This may be unrealistic when N is large. A weaker equilibrium requirement, akin to the notion of pairwise stability introduced by Jackson and Wolinsky (1996) for undirected networks, is to require agents to only consider the effects of adding or deleting a *single* link at time on their utility.

Under this weaker stability notion, which we call *single deviation stable* (SDS), we only require that the marginal utility of any link present in the network is non-negative, while that of any link not present is negative. This implies that the observed network $\mathbf{D}$ satisfies the system of $N(N-1)$ non-linear equations

$$D_{ij} = \mathbf{1}\left(A_i + B_j + W_{ij}'\lambda_0 + \gamma_0 s_{ij}(\mathbf{D}) \geq U_{ij}\right)$$

for $i, j = 1, \ldots, N$ and $j \neq i$.While we maintain the NE assumption in what follows, it turns out that our test is also valid if, instead, the observed network is only SDS. Although single deviation stability is a natural directed analog of pairwise stability, we are not aware of this equilibrium concept being considered before.

1.4 Likelihood

In the presence of multiple NE, Assumption 1.1 imposes no restrictions on which one is actually realized in the observed network. Our strategic network formation model is *incomplete*. Although we remain agnostic about equilibrium selection, it is nevertheless useful to develop a notation for, and establish some properties of, the unknown equilibrium selection rule. This allows us to write down a (well-defined) likelihood for the network, albeit abstractly.

Let $\mathcal{N}(\mathbf{d}, \mathbf{u}; \theta)$ be a function which assigns, for $\mathbf{U} = \mathbf{u}$, a probability weight to network $\mathbf{d}$:

$$\mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) : \mathbb{D}_N \times \mathbb{R}^n \to [0, 1] \tag{8}$$

In order for $\mathcal{N}(\mathbf{d}, \cdot; \theta)$ to be a valid NE selection function it must satisfy the conditions of Definition 1.1.

Definition 1.1. (Equilibrium Selection Function) For $\mathbf{U} = \mathbf{u}$ the realized vector of logistic link preference shocks and θ_0 the payoff function parameter, let $\mathbf{d}^* (\mathbf{u}; \theta_0)$ be a pure strategy NE or a network contained in the support of a mixed strategy NE and $\mathbb{D}_N^* (\mathbf{u}; \theta_0)$ be the set of all such networks. Function (8) is such that (i) $\mathcal{N} (\mathbf{d}, \mathbf{u}; \theta_0) \geq 0$ for all $\mathbf{d} \in \mathbb{D}_N^* (\mathbf{u}; \theta_0)$ (ii) $\sum_{\mathbf{d} \in \mathbb{D}_N^* (\mathbf{u}; \theta_0)} \mathcal{N} (\mathbf{d}, \mathbf{u}; \theta_0) = 1$ and (iii) $\mathcal{N} (\mathbf{d}, \mathbf{u}; \theta_0) = 0$ for all $\mathbf{d} \in \mathbb{D}_N \backslash \mathbb{D}_N^* (\mathbf{u}; \theta_0)$.

If $\mathcal{N}(\mathbf{d}, \cdot; \theta)$ satisfies the conditions of Definition 1.1, then the likelihood of observing network $\mathbf{D} = \mathbf{d}$ is

$$P(\mathbf{d}; \theta, \mathcal{N}) = \int_{\mathbf{u} \in \mathbb{R}^n} \mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) f_{\mathbf{u}}(\mathbf{u}) \mathrm{d}\mathbf{u}, \tag{9}$$

where $f_{\mathbf{u}}(\mathbf{u}) = \prod_{i \neq j} f_U(u_{ij})$ with $f_U(u) = e^u / [1 + e^u]^2$. Of course, for the likelihood (9) to be well-defined we require that $\mathcal{N}(\mathbf{d}, \cdot; \theta)$ is measurable.

Theorem 1.1. (Likelihood) *For any network* $\mathbf{d} \in \mathbb{D}_N$ *there exists a measurable function* $\mathcal{N}(\mathbf{d}, \cdot; \theta) : \mathbb{R}^n \to [0, 1]$, *which assigns to* $\mathbf{u} \in \mathbb{R}^n$ *a* NE *weight on the pure strategy combination corresponding to* $\mathbf{d}$.

The proof of Theorem 1.1 can be found in Appendix A.1. It follows directly from ideas in Beresteanu et al. (2011).

2 Testing for strategic interaction

In this section we introduce a model adequacy test for the dyadic logit null (the baseline model). We then show how to optimize the power of this test in certain directions of the alternative model space.

We begin by describing how to assess the adequacy of the baseline model as a description of the network in hand. Utilizing a conditioning argument we construct an *exact* test of the null of "correct specification". An alternative model is not explicitly formulated in this case, although researcher intuitions about plausible directions of mis-specification typically guides the choice of test statistic. As shown by Lehmann and Romano (2005), it is impossible to construct a test with power (greater than size) in all possible directions of mis-specification.

Next we consider applications where the analyst carefully specifies the alternative model (through an explicit choice of the network benefit function, $g_i(\mathbf{d})$). Here the researcher believes the true network formation model lies in either the null or the (specified) alternative model space; the purpose of testing is to determine which situation prevails. In this second application we seek to construct a test which rejects with high probability when the alternative is true, while continuing to control size under the null.

The mechanics of testing in both cases are the same, the difference lies solely in the choice of test statistic. This allows us to build up our results in a cumulative fashion. In sub-section 2.1 we show how the exponential family structure of the model under the null of no strategic interaction allows us to construct similar tests. Specifically we control size exactly by conditioning on the minimally sufficient statistic for the null model parameter, δ. With these basics established we show in sub-section 2.2 how to optimize power in the direction of a particular alternative. This step is non-trivial since under the alternative the network formation game may exhibit multiple equilibria.

Throughout, and crucially, we wish to remain agnostic about the distribution of any degree heterogeneity across agents as well as the form of any homophily and/or heterophily. Let Δ denote a subset of the K^2+2N dimensional Euclidean space in which $\delta_0 = (\lambda_0, \mathbf{A}_0, \mathbf{B}_0)$ is, a priori, known to lie. For technical reasons we assume that Δ contains a K^2+2N dimensional rectangle. The null model parameter space is

$$\Theta_0 = \{(\gamma, \delta') : \gamma = 0, \delta \in \Delta\}. \tag{10}$$

Our null hypothesis is the *composite* one:

$$H_0 : \theta \in \Theta_0 \tag{11}$$

since δ may range freely over $\Delta \subset \mathbb{R}^{K^2+2N}$ under the null.

Under the null the likelihood is $P_0(\mathbf{d};\delta) \stackrel{def}{\equiv} P(\mathbf{d};(0,\delta')',\mathcal{N}_0)$ with

$$\mathcal{N}_0(\mathbf{d},\mathbf{u};\theta) = \prod_i \prod_j \mathbf{1}\left(A_i + B_j + W_{ij}'\lambda \geq u_{ij}\right)^{d_{ij}} \times \mathbf{1}\left(A_i + B_j + W_{ij}'\lambda < u_{ij}\right)^{1-d_{ij}}.$$

Under the null the unique "equilibrium" network is the one where all links with positive marginal utility are present and those with negative marginal utility are not. Uniqueness follows from the fact that agents' best responses are constant in the strategy profiles of their peers when $\gamma = 0$. Consequently the marginal utility of edge ij is invariant to the presence or absence of links elsewhere in the network.

Evaluating the integral (9) under the null yields

$$P_0(\mathbf{d};\delta) = \prod_{i=1}^{N} \prod_{j\neq i} \left[\frac{\exp\left(W_{ij}'\lambda + R_i'\mathbf{A} + R_j'\mathbf{B}\right)}{1+\exp\left(W_{ij}'\lambda + R_i'\mathbf{A} + R_j'\mathbf{B}\right)}\right]^{d_{ij}} \times \left[\frac{1}{1+\exp\left(W_{ij}'\lambda + R_i'\mathbf{A} + R_j'\mathbf{B}\right)}\right]^{1-d_{ij}}, \tag{12}$$

where R_i is the $N \times 1$ vector with a 1 in its i^{th} element and zeros elsewhere.[16]

2.1 Testing without explicit specification of the alternative (i.e., baseline model adequacy analysis)

Under the null our likelihood, $P_0(\mathbf{d};\delta)$, is a member of the exponential family. To see this it is helpful to establish some additional notation. The *out-* and *in-degree sequences* equal:

$$\mathbf{S} = \begin{pmatrix} \mathbf{S}_{\text{out}} \\ \mathbf{S}_{\text{in}} \end{pmatrix} = \begin{pmatrix} D_{1+},\ldots,D_{N+} \\ D_{+1},\ldots,D_{+N} \end{pmatrix}. \tag{13}$$

Here $D_{+i} = \sum_j D_{ji}$ and $D_{i+} = \sum_j D_{ij}$ equal the in- and out-degree of agents $i = 1,\ldots,N$.

The $K \times K$ *cross-link matrix* equals

$$\mathbf{M} = \sum_i \sum_j D_{ij} X_i X_j'. \tag{14}$$

This matrix summarizes the inter-group link structure in the network (homophily). The kl^{th}

[16] Variants of this likelihood are analyzed by Chatterjee et al. (2011), Charbonneau (2017), Graham (2017), Jochmans (2018), Dzemski (2018) and Yan et al. (2018). See also Fernández-Val and Weidner (2016).

element of $\mathbf{M}$ records the number of links sent by type k agents to type l agents.

Let $\mathbf{S}, \mathbf{M}$ be a degree sequence and cross-link matrix. We say $\mathbf{S}, \mathbf{M}$ is *graphical* if there exists at least one arc set $\mathcal{A}$ such that $G(\mathcal{V}, \mathcal{A})$ is a simple directed graph with degree sequence $\mathbf{S}$ and cross link matrix $\mathbf{M}$. We call any such network a *realization* of $\mathbf{S}, \mathbf{M}$. The set of all possible realizations of $\mathbf{S}, \mathbf{M}$ is denoted by $\mathbb{G}_{\mathbf{S},\mathbf{M}}$. $\mathbb{D}_{\mathbf{S},\mathbf{M}}$ denotes the associated set of possible adjacency matrices:

$$\mathbb{D}_{\mathbf{S},\mathbf{M}} \overset{def}{\equiv} \Big\{\mathbf{d} \in \mathbb{D}_N \ : \ (d_{1+}, \dots, d_{N+}) = \mathbf{S}_{\text{out}}, \quad (d_{+1}, \dots, d_{+N}) = \mathbf{S}_{\text{in}}, \ \sum_i \sum_j d_{ij} X_i X_j' = \mathbf{M} \Big\}. \tag{15}$$

Let $\mathbf{T} = \left(\text{vec}\left(\mathbf{M}'\right)', \mathbf{S}_{\text{out}}, \mathbf{S}_{\text{in}}\right)'$. Note that associated with any graphical realization of $\mathbf{T}$ is a corresponding set of adjacency matrices $\mathbb{D}_{\mathbf{S},\mathbf{M}}$.

With this notation established it is easy to verify that the family of network formation models under the null of no strategic interaction constitutes an exponential family.

