

Causal Inference over Stochastic Networks

Duncan A. Clark

University of California – Los Angeles, Los Angeles, USA

E-mail: duncanclark@ucla.edu

Mark S. Handcock

University of California – Los Angeles, Los Angeles, USA

Summary. Claiming causal inferences in network settings necessitates careful consideration of the often complex dependency between outcomes for actors. Of particular importance are treatment spillover or outcome interference effects. We consider causal inference when the actors are connected via an underlying network structure. Our key contribution is a model for causality when the underlying network is endogenous; where the ties between actors and the actor covariates are statistically dependent. We develop a joint model for the relational and covariate generating process that avoids restrictive separability and fixed network assumptions, as these rarely hold in realistic social settings. While our framework can be used with general models, we develop the highly expressive class of Exponential-family Random Network models (ERNM) of which Markov Random Fields (MRF) and Exponential-family Random Graph models (ERGM) are special cases. We present potential outcome based inference within a Bayesian framework, and propose a modification to the exchange algorithm to allow for sampling from ERNM posteriors. We present results of a simulation study demonstrating the validity of the approach. Finally, we demonstrate the value of the framework in a case-study of smoking in the context of adolescent friendship networks.

Keywords: Causality, Social Networks, Network models, Spillover, Contagion, Interference, Gibbs measures

1. Introduction

Causal inference is difficult, especially in systems with partially known and likely complex structure. There is an extensive literature on causal inference methods for so called “network settings” from a variety of perspectives (Hudgens and Halloran, 2008; Shalizi and Thomas, 2011; Ogburn and VanderWeele, 2014; van der Laan, 2014; Sofrygin and van der Laan, 2017; DeAmour, 2016; Aronow and Samii, 2017; Shpitser et al., 2021; Tchetgen Tchetgen et al., 2020; Ogburn et al., 2020; Lee and Ogburn, 2021; Sävje et al., 2021). In these settings, we have a population of units, some of which receive an intervention. They are considered to be a networked population in the sense that the units can be thought of as nodes in a network and some of the units are connected with other units in some meaningful way (via “edges”).

This version of the article has been accepted for publication, after peer review and is subject to the publisher's terms of use, but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: <https://doi.org/>

There are recent empirical studies claiming strong causal results in network settings that have been controversial. For example, claims about the spread of characteristics through social settings, so-called social contagion, (Christakis and Fowler, 2007, 2008, 2010) with corresponding methodological criticism from others (e.g. Ogburn et al., 2022). Such results are controversial due not only to their surprising and perhaps provocative substantive nature, for example, statements such as “obesity is socially contagious”, but also due to the strong assumptions required to justify the methodology.

Much of the problem stems from the unknown underlying social processes. For example, as explicitly noted in Shalizi and Thomas (2011), contagion over time is often confounded with homophily when viewed statically. The former is the tendency of covariates on the nodes to change due to the influence of the covariate values of nodes they are tied to, while the latter is the tendency for tied nodes to have the same value of the covariate at a given time point. Shalizi and Thomas (2011) consider Directed Acyclic Graphs (DAGs) (Pearl, 1995; Spirtes et al., 2000) and demonstrate that contagion can be confounded when homophily on latent covariates is present.

Chain graphs have been posited as a possible representation of the causal structure of the equilibria of networks evolving over time (Tchetgen Tchetgen et al., 2020; Shpitser et al., 2021). By considering a DAG of a network generating process over time, Ogburn et al. (2020) and Lauritzen and Richardson (2002) suggest that estimating causal effects in this setting is not viable unless fine-grained temporal data is available and the causal structure is simple.

We consider the situation where we observe a single network, considered a realization of a random social process over both nodes and edges. Causal inference in this setting cannot be reduced to an independent and identically distributed (IID) problem in the same fashion as Sofrygin and van der Laan (2017). The chain graph approximation employed in Tchetgen Tchetgen et al. (2020) and Ogburn et al. (2020) cannot be used, as it does not allow for stochastic connections. There have been efforts to allow for network uncertainty (Toulis and Kao, 2013; Kao, 2017). However, these methods require unconfoundedness of the edges and the random nodal covariates, given enough fixed information about the nodes. In not making such an assumption, we require a joint probability model for the random edges and nodal covariates in the graph.

