

An optimal test for strategic interaction in network formation games

Andrin Pelican and Bryan S. Graham *

July 15, 2026

*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 Erdős, 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 Weerd, 2004). Jackson et al. (2017) present many other examples of networks in economics. Such data abound in 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 (beyond the one at hand).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), which depends only on the attributes of the sending (ego) and receiving (alter) agents. Another example is associated with “degree heterogeneity”: agents may vary systematically in the utility they generate by directing links (out-degree heterogeneity), or in their attractiveness as link targets for others (in-degree heterogeneity). 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.

However, in some situations, 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

¹In 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“Other links” include those directed by the sending agent to targets other than the one at hand.

detection of such externalities is the subject of this paper.

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. This combinatoric complexity, renders methods pioneered for the econometric analysis of discrete games with just a few players ill suited for 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.³ 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 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.⁴

Unlike this prior work, we do not consider set identification in this paper, focusing instead

³de Paula (2020) surveys work in this area and provides additional references.

⁴We 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 the same solution/equilibrium concept.

on the more modest goal of externality *detection*. With this target question in mind, we introduce an econometric model of strategic network formation which simultaneously allows for (i) agents to value both direct and indirect links (i.e., a freely specified network benefit function), (ii) the systematic returns to link formation to vary with observed dyad attributes, and (iii) 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 Weerd, 2004). Early formal econometric analyses include those by Charbonneau (2017), Jochmans (2018) and Yan et al. (2018); work that built upon large- N , large- T panel data research (e.g., Fernández-Val and Weidner, 2016) and 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, allowing 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, we make no auxiliary equilibrium selection assumptions.⁵

Let K be the (finite) number of support points in the distribution of observed agent attributes and N the number of agents.⁶ Our model includes (i) K^2 “homophily” parameters, collected in the $K \times K$ matrix $\Lambda \stackrel{\text{def}}{=} [\lambda_{kl}]$ for $k, l = 1 \dots K$, capturing how link returns vary systematically with ego and alter attributes (we define $\lambda \stackrel{\text{def}}{=} \text{vec}(\Lambda')$), (ii) two $N \times 1$ parameter vectors $\mathbf{A} \stackrel{\text{def}}{=} [A_i]$ and $\mathbf{B} \stackrel{\text{def}}{=} [B_i]$ for $i = 1 \dots N$, capturing, respectively, agent-specific out- and in-degree heterogeneity, and (iii) a scalar parameter, γ , measuring the extent

⁵More 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).

⁶That K is finite with $K \ll N$ is a strong assumption; one we return to below.

to which agents value indirect links. Our model also includes (iv) an “model implied” network selection function (This function assigns probabilities to all game outcomes consistent with NE equilibria for every possible realization of the $n = N(N - 1)$ link-specific random utility shocks, $\mathbf{U} \stackrel{def}{=} [U_{ij}]_{1 \leq i \neq j \leq N}$). 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 network selection mechanism as a nuisance function. This focuses our attention solely on the strategic parameter 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 Erdős-Rényi model to generate the reference distribution. This invariably results in “straw man” tests since few real world networks are well-described by the Erdős-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. 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 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 homophily on observables as well

⁷Our 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.

as unobserved agent-specific ego (sender) and alter (receiver) effects, therefore imposes a more stringent test for the no strategic interaction null (than an $\ddot{\text{E}}\text{rdos-R\'{e}nyi}$ benchmark).

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 is not interpreted as evidence in favor of any particular alternative.

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.
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 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”. 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.

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 practicality, 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 Weerd (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 A 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) $\mathcal{V} = \{1, \dots, N\}$ and a set of ordered pairs of nodes, respectively called *egos* and *alters*, $\mathcal{A} = \{(i, j), (k, l), \dots\}$ 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

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

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 $M \stackrel{\text{def}}{=} 2^{N-1}$ possible actions, corresponding to all possible configurations of links she may direct towards her peers. A mixed strategy for agent i , $\sigma_i = (\pi_{1i}, \pi_{2i}, \dots, \pi_{Mi})'$, is a probability distribution on these M possible pure strategy se-

lections; $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_N)'$ 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

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

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

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

where X_i is a $K \times 1$ vector of mutually exclusive group membership indicators that *is observed* by the econometrician and $\mathbf{U}_i = (U_{i1}, \dots, U_{ii-1}, U_{ii+1}, \dots, U_{iN})'$ is agent i 's vector of idiosyncratic logistic preference shocks over the $N - 1$ possible links she can direct (and $\mathbf{U} = (\mathbf{U}'_1, \dots, \mathbf{U}'_N)'$).⁸ 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, $\nu_i(\sigma_i, \sigma_{-i}; \theta, \mathbf{U}_i)$ to denote agent i 's *expected* utility under the mixed strategy profile $\sigma = (\sigma_i, \sigma_{-i})$.

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 $g_i(\mathbf{d})$ is often called the *network benefit* function, while $c_{ij}(X_i, X_j; \delta, U_{ij})$ 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.⁹ 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: $-c_{ij}(X_i, X_j; \delta, U_{ij})$ is simply the portion of the payoff i gets from directing an edge to j that

⁸More 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).

⁹Johnson and Gilles (2000) study the implications of cost heterogeneity on equilibrium network structure in the “connections” model.

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 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 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”.¹⁰

The $X_i' \Lambda_0 X_j \stackrel{def}{=} W_{ij}' \lambda_0$ term allows baseline utility to depend on whether agents assortatively match on their attributes (we define $W_{ij} \stackrel{def}{=} (X_i \otimes X_j)$ and recall that $\lambda \stackrel{def}{=} \text{vec}(\Lambda')$). 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.¹¹


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

¹⁰Alternatively 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.

¹¹An 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.

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 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 $g_i(\mathbf{d}) = \sum_{j \neq i} \phi(\ell_{ij}(\tilde{\mathbf{d}}))$ where $\tilde{\mathbf{d}}$ is the undirected network obtained from $\mathbf{d}$ (i.e., $\tilde{\mathbf{d}} = [\tilde{d}_{ij}]$ with $\tilde{d}_{ij} = 1 - (1 - d_{ij})(1 - d_{ji})$), $\phi : \{1, 2, \dots, N-1\} \rightarrow \mathbb{R}$ is a known function with $\phi(k) > \phi(k+1) > 0$ for any $k = 1, 2, \dots, N-1$, and $\ell_{ij}(\tilde{\mathbf{d}})$ 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 .

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. Kleinberg et al. (2008) emphasize the special benefits of lying on length two paths between disconnected agents. 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 $g_i(\mathbf{d}) = \sum_j \sum_{k \neq j} \phi(d_{ki}d_{ij}(1 - d_{kj}), \sum_l d_{kl}d_{lj}(1 - d_{kj}))$ with $\phi(0, k) \equiv 0$ and $\phi(1, k) > \phi(1, k+1) > 0$ for $k = 1, \dots, N-2$. See Goyal and Vega-Redondo (2007) for a related model.¹²

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 $g_i(\mathbf{d}) = \sum_j d_{ij}(\sum_k d_{ki}d_{kj})$. 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 $g_i(\mathbf{d}) = \sum_j d_{ij}(\sum_k d_{ik}d_{kj})$.

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

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

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

¹²We 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

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

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

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

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. 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 $i = 1, \dots, N$,

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

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.¹³

Assumption 1.1. (DATA GENERATING PROCESS)

1. (NON-STOCHASTIC $\mathbf{X}$) Let $\mathcal{V} = \{1, \dots, N\}$ 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 $(\mathcal{V}, \mathbb{D}_N, \{\nu_i(\cdot, \cdot; \theta_0, \mathbf{U}_i)\}_{i \in \mathcal{V}})$.

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 due solely to variation in $\mathbf{U}$ and any randomness in NE selection (see below).

¹³Observe 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

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

for $i, j = 1, \dots, N$ 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 as the observed network. Our strategic network formation model is *incomplete*. Although we remain agnostic about equilibrium selection, it is nevertheless useful to establish notation for how the unknown selection process assigns probability weights to different (model implied) network configurations. Doing so 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 any vector of logistic link preference shocks, $\mathbf{u} \in \mathbb{R}^n$, a probability weight to network $\mathbf{d}$:

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

With such a function, we can write a likelihood of observing network $\mathbf{D} = \mathbf{d}$ as

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

where $f_{\mathbf{u}}(\mathbf{u}) = \prod_{i \neq j} f_U(u_{ij})$ with $f_U(u) = e^u / [1 + e^u]^2$. For (9) to be *the* likelihood, we require that $\mathcal{N}(\mathbf{d}, \cdot; \theta)$ is (i) consistent with the assumption of NE play and (ii) measurable (so (9) is well-defined).

We use ideas from Beresteanu et al. (2011) to ensure both requirements. Begin with the first requirement: that $\mathcal{N}$ obeys all restrictions arising from the assumption of NE play. Fix θ and let $\sigma^* = (\sigma_1^*, \sigma_2^*, \dots, \sigma_N^*)' \in \Sigma$ denote a mixed strategy NE (which necessarily satisfies (7) above at $\mathbf{U} = \mathbf{u}$).¹⁴ Here $\sigma_i^* = (\sigma_{i1}^*, \dots, \sigma_{ii-1}^*, \sigma_{ii+1}^*, \dots, \sigma_{iN}^*)'$ is the vector of (equilibrium) probabilities with which i sends links to her $N - 1$ peers.

Let $h_{\mathbf{d}} : \Sigma \rightarrow [0, 1]$ denote the function which assigns, for a given mixed strategy NE, the probability of realizing network $\mathbf{d} \in \mathbb{D}_N$:

$$h_{\mathbf{d}}(\sigma^*; \theta) = \prod_i \prod_{j \neq i} [\sigma_{ij}^*]^{d_{ij}} [1 - \sigma_{ij}^*]^{1-d_{ij}}. \quad (10)$$

If σ^* corresponds to a pure strategy NE, such that $\sigma_i^* = \mathbf{d}_i^* \in \mathbb{D}_N$ for all $i = 1, \dots, N$, then $h_{\mathbf{d}^*}(\mathbf{d}^*; \theta) = 1$ and $h_{\mathbf{d}}(\mathbf{d}^*; \theta) = 0$ for all $\mathbf{d} \neq \mathbf{d}^*$. Otherwise, the available probability is spread across all networks in the support of σ^* according to (10). We call $h_{\mathbf{d}}(\sigma^*; \theta)$ the *NE weight* assigned to network $\mathbf{d}$ under equilibrium σ^* .

Next let $\xi(\mathbf{u})$ denote the NE actually selected (whether by agents or “nature”) when taste shock $\mathbf{U} = \mathbf{u}$ is realized. We define $\mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) = h_{\mathbf{d}}(\xi(\mathbf{u}); \theta)$. That is, $\mathcal{N}(\mathbf{d}, \cdot; \theta)$ is the

¹⁴Note that the set of available NE generally varies with $\mathbf{u}$.

composition of $h_{\mathbf{d}}$ and the (unknown) rule which picks out a NE from those available when $\mathbf{U} = \mathbf{u}$.