Lemma 2.1. *(i) $\mathcal{P}_{0,\delta} = \{P_0(\mathbf{d}; \delta) \ : \ \delta \in \Delta\}$ is an exponential family with $\mathbf{t}$ a boundedly complete sufficient statistic for δ,*[17] *(ii) its probability mass function (12) can be rewritten as:*

$$P_0(\mathbf{d}; \delta) = c(\delta) \exp\left(\mathbf{t}'\delta\right), \ \delta \in \Delta, \ \mathbf{d} \in \mathbb{D}_N, \tag{16}$$

with $c(\delta) \overset{def}{\equiv} \left[\sum_{d \in \mathbb{D}_N} \exp\left(\left[\sum_{i=1}^N \sum_{j \neq i} d_{ij} \left\{W_{ij}'\lambda + R_i'\mathbf{A} + R_j'\mathbf{B}\right\}\right]\right)\right]^{-1}$.

Proof. See Appendix A.2. □

The sufficient statistics for the $K^2 + N + N$ elements of the nuisance parameter δ, are (i) the cross link matrix, (ii) the out-degree sequence and (iii) the in-degree sequence.

Under H_0 the conditional likelihood of the event $\mathbf{D} = \mathbf{d}$ is

$$P_0\left(\mathbf{d} \middle| \mathbf{T} = \mathbf{t}\right) = \frac{P_0\left(\mathbf{d}; \delta\right)}{\sum_{\mathbf{v} \in \mathbb{D}_{\mathbf{s},\mathbf{m}}} P_0\left(\mathbf{v}; \delta\right)} = \frac{1}{\left|\mathbb{D}_{\mathbf{s},\mathbf{m}}\right|} \tag{17}$$

if $\mathbf{d} \in \mathbb{D}_{\mathbf{s},\mathbf{m}}$ and zero otherwise. Under the null of no strategic interaction all networks with the same in- and out-degree sequences and cross link structure are equally likely. Importantly this conditional likelihood is invariant to the actual value of the nuisance parameter δ.

By conditioning on $\mathbf{T}$, which is sufficient for δ, we isolate the information in the data that is relevant for assessing model adequacy (Barndorff-Nielsen and Cox, 1994). This follows

[17] A sufficient statistic is (boundedly) complete if, for all (bounded) functions $k(\mathbf{t})$, $\mathbb{E}_0\left[k(\mathbf{T})\right] = 0$ for all $P_0(\mathbf{d}; \delta) \in \mathcal{P}_{0,\delta}$ implies that $k(\mathbf{T}) = 0$ almost everywhere. See Section 3.6 of Ferguson (1967).

because conditional on $\mathbf{T}$, the null model *completely* specifies the distribution of $\mathbf{D}$. Consequently, the distribution of any statistic of the adjacency matrix, say $R(\mathbf{D})$, is also fully specified. Specifically the null distribution of $R(\mathbf{D})$ is the one induced by a discrete uniform distribution on $\mathbb{D}_{\mathbf{S},\mathbf{M}}$:

$$\Pr\left(R(\mathbf{D}) \leq r \middle| \mathbf{T}; \theta \in \Theta_0\right) = \frac{1}{|\mathbb{D}_{\mathbf{S},\mathbf{M}}|} \sum_{\mathbf{d} \in \mathbb{D}_{\mathbf{S},\mathbf{M}}} \mathbf{1}\left(R(\mathbf{d}) \leq r\right). \tag{18}$$

To test model goodness-of-fit/adequacy, we simply check whether the value of $R(\mathbf{D})$ in the network in hand is at an extreme quantile of this distribution. If it is, we take this as evidence against the baseline (null) model.

Similarity and conditioning

A test with critical function $\phi(\mathbf{D})$ will have size α if its null rejection probability (NRP) is less than or equal to α for *all* values of the nuisance parameter:

$$\sup_{\theta \in \Theta_0} \mathbb{E}_\theta\left[\phi(\mathbf{D})\right] = \sup_{\gamma=\gamma_0, \delta \in \triangle} \mathbb{E}_\theta\left[\phi(\mathbf{D})\right] = \alpha. \tag{19}$$

Since the nuisance parameter δ is very high dimensional, size control is *a priori* non-trivial. For some intuition as to why consider, as an example, the case where $s_{ij}(\mathbf{d}) = \sum_k d_{ki} d_{kj}$, such that agents' have a taste for supported links when $\gamma_0 > 0$. A natural test statistic in this case would be some function of $\mathbf{D}$ that is increasing in the number of supported links in the network.[18] The researcher would then reject the null of $\gamma_0 = 0$ when this statistic is large enough. Unfortunately, the expected number of supported links varies dramatically under the null depending on the value of δ. Certain configurations of $\mathbf{A}$, $\mathbf{B}$ and/or λ may result in a network with substantial link clustering (and hence support) *even when agents' have no taste for support per se*. If we choose a fixed critical value for rejection then, depending on the values of $\mathbf{A}$, $\mathbf{B}$ and/or λ, size may be very poor.

To avoid any size distortion induced by variation in δ over $\Delta \subset \mathbb{R}^{K^2+2N}$ we exploit the exponential family structure of our model (under the null). Let $\mathbb{T} = \{\mathbf{t} : \mathbf{s}, \mathbf{m} \text{ is graphical}\}$ be the set of possible sufficient statistics $\mathbf{T}$. Instead of choosing a fixed critical value, which may result in under- or over-rejection, depending on the value of δ, we proceed conditionally on $\mathbf{T} \in \mathbb{T}$, varying our critical value with $\mathbf{T}$. In this way we ensure good size control. By conditioning on $\mathbf{T}$ we can also remain agnostic about its marginal distribution (and hence the value of δ for which $\mathbf{T}$ is sufficient).

[18] Jackson et al. (2012) suggest the fraction of links in the network which are supported.

Formally, for each $\mathbf{t} \in \mathbb{T}$ we form a test with the property that, for all $\theta \in \Theta_0$,

$$\mathbb{E}_\theta \left[\phi \left(\mathbf{D} \right) \middle| \, \mathbf{T} = \mathbf{t} \right] = \alpha. \tag{20}$$

Such an approach ensures *similarity* of our test since, by iterated expectations,

$$\mathbb{E}_\theta \left[\phi \left(\mathbf{D} \right) \right] = \mathbb{E}_\theta \left[\mathbb{E}_\theta \left[\phi \left(\mathbf{D} \right) \middle| \, \mathbf{T} \right] \right] = \alpha \tag{21}$$

for any $\theta \in \Theta_0$ (Ferguson, 1967). By proceeding conditionally we ensure that the NRP is unaffected by the value of δ.

For any $\mathbf{t} \in \mathbb{T}$ we can construct an *exact* test, as is required by (20), because our model completely specifies the distribution of networks conditional on $\mathbf{T} = \mathbf{t}$ under the null. Condition (21) follows immediately. Using some well-known results from the theory of exponential families, we can make the stronger claim that similarity is only possible by conditioning.

Lemma 2.2. (Similarity) *Any similar test of $H_0 : \theta \in \Theta_0$ satisfies (20) for all $\mathbf{t} \in \mathbb{T}$.*

Proof. By Lemma 2.1 above, $\mathbf{T}$ is a boundedly complete sufficient statistic for θ under the null. The claim then follows from Ferguson (1967, Definition 4 and Theorem 2, Section 5.4). □

Implementation

For concreteness let $R(\mathbf{D})$ be the network reciprocity index (Newman, 2010):

$$R(\mathbf{D}) = \frac{2\hat{P}_{11}}{2\hat{P}_{11} + \hat{P}_{01}}, \tag{22}$$

where

$$\hat{P}_{01} = \frac{2}{N\left(N-1\right)} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \left[D_{ij} \left(1 - D_{ji}\right) + \left(1 - D_{ij}\right) D_{ji} \right] \tag{23}$$

equals the fraction of dyads which take an unreciprocated or "asymmetric" configuration and

$$\hat{P}_{11} = \frac{2}{N\left(N-1\right)} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} D_{ij} D_{ji} \tag{24}$$

the fraction which take a reciprocated or "mutual" configuration.

A conditional test based upon $R(\mathbf{D})$ will have a critical function of

$$\phi(\mathbf{d}) = \begin{cases} 1 & R(\mathbf{d}) > c_\alpha(\mathbf{t}) \\ g_\alpha(\mathbf{t}) & R(\mathbf{d}) = c_\alpha(\mathbf{t}) \\ 0 & R(\mathbf{d}) < c_\alpha(\mathbf{t}) \end{cases} \tag{25}$$

where the values of $c_\alpha(\mathbf{t})$ and $g_\alpha(\mathbf{t}) \in [0,1]$ are chosen to satisfy the requirement that $\mathbb{E}_\theta[\phi(\mathbf{D})|\,\mathbf{T} = \mathbf{t}] = \alpha$. Specifically, given the sufficient statistic $\mathbf{T}$ we first compute:

$$c_\alpha(\mathbf{T}) = \min\left\{ c \in \mathbb{R} \; : \; \frac{1}{|\mathbb{D}_{\mathbf{S},\mathbf{M}}|} \sum_{\mathbf{d} \in \mathbb{D}_{\mathbf{S},\mathbf{M}}} \mathbf{1}(R(\mathbf{d}) > c) \leq \alpha \right\}. \tag{26}$$

Second we set $g_\alpha(\mathbf{T})$ according to

$$g_\alpha(\mathbf{T}) = \frac{\alpha - \frac{1}{|\mathbb{D}_{\mathbf{S},\mathbf{M}}|} \sum_{\mathbf{d} \in \mathbb{D}_{\mathbf{S},\mathbf{M}}} \mathbf{1}(R(\mathbf{d}) > c_\alpha(\mathbf{T}))}{\frac{1}{|\mathbb{D}_{\mathbf{S},\mathbf{M}}|} \sum_{\mathbf{d} \in \mathbb{D}_{\mathbf{S},\mathbf{M}}} \mathbf{1}(R(\mathbf{d}) = c_\alpha(\mathbf{T}))}.$$

Observe that, formally, the test is randomized. Because the adjacency matrix is a discrete random variable, with finite support, it will generally not be possible to find a $c \in \mathbb{R}$ such that $|\mathbb{D}_{\mathbf{S},\mathbf{M}}|^{-1} \sum_{\mathbf{d} \in \mathbb{D}_{\mathbf{S},\mathbf{M}}} \mathbf{1}(R(\mathbf{d}) > c)$ *exactly* equals α. In such cases, the econometrician instead chooses the highest $c \in \mathbb{R}$ such that $|\mathbb{D}_{\mathbf{S},\mathbf{M}}|^{-1} \sum_{\mathbf{d} \in \mathbb{D}_{\mathbf{S},\mathbf{M}}} \mathbf{1}(R(\mathbf{d}) > c)$ is *strictly less than* α. She then probabilistically rejects when $R(\mathbf{D}) = c$ to ensure proper size. In typical uses cases, the cardinality of the set $\mathbb{D}_{\mathbf{S},\mathbf{M}}$ will generally be intractably large, such that a researcher can just use a non-randomized test in practice. That is, she will be able to find a $c \in \mathbb{R}$ such that $|\mathbb{D}_{\mathbf{S},\mathbf{M}}|^{-1} \sum_{\mathbf{d} \in \mathbb{D}_{\mathbf{S},\mathbf{M}}} \mathbf{1}(R(\mathbf{d}) > c) \approx \alpha$ to such a high level of accuracy that there will be little gained from using a randomized decision rule. In such settings, it would be difficult to accurately compute $g_\alpha(\mathbf{t})$ in any case since the event $R(\mathbf{d}) = c_\alpha(\mathbf{t})$ will occur with low probability. This will become clearer when we discuss simulation of the null distribution below.