1.1. Motivating Case

As our motivating example, we consider one of the school social networks from the National Longitudinal Study of Adolescent Health (Harris et al., 2007). The network we used represents all $n = 869$ students in the school. The social network is constructed by data collected from an in-school survey of all students in grades 7-12 of the school ($n = 869$). The instrument asked them to nominate up to five of their best friends of each sex within the school. Two students were friends if they each nominated each other. There were 462 were male and 407 female students, with 344 having reported trying a cigarette at least once (“smokers”). Intuitively, we believe that there is concurrently social selection (Robins et al., 2001), that smokers are more likely to be friends and social contagion (Robins et al., 2001), that smokers influence their non-smoking friends to smoke (or not). As noted in

Shalizi and Thomas (2011) these are confounded in the case of latent homophily. However, under the model introduced in this paper and assuming the lack of external confounding, we demonstrate estimates that can have a causal interpretation. This is very appealing in causal settings with networked populations, as we can then make statements regarding, for example, the difference in expected smoker status of a person, given a change in the number of smoking friends in the network.

1.2. *Related and Connecting Literature*

We note that the literature of network causal inference is at present sharply disconnected from the literature concerning generative models for social networks. As most approaches have considered the networks as fixed, there has been little interest in placing a probability distribution over the space of possible networks. Typically, generative models for social networks, for example the commonly used Exponential-family Random Graph Models (ERGM) (Frank and Strauss, 1986) consider the edges as the random variables to be modelled and nodal covariates as fixed. There has been much sophisticated work on understanding such models (Handcock, 2003; Robins et al., 2007; Schweinberger and Handcock, 2015; Lauritzen et al., 2018), and well developed Markov Chain Monte Carlo (MCMC) based fitting procedures (Snijders, 2002; Handcock, 2002) with associated complex software (Hunter et al., 2008). However, due to the assumption of fixed nodal covariates, these are of limited use for causal inference on nodal outcomes. Markov random field (MRF) models treat nodal covariates as stochastic but the connections between nodes as fixed. The autologistic actor attribute model (ALAAM) (Daraganova and Robins, 2012; Koskinen and Daraganova, 2022), is one particular instance of an MRF model, but we will retain the original name here. Exponential-family Random Network Models (ERNM) (Fellows and Handcock, 2012) encompass both model classes as special cases. ERNM allows for the edges and nodal covariates to be stochastic, thus the nodal covariates and the edges can be co-dependent. These are very expressive models that can represent complex network structures and allow mixed continuous and discrete covariates as well as valued tie variables.

Our framework allows other joint models of random attributes and random tie values. Fosdick and Hoff (2015) present a latent variable model that contains both additive and multiplicative latent effects and is able to represent complex network structure, including within-dyad correlation. This model can be motivated by the concept of an underlying social space, with the model a reduced dimension representation of the network structure (Hoff, 2005). It jointly models nodal covariates and ties variables via this latent space. Like the ERNM, it requires careful specification of the latent and manifest variable structure. Durante et al. (2017) propose joint modeling of categorical nodal covariates and network tie variables using a latent variable tensor factorization of the joint probability model. In the dynamic setting, Stochastic Actor-Oriented Models (SAOMs) have been developed to represent the co-evolution of networks and nodal covariates (behavior) (Steglich et al., 2010). SAOMs are very powerful and useful models for networks when network data are observed at multiple time points. Niezink et al. (2019) extend the framework to accommodate continuous nodal attributes. In this paper, we focus on the common situation where we do not have such data. We note that our approach to causal inference on networks can utilize

alternative models such as those of Fosdick and Hoff (2015) and Steglich et al. (2010) depending on the application and available data. The focus of this paper is on causal inference based on the plausible representation of complex social structure via ERNM.

1.3. Contributions

In this paper, we consider social structures represented by networks with stochastic links and nodal covariates that evolve over time. We present a generalized chain graph approximation to a credible social process DAG, which allows for a dependence structure that we believe to be compatible with such problems. We seek causal inference, that is, the effect on outcomes of a hypothetical intervention. We then frame causal inference in terms of network equilibrium potential outcomes, that is, potential outcomes derived from the chain graph structure. We utilize a Bayesian framework for the causal quantities and the network model. This allows for the incorporation of prior information, as well the automatic accounting of uncertainty in a theoretically consistent manner. The key contribution of our approach is to allow for uncertainty in the network structure and for codependency of edges and nodal covariates in the underlying social process.