Observe that $\mathcal{N}$ does not assign probabilities to the selection of particular NE, instead it assigns weights directly to possible networks. However, by construction, these weights are consistent with NE play.

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

Theorem 1.1 follows by verifying the existence of a measurable selection which assigns to each possible taste shock $\mathbf{u} \in \mathbb{R}^n$ a particular NE (from all those available). The proof, which follows directly from ideas in Beresteanu et al. (2011), can be found in Appendix A.1.

To understand $\mathcal{N}$, we can imagine the following ‘as if’ process. First, some rule selects an equilibrium from all the pure and mixed strategy NE available when $\mathbf{U} = \mathbf{u}$ (a “measurable selection”). Associated with the selected NE is a probability distribution over networks. If a NE in pure strategies is selected, this distribution is degenerate: all mass is placed on the equilibrium in pure strategies. Selection of mixed strategy NE, in contrast, induces a non-degenerate probability distribution across all networks in its support. Regardless Theorem 1.1 implies that the expectation of $\mathcal{N}(\mathbf{d}, \mathbf{U}; \theta)$, and hence the likelihood, is well-defined.

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, Theorem 14.6.1, p. 619), 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 decide 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

$$\Theta_0 = \{(\gamma, \delta') : \gamma = 0, \delta \in \Delta\}. \quad (11)$$

Our null hypothesis is the *composite* one:

$$H_0 : \theta \in \Theta_0 \quad (12)$$

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

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

$$\begin{aligned} \mathcal{N}_0(\mathbf{d}, \mathbf{u}; \theta) &= \prod_i \prod_j \mathbf{1}(A_i + B_j + W'_{ij}\lambda \geq u_{ij})^{d_{ij}} \\ &\quad \times \mathbf{1}(A_i + B_j + W'_{ij}\lambda < u_{ij})^{1-d_{ij}}. \end{aligned}$$

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

$$P_0(\mathbf{d}; \delta) = \prod_{i=1}^N \prod_{j \neq i} \left[\frac{\exp(W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B})}{1 + \exp(W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B})} \right]^{d_{ij}} \times \left[\frac{1}{1 + \exp(W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B})} \right]^{1-d_{ij}}, \quad (13)$$

where R_i is the $N \times 1$ vector with a 1 in its i^{th} element and zeros elsewhere.¹⁵

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:

$$\mathbf{S} = \begin{pmatrix} \mathbf{S}_{\text{out}} \\ \mathbf{S}_{\text{in}} \end{pmatrix} = \begin{pmatrix} D_{1+}, \dots, D_{N+} \\ D_{+1}, \dots, D_{+N} \end{pmatrix}. \quad (14)$$

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

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

$$\mathbf{M} = \sum_i \sum_j D_{ij} X_i X'_j. \quad (15)$$

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

$$\mathbb{D}_{\mathbf{S}, \mathbf{M}} \stackrel{\text{def}}{=} \left\{ \mathbf{d} \in \mathbb{D}_N : (d_{1+}, \dots, d_{N+}) = \mathbf{S}_{\text{out}}, \right. \\ \left. (d_{+1}, \dots, d_{+N}) = \mathbf{S}_{\text{in}}, \sum_i \sum_j d_{ij} X_i X'_j = \mathbf{M} \right\}. \quad (16)$$

¹⁵Variants of this likelihood are analyzed by Chatterjee et al. (2011), Charbonneau (2017), Graham (2017), Jochmans (2018), Dzinski (2018) and Yan et al. (2018). See also Fernández-Val and Weidner (2016).

Let $\mathbf{T} = (\text{vec}(\mathbf{M}')', \mathbf{S}_{\text{out}}, \mathbf{S}_{\text{in}})'$. 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 ,¹⁶ (ii) its probability mass function (13) can be rewritten as:

$$P_0(\mathbf{d}; \delta) = c(\delta) \exp(\mathbf{t}'\delta), \quad \delta \in \Delta, \quad \mathbf{d} \in \mathbb{D}_N, \quad (17)$$

$$\text{with } c(\delta) \stackrel{\text{def}}{=} \left[\sum_{\mathbf{d} \in \mathbb{D}_N} \exp \left(\left[\sum_{i=1}^N \sum_{j \neq i} d_{ij} \{W'_{ij}\lambda + R'_i \mathbf{A} + R'_j \mathbf{B}\} \right] \right) \right]^{-1}.$$

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

$$P_0(\mathbf{d} | \mathbf{T} = \mathbf{t}) = \frac{P_0(\mathbf{d}; \delta)}{\sum_{\mathbf{v} \in \mathbb{D}_{\mathbf{S}, \mathbf{M}}} P_0(\mathbf{v}; \delta)} = \frac{1}{|\mathbb{D}_{\mathbf{S}, \mathbf{M}}|} \quad (18)$$

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 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 $\mathbb{D}_{\mathbf{S}, \mathbf{M}}$:

$$\Pr(R(\mathbf{D}) \leq r | \mathbf{T}; \theta \in \Theta_0) = \frac{1}{|\mathbb{D}_{\mathbf{S}, \mathbf{M}}|} \sum_{\mathbf{d} \in \mathbb{D}_{\mathbf{S}, \mathbf{M}}} \mathbf{1}(R(\mathbf{d}) \leq r). \quad (19)$$

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.

¹⁶A sufficient statistic is (boundedly) complete if, for all (bounded) functions $k(\mathbf{t})$, $\mathbb{E}_0[k(\mathbf{T})] = 0$ 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).

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:

$$\sup_{\theta \in \Theta_0} \mathbb{E}_\theta [\phi(\mathbf{D})] = \sup_{\gamma=\gamma_0, \delta \in \Delta} \mathbb{E}_\theta [\phi(\mathbf{D})] = \alpha. \quad (20)$$

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.¹⁷ 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 size distortion as δ varies 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).

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

$$\mathbb{E}_\theta [\phi(\mathbf{D}) | \mathbf{T} = \mathbf{t}] = \alpha. \quad (21)$$

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

$$\mathbb{E}_\theta [\phi(\mathbf{D})] = \mathbb{E}_\theta [\mathbb{E}_\theta [\phi(\mathbf{D}) | \mathbf{T}]] = \alpha \quad (22)$$

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 (21), because our model completely specifies the distribution of networks conditional on $\mathbf{T} = \mathbf{t}$ under the null. Condi-

¹⁷Jackson et al. (2012) suggest the fraction of links in the network which are supported.

tion (22) 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 (21) 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). $\square$

Implementation

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

$$R(\mathbf{D}) = \frac{2\hat{P}_{11}}{2\hat{P}_{11} + \hat{P}_{01}}, \quad (23)$$

where

$$\hat{P}_{01} = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^N [D_{ij}(1 - D_{ji}) + (1 - D_{ij})D_{ji}] \quad (24)$$

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

$$\hat{P}_{11} = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^N D_{ij}D_{ji} \quad (25)$$

the fraction which take a reciprocated or “mutual” configuration.

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

$$\phi(\mathbf{d}) = \begin{cases} 1 & R(\mathbf{d}) > c_\alpha(\mathbf{t}) \\ g_\alpha(\mathbf{t}) & R(\mathbf{d}) = c_\alpha(\mathbf{t}) \\ 0 & R(\mathbf{d}) < c_\alpha(\mathbf{t}) \end{cases} \quad (26)$$

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:

$$c_\alpha(\mathbf{T}) = \min \left\{ c \in \mathbb{R} : \frac{1}{|\mathbb{D}_{\mathbf{S}, \mathbf{M}}|} \sum_{\mathbf{d} \in \mathbb{D}_{\mathbf{S}, \mathbf{M}}} \mathbf{1}(R(\mathbf{d}) > c) \leq \alpha \right\}. \quad (27)$$

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 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.¹⁸ Therefore a method of approximating the exact null distribution is required. The simulation algorithm introduced in Section 3 below provides 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

¹⁸In 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, as well as some examples, from a related setting.

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 $\mathbb{E}_\theta[\phi(\mathbf{D}) | \mathbf{T} = \mathbf{t}] = \alpha$ for all $\theta \in \Theta_0$. 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 get¹⁹

$$\left. \frac{\partial \beta(\gamma, \mathbf{t})}{\partial \gamma} \right|_{\gamma=0} = \mathbb{E}[\phi(\mathbf{D}) \mathbb{S}_\gamma(\mathbf{D} | \mathbf{T}; \theta) | \mathbf{T} = \mathbf{t}] \quad (28)$$

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 the population value δ_0 .) By the Neyman-Pearson lemma, the test with the critical function given by equation (26) above, where the test statistic, $R(\mathbf{d})$, is set equal to the log-likelihood gradient, $\left. \frac{1}{P_0(\mathbf{d}; \delta)} \frac{\partial P(\mathbf{d}; \theta)}{\partial \gamma} \right|_{\gamma=0}$, 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 $\left. \frac{1}{P_0(\mathbf{d}; \delta)} \frac{\partial P(\mathbf{d}; \theta)}{\partial \gamma} \right|_{\gamma=0}$. 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.

¹⁹Differentiability of the likelihood function is formally established by Theorem 2.3 below.

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

$$\left. \frac{\partial P(\mathbf{d}; \theta, \mathcal{N})}{\partial \gamma} \right|_{\gamma=0} = P_0(\mathbf{d}; \delta) \times \left[\sum_{i \neq j} s_{ij}(\mathbf{d}) \left\{ d_{ij} \frac{f_U(\mu_{ij})}{\int_{-\infty}^{v_{ij}} f_U(u) du} - (1 - d_{ij}) \frac{f_U(\mu_{ij})}{\int_{v_{ij}}^{\infty} f_U(u) du} \right\} \right], \quad (29)$$

where $\mu_{ij} = A_i + B_j + X_j' \Lambda_0 X_i$ equals the systematic, non-strategic, component of utility generated by arc ij and 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 $\mathcal{N}$, and hence 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 still unique in this scenario so 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 (29) to:

$$\frac{1}{P_0(\mathbf{d}; \delta)} \left. \frac{\partial P(\mathbf{d}; \theta)}{\partial \gamma} \right|_{\gamma=0} = \sum_{i \neq j} [d_{ij} - F_U(\mu_{ij})] s_{ij}(\mathbf{d}) \quad (30)$$

where $F_U(u) = e^u / [1 + e^u]$ 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 $s_{ij}(\mathbf{d}) = d_{ji}$, as would be true when agents’ have a taste for reciprocated links. Observe that $F_U(\mu_{ij})$ 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, give 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 does not have to 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.²⁰ 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 than tests based on ad hoc statistics. This is true even in the “sparse” designs we consider.