Under the null all adjacency matrices with the $\mathbf{S} = \mathbf{s}$ and $\mathbf{M} = \mathbf{m}$ are equally probable. By enumerating all adjacency matrices in $\mathbb{D}_{\mathbf{s},\mathbf{m}}$ we could exactly compute the null distribution of $R(\mathbf{D})$ and hence the critical values $c_\alpha(\mathbf{t})$ and $g_\alpha(\mathbf{t})$ defined above. In general such a brute force approach will be infeasible.[19] Therefore a method of approximating the exact null distribution is required. The simulation algorithm introduced in Section 3 below provides

[19] In fact very little is known about the set $\mathbb{D}_{\mathbf{s},\mathbf{m}}$; for example we are aware of no method for checking whether a given $\mathbf{s}, \mathbf{m}$ pair is graphic. From related settings we believe that the cardinality of $\mathbb{D}_{\mathbf{s},\mathbf{m}}$ will typically be intractably huge even for modestly-sized networks. See Blitzstein and Diaconis (2011) for discussion of this point and examples from a related setting.

such a method.

The intuition behind this test is straightforward. If the network in hand has an "unusually" large value of $R(\mathbf{D})$ relative to the set of all networks with same in- and out-degree sequences and cross-link matrices, then we reject the null that the baseline model is correctly specified. A rejection is not interpreted as evidence in favor of a particular alternative model. Relatedly, a feature of goodness-of-fit tests, including this one, is that we have may have low, or even, power equal to size in certain directions (Lehmann and Romano, 2005).

Observe that the test is exact, involving no appeal to approximations associated with an asymptotic thought experiment.

2.2 Optimal testing with an explicit alternative

In this subsection we discuss how to optimize our test when the alternative model space is explicitly specified. That is, when the researcher explicitly specifies the network benefit function in (2) and proceeds under the premise that the true network generating process lies either in the null or the (explicitly specified) alternative model space. In such settings a rejection provides evidence that $\gamma_0 > 0$ (in the context of a specific network benefit function). Naturally the researcher would like to maximize her power to reject, while continuing to maintain similarity. To accomplish this requires choosing the right test statistic.

Because an equilibrium selection mechanism is not explicitly specified under the alternative, likelihood ratio (LR) testing is not feasible (cf., Chen et al., 2018). As an alternative to a LR test, we instead choose, for each $\mathbf{t} \in \mathbb{T}$, the critical function, $\phi(\mathbf{D})$ to maximize the *derivative* of the (conditional) power function $\beta(\gamma, \mathbf{t}) = \mathbb{E}[\phi(\mathbf{D})|\, \mathbf{T} = \mathbf{t}]$ evaluated at $\gamma = 0$ subject to the (conditional) size constraint $\mathbb{E}_\theta[\phi(\mathbf{D})|\, \mathbf{T} = \mathbf{t}] = \alpha$. Such a $\phi(\mathbf{D})$ is *locally best* (Ferguson, 1967, Lemma 1, Section 5.5). Remarkably we show that the locally best test does not depend upon the form of the equilibrium selection mechanism $\mathcal{N}(\mathbf{d}, \mathbf{u}; \theta)$.

Differentiating the power function we get[20]

$$\left.\frac{\partial \beta(\gamma, \mathbf{t})}{\partial \gamma}\right|_{\gamma=0} = \mathbb{E}\left[\phi(\mathbf{D})\, \mathbb{S}_\gamma(\mathbf{D}|\, \mathbf{T}; \theta)\right|\, \mathbf{T} = \mathbf{t}] \tag{27}$$

[20]Differentiability of the likelihood function is formally established by Theorem 2.3 below.

with $\mathbb{S}_\gamma(\mathbf{d}|\,\mathbf{t};\theta)$ denoting the conditional score function

$$\begin{aligned}\mathbb{S}_\gamma(\mathbf{d}|\,\mathbf{t};\theta) &= \frac{1}{P_0(\mathbf{d};\delta)}\left.\frac{\partial P(\mathbf{d};\theta)}{\partial\gamma}\right|_{\gamma=0} - \sum_{\mathbf{v}\in\mathbb{D}_{\mathbf{s},\mathbf{m}}}\left.\frac{\partial P(\mathbf{v};\theta)}{\partial\gamma}\right|_{\gamma=0} \\ &= \frac{1}{P_0(\mathbf{d};\delta)}\left.\frac{\partial P(\mathbf{d};\theta)}{\partial\gamma}\right|_{\gamma=0} + k(\mathbf{t})\end{aligned}$$

and $k(\mathbf{t})$ only depending on the data through $\mathbf{T}=\mathbf{t}$ (Here, and in the balance of this section, it is understood that δ is evaluated at is population value δ_0). By the Neyman-Pearson lemma, the test with the critical function given by equation (25) above, where the test statistic, $R(\mathbf{d})$, is set equal to the log-likelihood gradient, $\frac{1}{P_0(\mathbf{d};\delta)}\left.\frac{\partial P(\mathbf{d};\theta)}{\partial\gamma}\right|_{\gamma=0}$, will be locally best within the class of similar tests.

The idea behind the locally best test is as follows. If the likelihood increases sharply as we move away from the null *in the direction of the alternative of interest*, then we take this as evidence against the null. Intuitively if the likelihood gradient in the neighborhood of the null is large, then the likelihood ratio will also be large for simple alternatives close to the null (i.e., when $\gamma\in(0,\epsilon]$).

Constructing the locally best critical function requires calculating $\frac{1}{P_0(\mathbf{d};\delta)}\left.\frac{\partial P(\mathbf{d};\theta)}{\partial\gamma}\right|_{\gamma=0}$. This is not straightforward since it depends on properties of the likelihood under the alternative (and consequently the equilibrium selection function). Nevertheless, we are able to derive the form of this derivative.

Theorem 2.3. (Locally Best Test) *(i) $P(\mathbf{d};\theta,\mathcal{N})$ is twice differentiable with respect to γ at $\gamma=0$. Its first derivative at $\gamma=0$ is*

$$\left.\frac{\partial P(\mathbf{d};\theta,\mathcal{N})}{\partial\gamma}\right|_{\gamma=0} = P_0(\mathbf{d};\delta) \times\left[\sum_{i\neq j}s_{ij}(\mathbf{d})\left\{d_{ij}\frac{f_U(\mu_{ij})}{\int_{-\infty}^{v_{ij}}f_U(u)\,\mathrm{d}u}-(1-d_{ij})\frac{f_U(\mu_{ij})}{\int_{v_{ij}}^{\infty}f_U(u)\,\mathrm{d}u}\right\}\right], \quad (28)$$

where $\mu_{ij}=A_i+B_j+X_j'\Lambda_0X_i$ equals the systematic, non-strategic, component of utility generated by arc ij and that f_U is the logistic density; (ii) the test statistic $R(\mathbf{d})=\frac{1}{P_0(\mathbf{d};\delta)}\left.\frac{\partial P(\mathbf{d};\theta)}{\partial\gamma}\right|_{\gamma=0}$ yields the locally best test in the direction of the specified alternative within the class of similar tests.

The proof of Theorem 2.3, along with some additional commentary, can be found in Section A.3 of the Supplemental Web Appendix. A key implication of Theorem 2.3 is that the form of the locally best test statistics does *not* depend upon $\mathcal{N}$, the equilibrium selection mechanism. This is essential, since optimal testing would not be feasible otherwise (at least without additional assumptions). One intuition for this finding is that equilibrium is unique with high probability when γ is close to zero. This means we can effectively ignore draws of $\mathbf{U}$ which lead to multiple equilibria when differentiating the likelihood.

Indeed, when γ is close to zero most players will have a strictly dominant strategy (that is the optimal set of links for them to send will be invariant to the play of their peers). Of course we need more information to recover the gradient with respect to γ, since this parameter measures the responsiveness of agents to their peers' actions. It turns out that a key scenario used in the derivative calculation involves considering draws of $\mathbf{U}$ where all players *except one* have strictly dominant strategies. The one player without a strictly dominant strategy provides the needed gradient information. This player's actions are sensitive to the play of their peers' – this delivers non-trivial gradient information – but NE is unique in this scenario such that the details of equilibrium selection do not matter. In the proof we show that the effect of realizations of $\mathbf{U}$ associated with multiple NE is negligible when γ is small enough.

Locally best vs. heuristic test statistics

With a little manipulation we can simplify (28) to:

$$\frac{1}{P_0\left(\mathbf{d};\delta\right)}\left.\frac{\partial P\left(\mathbf{d};\theta\right)}{\partial\gamma}\right|_{\gamma=0}=\sum_{i\neq j}\left[d_{ij}-F_U\left(\mu_{ij}\right)\right]s_{ij}\left(\mathbf{d}\right) \tag{29}$$

where $F_U\left(u\right)=e^u/\left[1+e^u\right]$ is the logistic CDF. This form of the statistic provides insight into how our test accumulates evidence against the null in practice. Consider the case where $s_{ij}\left(\mathbf{d}\right)=d_{ji}$, as would be true in agents' have a taste for reciprocated links. Observe that $F_U\left(\mu_{ij}\right)$ corresponds to the probability of the edge ij under the null. Therefore the optimal test statistic is large if we observe that many ij links *with low probability under the null* are reciprocated. It is not many reciprocated links that drives rejection per se, but the presence of many "unexpected" reciprocated links.

Consider a network of boys and girls with agents exhibiting a strong taste for gender-based homophily. The optimal test statistic in this case is the *conditional* sample covariance of D_{ij} and D_{ji} given (A_i, B_i, X_i) and (A_j, B_j, X_j). The test based upon the reciprocity index is – essentially – based upon the *unconditional* covariance. The effect of conditioning is to, for example, given more weight to heterophilous reciprocated links than to homophilous ones. Similarly we give more weight to reciprocated links across low degree agents, than to those across high degree agents.

We close by observing that the locally best property, like similarity, is a *non-asymptotic* one. In our "finite sample" setting we can say nothing about test consistency. We also note that optimality is in the region of $\gamma_0 = 0$. Test power need not be monotonic in γ_0 as we move far from zero. This is a feature of many other score tests.

Implementation

Two practical issues remain. The first, how to simulate the null distribution of the test statistic, is covered in the next section. Second, although the locally best test statistic does not depend on the details of equilibrium selection, it does depend on δ_0. Although the test will remain admissible when δ_0 is replaced by some other, perhaps arbitrary, δ, it will not be locally best.

A practical solution to this problem is to replace δ_0 with its joint maximum likelihood estimate (MLE) computed under the null. This particular MLE is elegantly studied by Fernández-Val and Weidner (2016) (see also Yan et al. (2018)). We emphasize that the feasible test based on the MLE is no longer locally best. When δ_0 is poorly estimated, as may occur when $\mathbf{D}$ is sparse, the power of the feasible test could be appreciably lower than that of the oracle.[21] The feasible test remains admissible.

Nevertheless, in our Monte Carlo experiments, some of which are reported in the Supplemental Web Appendix, we have found that replacing δ_0 with its MLE, results in a test which is nearly as powerful as the infeasible oracle test based on δ_0, and far more powerful that tests based on ad hoc statistics. This is true even in the "sparse" designs we consider. Of course this finding is specific to the particular experiments we considered. It would be interesting to study how to choose $\delta \in \Delta$ rigorously. While replacing δ_0 with a "good guess" seems sensible, what constitutes a "good guess" in this setting is a a topic we leave for future research.