The Bayesian approach does not require appeals to asymptotic arguments for its validity. Indeed, asymptotics for causal quantities in our setting are conceptually difficult as there is no single asymptotic framework that is compelling. In particular, the number of nodes, N , is a fundamental characteristic of the social process and not a sampling design characteristic, as it is in most of Statistics. For example, the interactions of a class with 5 students will be quite different from a class of size 75. Hence, asymptotic approximations must identify credible invariant parametrisations (Krivitsky et al., 2011). Different values of N change the fundamental structure of the social network, as dependent edge behaviour is strongly related to the number of nodes in a network.

We use an augmented variable MCMC algorithm to sample from their posterior distributions. We develop a modification to the exchange algorithm (Murray et al., 2006; Liang et al., 2016) to allow for sampling from ERNM posteriors, which we use to infer the posterior distributions of potential outcomes and, hence, estimate causal estimands of interest.

Fellows and Handcock (2012) introduced the ERNM, developed its conceptual features and discussed maximum likelihood-based inference for it. This paper develops a model-based approach to causal inference for networked populations, utilizing the strengths of the ERNM in jointly modeling network tie and nodal covariates.

Our approach has the advantage of capturing uncertainty about causal inference due to the network generation process. It has an additional advantage: the ability to generalize beyond the population network at hand, to the super population of population networks from which the population network is taken. Our inferences hold for the given population of units (“node set”) and social process, whereas assuming a fixed network narrows the scope to that observed network only.

This paper is structured as follows. Section 2 introduces our general network setting and our notation. Section 3 defines causal quantities of interest in terms of equilibrium potential outcomes. Section 4 considers the DAG of a network process over time allowing for network

uncertainty with a chain graph approximation. Section 5 discusses identifiability and inference. Section 6 presents a simulation study, the details of which are contained in the supplement. Section 7 considers a case-study of a network from the National Longitudinal Study of Adolescent Health (Harris et al., 2007) and gives estimates of unknown causal quantities. Section 8 provides general discussion of the method, and its ability to generate credible causal inference.

2. Notation and Setting

We consider a known fixed population (set) of N units (nodes) with $N \geq 3$. Each node has a random nodal outcome Y_i , thus the whole network outcome is $Y = (Y_1, \dots, Y_N)$. Realizations of the random nodal outcomes Y_i and Y are denoted with lower case y_i and y . For this paper, we only consider binary outcomes (although the ideas are easily extended to non-binary outcomes). Each node is also permitted to have further multivariate nodal covariates, similarly denoted $X = (X_1, \dots, X_N)$ with $X_i \in \mathbb{R}^p$ for some $p \geq 0$.

We denote the random edges between nodes as the random variable $A = \{A_{i,j}\}_{i=1,j=1}^N$, with realizations a . A can be considered a random adjacency matrix. We also restrict $A_{i,j}$ to be binary with 1 indicating a connection and 0 representing the absence of a connection. For this paper, we make the restriction that our networks are undirected, i.e. $A_{i,j} = A_{j,i} \quad \forall i, j$.

A network realization is defined to be a set $\{y, a, x\}$. When considering the dynamic network of the process over time we indicate the outcomes at time t with superscript, for example, the outcome random variable for node i at time t is Y_i^t and the whole network random variable at time t is $\{Y^t, A^t, X^t\}$. If temporal dynamics with distribution P_t result in an equilibrium distribution, we denote it by $\lim_{t \rightarrow \infty} P_t(Y^t, A^t, X^t) = P(Y, A, X)$.

As the node set is fixed, the nodal covariates X are often in practice fixed throughout the evolution of the social process. Going forward we omit X from our notation, that is for clarity, we consider our networks as realizations of $(Y, A) | X = x_{\text{observed}}$ but write (Y, A) .

We represent the treatment of nodes via the treatment vector $Z = (z_1, \dots, z_N)$ with realizations z . For the purposes of this paper we consider the treatment to be applied prior to the evolution of the network process, though perhaps conditional on the fixed nodal covariates. We leave allowing for time varying treatments assignments and outcome evolution to future research, though we believe it to be compatible with our approach.

In the following section, we will introduce a DAG to represent the dependence structure of the social network. We emphasize that the nodes in the DAGs represent random variables in the stochastic social process underlying the above described social network setting (rather than actors in the social network). That is, nodes in the DAG represent random variables in the social process generating the network, they may either be treatments, random nodal covariates, or random edges.

3. Network Potential Outcomes and Causal Estimands of Interest

Hudgens and Halloran (2008), Toulis and Kao (2013) and Aronow and Samii (2017) all considered potential outcome-based frameworks of assumptions and definitions as a basis

for causal inference for nodal covariates. We consider network potential outcomes as realizations of an equilibrium distribution of a social process that evolved over time.