²⁰We thank a referee for raising this concern.

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 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 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.²¹

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. Czaparka 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 $\text{NP} \neq \text{P}$), using a Blitzstein and Diaconis (2011) type algorithm to draw from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$ is not feasible.

²¹See 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.

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.

3.1 Notation and definitions

We start by defining an alternating walk.

Definition 3.1. (ALTERNATING WALK) An alternating walk H is a sequence of (ordered) dyads of the form

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

or

$$H := (i_2, i_1), (i_2, i_3), (i_4, i_3), \dots, (i_{l-1}, i_l) \quad (32)$$

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 $k = 2, \dots, l-1$.

For brevity we will often refer to a walk simply by its node sequence, writing $H := i_1 i_2, \dots, i_l$.

To unpack Definition 3.1 it is easiest to consider an example. Observe that for $H := i_1 i_2, \dots, i_l$, the adjacency matrix entries $D_{i_1 i_2}, D_{i_3 i_2}, \dots, D_{i_l i_{l-1}}$ 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.²² 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.²³

²²This description is essentially due to (Tao, 2016, p. 124).

²³Figure 6, Panel B, in the Supplemental Web Appendix provides some additional examples of alternating walks.

Figure 2: Constructing an alternating walk

A: Alternating Walk										B: Degree Sequence		
	a	b	c	d	e	f	g	h	i	j	Indegree	Outdegree
a	0	1	0	0	1	0	0	0	0	0	0	2
b	0	0	0	0	0	0	0	0	0	0	1	0
c	0	0	0	1	0	0	0	1	0	0	2	2
d	0	0	0	0	0	0	0	0	0	0	1	0
e	0	0	1	0	0	0	0	0	0	0	1	1
f	0	0	0	0	0	0	0	1	1	0	0	2
g	0	0	0	0	0	0	0	0	0	0	2	0
h	0	0	1	0	0	0	1	0	0	0	2	2
i	0	0	0	0	0	0	0	0	0	0	2	0
j	0	0	0	0	0	0	1	0	1	0	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 6 of the Supplemental Materials. 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 in-degree and out-degree of each agent in the network. Observe that in active steps the probability of any feasible choice equals the inverse of the out-degree of the current agent. In passive steps the probability of any feasible choice equals the inverse of the number of nodes minus the in-degree 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.

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 $D_{i_1 i_2}, D_{i_3 i_1}, \dots, D_{i_l i_{l-1}}$ 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).

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” (unused, see Algorithm 2) 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*, denoted - in a small abuse of notation - by say 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,

²⁴The 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 6 in the Supplemental Materials for an example.

$C := abcd$, 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.

We call a sequence of schlaufen $\mathcal{R} = (R_1, \dots, R_k)$ 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.

Appendix A.4, provides additional discussion of the MCMC algorithm as well as proof of its correctness.

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.

Theorem 3.1. *Algorithm 1 is a random walk on the state graph Φ which samples uniformly a network from $\mathbb{D}_{\mathbf{s},\mathbf{m}}$ for $\tau \rightarrow \infty$.*

Proof. See Appendix A.4. □

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

De Weerd (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 3.

Modeling the configuration of links shown in Figure 3 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 Weerd (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

$$g_i(\mathbf{d}) = \sum_j \sum_{k \neq j} \frac{D_{ki} D_{ij} (1 - D_{kj})}{\max(1, \sum_l D_{kl} D_{lj} (1 - D_{kj}))}. \quad (33)$$

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

²⁵The prompt used by De Weerd (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.

between j and k , they each get half the benefit and so on. The marginal network benefit of edge ij is thus

$$s_{ij}(\mathbf{d}) = \sum_{k \neq j} \frac{D_{ki} D_{ij} (1 - D_{kj})}{\max(1, \sum_l D_{kl} D_{lj} (1 - D_{kj}))}, \quad (34)$$

from which the form of the locally best test follows.

From De Weerd (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:

$$C_i^{\text{BC}}(\tilde{\mathbf{d}}) = \sum_{j < k} \frac{\phi_i^{jk}(\tilde{\mathbf{d}})}{\phi^{jk}(\tilde{\mathbf{d}})}. \quad (35)$$

Equation (35) 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 $\binom{N-1}{2} = \frac{(N-1)(N-2)}{2}$.

Figure 3: Nyakatoke Village Risk-Sharing Network



Source: De Weerd (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.

The “90-50 gap” in the empirical distribution of $C_i^{\text{BC}}(\tilde{\mathbf{d}})$ 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 it gives precision to the alternative of interest (in turn suggesting a suitable, in fact, optimal test statistic).²⁶


The left panel of Figure 4 plots simulation estimates of the distribution of the 90 – 50 betweenness-centrality gap across three reference sets of networks: (i) $\ddot{\text{E}}$ rdos-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 $\ddot{\text{E}}$ rdos-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

²⁶Of 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.

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

Figure 4: Testing for bridging/brokerage preferences



Source: De Weerd (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.

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 4 plots the null distribution of the locally best test statistic for the alternative that households gain utility by bridging disconnected pairs of agents (as 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 4, 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

²⁸The 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.

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 Jupyter 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 unknown NE selection process, 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.

Our set-up assumes that the 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

²⁹We thank the referees for this observation.

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 structure 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 in-degree of agent j . Hence the distribution 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

³⁰For 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.

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

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’ tool-kits (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.

Data Availability Statement

The data used in this paper were extracted from a larger dataset collected by Joachim de Weerd. A basic reference is De Weerd (2004). The raw data files used to construct our “empirical dataset”, as well as processing code, are included in the replication package associated with this paper. The exact files used to produce Figure 4 above, as well as all associated computer code, may also be found in the replication archive: <https://doi.org/10.5281/zenodo.19598770>. We certify that the author(s) of the manuscript have legitimate access to and permission to use the data used in this manuscript.

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.
- 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.
- 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*, 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.

- 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.

- 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*.

A Supplemental Web Appendix

The appendix includes proofs of the theorems stated in the main text, as well as statements and proofs of supplemental lemmata. All notation is as established in the main text, unless stated otherwise. Equation numbering continues in sequence with that established in the main text.

In addition to proofs, Section A.5 of the Appendix summarizes the results of a small set of Monte Carlo experiments and Section A.6 discusses additional applications of our MCMC algorithm.

A.1 Measurability of the likelihood

For the equation (9) to be well-defined we must show that $\mathcal{N}(\mathbf{d}, \cdot; \theta)$ is measurable. For any network $\mathbf{d}$ we can define a function $\mathcal{N}(\mathbf{d}, \cdot; \theta)$, which assigns to the realization $\mathbf{U} = \mathbf{u}$ a probability weight for the pure strategy combination which corresponds to $\mathbf{d}$. We now show that there is a measurable function $\mathcal{N}(\mathbf{d}, \cdot; \theta)$ satisfying these conditions.

In order to show measurability, we use random set theory. Let X be the set of NE in the game. This set depends on the random realization of the taste shock. We now use Proposition 3.1 of Beresteanu et al. (2011) to show that X is a random closed set. We verify that assumptions (i) to (iii) of their Proposition 3.1 are fulfilled.

- **(i)** The game is a network formation game, which implies the finiteness of the outcome set and that each player has only a finite number of strategies. This ensures the first assumption.
- **(ii)** The second assumption states that the observed outcome of the game results from static, simultaneous-move Nash play, which is ensured by the structure of our game's setup.
- **(iii)** The payoff functions are known for a given strategy and, without loss of generality, can be assumed to be normalized. The payoff functions are continuous in covariates and the taste shocks, which follows directly from equation (2). Without loss of generality, we can assume the parameters belong to a compact subset of the real cube, as the parameters are indicator vectors, constants, and fixed for the test.

With these assumptions, Proposition 3.1 states that X is a random closed set.

We know from Nash's existence theorem that for each game there exists a NE. Therefore X is nonempty. The set of NE is a closed set. We therefore can apply Molchanov's (2017) Fundamental Selection Theorem (Theorem 1.4.1 on p. 77) and find a measurable selection $\xi : \mathbb{R}^n \rightarrow \Sigma$ which assigns to each taste shock $\mathbf{u}$, a NE.

Let $h_{\mathbf{d}} : \Sigma \rightarrow [0, 1]$ be the function which assigns to every mixed strategy the probability of $\mathbf{d}$ by multiplying the mixed strategies weights corresponding to $\mathbf{d}$. Since multiplication is a measurable operation $h_{\mathbf{d}}$ is measurable and $\mathcal{N}(\mathbf{d}, \cdot; \theta) := h_{\mathbf{d}} \circ \xi$ satisfies the desired properties.

A.2 Proof of Lemma 2.1

Let $c(\delta) = \left[\prod_{i=1}^N \prod_{j \neq i} (1 + \exp(W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B})) \right]^{-1}$; algebraic manipulation of (13) in the main text yields

$$\begin{aligned}
P_0(\mathbf{d}; \delta) &= \prod_{i=1}^N \prod_{j \neq i} \left[\frac{\exp(W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B})}{1 + \exp(W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B})} \right]^{d_{ij}} \times \left[\frac{1}{1 + \exp(W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B})} \right]^{1-d_{ij}} \\
&= \prod_{i=1}^N \prod_{j \neq i} \frac{\exp(W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B})^{d_{ij}}}{1 + \exp(W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B})} \\
&= c(\delta) \exp \left(\left[\sum_{i=1}^N \sum_{j \neq i} d_{ij} W_{ij} \right]' \lambda \right) \\
&\quad \times \exp \left(\left[\sum_{i=1}^N \sum_{j \neq i} d_{ij} R_i \right]' \mathbf{A} \right) \exp \left(\left[\sum_{i=1}^N \sum_{j \neq i} d_{ij} R_j \right]' \mathbf{B} \right) \\
&= c(\delta) \exp(\text{vec}(\mathbf{m}')' \lambda) \exp(\mathbf{s}'_{\text{out}} \mathbf{A}) \exp(\mathbf{s}'_{\text{in}} \mathbf{B}) \\
&= c(\delta) \exp(\mathbf{t}' \delta),
\end{aligned}$$

which belongs to the exponential family of distributions with a sufficient statistic for δ of $\mathbf{T}$ (see, for example, Definition 1 in Ferguson (1967, p. 126)).

That $\mathbf{T}$ is a boundedly complete sufficient statistic for δ follows from the observation that, if $k(\mathbf{t})$ is a bounded real-valued function for which $\mathbb{E}_0[k(\mathbf{T})] = 0$ for all $\delta \in \Delta$, then

$$c(\delta) \sum_{\mathbf{t} \in \mathbf{T}} k(\mathbf{t}) \exp(\mathbf{t}' \delta) \equiv 0, \text{ for all } \delta \in \Delta. \quad (36)$$

Since the range of the exponential function is the set of all positive real numbers and $c(\delta) > 0$ for all $\delta \in \Delta$, we have $\mathbb{E}_0[k(\mathbf{T})] = 0$ for all $\delta \in \Delta$ implies that $\Pr_0(k(\mathbf{T}) = 0) = 1$ (recall

that Δ includes an open rectangle, hence if the summands in (36) exactly sum to zero for a particular δ in the open rectangle, they will not do so for adjacent values of δ by continuity). We have that $\mathbf{T}$ is a complete sufficient statistic for δ . Minimal sufficiency follows from completeness. See Ferguson (1967, Section 1, pp. 132 - 136) for additional background and related results.