3 Simulation: drawing uniformly at random from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$

Because a complete enumeration of $\mathbb{D}_{\mathbf{s},\mathbf{m}}$ is not feasible unless N is very small, making our test practical requires a method of constructing uniform random draws from this set. Such draws can be used to simulate the null distribution of any test statistic of interest.

The problem of simulating networks with fixed degree sequences is well-studied; with many domain specific applications (e.g., Sinclair, 1993). We add to this problem the additional requirement that the simulated network satisfies the cross-link matrix constraint.

Prior work on network simulation adopts one of two basic approaches. The first approach begins with an empty graph and randomly adds links. Links need to be added such that the

[21] We thank a referee for raising this concern.

end graph satisfies the degree sequence constraint. Blitzstein and Diaconis (2011) develop an algorithm along these lines. They cleverly use checks for graphicality of a degree sequence, available in the discrete math literature, to add links in a way which constrains the end graph to be in the target set.[22]

The second approach, to which our new method belongs, uses Markov Chain Monte Carlo (MCMC). Specifically an initial graph, satisfying the target constraints, is randomly rewired many times to create a new graph from the target set. Key to this approach is ensuring that each rewiring is compatible with the target constraints (e.g., maintains the network's degree sequence). The algorithm also needs to be constructed carefully to ensure that the end graph is a *uniform* random draw from the target set. Sinclair (1993), Rao et al. (1996), McDonald et al. (2007), Berger and Müller-Hannemann (2009) and Tao (2016) all developed MCMC methods for simulating graphs (or digraphs) with given degree sequences.

We are aware of no method of generating adjacency matrix draws from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$. The novelty of this problem, relative to the work described above, is the presence of the additional cross link matrix constraint, $\mathbf{M}$. In the discrete math literature the cross link matrix constraint corresponds to what is called a partition adjacency matrix (PAM) constraint. Czabarka et al. (2021) conjecture that determining whether a given $\mathbf{s}, \mathbf{m}$ pair is graphical, the PAM realization problem, is NP-complete. If their conjecture is correct (and NP $\neq$ P), using a Blitzstein and Diaconis (2011) type algorithm to draw from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$ is not feasible.

This leaves MCMC methods. Erdős et al. (2017) showed that naively incorporating a PAM constraint into existing MCMC algorithms destroys their correctness. In this section we introduce a new MCMC algorithm that *does* generate uniform random draws from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$. This algorithm is of independent interest. Before describing the algorithm we introduce some additional definitions and notation.

[22] See also Del Genio et al. (2010) and Kim et al. (2012). Graham and Pelican (2020) provide a textbook discussion of the Blitzstein and Diaconis (2011) algorithm.

3.1 Notation and definitions

We start by defining an alternating walk.

Definition 3.1. (Alternating Walk) An alternating walk H is sequence of (ordered) dyads of the form

$$H := (i_1, i_2), (i_3, i_2), (i_3, i_4), \ldots, (i_l, i_{l-1}) \tag{30}$$

or

$$H := (i_2, i_1), (i_2, i_3), (i_4, i_3), \ldots, (i_{l-1}, i_l) \tag{31}$$

with $i_k \in \mathcal{V}(G)$, $i_k \neq i_{k+1}$, $i_k \neq i_{k-1}$and
(i) if $(i_k, i_{k-1}) \in \mathcal{A}(G)$, then $(i_k, i_{k+1}) \notin \mathcal{A}(G)$
(ii) if $(i_k, i_{k-1}) \notin \mathcal{A}(G)$, then $(i_k, i_{k+1}) \in \mathcal{A}(G)$
(ii) if $(i_{k-1}, i_k) \in \mathcal{A}(G)$, then $(i_{k+1}, i_k) \notin \mathcal{A}(G)$
(iv) if $(i_{k-1}, i_k) \notin \mathcal{A}(G)$, then $(i_{k+1}, i_k) \in \mathcal{A}(G)$
for all $k = 2, \ldots, l-1$.

For brevity we will often refer to a walk simply by its node sequence, writing $H := i_1 i_2, \ldots, i_l$. To unpack Definition 3.1 it is easiest to consider an example. In Figure 3, Panel B, three altering walks are shown (the links not present are depicted as dotted arrows).

Observe that for $H := i_1 i_2, \ldots, i_l$, the adjacency matrix entries $D_{i_1 i_2}, D_{i_3 i_2}, \ldots, D_{i_l i_{l-1}}$ alternate between ones and zeros (or zeros and ones). This observation suggests a method of constructing an alternating walk via a sequence of "hops" across the adjacency matrix: pick row i_1 of the adjacency matrix and move horizontally to column i_2, where i_2 corresponds to one of the agents to which i_1 directs a link, next move vertically to row i_3, where i_3 is an agent which does not direct a link to i_2, and so on.[23] We call the horizontal moves *active steps* and vertical moves *passive steps*. Figure 2 provides an example construction. The different cases in Definition 3.1 correspond to walks beginning/ending with passive/active steps.

The length of an alternating walk equals the number of ordered dyads used to define it. An important type of alternating walk, which following Tao (2016), we call an *alternating cycle*, is central to our algorithm.

Definition 3.2. (Alternating Cycle) The alternating walk C is an alternating cycle if $i_1 = i_l$ and C has even length.

The length of an alternating cycle is at least four. Let $D_{i_1 i_2}, D_{i_3 i_1}, \ldots, D_{i_l i_{l-1}}$ be the sequence of adjacency matrix entries associated with alternating cycle C in $\mathbf{D}$. These entries necessarily form a sequence of zeros and ones (or ones and zeros).

[23] This description is essentially due to (Tao, 2016, p. 124).

Figure 2: Constructing an alternating walk

A: Alternating Walk

	a	b	c	d	e	f	g	h	i	j
a	0	1	0	0	1	0	0	0	0	0
b	0	0	0	0	0	0	0	0	0	0
c	0	0	0	1	0	0	0	1	0	0
d	0	0	0	0	0	0	0	0	0	0
e	0	0	1	0	0	0	0	0	0	0
f	0	0	0	0	0	0	0	1	1	0
g	0	0	0	0	0	0	0	0	0	0
h	0	0	1	0	0	0	1	0	0	0
i	0	0	0	0	0	0	0	0	0	0
j	0	0	0	0	0	0	1	0	1	0

B: Degree Sequence

	Indegree	Outdegree
a	0	2
b	1	0
c	2	2
d	1	0
e	1	1
f	0	2
g	2	0
h	2	2
i	2	0
j	0	2

Source: Authors' calculations.
Notes: Panel A depicts an alternating walk $j, g, a, b, c, d, e, c, a$ constructed using the adjacency matrix. The same altering walk is colored blue in Figure 2. Agent labels are given in the first column and row of the table. To construct such a walk randomly we begin by choosing an agent at random. Here agent j is chosen, with an ex ante probability of $\frac{1}{10}$ since there are ten agents in the network. Next we take an active step where one of agent j's outlinks is chosen at random. Here we choose the outlink to agent g, an event with an ex ante probability of $\frac{1}{2}$ since agent j has just two outlinks. Following the active step comes a passive step. In a passive step we move vertically to the row of an agent which does not direct a link to the current agent. Here we choose a from the set $\{a, b, c, d, e, f, i\}$ uniformly at random (i.e., with an ex ante probability of $\frac{1}{7}$). We continue with active and passive steps until we choose to stop or can proceed no further. Panel B reports the indegree and outdegree of each agent in the network. Observe that in active steps the probability of any feasible choice equals the inverse of the outdegree of the current agent. In passive steps the probability of any feasible choice equals the inverse of the number of nodes minus the indegree of the node chosen in the prior step minus 1 (since $i_k \neq i_{k+1}$). We can also construct alternating walks by the above procedure, but instead starting with a passive step. The shaded cells in the table shows which edges (ones) and non-edges (zeros) are in the walk.

Consider constructing an alternative digraph, say $\mathbf{D}'$, by replacing all the "ones" in the alternating cycle C with "zeros" and all "zeros" with "ones". Rewiring $\mathbf{D}$ in this way is degree preserving: $\mathbf{D}'$ has the same in- and out-degree sequence as $\mathbf{D}$. We refer to such operations as switching the cycle (since we switch the zeros and ones).

We use random alternating walks on the network in order to find alternating cycles. We then use these alternating cycles to rewire the network. This motivates the definition of what we call a schlaufe. A schlaufe is either an alternating walk which contains an alternating cycle (as the last part of the walk) or it is an alternating walk which cannot be continued. More precisely

Definition 3.3. (SCHLAUFE) An alternating walk $H := i_1 i_2 \dots i_l$ is a *schlaufe* if either
(i) There is a node $i_k \in \{i_1 i_2 \dots i_l\}$ with $k \neq l$ such that $i_k = i_l$ and $(k - l) \bmod 2 = 0$. Furthermore for any two nodes i_j and i_h in $\{i_1 i_2 \dots i_{l-1}\}$ with $i_j = i_h$ and $j \neq h$ it holds that $(j - h) \bmod 2 = 1$.
(ii) At node i_l there is no other node i_{l+1} such that the alternating walk could be extended with the unmarked link (i_l, i_{l+1}).

In German schlaufe corresponds to "loop", "bow" or "ribbon" (its plural is schlaufen); the latter translation is evocative of our meaning here. In the first case the schlaufe will coincide with an alternating walk which includes exactly one alternating cycle.[24] Visually schlaufen of the first type, with the nodes appropriately placed, will look like loops and ribbons. In the second case the schlaufe does not include an alternating cycle.

Associated with a schlaufe, R, is a $K \times K$ violation matrix which records the number of extra links from group k to group l generated by switching the alternating cycle in R (if there is one). Consider an alternating rectangle consisting of two boys and two girls. If initially one boy directs a link to the other and one girl directs a link to the other, then after switching the cycle the violation matrix will equal:

Ego \Alter	Boy	Girl
Boy	-1	1
Girl	1	-1

After switching the cycle there are too few same gender links and too many mixed gender ones.

[24] The requirement that $i_k = i_l$ and $(k - l) \bmod 2 = 0$ ensures that $C = i_k i_{k+1} \dots i_l$ is an alternating cycle (imposing even length). The "furthermore..." requirement ensures that if another node is visited multiple times it does not form an alternative cycle (imposing non-even length). See Figure 3 for an example.

We call a sequence of schlaufen $\mathcal{R} = (R_1, \ldots, R_k)$ feasible if (i) the cycles of the schlaufen are link disjoint and (ii) the sum of their violation matrices is zero (and for $i < k$ the sum of their violation matrices is not zero).

Conventional MCMC adjacency matrix re-wiring algorithms work by switching short cycles (e.g., alternating rectangles and compact alternating hexagons as in Rao et al. (1996)). Switches of this type, while preserving the in- and out-degree sequence of the network will typically generate networks with the wrong inter-group link structure (i.e., non-zero link violation matrices). Our approach to solving this problem involves switching many alternating cycles simultaneously such that their individual link violation matrices sum to zero.