Our causal estimands should be interpreted as the effect of an intervention on the equilibrium distribution, this is implicit in Toulis and Kao (2013) and Aronow and Samii (2017). Estimands in Tchetgen Tchetgen et al. (2020) and Ogburn et al. (2020) are based on interventions on nodal statuses prior to the evolution of the social process, and estimate network effects, rather than nodal effects. We note that these are whole network direct or spillover treatment outcomes for pre-social process interventions. The ERNM (Fellows and Handcock, 2012) presented in Section 5.2 is still fully compatible with such regimes, and indeed a generalization of the MRF models used in Tchetgen Tchetgen et al. (2020) and Ogburn et al. (2020).

We set $Y_i(Z, A, Y_{-i})$ to be the equilibrium potential outcome of node i , given the treatment vector of all nodes Z , the network A , and the outcomes of all the other nodes Y_{-i} . In this definition, each potential outcome i depends fully on the entire rest of the edges and nodal potential outcomes in the network.

We introduce the k -step neighborhood of a node as follows.

DEFINITION 1. *k -step neighbourhood*

For a network A :

$$\mathcal{N}_i^k(A) = \{j \mid \exists \text{ a path of length } k \text{ or less between node } i \text{ and } j\} \quad (1)$$

The k -step neighbourhood of node i is the subset of nodes in A such that, there is a path of length less than or equal to k between the two nodes.

For example, the one-step neighbourhood of a node is the set of nodes directly tied to that node. The k -step neighbourhood of a subset of nodes in the network, is defined to be the union of the neighbourhoods of each of the nodes. We define the subnetwork induced by a subset of nodes, to be the network, with all nodes outside of the subset removed, and all edges not between nodes in the subset removed.

Our estimands of interest relate to one-step neighbourhoods of nodes, which dropping k we denote as $\mathcal{N}_i(A)$. We next make the one-step neighbourhood assumption on the nodal potential outcomes, within which we assume the absence of outside confounders. Using the one-step neighbourhood assumption allows us to dramatically reduce the required number of simulations to estimate missing potential outcomes in Section 5.

ASSUMPTION 1. *One-Step Neighbourhood*

$$\forall i \quad Y_i(Z, A, Y_{-i}, C) = Y_i(Z_i, Z_{\mathcal{N}_i(A)}, \mathcal{N}_i(A), Y_{\mathcal{N}_i(A)}) \quad (2)$$

where $\mathcal{N}_i(A)$ is the neighbourhood of node i in the network A and C represents the context of the networked population. That is, the individual outcomes only depend on treatments, edges and outcomes in their own one-step neighbourhood. The assumption of no unaccounted for outside confounders is almost universal to casual settings and social network models. We suppress reference to C from this point on.

We next state the causal estimands that we will pursue in their full generality. In section 5 we will consider models, and simplifying assumptions that allow for inference in practice.

DEFINITION 2. Primary Effects

Let ξ_i be the i^{th} primary effect and ξ be the average primary effect.

$$\xi_i = \frac{1}{|(Z_{-i}, A, Y_{-i})|} \sum_{(z, a, y) \in (Z_{-i}, A, Y_{-i})} \left(Y_i(z_i = 1, z, a, y) - Y_i(z_i = 0, z, a, y) \right) \quad (3)$$

$$\xi = \frac{1}{N} \sum_{i=1}^N \xi_i \quad (4)$$

where (Z_{-i}, A, Y_{-i}) are all possible combinations of edges as well as the treatments and outcomes of nodes other than i

The primary effect can be interpreted as the average difference between nodal potential outcomes under treatment and not under treatment.

DEFINITION 3. k -peer Treatment Effect

Let δ_i^k be the i^{th} k -peer treatment effect and δ^k be the average k -peer treatment effect.

$$\delta_i^k = \frac{1}{|(Z, A)_{\mathcal{N}_i}^k|} \sum_{(z, a, y) \in (Z, A)_{\mathcal{N}_i}^k} \left(Y_i(z_i = 0, z, a, y) - Y_i(z_i = 0, (z, a) \in (Z, A)_{\mathcal{N}_i}^0, y) \right) \quad (5)$$

$$\delta^k = \frac{1}{N} \sum_{i=1}^N \delta_i^k \quad (6)$$

where $(Z, A)_{\mathcal{N}_i}^k$ is the set of all possible combinations of edge realizations and treatment assignments such that node i is exposed to k other treated nodes.