To get the form of $c(\delta)$ given in the statement of the Lemma observe that

$$\begin{aligned} c(\delta)^{-1} &= \left[\prod_{i=1}^N \prod_{j \neq i} (1 + \exp(W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B})) \right] \\ &= \left[\prod_{i=1}^N \prod_{j \neq i} (\exp(0) + \exp(W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B})) \right] \\ &= \sum_{d \in \mathbb{D}_N} \exp \left(\left[\sum_{i=1}^N \sum_{j \neq i} d_{ij} \{W'_{ij}\lambda + R'_i\mathbf{A} + R'_j\mathbf{B}\} \right] \right). \end{aligned}$$

A.3 Derivation of locally best test statistic

We begin with a high level overview of our argument; with a formal proof immediately following.

One feature of $s_{ij}(\mathbf{d})$, which will prove central to our analysis, is that it has finite range. To see this observe that since the set of all networks $\mathbb{D}_N$ is finite, $s_{ij}(\mathbf{d})$ also takes only a finite number of values. Let $\mathbb{S} = \{\underline{s}, s_1, \dots, s_M, \bar{s}\}$ be the set of possible values for $s_{ij}(\mathbf{d})$, ordered from smallest to largest.

An example illustrates. If $s_{ij}(\mathbf{d}) = d_{ji}$, as occurs when agents prefer reciprocated links, we have $\mathbb{S} = \{0, 1\}$. If $s_{ij}(\mathbf{d}) = \sum_k d_{ki}d_{kj}$, as when agents prefer supported links, we have $\mathbb{S} = \{0, 1, \dots, N-2\}$. Finiteness of the cardinality of $\mathbb{S}$ (for a given N) plays an important role in our analysis, as will become apparent below.

To understand the likelihood (9) it is helpful to consider a (relatively) simple example. This example will also help in understanding our derivation of the optimal test statistic below. Assume that $s_{ij}(\mathbf{d}) = d_{ji}$ such that agents prefer reciprocated links when $\gamma > 0$. In this example $s_{ij}(\mathbf{d})$ equals either zero (j does not reciprocate) or one (j does reciprocate). We can use the two elements of $\mathbb{S}$ to partition the real line into what we will call *buckets*:

$$\mathbb{R} = (-\infty, \mu_{ij}] \cup (\mu_{ij}, \mu_{ij} + \gamma] \cup (\mu_{ij} + \gamma, \infty). \quad (37)$$

Here $\mu_{ij} = A_i + B_j + X'_j\Lambda_0X_i$ equals the systematic component of baseline utility generated by arc ij . Next consider the realization of U_{ij} , the idiosyncratic utility agent i gets when she directs a link to j . If U_{ij} falls into the first bucket in (37), then agent i will always direct a

link to j ; irrespective of whether j chooses to direct a link to i or not. If U_{ij} falls into the middle or *inner* bucket, however, then i will direct a link to j only if j reciprocates. Finally, if U_{ij} falls into the last bucket, then i will never direct a link to j regardless of whether j directs a link to i or not. We will call the first and last buckets in (37) *outer* buckets.

If both U_{ij} and U_{ji} fall in their respective *inner* buckets, then the $\{i, j\}$ dyad can either take the empty ($D_{ij} = D_{ji} = 0$) or reciprocated ($D_{ij} = D_{ji} = 1$) configuration in equilibrium. In contrast, if *either* U_{ij} or U_{ji} falls into an outer bucket, then the $\{i, j\}$ dyad's wiring is uniquely determined. For example if U_{ij} is in the first outer bucket and U_{ji} is in the inner bucket, then the $\{i, j\}$ dyad will take the reciprocated form with probability one. It is a strictly dominant strategy for i to direct an link to j in this case and a best response for j to reciprocate.

For $\mathbf{U} = \mathbf{u}$ and $\theta = \theta_0$, let $J(\mathbf{u}; \theta_0) \leq \binom{N}{2}$ equal the number of dyads $\{i, j\}$, where both u_{ij} and u_{ji} fall into their inner bucket. For each of these dyads both the empty and reciprocated configuration is an equilibrium outcome. There are therefore $2^{J(\mathbf{u}; \theta_0)}$ equilibrium networks in this case; the $\mathcal{N}(\mathbf{d}, \mathbf{u}; \theta_0)$ function would assign some probability between zero and one to each of these $2^{J(\mathbf{u}; \theta_0)}$ networks (summing to one in total).

Recall that $\mathbb{S} = \{\underline{s}, s_1, \dots, s_M, \bar{s}\}$ equals the possible values of $s_{ij}(\mathbf{d})$, arranged from smallest to largest. We can use these support points to partition $\mathbb{R}$ into a set of intervals $\mathbb{B}$:

$$\mathbb{R} = (-\infty, \mu_{ij} + \gamma \underline{s}] \cup (\mu_{ij} + \gamma \underline{s}, \mu_{ij} + \gamma s_1] \cup \dots \cup (\mu_{ij} + \gamma s_M, \mu_{ij} + \gamma \bar{s}] \cup (\mu_{ij} + \gamma \bar{s}, \infty). \quad (38)$$

The elements of $\mathbb{B}$, called *buckets*, correspond to the intervals listed in (38). In principle we should write $\mathbb{B}_{ij}$ instead of $\mathbb{B}$, reflecting the dependence of the bucket definitions on the value of μ_{ij} , the systematic non-strategic utility associated with an i -to- j link. However, since this dependence is not essential to any of the arguments that follow we leave it implicit. Note that the cardinality of $\mathbb{B}$ does not depend on μ_{ij} , but instead equals $|\mathbb{S}| + 1$.

Agent i 's linking behavior vis-a-vis j depends on which bucket U_{ij} falls into. For $B \in \mathbb{B}$, if $U_{ij} \in B$, then we say U_{ij} is in, or falls into, bucket B . The first and last buckets, respectively $(-\infty, \mu_{ij} + \gamma \underline{s}]$ and $(\mu_{ij} + \gamma \bar{s}, \infty)$, play an important role in our argument. We call these two buckets *outer buckets*. The rest of the buckets we call *inner buckets*.

If U_{ij} falls into one of these outer buckets then player i has a pure strategy for d_{ij} which is strictly dominating. Specifically if U_{ij} falls into the lowest bucket, then i will direct an link to j regardless of what actions are taken by the other agents in the network. The marginal utility generated by link ij is so large that it remains positive across all possible configurations of the rest of the network; hence i always chooses to direct an link to j .

If, instead, U_{ij} falls into the highest bucket, then i will never direct an link to j . In this case the marginal utility associated with link ij is so low that it remains negative across all possible configurations of the rest of the network; hence i never chooses to direct a link to j .

Finally, if U_{ij} falls into an inner bucket, say $(\mu_{ij} + \gamma s_m, \mu_{ij} + \gamma s_{m+1}]$, then agent i 's optimal choice for d_{ij} is contingent upon the linking behavior of other agents. If other agents' link actions are such that $s_{ij}(\mathbf{d}) \geq s_m$, then it is a best response for i to link with j , but not otherwise.

The vector of idiosyncratic taste shocks, $\mathbf{U}$ contains $n = N(N-1)$ elements; one for each possible arc. Let the boldface subscripts $\mathbf{i} = \mathbf{1}, \mathbf{2}, \dots$ index these potential arcs in arbitrary order (e.g., $\mathbf{i}$ maps to some ij and vice-versa). Let $\mathbf{b} \in \mathbb{B}^n \stackrel{\text{def}}{=} \mathbb{B} \times \dots \times \mathbb{B}$ and $\mathbf{U} = (U_1, \dots, U_n)'$; we have that $\mathbf{U} \in \mathbf{b}$ for $\mathbf{b} \in \mathbb{B}^n$ so that each element of $\mathbf{u}$ falls into a bucket.

With the above notation established we can rewrite the likelihood (9) as:

$$P(\mathbf{d}; \theta, \mathcal{N}) = \sum_{\mathbf{b} \in \mathbb{B}^n} \int_{\mathbf{u} \in \mathbf{b}} \mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} \quad (39)$$

Expression (39) suggests a derivation by cases approach to finding $\left. \frac{\partial P(\mathbf{d}; \theta, \mathcal{N})}{\partial \gamma} \right|_{\gamma=0}$. Fortunately a brute force exhaustive approach is not required because it is possible to show that most of the summands in (39) do not influence the derivative at $\gamma = 0$.

Let $\tilde{\mathbb{B}}^n$ be the set of bucket configurations with at least two inner buckets. If at least two elements of $\mathbf{U}$ fall in inner buckets, then we have that $\mathbf{U} \in \mathbf{b}$ with $\mathbf{b} \in \tilde{\mathbb{B}}^n$. If, instead, at most one element of $\mathbf{U}$ falls in an inner bucket, then we have that $\mathbf{U} \in \mathbf{b}$ with $\mathbf{b} \in \mathbb{B}^n \setminus \tilde{\mathbb{B}}^n$. This set-up gives the likelihood decomposition:

$$P(\mathbf{d}; \theta, \mathcal{N}) = \tilde{P}(\mathbf{d}; \theta, \mathcal{N}) + Q(\mathbf{d}; \theta, \mathcal{N}), \quad (40)$$

with

$$\tilde{P}(\mathbf{d}; \theta, \mathcal{N}) = \sum_{\mathbf{b} \in \mathbb{B}^n \setminus \tilde{\mathbb{B}}^n} \int_{\mathbf{u} \in \mathbf{b}} \mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} \quad (41)$$

$$Q(\mathbf{d}; \theta, \mathcal{N}) = \sum_{\mathbf{b} \in \tilde{\mathbb{B}}^n} \int_{\mathbf{u} \in \mathbf{b}} \mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u}. \quad (42)$$

To prove Theorem 2.3 we show that for $\gamma \rightarrow 0$

$$P(\mathbf{d}; \theta, \mathcal{N}) = \tilde{P}(\mathbf{d}; \theta, \mathcal{N}) + \mathcal{O}(\gamma^2). \quad (43)$$

Intuitively, this follows from the fact that the chance that two or more elements of $\mathbf{U}$ fall in

inner buckets is negligible when γ is close to zero (because most of the probability mass for U_{ij} is contained in the two outer buckets when strategic interactions are small). Hence when calculating the optimal test statistic we are free to focus on the cases where either all, or all but one, of the elements of $\mathbf{U}$ fall in outer buckets. We can then show that

$$\left. \frac{\partial P(\mathbf{d}; \theta, \mathcal{N})}{\partial \gamma} \right|_{\gamma=0} = \left. \frac{\partial \tilde{P}(\mathbf{d}; \theta, \mathcal{N})}{\partial \gamma} \right|_{\gamma=0}. \quad (44)$$

Hence to derive the form of $\left. \frac{\partial P(\mathbf{d}; \theta, \mathcal{N})}{\partial \gamma} \right|_{\gamma=0}$ we need only calculate $\left. \frac{\partial \tilde{P}(\mathbf{d}; \theta, \mathcal{N})}{\partial \gamma} \right|_{\gamma=0}$. This calculation is non-trivial, but doable. Details of this calculation are provided in the proof.