3.2 The MCMC algorithm

Let $\mathbf{S} = \mathbf{s}$ and $\mathbf{M} = \mathbf{m}$ be the degree sequence and cross link matrix of the network in hand. In order to a draw, say $\mathbf{D}'$, from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$ we (i) start with a realization of $(\mathbf{s}, \mathbf{m})$, say $\mathbf{D}$, (ii) randomly construct (link disjoint) schlaufen, and (iii) switch any alternating cycles in them. While switching cycles will preserve the degree sequence, it may – as discussed earlier – result in a graph without the appropriate cross link matrix. In order to ensure that $\mathbf{D}'$ has the appropriate cross link matrix, we construct schlaufen until either the sum of their violation matrices equals zero or we stop randomly. If the sum of the schlaufen violation matrices is zero we move to $\mathbf{D}'$ from $\mathbf{D}$ by switching the cycles, otherwise we set $\mathbf{D}' = \mathbf{D}$. Proceeding in this way ensures that $\mathbf{D}'$ is, in fact, a random draw from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$. After sufficiently many iterations of this process we show that a graph constructed in this way corresponds to uniform random draw from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$. A formal statement of the procedure is provided by Algorithm 1.

Algorithm 1 uses a subroutine to find schlaufen. This subroutine, described in Algorithm 2, finds and marks a schlaufe in the graph.

To illustrate our method in more detail consider the network depicted in Panel A of Figure 3. This network consists of two types of agents: gold (light) and blue (dark). The cross link matrix for the graph is given in Panel D. In Panels B and C a sequence of three schlaufen is shown. The first schlaufen is $R_1 = jgabcdeca$. It is constructed through a sequence of active and passive steps as described earlier (see also the notes to Figure 2 above). We begin by choosing agent j randomly with a probability of $\frac{1}{10}$ (since there are ten agents in the network). We then take an active step, randomly choosing one of the two agents to which j directs a link (i.e., either agent g or i). Here we choose agent g. Next we take a passive step. Specifically we choose an agent at random from the set of agents that *do not* direct a link to g (the agent chosen in the previous active step). The probability associated with our choice in this passive step is $\frac{1}{7}$; this corresponds to the reciprocal number of agents in the network (i.e., 10) minus the indegree the current agent (i.e., 2) minus one (since self-loops

Algorithm 1 MARKOV DRAW ALGORITHM

Inputs: An adjacency matrix $\mathbf{d} \in \mathbb{D}_{\mathbf{s},\mathbf{m}}$; a mixing time τ

Procedure:

1. Set $t = 0$.
2. With probability $1 - q$ go to step 3, with probability q go to step 4.
3. find and mark a schlaufe (see Algorithm 2):
   (a) **if** the sum of the schlaufen violation matrices is zero, then
      i. switch the cycles in the schlaufen (changing the adjacency matrix $\mathbf{d}$),
      ii. unmark all links,
      iii. go to step 4.
   (b) **else**
      i. with probability $\frac{1}{2}$, go to step 3 or
      ii. with probability $\frac{1}{2}$, unmark all links and go to step 4.
4. Set $t = t + 1$
   (a) **if** $t = \tau$ then return $\mathbf{d}$
   (b) **else** go to step 2

Output: A uniform random draw $\mathbf{d}$ from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$

Algorithm 2 SCHLAUFEN DETECTION ALGORITHM

Inputs: An adjacency matrix $\mathbf{d} \in \mathbb{D}_{\mathbf{s},\mathbf{m}}$ (this network may have marked links in it)

Procedure:

1. Choose an agent/node, say i, at random.
2. Mark agent i as active and
   (a) **if** feasible, randomly choose one of $i's$ (unmarked) outlinks, say to j, and go to step 3;
   (b) **else** (i.e., no unmarked outlinks available) go to step 6.
3. Mark edge ij, chosen in step 2 and
   (a) **if** agent j is already marked passive, then go to step 6;
   (b) **else** go to step 4.
4. Mark agent j, chosen in step 3, as passive and
   (a) **if** feasible, randomly choose an agent, say k, from among those who *do not* direct links to j, and go to step 5,
   (b) **else** go to step 6.
5. Mark edge kj, with k the agent chosen in step 4, as passive and
   (a) **if** agent k is already marked active, then go to step 6;
   (b) **else** go to step 2.
6. return the (marked) adjacency matrix, the constructed schlaufe and its violation matrix.

Output: A schlaufe, its violation matrix and a marked adjacency matrix.

are not allowed). We continue taking active and passive steps in this way until we visit a for the second time. At this point we stop since our schlaufe now includes the alternating cycle $C_1 = abcdeca$. Note that c is also visited twice, but also that $cdec$ is not an alternating cycle since it is not of even length (see Definition 3.2).

As seen in the example we can calculate the probability of a schlaufe R as we go through the algorithm (see Panel E). In Step 1 of Algorithm 2 an agent is chosen with probability $\frac{1}{N}$. Next let $r^a_D(i)$ be the cardinality of the set of feasible out links in an active step. This set consists of all the out links of node i, which are not already marked in $\mathbf{D}$. Similarly, let $r^p_D(i)$ be the cardinality of the set of feasible outlinks in an passive step. That set consists of all the links ij for which ji is not in $\mathbf{D}$ and which are not already marked. The probability of $R = (i_1, .., i_l)$ can now be written as

$$p_{\mathbf{D}}(R) = \frac{1}{N} \prod_{k=1}^{l-1} \left(\frac{1}{r^a_{\mathbf{D}}(i_k)} \left[k \bmod 2 \right] + \frac{1}{r^p_{\mathbf{D}}(i_k)} \left[(k-1) \bmod 2 \right] \right) \tag{32}$$

In step 2 of Algorithm 1 we attempt to find a sequence of schlaufen with probability $1 - q$ and do not change the adjacency matrix otherwise. In step 3, a schlaufen sequence $\mathcal{R} = (R_1, .., R_h)$ is constructed/found. After each detected schlaufe in this sequence, say R_k, any cycle in it is marked. Let $\mathbf{D}_k$ be the graph with the cycles of $R_1, .., R_{k-1}$ marked. After each schlaufe added the construction is stopped with probability $\frac{1}{2}$. The probability of finding a cycle R_k is $p_{\mathbf{D}_k}(R_k)$ as given in equation (32) above. The total probability of a feasible schlaufen sequence $\mathcal{R}$ is therefore

$$p_{\mathbf{D}}(\mathcal{R}) = (1 - q) \frac{1}{2^{(h-1)}} \prod_{i=1}^{h} p_{\mathbf{D}_k}(R_k). \tag{33}$$

3.3 Correctness

To show that our algorithm does indeed generate a uniform random draw from the set $\mathbb{D}_{\mathbf{s},\mathbf{m}}$ we use standard Markov chain theory (e.g., Chapters 7 and 10 of Mitzenmacher and Upfal (2005)).

The random rewiring of the network implemented by Algorithm 1 can be described as a Markov chain. To show that, for τ large enough, it returns a uniform random draw from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$ we prove that the stationary distribution of the Markov chain generated by Algorithm 1 is uniform on $\mathbb{D}_{\mathbf{s},\mathbf{m}}$. To show this it is helpful to develop a graphical representation of the Markov chain.

We denote the state graph of the Markov chain by $\Phi = (\mathcal{V}_\phi, \mathcal{A}_\phi)$. Its underlying vertex set $\mathcal{V}_\phi$ is the set of all elements in $\mathbb{D}_{\mathbf{s},\mathbf{m}}$: each node in our state graph is a network with degree

Figure 3: A feasible schlaufen sequence

A: Network prior to edge swaps

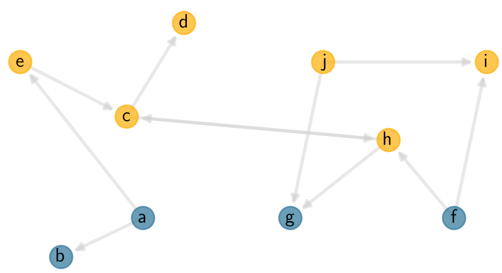

D: Cross-link (PAM) matrix

	Blue	Gold
Blue	1	3
Gold	2	5

B: Network with
three schlaufen shown

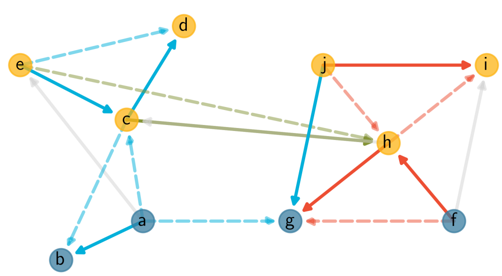

E: Three schlaufen

	j	g	a	b	c	d	e	c	a
i	1	2	3	4	5	6	7	8	9
sel. Pr	$\frac{1}{10}$	$\frac{1}{2}$	$\frac{1}{7}$	$\frac{1}{2}$	$\frac{1}{8}$	$\frac{1}{2}$	$\frac{1}{8}$	$\frac{1}{1}$	$\frac{1}{8}$

	c	h	e
i	1	2	3
sel. Pr	$\frac{1}{10}$	$\frac{1}{1}$	$\frac{1}{7}$

	f	h	j	i	h	g	f
i	1	2	3	4	5	6	7
sel. Pr	$\frac{1}{10}$	$\frac{1}{2}$	$\frac{1}{6}$	$\frac{1}{1}$	$\frac{1}{7}$	$\frac{1}{1}$	$\frac{1}{6}$

C: Network after
edge swaps

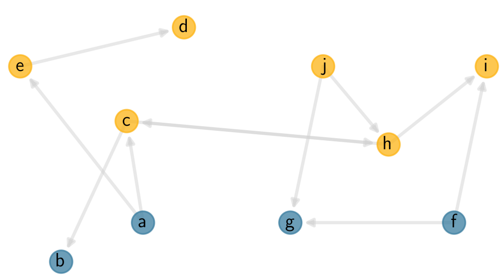

F: Violation matrices for
the three schlaufen

	Blue	Gold
Blue	-1	+1
Gold	+1	-1

	Blue	Gold
Blue	0	0
Gold	0	0

	Blue	Gold
Blue	+1	-1
Gold	-1	+1

Source: Authors' calculations.
Notes: See the discussion in the main text. The figure depicts three link disjoint schlaufen with violation matrices which sum to zero. Panel E reports the (ex ante) probability that a given node was selected as the schlaufe was constructed. See equation (32).

sequence $\mathbf{S} = \mathbf{s}$ and cross link matrix $\mathbf{M} = \mathbf{m}$. For network $\mathbf{D}$ in $\mathbb{D}_{\mathbf{s},\mathbf{m}}$, we denote by $v_{\mathbf{D}}$ the corresponding vertex in $\mathcal{V}_\phi$. The arc set $\mathcal{A}_\phi$ is defined as follows.

1. For all vertices we add the self loop $(v_{\mathbf{D}}, v_{\mathbf{D}})$ with (probability) weight q (see Step 2 of Algorithm 1).

2. Let $\mathbf{D}$ and $\mathbf{D}'$ be two different networks in $\mathbb{D}_{\mathbf{s},\mathbf{m}}$. Let $\mathbf{D}\Delta\mathbf{D}'$ equal the union of the set of edges in $\mathbf{D}$, but not in $\mathbf{D}'$ and the set of edges in $\mathbf{D}'$, but not in $\mathbf{D}$. For each feasible schlaufen-sequence $\mathcal{R}$, with cycle edge set equal to $\mathbf{D}\Delta\mathbf{D}'$ we add the edge $(v_{\mathbf{D}}, v_{\mathbf{D}'})$ and assign to it probability weight $p_{\mathbf{D}}(\mathcal{R})$.