Preliminary results

Lemma A.1. *Any differentiable function $f \in \mathcal{O}(\gamma^2)$ with $f(0) = 0$ has a derivative of zero at point zero.*

Proof. For $f \in \mathcal{O}(\gamma^2)$ we have, for some $C > 0$ and $\epsilon > 0$, that

$$|f(\gamma)| < C\gamma^2 \quad (45)$$

for all $\gamma \in [-\epsilon, \epsilon]$. The derivative of f at $\gamma = 0$ equals

$$f'(0) = \lim_{\gamma \rightarrow 0} \frac{f(\gamma) - f(0)}{\gamma} = \lim_{\gamma \rightarrow 0} \frac{f(\gamma)}{\gamma}, \quad (46)$$

with the second equality because $f(0) = 0$. As $\gamma \rightarrow 0$, we will have $\gamma < \epsilon$ so that

$$f'(0) = \lim_{\gamma \rightarrow 0} \frac{f(\gamma)}{\gamma} \leq \lim_{\gamma \rightarrow 0} \frac{C\gamma^2}{\gamma} = \lim_{\gamma \rightarrow 0} C\gamma \quad (47)$$

which goes to zero as $\gamma \rightarrow 0$ as claimed. $\square$

Proof of Theorem 2.3

We begin with the likelihood decomposition (39) given above. The number of summands in (39) depends on the partition that $s_{ij}(\mathbf{d})$ induces on $\mathbb{R}$. For a positive γ , the number neither depends on the exact value of γ , nor on the other covariates and parameters. Intuitively, as long as γ is positive, there is a positive probability that $\mathbf{U}$ falls in any combination of buckets. The number of summands in (39) is typically large. The buckets $\mathbf{b}$ of $\mathbb{B}^n$ and the function $\mathcal{N}$ depend on γ .

We have that

$$\begin{aligned}\frac{\partial P(\mathbf{d}; \theta, \mathcal{N})}{\partial \gamma} &= \frac{\partial}{\partial \gamma} \left\{ \sum_{\mathbf{b} \in \mathbb{B}^n} \int_{\mathbf{u} \in \mathbf{b}} \mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} \right\} \\ &= \sum_{\mathbf{b} \in \mathbb{B}^n} \frac{\partial}{\partial \gamma} \int_{\mathbf{u} \in \mathbf{b}} \mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u},\end{aligned}$$

The switching of summation and derivative operator is possible because the number of summands does not depend on γ . We could try to take the derivative of each summands integral boundaries and of $\mathcal{N}(\mathbf{d}, \cdot; \theta)$. But there is no need to boil the ocean, because regardless of $\mathcal{N}(\mathbf{d}, \cdot; \theta)$ most of the summands are 0. To show this we consider three sets of summands.

Case 1: more than two buckets in $\mathbf{B}$ are inner buckets

Recall that the boldface subscripts $\mathbf{i} = \mathbf{1}, \mathbf{2}, \dots$ index the $n = N(N-1)$ directed dyads in arbitrary order. Consider a set of buckets $\mathbf{b}$ where two or more of them are inner buckets. Without loss of generality assume that the $L \geq 2$ inner buckets correspond to $b_1, \dots, b_L$ of $\mathbf{b} = (b_1, \dots, b_n)$. The shape of the $\mathbf{l}^{th}$ bucket is $(\gamma \underline{s}_1, \gamma \bar{s}_1]$ with $\underline{s}_1 < \bar{s}_1$ coinciding with the bucket borders induced by the precise form of strategic interaction specified under the alternative. We normalize the dyad-specific systematic utility component $\mu_{ij} = 0$ without loss of generality.

Recall that $\tilde{\mathbb{B}}^n$ is the set of bucket configurations with two or more inner buckets. For any $\mathbf{b} \in \tilde{\mathbb{B}}^n$ we can derive the upper bound:

$$\begin{aligned}\int_{\mathbf{u} \in \mathbf{b}} \mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} &= \int_{\mathbf{u} \in \mathbf{b}} \mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) \left[\prod_{\mathbf{i}} f_U(u_{\mathbf{i}}) \right] d\mathbf{u} \\ &\leq \int_{\gamma \underline{s}_1}^{\gamma \bar{s}_1} f_U(u_1) \times \dots \times \int_{\gamma \underline{s}_L}^{\gamma \bar{s}_L} f_U(u_L) \int_{\mathbf{u}_{-L} \in \mathbf{b}_{-L}} f_{\mathbf{U}_{-L}}(\mathbf{u}_{-L}) d\mathbf{u} \\ &< \int_{\gamma \underline{s}_1}^{\gamma \bar{s}_1} f_U(u_1) \times \dots \times \int_{\gamma \underline{s}_L}^{\gamma \bar{s}_L} f_U(u_L) du_1 \dots du_L \\ &< \int_{\gamma \underline{s}_1}^{\gamma \bar{s}_1} 1 \times \dots \times \int_{\gamma \underline{s}_L}^{\gamma \bar{s}_L} 1 du_1 \dots du_L \\ &= \gamma^L (\bar{s}_1 - \underline{s}_1) \times \dots \times (\bar{s}_L - \underline{s}_L)\end{aligned}$$

where $\mathbf{u}_{-L}$ denotes the vector $\mathbf{u}$ after removal of its first L components and similarly for $\mathbf{b}_{-L}$. The first equality follows from independence of the components of $\mathbf{u}$, the second (weak) inequality from the fact that $\mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) \leq 1$ for all $\mathbf{u} \in \mathbb{U}$. The third (strict) inequality follows because $f_{\mathbf{U}_{-L}}(\mathbf{u}_{-L})$ is a density and the integration is not over all of $\mathbb{R}^{n-L}$. The

fourth (strict) inequality arises because when $f_U(u)$ is the logistic density we have that $f_U(u) = F_U(u)[1 - F_U(u)] < 1$ for all u on a compact interval of the real line. We conclude that any summand where $\mathbf{b}$ has two or more inner buckets is $\mathcal{O}(\gamma^2)$ for $\gamma \rightarrow 0$.

We have, directly from this argument, that $Q(\mathbf{d}; \theta, \mathcal{N}) \in \mathcal{O}(\gamma^2)$ and furthermore that $Q(\mathbf{d}; (0, \delta)', \mathcal{N}) = 0$ (since inner buckets have zero probability when $\gamma = 0$). Hence, by Lemma A.1, we have that

$$\left. \frac{\partial Q(\mathbf{d}; \theta, \mathcal{N})}{\partial \gamma} \right|_{\gamma=0} = 0.$$

This is enough to show equation (44) above. This simplification is essential to the overall result, as it allows us to proceed without knowing any details about the form of the weight function $\mathcal{N}$ when $\mathbf{U}$ takes values which admit multiple NE networks.

Case 2: No bucket in $\mathbf{b}$ is an inner bucket (i.e., all buckets are outer buckets)

If all components of $\mathbf{u}$ fall in either their first or last buckets, then the network is uniquely defined. This occurs because agent-level preferences for forming (or not forming) a link are so strong that they do not depend on the presence or absence of other links in the network. Each agent i either prefers to send a link to j , regardless of the actions taken by others, or does not wish to send a link. Put differently, each agent has a pure link formation strategy which is strictly dominating in such games; therefore $\mathcal{N}(\mathbf{d}, \mathbf{u}; \theta)$ is either zero or one.

For a particular network $\mathbf{d}$, $\mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) = 1$ if, for all (directed) dyads ij such that $d_{ij} = 1$, we have that u_{ij} falls in the first bucket and for all dyads ij such that $d_{ij} = 0$ we have that u_{ij} falls in the last bucket. These considerations give the equality

$$\int_{\mathbf{u} \in \mathbf{b}} \mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} = \prod_{i \neq j} \left[\int_{-\infty}^{\mu_{ij} + \gamma \underline{s}} f_U(u_{ij}) du_{ij} \right]^{d_{ij}} \left[\int_{\mu_{ij} + \gamma \bar{s}}^{\infty} f_U(u_{ij}) du_{ij} \right]^{1-d_{ij}} \quad (48)$$

$$= \prod_{i \neq j} [F_U(\mu_{ij} + \gamma \underline{s})]^{d_{ij}} [1 - F_U(\mu_{ij} + \gamma \bar{s})]^{1-d_{ij}} \quad (49)$$

Taking logarithms of the expression above, differentiating with respect to γ , evaluating at $\gamma = 0$, and multiplying by $P_0(\mathbf{d}; \delta)$ yields a derivative for summands where all buckets in $\mathbf{b}$ are outer buckets of

$$P_0(\mathbf{d}; \delta) \sum_{i \neq j} \left[d_{ij} \underline{s} \frac{f_U(\mu_{ij})}{F_U(\mu_{ij})} - (1 - d_{ij}) \bar{s} \frac{f_U(\mu_{ij})}{1 - F_U(\mu_{ij})} \right]. \quad (50)$$

Case 3: Exactly one bucket in $\mathbf{b}$ is an inner bucket

If all but one component of $\mathbf{u}$ falls into its first or last bucket, then the resulting network is uniquely defined except for the presence or absence of one arc, say, ij . For any such draw of $\mathbf{u}$, since all other links are formed according to a strictly dominating strategy, player i will either benefit from forming the ij arc or not. Hence $\mathcal{N}(\mathbf{d}, \mathbf{u}; \theta)$ is also either zero or one in this case.

For a particular network $\mathbf{d}$, $\mathcal{N}(\mathbf{d}, \mathbf{u}; \theta)$ will equal one if two conditions hold. First, for all directed dyads $kl \neq ij$ such that $d_{kl} = 1$ we have that u_{kl} falls in the first bucket and for all dyads $kl \neq ij$ such that $d_{kl} = 0$ we have that u_{kl} falls in the last bucket. Second, for the dyad ij with u_{ij} falling in an inner bucket, we require that if $u_{ij} \in [\mu_{ij} + \gamma \underline{s}, \mu_{kl} + \gamma s_{ij}(\mathbf{d})]$ that $d_{ij} = 1$, while if $u_{ij} \in [\mu_{kl} + \gamma s_{ij}(\mathbf{d}), \mu_{ij} + \gamma \bar{s}]$ we require that $d_{ij} = 0$. The overall likelihood contribution for this case therefore equals:

$$\begin{aligned} \int_{\mathbf{u} \in \mathbf{b}} \mathcal{N}(\mathbf{d}, \mathbf{u}; \theta) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} &= \prod_{kl \neq ij} \left[\int_{-\infty}^{\mu_{kl} + \gamma \underline{s}} f_U(u_{kl}) du_{kl} \right]^{d_{kl}} \left[\int_{\mu_{kl} + \gamma \bar{s}}^{\infty} f_U(u_{kl}) du_{kl} \right]^{1-d_{kl}} \\ &\quad \times \left[\int_{\mu_{ij} + \gamma \underline{s}}^{\mu_{ij} + \gamma s_{ij}(\mathbf{d})} f_U(u_{ij}) du_{ij} \right]^{d_{ij}} \left[\int_{\mu_{ij} + \gamma s_{ij}(\mathbf{d})}^{\mu_{ij} + \gamma \bar{s}} f_U(u_{ij}) du_{ij} \right]^{1-d_{ij}} \\ &= \prod_{kl \neq ij} [F_U(\mu_{kl} + \gamma \underline{s})]^{d_{kl}} [1 - F_U(\mu_{kl} + \gamma \bar{s})]^{1-d_{kl}} \\ &\quad \times [F_U(\mu_{ij} + \gamma s_{ij}(\mathbf{d})) - F_U(\mu_{ij} + \gamma \underline{s})]^{d_{ij}} \\ &\quad \times [F_U(\mu_{ij} + \gamma \bar{s}) - F_U(\mu_{ij} + \gamma s_{ij}(\mathbf{d}))]^{1-d_{ij}}. \end{aligned}$$

Recall that $s_{ij}(\mathbf{d})$ is the “strategic” part of the marginal utility agent i gets if he forms ij . Because the last two terms in $[\cdot]$ in the expression above are zero at $\gamma = 0$ we only need to consider their derivative (by the product rule the other term equals zero at $\gamma = 0$). Differentiating the last two terms with respect to γ (and multiplying by the balance of

preceding terms) yields

$$\begin{aligned}
& \prod_{kl \neq ij} [F_U(\mu_{kl} + \gamma \underline{s})]^{d_{kl}} [1 - F_U(\mu_{kl} + \gamma \bar{s})]^{1-d_{kl}} \\
& \times [s_{ij}(\mathbf{d}) f_U(\mu_{ij} + \gamma s_{ij}(\mathbf{d})) - \underline{s} f_U(\mu_{ij} + \gamma \underline{s})]^{d_{ij}} \\
& \times [\bar{s} f_U(\mu_{ij} + \gamma \bar{s}) - s_{ij}(\mathbf{d} + ij) f_U(\mu_{ij} + \gamma s_{ij}(\mathbf{d} + ij))]^{1-d_{ij}} \\
& = \prod_{i \neq j} [F_U(\mu_{ij} + \gamma \underline{s})]^{d_{ij}} [1 - F_U(\mu_{ij} + \gamma \bar{s})]^{1-d_{ij}} \\
& \times \left[s_{ij}(\mathbf{d}) \frac{f_U(\mu_{ij} + \gamma s_{ij}(\mathbf{d}))}{F_U(\mu_{ij} + \gamma \underline{s})} - \underline{s} \frac{f_U(\mu_{ij} + \gamma \underline{s})}{F_U(\mu_{ij} + \gamma \underline{s})} \right]^{d_{ij}} \\
& \times \left[\bar{s} \frac{f_U(\mu_{ij} + \gamma \bar{s})}{F_U(\mu_{ij} + \gamma \bar{s})} - s_{ij}(\mathbf{d}) \frac{f_U(\mu_{ij} + \gamma s_{ij}(\mathbf{d}))}{F_U(\mu_{ij} + \gamma \bar{s})} \right]^{1-d_{ij}}.
\end{aligned}$$

Summing this expression over all potential arcs (and evaluating at $\gamma = 0$) gives a total contribution of “one inner bucket in $\mathbf{b}$ ” summands to the derivative of:

$$\begin{aligned}
P_0(\mathbf{d}; \delta) \sum_{i \neq j} d_{ij} \left[s_{ij}(\mathbf{d}) \frac{f_U(\mu_{ij})}{F_U(\mu_{ij})} - \underline{s} \frac{f_U(\mu_{ij})}{F_U(\mu_{ij})} \right] \\
+ (1 - d_{ij}) \left[\bar{s} \frac{f_U(\mu_{ij})}{F_U(\mu_{ij})} - s_{ij}(\mathbf{d}) \frac{f_U(\mu_{ij})}{F_U(\mu_{ij})} \right]. \quad (51)
\end{aligned}$$

Summing (48) and (51) then gives the expression in the statement of Theorem 2.3. Using similar methods we can show that $P(\mathbf{d}; \theta)$ can be differentiated with respect to γ twice as claimed.


A.4 MCMC Proofs

Before proving Theorem 3.1, we provide some additional discussion of our MCMC algorithm. This also allows us to introduce additional notation required for the proof.

Our algorithm involves finding sequences of *schlaufen* such that when we “swap their edges” the in-degree and out-degree sequence of the network remains unchanged. Individually, such swaps may induce cross link matrix violations, but by swapping the edges in many *schlaufen* simultaneously, we can ensure such violations net to zero.

First consider edge swaps which leave the in- and out-degree sequence of a network unchanged. Figure 5 depicts two canonical alternating cycles (*schlaufen*) (Rao et al., 1996). The first, $C := abcd a$, is a so called *alternating rectangle*. In the configuration to the left a and d each have a single out-link and b and c a single in-link; this is also true in the configuration to the right, which corresponds to the the network generated by switching C . The second

Figure 5: Examples of alternating cycles



Source: Authors' calculations.

Notes: The first row depicts an alternating rectangle before and after switching. The second row depicts a compact alternating hexagon before and after switching. The final column of the figure shows the probability with which each node is chosen, resulting in the total probability of the schlaufe.

cycle, called a *compact alternating hexagon*, can be constructed from a so-called 030C triad.

Let $\mathbb{G}_s$ denote the set of all digraphs with degree sequence s . Rao et al. (1996) showed that, for $\mathbf{G}$ and $\mathbf{G}'$ distinct and belonging to $\mathbb{G}_s$, it is possible to obtain $\mathbf{G}'$ from $\mathbf{G}$ by switching a sequence of alternating rectangles and compact alternating hexagons. In practice it may be useful, in the sense that one can move from $\mathbf{G}$ to $\mathbf{G}'$ with fewer cycle switches, if longer alternating cycles are used (e.g., McDonald et al., 2007; Tao, 2016).

These concepts are further illustrated in Figure 6, which also illustrates the important idea of the cross-link violation matrices, associated with a sequence of schlaufen, summing to zero. Consider the network depicted in Panel A of Figure 6. 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 schlaufe is $R_1 = jgabcdeca$. It is constructed through a sequence of active and passive steps described in the main text (see also the notes to Figure 2 in the main text). 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 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 $r_D^a(i)$ 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 $r_D^p(i)$ be the cardinality of the set of feasible outlinks in a 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 $R = (i_1, \dots, i_l)$ can now be written as

$$p_{\mathbf{D}}(R) = \frac{1}{N} \prod_{k=1}^{l-1} \left(\frac{1}{r_{\mathbf{D}}^a(i_k)} [k \bmod 2] + \frac{1}{r_{\mathbf{D}}^p(i_k)} [(k-1) \bmod 2] \right) \quad (52)$$

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 $\mathcal{R} = (R_1, \dots, R_h)$ 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 $R_1, \dots, R_{k-1}$ 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 (52) above. The total probability of a feasible schlaufen sequence $\mathcal{R}$ is therefore


$$p_{\mathbf{D}}(\mathcal{R}) = (1 - q) \frac{1}{2^{(h-1)}} \prod_{i=1}^h p_{\mathbf{D}_i}(R_i). \quad (53)$$

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 6: A feasible schlaufen sequence



Source: Authors' calculations.

Notes: See the discussion in the 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 (52).

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 A.2. *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. Let $A_{\mathbf{D}, \mathbf{D}'}$ be the set of arcs from the node $v_{\mathbf{D}}$ to the node $v_{\mathbf{D}'}$. We construct a bijection $\varphi : A_{\mathbf{D}, \mathbf{D}'} \rightarrow A_{\mathbf{D}', \mathbf{D}}$. Then we show that the probability of an arc $p(a)$ is equal to $p(\varphi(a))$. If that is proven, the probability of a transition from $v_{\mathbf{D}}$ to $v_{\mathbf{D}'}$ is

$$\begin{aligned}
\sum_{a \in A_{\mathbf{D}', \mathbf{D}}} p(a) &= \sum_{a \in A_{\mathbf{D}', \mathbf{D}}} p(\varphi(a)) \\
&= \sum_{\varphi^{-1}(a') \in A_{\mathbf{D}', \mathbf{D}}} p(a') \\
&= \sum_{a' \in \varphi(A_{\mathbf{D}', \mathbf{D}})} p(a') \\
&= \sum_{a' \in A_{\mathbf{D}, \mathbf{D}'}} p(a')
\end{aligned}$$

which is the probability for a transition from $v_{\mathbf{D}'}$ to $v_{\mathbf{D}}$.

For the construction of the bijection consider that every arc $A_{\mathbf{D}, \mathbf{D}'}$ corresponds uniquely to a schlaufen-sequence $\mathcal{R} = (R_1, \dots, R_h)$. Let $R_k = (i_1, \dots, i_m, \dots, i_l)$ with i_m the start of the

cycle (if there is no cycle in R , we set $R = \bar{R}$). We define $\bar{R}_k = (i_1, \dots, i_m, i_{l-1}, \dots, i_{m+1}, i_l)$ and $\bar{\mathcal{R}} = (\bar{R}_1, \dots, \bar{R}_h)$.

Note that the $R_1, \dots, R_h$ are link disjoint and as soon as the cycle of R_k is switched $\bar{R}_k$ is a schlaufe. The violation matrix of $\bar{R}_k$ is the negative violation matrix of R_k . This implies that if $\mathcal{R}$ is a feasible schlaufen sequence for G which defined an arc in $A_{\mathbf{D}, \mathbf{D}'}$ then $\bar{\mathcal{R}}$ is a feasible schlaufen-sequence for $\mathbf{D}'$ and defines an arc $A_{\mathbf{D}', \mathbf{D}}$.