3. Finally we add a directed loop $(v_{\mathbf{D}}, v_{\mathbf{D}})$ if the probability of all arrows leaving $v_{\mathbf{D}}$, introduced in points 1 and 2 immediately above, do not sum to 1. The probability of this loop is 1 minus the sum of the probability of all other outward arrows.

The probability of any arc $a \in \mathcal{A}_\phi$ is denoted by $p(a)$. Note, by definition, the state graph can have parallel arcs and loops.

With these definitions in place we can prove correctness of the algorithm. First we show that the probability of the algorithm moving from graph $\mathbf{D}$ to $\mathbf{D}'$ coincides with the probability of moving in the reverse direction.

Lemma 3.1. *For any two vertexes $v_{\mathbf{D}}, v_{\mathbf{D}'}$ the transition probability attached to $(v_{\mathbf{D}}, v_{\mathbf{D}'})$ equals that attached to $(v_{\mathbf{D}'}, v_{\mathbf{D}})$.*

Proof. See appendix A.4. □

Next we show the state graph is strongly connected. This means our Algorithm moves from any $\mathbf{D} \in \mathbb{D}_{\mathbf{s},\mathbf{m}}$ to any other $\mathbf{D}' \in \mathbb{D}_{\mathbf{s},\mathbf{m}}$ with positive probability.

Lemma 3.2. *The state graph Φ is strongly connected.*

Proof. See appendix A.4. □

With these two lemmata it is easy to show that the stationary distribution is uniform on $\mathbb{D}_{\mathbf{s},\mathbf{m}}$. This gives us the main result of the section.

Theorem 3.3. *Algorithm 1 is a random walk on the state graph Φ which samples uniformly a network from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$ for $\tau \to \infty$.*

Proof. See appendix A.4. □

Figure 4: Nyakatoke Village Risk-Sharing Network

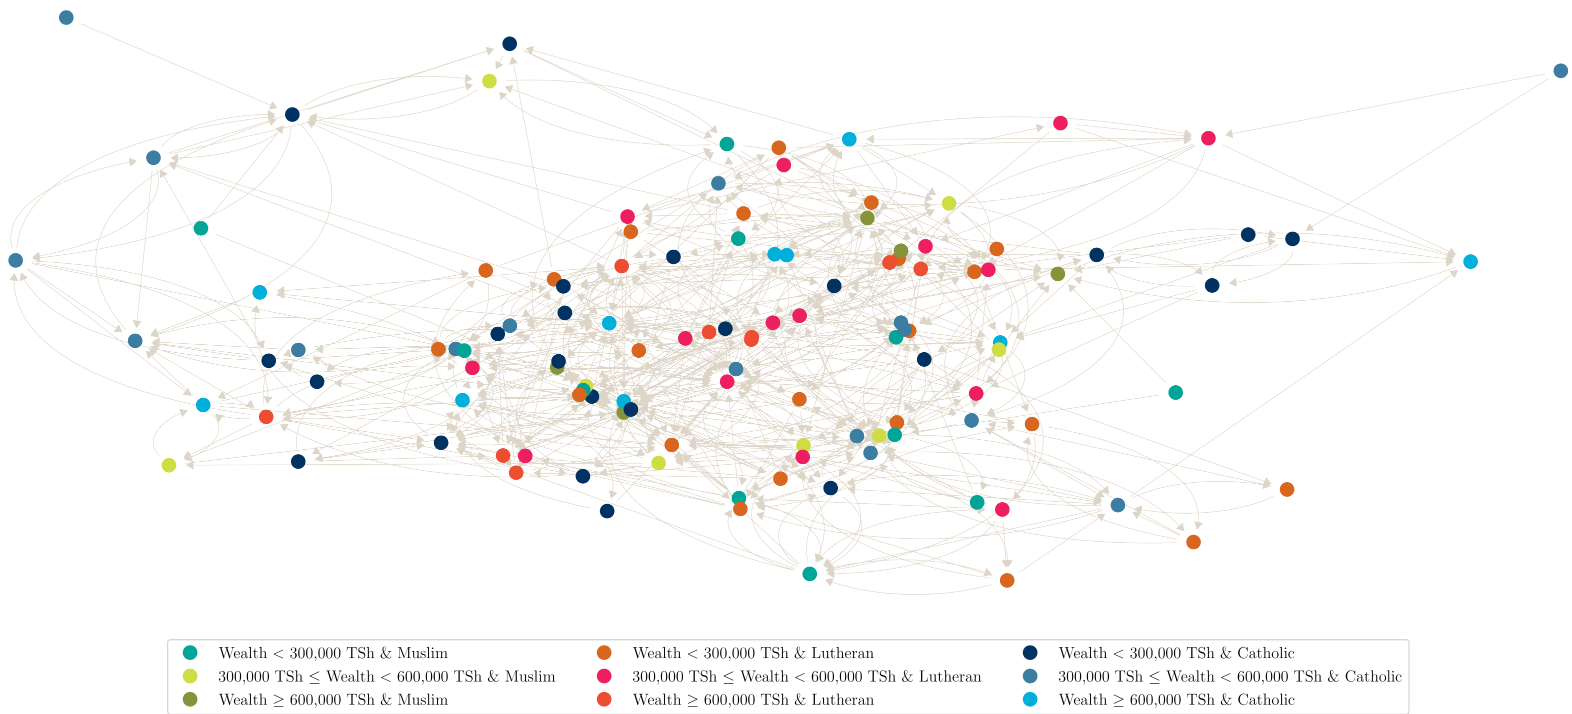

Source: De Weerdt (2004) and authors' calculations.
Notes: Each household is colored according to their land and livestock wealth (measured in Tanzanian Shillings) and religion. The arrow head on the edges points to the "alter" household with the link being sent by the tail "ego" household.

4 Empirical illustration: risk-sharing links when agents value bridging capital

De Weerdt (2004) studied the formation of risk-sharing links across 119 households in the rural village of Nyakatoke (located in Tanzania). He asked all adult individuals in the village who they could rely upon for help and, from their responses, constructed a network of directed links across households.[25]. The resulting set of links is shown in Figure 4.

Modeling the configuration of links shown in Figure 4 as a NE of a complete information network formation games is reasonable in our setting. For example, the observed network may correspond to the long-run rest point associated with an un-modeled dynamic network formation process. The small village rural setting of Nyakatoke is consistent with agents having high levels of information about their own and others' payoffs. Finally we interpret the De Weerdt (2004) prompt at face value: households report who they would – in fact – turn to in a time of need.

Here we assess whether households value "bridging capital", as suggested by Burt (1995) and formalized in game-theoretic terms by Kleinberg et al. (2008) and others. If k directs a link to i but not to j, then i, by directing a link to j, may position herself to serve as a "bridge" or "broker" between k and j. See Figure 1.2 above.

In the formal model of Kleinberg et al. (2008) agents gain utility from positioning themselves on length two paths connecting agents not directly connected themselves; however such utility gains are decreasing in the number of "rival" length two paths (i.e., those with other agents in the center). This suggest, for example, a network benefit function of

$$g_i(\mathbf{d}) = \sum_{j\neq i,k,j} \sum_{k\neq i,j} \frac{D_{ki}D_{ij}(1-D_{kj})}{\max\left(1, \sum_{l\neq j,k} D_{kl}D_{lj}(1-D_{kj})\right)}. \tag{34}$$

In this formulation any "bridging" capital is shared equally across all agents l on length two paths from j to k (with arc jk absent). For example, if there are two bridging agents situated between j and k, they each get half the benefit and so on. The marginal network benefit of edge ij is thus

$$s_{ij}(\mathbf{d}) = \sum_{k\neq i,j} \frac{D_{ki}D_{ij}(1-D_{kj})}{\max\left(1, \sum_{l\neq j,k} D_{kl}D_{lj}(1-D_{kj})\right)}, \tag{35}$$

[25] The prompt used by De Weerdt (2004) is suggestive of both mutuality and directionality, leading to some ambiguity in whether to interpret the collected edges as undirected or directed. Comola and Fafchamps (2014) present evidence suggesting that the links given by households are directed. Specifically that they indicate to which other households they would turn to in the event of need. It is this interpretation that we give the links here.

from which the form of the locally best test follows.

From De Weerdt (2004) we also know that household land and livestock wealth, as well as religion (Catholic, Lutheran or Muslim), are important drivers of link formation in Nyakatoke. We divide households into three wealth bins, which in conjunction with religion, partitions households into nine groups; X_i consists of the nine resulting group membership dummies with the 81 elements of Λ parameterizing any homophily/heterophily across these groups. The remaining null model parameters are the $238 = 119 \times 2$ household-specific in- and out-degree heterogeneity parameters. This gives $\dim(\delta) = 2 \times 119 + 9 \times 9 = 319$ null model "nuisance" parameters. It is hard to imagine a testing approach with good properties in this setting which would not involve "conditioning away" the null model parameter.

One aim of our empirical illustration is to compare the performance of the locally best test statistic, which follows naturally from the form of the Kleinberg et al. (2008) network benefit function, to that of heuristically motivated *ad hoc* test statistic.

It is not entirely clear how to form a heuristic test with power to detect the more qualitative alternative "agents like to bridge disconnected groups". Indeed, this lack of clarity is one argument for using a locally best test. Such tests proceed in a principled way from an explicit network benefit function.

After some experimentation we settled on the difference between the 90th and 50th percentiles of the empirical distribution of *betweenness-centrality* across agents in the network as a suitable *ad hoc* test statistic (other measures of dispersion give similar results).

The reasoning behind this choice is as follows. As before, let $\tilde{\mathbf{d}}$ denote the undirected network obtained from $\mathbf{d}$. Next let $\phi_i^{jk}(\tilde{\mathbf{d}})$ denote the number of paths between j and k in $\tilde{\mathbf{d}}$ which pass through i and $\phi^{jk}(\tilde{\mathbf{d}})$ the total number of paths connecting j and k (whether they pass through i or not).

Agent i's *betweenness centrality* (Wasserman and Faust, 1994, p. 190) equals:

$$C_i^{\mathrm{BC}}\left(\tilde{\mathbf{d}}\right) = \sum_{j<k} \frac{\phi_i^{jk}(\tilde{\mathbf{d}})}{\phi^{jk}(\tilde{\mathbf{d}})}. \tag{36}$$

Equation (36) is maximal when i is situated on every path between every pair of connected agents. When $\tilde{\mathbf{d}}$ is connected, this maximal value equals the number of dyads excluding i or $\binom{N-1}{2} = \frac{(N-1)(N-2)}{2}$.

The "90-50 gap" in the empirical distribution of $C_i^{\mathrm{BC}}\left(\tilde{\mathbf{d}}\right)$ measures right-tail inequality in betweenness centrality. If agents value bridging capital, then it is plausible that the top 10 percent of agents in the network will acquire substantially more of such capital – as proxied by betweenness centrality – than the typical (or median) agent.

The intuition behind this claim is that acquiring bridging capital is inherently rivalrous;

the addition of links by other agents may reduce one's own betweenness centrality. Competition to accumulate bridging capital should therefore lead to more dispersion in betweenness-centrality across agents (than in a reference set of null model graphs). Winners of this competition (the 90th percentile) will have more bridging capital than the typical agent (the 50th percentile) in the network.

We wish to emphasize that the "ad hoc" descriptor of this statistic is apt. The reasoning outlined above is both meandering and speculative; we provide it simply as an example of how one might select a test statistic heuristically. In contrast, an advantage of the formalism of an explicit network benefit function is that gives precision to the alternative of interest (in turn suggesting a suitable, in fact, optimal test statistic).[26]

The left panel of Figure 5 plots simulation estimates of the distribution of the $90-50$ betweenness-centrality gap across three reference sets of networks: (i) Ërdos-Rényi graphs with the same number of links as observed in Nyakatoke, (ii) the set of all graphs with the same in- and out-degree sequences as observed in Nyakatoke, and (iii) the set of all graphs which additionally constrain the number of cross-group links to be the same as observed in Nyakatoke. The vertical line in the figure marks the value of the actual $90-50$ betweenness-centrality gap in Nyakatoke.

The three reference distributions in Panel A allow us to undertake three model adequacy tests: is Nyakatoke well-described by (i) the Ërdos-Rényi model, (ii) a directed β-model which places equal probability on all networks with the same in- and out-degree sequence as in Nyakatoke, or (iii) by the full baseline model described above (which additionally accommodates homophily)? In all three cases we reject, but notice that as we enrich the null model the simulated reference distributions shift to the right.[27] Put differently a portion of the dispersion in betweenness-centrality across households observed in Nyakatoke is likely a by-product of degree heterogeneity and homophily. The rightward shifts in the reference distributions as we enrich the null model is indicative of how using a realistic null model may be important for avoiding spurious rejections in practice. That said, our decisive rejection of even the 319 parameter baseline model model indicates that degree heterogeneity and wealth/religion homophily cannot explain *all* of the inequality in betweenness-centrality we observe across agents in Nyakatoke.

The right panel of Figure 5 plots the null distribution of the locally best test statistic for the alternative that households gain utility by bridging disconnected pairs of agents (as

[26] Of course one can also start with a specific test statistics and then reverse engineer a network benefit function for which it is locally best. We do not advocate proceeding in this way in practice, but such an exercise can be useful for understanding the game theoretic implications of test statistics initially proposed in other settings. We thank the co-editor for this observation.

[27] The incremental effect of additionally controlling for homophily is modest.

Figure 5: Testing for bridging/brokerage preferences

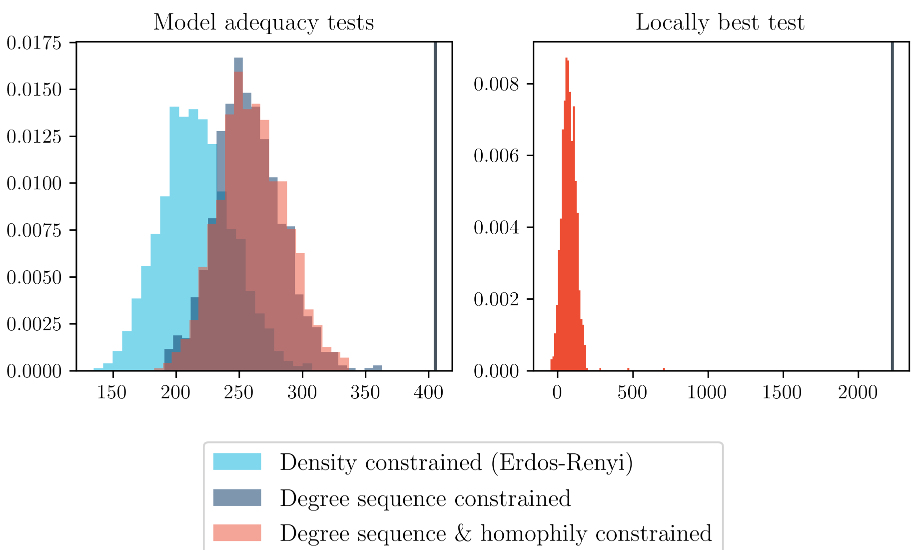

Source: De Weerdt (2004) and authors' calculations.
Notes: Panel A presents MCMC estimates of the *distribution* of the $90-50$ betweenness-centrality gap across agents for three reference sets of networks (as listed in the legend). Panel B shows the null distribution of the locally best test described in the main test. In this panel the reference set is all networks with the same in- and out-degree sequences and cross-link matrix as observed in Nyakatoke.

formalized by Kleinberg et al. (2008)). If we are willing to maintain that the true data generating process is either in the null or specified alternative model space, we can interpret a rejection as evidence *for* γ_0 being positive. To implement this test we replace δ_0 with its maximum likelihood estimate (MLE) computed under the null.[28] As is clear from Panel B of Figure 5, we decisively reject the null.

Panel B is also suggestive of the power gains associated with the locally best test. If we were to standardize each of our test statistics using their respective reference distribution's mean and standard deviation, it is obvious that the locally best test statistic is more extremely positioned in the right tail of its null distribution (the Monte Carlo experiments reported in the Supplemental Web Appendix confirm the power advantages of the locally best test).

Using Algorithm 1 requires a choice of the mixing time parameter τ. Although the mixing properties of our MCMC procedure are largely unexplored, we have found - by Monte Carlo experimentation – that choosing τ such that each edge in the input graph is, on average, swapped at least once before the resulting output is considered a uniform random draw from the target set to yield acceptable results in practice. We use this approach here (also see the Python Juypter Notebook in the Supplemental Materials). The required value for τ is increasing in the dimension of the nuisance parameter δ and especially in the dimension of Λ. Hence the speed of the simulation algorithm declines in both N and K.

5 Limitations and future research

The analysis in this paper, like much of the wider econometrics literature on games, is likelihood based. Our null model is fully parametric (albeit flexibly-so), while the alternative, due to the unmodeled NE selection function, is semiparametric. Under correct specification our test reveals whether $\gamma_0 = 0$ or $\gamma_0 > 0$ (with a researcher-specified *exact* Type I error rate, and a locally best Type II error rate). That is, we present a method for detecting whether agents form links "strategically" in the presence of any pattern of homophily and degree heterogeneity allowed by the baseline null model.

It would be interesting to know whether detecting strategic interaction in the presence of *arbitrary* homophily on observables and degree heterogeneity is possible. We know from the panel data literature that detecting state-dependence in the presence of heterogeneity is non-trivial and that modeling details matter (e.g., Chamberlain, 1985). Analogous questions arise here.

[28] The computation of this MLE is described in detail by Yan et al. (2018) and implemented in our Python package **ugd** for "uniform graph draw". Additional discussion can be found in the Python Jupyter notebooks included in the Supplemental Materials.

Our set-up assumes that researcher is able to a priori partition the support of agents' covariates into K regions along which all homophilous sorting occurs. In practice this is an approximation. Developing data-based discretization rules (e.g., using clustering algorithms) and formalizing the nature of the approximations involved would be useful. It is possible that recent results on randomization inference by Canay et al. (2017) could be useful for such an analysis.

We conjecture that, for K sufficiently large, further increases in it will reduce power. In contrast, too coarse of a covariate partition could lead the researcher to reject not because of any strategic interaction, but simply because the baseline model is misspecified.[29] Note our test correctly rejects the null in this case, the subtlety centers upon interpretation. Such concerns arises in other specification testing problems.[30] Our test could also reject in the presence of homophily on *latent* attributes.

Key to our set-up is the exponential family structure (under the null) induced by the assumption of logistic random link-specific utility. While this is a strong assumption, it comes with considerable pay-off: we are (i) able to *exactly* control size in (ii) the presence of a high dimensional nuisance parameter while (iii) also making no assumptions about equilibrium selection. Exponential family structures has proved highly fruitful in other areas of econometrics; applications in panel data being most closely connected to the present setting. Our similarity and local optimality results build on classic results in the theory of testing in exponential families (e.g., Ferguson (1967) and Lehmann and Romano (2005)).

Recent work studies dyadic regression in settings richer than our baseline model (e.g., Gao, 2020). Adapting such work to our testing problem would be an interesting area for future research. The loss of exponential family structure would mean a loss of exact size control and optimality. However, the insight that score tests in the direction of certain alternatives can be constructed without modeling the details of NE selection should still hold. Any such extension would require difficult asymptotic arguments; but we expect insights from the large-N, large-T panel data literature (e.g., Fernández-Val and Weidner, 2016) as well as the large games literature (e.g., Menzel, 2016) to be useful in any such extension.

While obvious, and generic to most testing problems, it is important to understand that our test may have low power in some directions (in extreme cases even power equal to size). As an example imagine agents gain utility from linking with popular agents (as in preferential attachment models), such that $g_i(\mathbf{d}) = \sum_{j \neq i} d_{ij} \left[\sum_{k \neq i} d_{kj}\right]$. This model yields $s_{ij}(\mathbf{d}) = \sum_{k \neq i} d_{kj}$, which is almost equal to the indegree of agent j. Hence the distribution

[29] We thank the referees for this observation.

[30] For example a Sargan-Hansen test of over-identifying restrictions might reject in the linear instrumental variables setting because one of the instrument exclusion restrictions is violated or because of treatment effect heterogeneity; with rather different implications for how to proceed in each case.

of $s_{ij}(\mathbf{D})$ across $\mathbb{D}_{\mathbf{s},\mathbf{m}}$ will be nearly degenerate. Examples of this type are not unique to our setting. See Lehmann and Romano (2005) for general impossibility results.

Finally, while we are able to prove that our simulation algorithm works for τ "large enough", we don't currently have a formal handle on the mixing properties of our MCMC algorithm. This is not just a limitation of our work, but of much of the related work in the discrete math and computer science literature (e.g., Cooper et al. (2007) and Erdos et al. (2018)). Our limited simulation experiments suggest relatively fast mixing. [31]

These limitations notwithstanding, we nevertheless see potential for the widespread use of the methods presented in this paper in empirical social and economic network research (and, with modification, in other settings where strategic interaction is important). We hope that the ability to easily embed formal game-theoretic models of network formation of the type surveyed by, for example, Jackson (2008) and Goyal (2023) into heterogeneity-rich dyadic linking models will be attractive to empirical researchers. While not emphasized here, we also expect our simulation algorithm to find use in other settings where binary matrix simulation is an important part of researchers' toolkits (e.g., Gotelli, 2000). Finally our focus on score type tests may represent a fruitful direction for further research on testing in incomplete models (e.g., Chen and Kaido, 2021).

The Supplemental Web Appendix shows how to adapt our results to bi-partite networks. There we show how ideas in this paper might be used to, for example, study airline entry into different routes as in Ciliberto and Tamer (2009). The set-up allows for complex airline preferences over their own route map as well as the route maps chosen by their competitors. We also shows how our simulation algorithm can be used for more traditional conditional likelihood estimation and inference problems. A carefully annotated Python Jupyter, Notebook illustrating how the methods in this paper work in practice, is available in the Supplemental Materials.

References

Amelkin, V. and Vohra, R. (2023). Yield uncertainty and strategic formation of supply chain networks. *Networks*, 83(1):131 – 152.

Andrews, D. W. and Ploberger, W. (1994). Optimal tests when a nuisance parameter is present only under the alternative. *Econometrica*, 62(6):1383 – 1414.

Apicella, C. L., Marlowe, F. W., Fowler, J. H., and Christakis, N. A. (2012). Social networks and cooperation in hunter-gatherers. *Nature*, 481(7382):497 – 501.

[31] A simple heuristic is to increase τ until one's results are not sensitive to further increases in it.