We define now φ as the function which maps the arc in $A_{\mathbf{D}, \mathbf{D}'}$ with schlaufen sequence $\mathcal{R}$ to the arc in $A_{\mathbf{D}', \mathbf{D}}$ with schlaufen sequence $\bar{\mathcal{R}}$. By construction φ is injective, which implies $|A_{\mathbf{D}, \mathbf{D}'}| \leq |A_{\mathbf{D}', \mathbf{D}}|$. By symmetry we conclude $|A_{\mathbf{D}', \mathbf{D}}| \geq |A_{\mathbf{D}, \mathbf{D}'}|$, which implies $|A_{\mathbf{D}', \mathbf{D}}| = |A_{\mathbf{D}, \mathbf{D}'}|$ and that φ is bijective.

It remains to show that the probability of an arc $p(a)$ is equal to $p(\varphi(a))$. For any node there are equally many feasible active / passive outlinks in $\mathbf{D}$ as in $\mathbf{D}'$. If for a node one outlink is marked due to an link in R_k then for the same node one outlink is marked in $\bar{R}_k$. Therefore $r_{G'_k}(i)$ is equal to $r_{G_k}(i)$ for an active as well as a passive step. Looking at equation (52) the $p_{G_k}(R_k)$ is only different from $p_{G'_k}(\bar{R}_k)$ in the numbering of the factors. But in a cycle of a schlaufe the start node i_m and the end node i_l are such that $m - l \bmod 2 = 0$. The reordering leaves even indexes even and odd indexes odd. Therefore $p_G(R_k) = p'_G(\bar{R}_k)$. From equation (53) it follows directly that $p_G(\mathcal{R}) = p_{G'}(\bar{\mathcal{R}})$ which completes the proof. $\square$

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

Lemma A.3. *The state graph Φ is strongly connected.*

Proof. The symmetric difference of two realizations of $\mathbb{D}_{s, \mathbf{m}}$, which we denote by $\mathbf{D}$ and $\mathbf{D}'$ is a set of alternating cycles. Cycles are in particular schlaufen. We order them arbitrarily as $(R_1, \dots, R_h)$. The sum of the violation matrices is 0. Therefore $(R_1, \dots, R_h)$ is either a feasible schlaufen-sequence or a concatenation of feasible schlaufen-sequences. In the first case there is an arc from $v_{\mathbf{D}}$ to $v_{\mathbf{D}'}$. In the second case, all the feasible schlaufen-sequence define an arc to a new node, resulting in a directed path starting at $v_{\mathbf{D}}$ and ending in $v_{\mathbf{D}'}$. Thus between any two vertexes in Φ there is a directed path. $\square$

With Lemmata A.2 and A.3 in hand, proving Theorem 3.1 is straightforward.

Proof of Theorem 3.1

Every time the Algorithm 1 arrives at step 2 a new arc of the state graph is crossed. At step 2 the algorithm follows a loop arc of type 1 with probability q . Otherwise it proceeds

to step 3. In step 3 a schlaufen-sequence $\mathcal{R}$ is constructed. If the violation matrices of this schlaufen-sequence sum up to 0, the its cycles are switched and an arc of type 2 is followed with probability $p_G(\mathcal{R})$. If the violation matrices do not sum up to 0, then an arc of type 3 is followed. All the cases in which the violation matrices don't sum up to 0 correspond to the residual probability. Therefore Algorithm 1 is a random walk on the state graph Φ .

According to Lemma A.2 Φ is (weighted) symmetric and according to Lemma A.3 it is strongly connected. Due to the self-loops, Φ is not bipartite. Therefore the limit distribution is uniform.

A.5 Monte Carlo experiments

In this appendix we summarize the results of a small number of Monte Carlo experiments. These experiments illustrate the excellent size control and good power properties of our tests. For the Monte Carlo experiments we work with the general utility function introduced in the main paper. We assume that $A_i \in \mathbb{A} \stackrel{\text{def}}{=} \{\alpha_L, \alpha_H\}$, $B_i \in \mathbb{B} \stackrel{\text{def}}{=} \{\beta_L, \beta_H\}$ and $X_i \in \mathbb{X} \stackrel{\text{def}}{=} \{0, 1\}$. We assume that each support point in $\mathbb{A} \times \mathbb{B} \times \mathbb{X}$ occurs with equal probability (i.e., with probability equal to $\frac{1}{8}$).

Observe that there are four types of sending agents: $(A_i = \alpha_L, X_i = 0)$, $(A_i = \alpha_H, X_i = 0)$, $(A_i = \alpha_L, X_i = 1)$ and $(A_i = \alpha_H, X_i = 1)$. Similarly there are four types of receiving agents. The null model is therefore fully described by $16 = 4 \times 4$ linking probabilities. These probabilities are, in turn, a function of the 8 model parameters. We set these parameters as follows: $-\alpha_L = \alpha_H = 0.7$, $-\beta_L = \beta_H = 0.5$, $\lambda_{00} = \lambda_{11} = -2$ and $\lambda_{10} = \lambda_{01} = -4$. This yields a null model expected graph density of about 0.10.


Our parameter choices generate meaningful degree heterogeneity and homophily under the null. Across 1,000 Monte Carlo simulations with $N = 100$, average network density was 0.1, average reciprocity was 0.1, and the *average* standard deviation of in- and out-degree, was 7.5.

We set the network benefit function to $g_i(\mathbf{d}) = \sum_j d_{ij} (\sum_k d_{ik} d_{kj})$ as is appropriate when agents prefer transitive ties. To simulate a network under the alternative we draw $\mathbf{U}$ and then, starting with an empty adjacency matrix, iterate until all links with positive marginal utility are present, and all those with negative marginal utility are absent. By Tarski's Theorem this finds us the least dense pure strategy Nash Equilibrium.

We compare the performance of three tests: (i) the infeasible locally best test that is based upon the true value of $\delta = \delta_0$; (ii) the feasible version of this test which replaces δ with its maximum likelihood estimate computed under the null; finally, (iii) we construct an *ad hoc* test based upon the transitivity index. This last test is the one most often used in practice (where it is compared to a reference value derived from a simple random graph null).

Figure 7 summarizes our findings. The horizontal axis of the figure correspond to different values of the strategic interaction parameter, γ_0 ; the vertical axis to the rejection frequency. With 1000 Monte Carlo replications the standard error of our simulation estimate of size is $\sqrt{(0.05(1 - 0.05)/1000)} \approx 0.007$.

Figure 7: Power Analysis



Source: Authors' calculations.

Notes: The figures plot the frequency with which $H_0 : \gamma_0 = 0$ is rejected across 1,000 Monte Carlo replications for networks with $N = 100$ agents. The y-axis reports the estimated rejection frequency, the x-axis gives the value of the strategic interaction parameter, γ . The minimal pure strategy NE is used to simulate each network. For each simulation a total of 100 MCMC draws from $\mathbb{D}_{s,m}$ were used to compute critical values. The mixing time was chosen such that (approximately) every edge is twice randomly modified before the network is considered a uniform draw. The marginal utility function equals $A_i + B_j + X_i' \Lambda_0 X_j + \gamma_0 s_{ij}(\mathbf{d}) - U_{ij}$ with $s_{ij}(\mathbf{d}) = \sum_k d_{ik} d_{kj} + \sum_{k \neq j} d_{ik} d_{jk}$. The distribution of (A_i, B_i, X_i) and the model parameters are as described in the main text.

As expected, the actual size of our test is indistinguishable (i.e., equal up to simulation error) from its nominal size. For the designs considered here the power gains associated with using the locally best test statistic derived in Section 2 are considerable. Furthermore the

feasible locally best test, which replaces δ_0 with its MLE (computed under the null), performs almost as well as the infeasible locally best test based on the actual value of δ_0 .

The Monte Carlo experiments highlight that the locally best test, which upweights episodes of “unexpected” transitivity, is more powerful than the ad hoc test based on comparing the transitivity index with its null distribution. Note both tests are valid and correctly-sized.

Next we consider the behavior of our test under mis-specification. Specifically we consider a data generating process where the link-specific random utility shocks are Gaussian instead of logistic. We set the variance of the Gaussian distribution to $\pi^2/3$ so that its scale is the same as in the logistic case. All other model parameters remain as defined above. We simulate 1,000 networks with Gaussian errors, but then proceed “as if” our logistic assumption were true. Note our sufficiency, conditioning, and similarity arguments are no longer valid.

The results of these experiments are summarized in Figure 8. Our tests are conservative in these experiments, with actual sizes below their nominal 0.05 level. The power properties of our tests remain good. Note the “oracle” test is based on an evaluation of a logit probability function at the true utility function coefficients (from the the Gaussian model). The feasible test is based upon quasi maximum likelihood estimates of the pseudo-true utility function coefficients (under the null). There is no a priori reason to expect one test to perform better than the other in these designs (indeed one might expect the feasible test to do better under misspecification). In this example the two power curves cross.

This is just one experiment, but it provides some suggestive evidence that our test may still be useful in settings with modest departures from the logistic assumption.


A.5.1 A note on computation

In our simulation experiments $\hat{\delta}$ only enters the locally best test statistic through the null linking probabilities $P(\hat{\mu}_{ij})$. This is relevant because in some settings the MLE of δ_0 may not exist, while that of $P(\mu_{ij})$ does (c.f., Chamberlain, 1980). This can be especially true if $\mathbf{D}$ is “sparse”.

Consider a simple case where one agent sends no links in $\mathbf{D}$. The MLE of her out-degree parameter will be negative infinity. The MLEs of her dyadic out-link probabilities will all equal zero (which is finite). If we consider the form of the optimal test statistic, we see that any dyads ij , where i sends no links in the network in hand, makes a contribution of zero to the feasible test statistic. An analogous argument can be used to show that any ego which directs a link to all alters will also not contribute to the locally best test statistic.

In general, the MLEs of the baseline model choice probabilities are well-defined, although in a “sparse” graph some of these fitted probabilities may be zero or one. In such cases these

Figure 8: Test Behavior Under a Gaussian Random Utility Distribution



Source: Authors' calculations.

Notes: All features of these experiments are as reported in the notes to Figure 7 above, with the exception that the link-specific random utility shocks are Gaussian (with a variance of $\pi^2/3$). Estimation proceeds “as if” the logistic assumption nevertheless holds.

dyads simply make no contribution to the test statistic. Our code automatically handles such cases.

Recall that our test remains admissible for any $\delta \in \Delta$.

A.6 Additional applications

Bi-partite networks

Let $i = 1, \dots, N_1$ index a set of firms deciding which of $j = 1, \dots, N_2$ markets to enter or not. To be concrete consider the problem of airlines deciding which routes they will operate in. The $K_1 \times 1$ binary vector X_i^f indicates what type firm i is (e.g., a legacy carrier or a low cost airline); the binary vector X_j^m indicates what type market j is (e.g., small, medium or large). Let $\mathbf{E} = [E_{ij}]_{i=1, \dots, N_1, j=1, \dots, N_2}$ be the $N_1 \times N_2$ matrix which records which markets each firm chooses to enter (“ $\mathbf{E}$ ” for market entry matrix).