Atalay, E., Hortaçsu, A., Roberts, J., and Syverson, C. (2011). Network structure of production. *Proceedings of the National Academy of Sciences*, 108(13):5199 – 5202.

Bala, V. and Goyal, S. (2000). A noncooperative model of network formation. *Econometrica*, 68(5):1181 – 1229.

Barabási, A.-L. (2016). *Network Science*. Cambridge University Press, Cambridge.

Barndorff-Nielsen, O. E. and Cox, D. R. (1994). *Inference and Asymptotics*. Chapman and Hall/CRC, Boca Raton.

Bartolucci, F. (2024). Mcmc conditional maximum likelihood for the two-way fixed-effects logit. *Econometric Reviews*, 43(6):379 – 404.

Bennett, D. S. and Stam, A. C. (2000). Research design and estimator choices in the analysis of interstate dyads: when decisions matter. *Journal of Conflict Resolution*, 44(5):653 – 685.

Beresteanu, A., Molchanov, I., and Molinari, F. (2011). Sharp identification regions in models with convex moment predictions. *Econometrica*, 79(6):1785 – 1821.

Berger, A. and Müller-Hannemann, M. (2009). Uniform sampling of undirected and directed graphs with a fixed degree sequence. *CoRR*, abs/0912.0685.

Blitzstein, J. and Diaconis, P. (2011). A sequential importance sampling algorithm for generating random graphs with prescribed degrees. *Internet Mathematics*, 6(4):489 – 522.

Bloch, F. and Jackson, M. O. (2007). The formation of networks with transfers among players. *Journal of Economic Theory*, 113(1):83 – 110.

Burt, R. S. (1995). *Structural Holes: The Social Structure of Competition*. Harvard University Press, Cambridge, MA.

Canay, I. A., Romano, J. P., and Shaikh, A. M. (2017). Randomization tests under an approximate symmetry assumption. *Econometrica*, 85(3):1013 – 1030.

Chamberlain, G. (1980). Analysis of covariance with qualitative data. *Review of Economic Studies*, 47(1):225 – 238.

Chamberlain, G. (1985). *Longitudinal Analysis of Labor Market Data*, chapter Heterogeneity, omitted variable bias, and duration dependence, pages 3 – 38. Cambridge University Press, Cambridge.

Chamberlain, G. (2010). Binary response models for panel data: identification and information. *Econometrica*, 78(1):159 – 168.

Charbonneau, K. B. (2017). Multiple fixed effects in binary response panel data models. *Econometrics Journal*, 20(3):S1 – S13.

Chatterjee, S., Diaconis, P., and Sly, A. (2011). Random graphs with a given degree sequence. *Annals of Applied Probability*, 21(4):1400 – 1435.

Chen, M., Fernández-Val, I., and Weidner, M. (2021). Nonlinear factor models for network and panel data - journal of econometrics, 2021. *Journal of Econometrics*, 220(2):296 – 324.

Chen, S. and Kaido, H. (2021). Robust tests of model incompleteness in the presence of nuisance parameters. Technical report, Boston University.

Chen, X., Christensen, T. M., and Tamer, E. (2018). Monte carlo confidence sets for identified sets. *Econometrica*, 86(6):1965 – 2018.

Christakis, N. A., Fowler, J. H., Imbens, G. W., and Kalyanaraman, K. (2020). *The Econometric Analysis of Network Data*, chapter An empirical model for strategic network formation, pages 123 – 148. Academic Press.

Ciliberto, F. and Tamer, E. (2009). Market structure and multiple equilibria in airline markets. *Econometrica*, 77(6):1791 – 1828.

Comola, M. and Fafchamps, M. (2014). Testing unilateral and bilateral link formation. *Economic Journal*, 124(579):954 – 975.

Cooper, C., Dyer, M., and Greenhill, C. (2007). Sampling regular graphs and a peer-to-peer network. *Combinatorics, Probability and Computing*, 16(4):557–593.

Czabarka, E., Szekely, L. A., Toroczkai, Z., and Walker, S. (2021). An algebraic monte-carlo algorithm for the partition adjacency matrix realization problem. *Algebraic Statistics*, 12(2):115–124.

de Paula, Á. (2020). Econometric models of network formation. *Annual Review of Economics*, 12:775 – 799.

de Paula, Á., Richards-Shubik, S., and Tamer, E. (2018). Identifying preferences in networks with bounded degree. *Econometrica*, 86(1):263 – 288.

De Weerdt, J. (2004). *Insurance Against Poverty*, chapter Risk-sharing and endogenous network formation, pages 197 – 216. Oxford University Press, Oxford.

Del Genio, C. I., Kim, H., Toroczkai, Z., and Bassler, K. E. (2010). Efficient and exact sampling of simple graphs with given arbitrary degree sequence. *Plos One*, 5(4):1–7.

Dzemski, A. (2018). An empirical model of dyadic link formation in a network with unobserved heterogeneity. *Review of Economics and Statistics*. University of Mannheim.

Erdos, P. L., Miklos, I., and Toroczkai, Z. (2018). New classes of degree sequences with fast mixing swap markov chain sampling. *Combinatorics, Probability and Computing*, 27(2):186 – 207.

Erdős, P., Hartke, S. G., van Iersel, L., and Miklós, I. (2017). Graph realizations constrained by skeleton graphs. *Electronic Journal of Combinatorics*, 24(2):1 – 18.

Ferguson, T. S. (1967). *Mathematical Statistics: A Decision Theoretic Approach*. Academic Press, New York.

Fernández-Val, I. and Weidner, M. (2016). Individual and time effects in nonlinear panel data models with large n, t. *Journal of Econometrics*, 192(1):291 – 312.

Freeman, L. (1977). A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35 – 41.

Fudenberg, D. and Tirole, J. (1998). *Game Theory*. The MIT Press, Cambridge, MA.

Gao, W. Y. (2020). Nonparametric identification in index models of link formation. *Journal of Econometrics*, 215(2):399 – 413.

Gotelli, N. A. (2000). Null model analysis of species co-occurrence patterns. *Ecology*, 81(9):2606 –2621.

Goyal, S. (2023). *Networks: An Economics Approach*. The MIT Press.

Goyal, S. and Vega-Redondo, F. (2007). Structural holes in social networks. *Journal of Economic Theory*, 137(1):460 – 492.

Graham, B. S. (2017). An econometric model of network formation with degree heterogeneity. *Econometrica*, 85(4):1033 – 1063.

Graham, B. S. (2020). *Handbook of Econometrics*, volume 7A, chapter Network data. North-Holland, Amsterdam.

Graham, B. S. and Pelican, A. (2020). *The Econometric Analysis of Network Data*, chapter Testing for externalities in network formation using simulation, pages 65 – 82. Academic Press, London.

Grandjean, M. (2016). A social network analysis of twitter: Mapping the digital humanities community. *Cogent Arts & Humanities*, 3(1):1171458.

Hoff, P. (2005). Bilinear mixed-effects models for dyadic data. *Journal of the American Statistical Association*, 100:286–295.

Holland, P. W. and Leinhardt, S. (1976). Local structure in social networks. *Sociological Methodology*, 7:1 – 45.

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press, Princeton.

Jackson, M. O., Rodriguez-Barraquer, T., and Tan, X. (2012). Social capital and social quilts: network patterns of favor exchange. *American Economic Review*, 102(5):1857–1897.

Jackson, M. O., Rogers, B. W., and Zenou, Y. (2017). The economic consequences of social-network structure. *Journal of Economic Literature*, 55(1):49 – 95.

Jackson, M. O. and Wolinsky, A. (1996). A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44 – 74.

Jia, P. (2008). What happens when wal-mart comes to town: an empirical analysis of the discount retailing industry. *Econometrica*, 76(6):1263 – 1316.

Jochmans, K. (2018). Semiparametric analysis of network formation. *Journal of Business and Economic Statistics*, 36(4):705 – 713.

Johnson, C. and Gilles, R. P. (2000). Spatial social networks. *Review of Economic Design*, 5(3):273 – 299.

Kaido, H. and Zhang, Y. (2019). Robust likelihood ratio tests for incomplete economic models. arXiv Working Paper arXiv:1910.04610v2 [econ.EM], Boston Univerity and Jinan University.

Kim, H., Genio, C. I. D., Bassler, K. E., and Toroczkai, Z. (2012). Constructing and sampling directed graphs with given degree sequences. *New Journal of Physics*, 14(2):023012.

Kleinberg, J., Suri, S., Tardos, É., and Wexler, T. (2008). Strategic network formation with structural holes. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 284 – 293.

Lehmann, E. L. and Romano, J. P. (2005). *Testing Statistical Hypotheses*. Springer, New York, 3rd edition.

Marotta, L., Miccichè, S., Fujiwara, Y., Iyetomi, H., Aoyama, H., Gallegati, M., and Mantegna, R. N. (2015). Bank-firm credit network in japan: an analysis of a bipartite network. *Plos One*, 10(5):e0123079.

McDonald, J. W., Smith, P. W. F., and Forster, J. J. (2007). Markov chain monte carlo exact inference for social networks. *Social Networks*, 29(1):127 – 136.

McFadden, D. (1974). *Frontiers in Econometrics*, chapter Conditional logit analysis of qualitative choice behavior, pages 105 – 142. Academic Press, New York.

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.

Mele, A. (2017). A structural model of dense network formation. *Econometrica*, 85(3):825 – 850. John Hopkins University.

Menzel, K. (2016). Strategic network formation with many agents. New York University.

Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D., and Alon, U. (2002). Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824 – 827.

Mitzenmacher, M. and Upfal, E. (2005). *Probability and Computing*. Cambridge University Press, Cambridge.

Miyauchi, Y. (2016). Structural estimation of a pairwise stable network with nonnegative externality. *Journal of Econometrics*, 195(2):224 – 235.

Molchanov, I. (2017). *Theory of Random Sets*, volume 87 of *Probability Theory and Stochastic Modelling*. Springer, 2nd edition.

Moreira, M. J. (2009). Tests with correct size when instruments can be arbitrarily weak. *Journal of Econometrics, 152(2), 131-140*, 152(2):131 – 140.

Newman, M. E. J. (2010). *Networks: An Introduction*. Oxford University Press, Oxford.

Oneal, J. R. and Russett, B. (1999). The kantian peace: the pacific benefits of democracy, interdependence, and international organizations. *World Politics*, 52(1):1 – 37.

Rao, A. R., Jana, R., and Bandyopadhyay, S. (1996). A markov chain monte carlo method for generating random (0,1)-matrices with given marginals. *Sankhya*, 58(2):225 – 242.

Sheng, S. (2020). A structural econometric analysis of network formation games through subnetworks. *Econometrica*, 88(5):1829 – 1858.

Sinclair, A. (1993). *Algorithms for Random Generation and Counting: A Markov Chain Approach*. Progress in Theoretical Computer Science. Birkhäuser, Boston.

Tao, T. (2016). An improved mcmc algorithm for generating random graphs from constrained distributions. *Network Science*, 4(1):117 – 139.

Tinbergen, J. (1962). *Shaping the World Economy: Suggestions for an International Economic Policy*. Twentieth Century Fund, New York.

Wasserman, S. and Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge.

Yan, T., Jiang, B., Fienberg, S. E., and Leng, C. (2018). Statistical inference in a directed network model with covariates. *Journal of the American Statistical Association*.