The payoff firm $i = 1, \dots, N_1$ gets from a given industry-wide pattern of market entry $\mathbf{E} = \mathbf{e}$ equals

$$\nu_i(\mathbf{e}_i, \mathbf{e}_{-i}; \theta, \mathbf{U}_i) = \gamma_0 g_i(\mathbf{e}) + \sum_{j=1}^{N_2} e_{ij} \left(A_i + B_j + X_i^{f'} \Delta X_j^m - U_{ij} \right). \quad (54)$$

Here A_i captures unobserved heterogeneity across firms; for example some firms may be able to operate at systematically lower cost and hence profitably enter more markets (corresponding to high values of A_i). Likewise B_j captures heterogeneity across markets; some markets may be intrinsically more profitable than others and hence many firms operate in them (corresponding to high values of B_j). The term $X_i^{f'} \Delta X_j^m$ allows certain types of markets to be systematically more attractive to certain types of firms. Finally U_{ij} is a firm-by-market specific logistic profit shock.

In a simple entry game, we might set $g_i(\mathbf{e}) = -\sum_{j=1}^{N_2} e_{ij} \left[\sum_{k=1}^{N_1} e_{kj} \right]$, such that entry into market j is less attractive to firm i when many other firms k also enter market j (e.g., Ciliberto and Tamer, 2009). This payoff function implies independence of entry decisions across markets.

An advantage of thinking about entry decisions as resulting in a bi-partite network, with arcs from firms to markets, is that it makes it easy to consider more complex preference structures. These preference structures can allow for interdependence in entry decisions across markets.

For example, in the model of Jia (2008) entry into market j is more attractive if the firm

also enters other nearby markets. In this case we might set

$$g_i(\mathbf{e}) = - \sum_{j=1}^{N_2} e_{ij} \left[\sum_{k=1}^{N_1} e_{kj} \right] + \lambda(\mathbf{e}_i) \quad (55)$$

where $\lambda(\mathbf{e}_i)$ varies inversely with some measure of the spatial spread of those markets i enters (calibrated to measure how the operating costs of the firm vary with the geographic dispersion of the markets entered). In the airline example it may be more costly, or less profitable, to operate across markets that are disconnected. That is an airline may prefer route maps which allow a customer to travel to all other airports in their network without having to fly with competitor. In this example $\lambda(\mathbf{e}_i)$ might equal a decreasing function of the number of connected components in i 's route network. Alternatively it could vary with the number of location pairs which can be reached on either a direct flight or by making just one connection.

We can also allow for more complex forms of spatial competition. If j indexes airline routes, then we might set

$$g_i(\mathbf{e}) = - \sum_{j=1}^{N_2} e_{ij} \phi_j(e_{-i}) \quad (56)$$

with $\phi_j(\mathbf{e}_{-i})$ returning how many competitor airlines operate in routes with origin and destination airports both within a one hour drive of the corresponding route j airports (e.g., entry into the SFO-LAX market may depend on how many competitors operate on the OAK-BUR, SFO-BUR, and OAK-LAX routes as well as the number which operate on the SFO-LAX route). One apparent advantage of this “network” perspective is that it allows for the incorporation of complex cross-market complementarities as well as rich forms of spatial competition.

To connect this entry model to the directed network problem defined in the main text we define the $N \times N$ matrix:

$$\mathbf{D} = \begin{bmatrix} \mathbf{0} & \mathbf{E} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \quad (57)$$

and then proceed as described in the paper.

The network $\mathbf{D}$ is bipartite. There are no firm-to-firm or market-to-market links. Furthermore, only firms may direct links (with arcs denoting market entry decisions). These features of the problem induce the special structure of $\mathbf{D}$ above. The cross link matrix will also have a structure analogous to that of the adjacency matrix.

In this example the null reference distribution will assign a zero to many “links” with probability one; this is not a problem for our MCMC simulation algorithm. The imple-

mentation available in the Python **ugd** package can handle degree sequences and cross link matrices with zero elements.

The null set of networks corresponds to all networks where the same number of firms enter each market as observed in the network in hand and the observed patterns of entry are the same as in the network in hand. For example, the number of low cost airlines entering small, medium and large markets is held fixed and so on.

A natural test statistic would be

$$R(\mathbf{D}) = \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \left(E_{ij} - p_{ij}(\hat{\delta}) \right) s_{ij}(E_{ij}) \quad (58)$$

with $p_{ij}(\hat{\delta})$ an estimate of the probability that firm i enters market j under the null that $\gamma_0 = 0$ and $s_{ij}(\mathbf{E})$ the marginal network benefit of entering market j for firm i holding all other own and competitor route choices fixed.

As the above example makes clear, the application of the methods proposed in the main text to bipartite networks is conceptually straightforward. This, in turn, expands the class of many player games to which our methods apply. Note that questions of test power are game specific.

Conditional inference in dyadic models

Following a suggestion in Graham (2017), our MCMC algorithm can also be used for conditional maximum likelihood estimation (CMLE) of non-strategic dyadic logit models (see also Bartolucci (2024)). Let $D_{ij} = 1$ if country i attacks country j within some researcher-defined time period. Let X_i be a $K \times 1$ vector of country types (e.g., a partition of countries into broad geographic regions), finally let $Z_{ij} = 1$ if both i and j are democracies and zero otherwise. We posit the following model for the initiation of conflict by i against j

$$D_{ij} = \mathbf{1} \left(A_i + B_j + X_i' \Lambda_0 X_j + Z_{ij}' \beta_0 - U_{ij} \geq 0 \right), \quad (59)$$

for $i \neq j$, $i, j = 1, \dots, N$ and U_{ij} logistic.

According to democratic peace theory, democracies are less likely to engage in conflict with other democracies such that $\beta_0 < 0$ (e.g. Oneal and Russett, 1999). Note that any level (or monadic) effect of democracy on the propensity to initiate conflict generally is absorbed into the ego effects $\{A_i\}_{i=1}^N$ (out-degree effects), while any level effect on the propensity to be militarily targeted by others is absorbed into the alter effects $\{B_j\}_{j=1}^N$ (in-degree effects). Systematic cross-regional patterns in the costs and benefits of conflict are controlled for by

the “homophily” term $X_i' \Lambda_0 X_j = W_{ij}' \lambda_0$.

The conditional likelihood of the network in hand, $\mathbf{D}$, here the observed pattern of conflict among nations, is

$$\begin{aligned} L(\mathbf{D}; \delta_0, \beta_0) &= \prod_{i \neq j} \left[\frac{\exp(W_{ij}' \lambda_0 + R_i' \mathbf{A} + R_j' \mathbf{B} + Z_{ij}' \beta_0)}{1 + \exp(W_{ij}' \lambda_0 + R_i' \mathbf{A} + R_j' \mathbf{B} + Z_{ij}' \beta_0)} \right]^{D_{ij}} \\ &\quad \times \left[\frac{1}{1 + \exp(W_{ij}' \lambda_0 + R_i' \mathbf{A} + R_j' \mathbf{B} + Z_{ij}' \beta_0)} \right]^{1-D_{ij}} \\ &= c(\mathbf{X}, \mathbf{Z}; \delta_0, \beta_0) \prod_{i \neq j} \exp(\mathbf{T}' \delta_0) \exp \left(\left[\sum_{i \neq j} D_{ij} Z_{ij} \right]' \beta_0 \right) \end{aligned}$$

with R_i , as earlier, a $N \times 1$ vector with a 1 in its i^{th} element and zeros elsewhere and $c(\mathbf{X}, \mathbf{Z}; \delta_0, \beta_0)$ not varying with $\mathbf{D}$. Here $\mathbf{T}$ includes the vectorized cross-link matrix as well as the out- and in-degree sequences as discussed in the main text.

Let $\tau(\mathbf{v})$ be the sufficient statistics for δ in network $\mathbf{v}$. Conditioning on $\mathbf{T} = \mathbf{t}$ yields

$$\begin{aligned} \Pr(\mathbf{D} = \mathbf{d} | \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}, \mathbf{T} = \mathbf{t}; \delta_0, \beta_0) &= \frac{\prod_{i \neq j} \exp(\mathbf{t}' \delta_0) \exp \left(\left[\sum_{i \neq j} d_{ij} z_{ij} \right]' \beta_0 \right)}{\sum_{\mathbf{v} \in \mathbb{D}_{\mathbf{s}, \mathbf{m}}} \prod_{i \neq j} \exp(\tau(\mathbf{v})' \delta_0) \exp \left(\left[\sum_{i \neq j} v_{ij} z_{ij} \right]' \beta_0 \right)} \\ &= \frac{\exp \left(\left[\sum_{i \neq j} d_{ij} z_{ij} \right]' \beta_0 \right)}{\sum_{\mathbf{v} \in \mathbb{D}_{\mathbf{s}, \mathbf{m}}} \exp \left(\left[\sum_{i \neq j} v_{ij} z_{ij} \right]' \beta_0 \right)} \\ &= \Pr(\mathbf{D} = \mathbf{d} | \mathbf{Z} = \mathbf{z}, \mathbf{T} = \mathbf{t}, \beta_0), \end{aligned}$$

where the second equality follows from the fact that $\tau(\mathbf{v}) = \mathbf{t}$ for all $\mathbf{v}$ in $\mathbb{D}_{\mathbf{s}, \mathbf{m}}$.

Taking logs yields, after some manipulation,

$$\begin{aligned} \ln \Pr(\mathbf{D} = \mathbf{d} | \mathbf{Z} = \mathbf{z}, \mathbf{T} = \mathbf{t}, \beta_0) &= \left[\sum_{i \neq j} d_{ij} z_{ij} \right]' \beta_0 \\ &\quad - \ln \left(\mathbb{E} \left[\exp \left(\left[\sum_{i \neq j} D_{ij} z_{ij} \right]' \beta_0 \right) \right] \right) - \ln |\mathbb{D}_{\mathbf{s}, \mathbf{m}}|, \quad (60) \end{aligned}$$

where the expectation in the second term to the right of the equality is taken with respect to the discrete uniform distribution on $\mathbb{D}_{\mathbf{s}, \mathbf{m}}$. This expectation can be estimated using our simulation algorithm. The third term is invariant to β_0 and can consequently be ignored. Basing

estimation and inference upon (60) allows a researcher to learn about β_0 in the presence of out- and in-degree heterogeneity as well as potentially complex patterns of homophily.

In our democratic peace theory example, Z_{ij} is binary. In this case $\sum_{i \neq j} D_{ij} Z_{ij}$ is simply a count of how many wars are initiated by democracies against other democracies. To compute the expectation in (60) we need an estimate of the distribution of this count induced by the discrete uniform distribution on $\mathbb{D}_{\mathbf{s}, \mathbf{m}}$.