The k -peer treatment effect can be interpreted as the average difference between nodal potential outcomes (that are not treated) with k treated neighbours and 0 treated neighbours.

DEFINITION 4. k -peer Outcome Effect For Binary Outcomes

Let δ_i^k be the i^{th} k -peer outcome effect and δ^k be the average k -peer outcome effect.

$$\delta_i^k = \frac{1}{|(A, Y)_{\mathcal{N}_i}^k|} \sum_{(z, a, y) \in (A, Y)_{\mathcal{N}_i}^k} \left(Y_i(z_i = 0, z, a, y) - Y_i(z_i = 0, z, (a, y) \in (A, Y)_{\mathcal{N}_i}^0) \right) \quad (7)$$

$$\delta^k = \frac{1}{N} \sum_{i=1}^N \delta_i^k \quad (8)$$

where $(A, Y)_{\mathcal{N}_i}^k$ is the set of all possible combinations of edge and outcome realizations such that node i is exposed to k other nodes with outcome 1.

The k -peer outcome effect can be interpreted as the average difference between nodal potential outcomes (that are not treated) with k neighbour outcomes equal to 1 and 0 neighbour outcomes equal to 1.

We note that the k -peer outcome effect could be considered as a special case of a k -peer treatment effect where for a treatment such that $P(Y_i = 1|Z_i = 1) = 1$. However, we find it convenient to express the estimand separately, as it is often of substantive interest in networks where some other treatment is administered, but the k -peer outcome effect is important.

4. Causal Framework

DAGs provide a means for precisely specifying the structure of relationships between random variables (see Pearl, 2009, for an introduction). Pearl (1995) developed strict criteria for identification of causal effects. Formal equivalence with the potential outcome framework was shown in Richardson and Robins (2013). We note that there is no contradiction between using an undirected model for the adjacency matrix A and a DAG for the relationships between the random variables in A as they are separate networks. The adjacency matrix A represents data structure (e.g. an observed friendship network) and the DAG represents the conditional independence structure of the system.

Nodes in the DAG represent random variables, with edges drawn between nodes being strictly uni-directional, and cycles of edges prohibited, so that so called “feedback” loops are not permitted. The interpretation of a directed line from node i to node j is that i causally effects j . For a node indexing set V , denote the corresponding random variables $\{x_v\}_{v \in V}$. This graphical structure encodes the following factorization of their joint probability distribution $P(x) = \prod_{v \in V} P(x_v|x_{pa(v)})$, where $pa(v)$ are the parents of v , that is the nodes in the DAG with edges into v .

The causal effect of X on Y is represented by $P(y|do(X = x))$, the distribution of the random variable Y when X has been externally set to x . Pearl (2009) provides transformations for expressing this new distribution $P(y|do(X = x))$ in terms of distribution on observed random variables, e.g. $P(y|x)$ or $P(y|x, z)$ for z some other variable in the system.

Chain graphs permit undirected as well as directed edges, which allow for different Markov properties from DAGs. In fact, chain graphs can represent dependence structures that are not possible under a DAG. We include some discussion in the supplement, however we omit subtleties of their Markov properties discussed in Frydenberg (1990) and utilized for causal analysis in Lauritzen and Richardson (2002). In their full generality, chain graphs can express complex dependence structures. However, in our case, our example chain graph only has one chain component which results in none of the outcomes being rendered conditionally independent, thus we do not require an in-depth review for the purposes of this paper. Practically, one possible interpretation of the undirected edges in a chain graph is that the two variables interact with a causal feedback sense over time.

DAGs representing the causal structure of outcomes of nodes in networks under both interference and contagion are given a clear and detailed treatment in Ogburn and VanderWeele (2014). We generalize the related conjecture in Ogburn et al. (2020) for a chain graph approximation of causal structure of a social network, which slowly evolves over time. This is based on an interpretation of causality with feedback relationships over time, which

chain graphs can be used to explain (Lauritzen and Richardson, 2002). Our generalization is to consider a social network where the connections are not fixed and are motivated by the empirical observation that social connections are often strongly dependent on other nodes' connections as well as other nodes' covariates, treatments and outcomes.

As an illustration for our chain graph approximation, Figure 1 represents the structure of a network evolving over time for a three node network similar to that of Ogburn et al. (2020) and Tchetgen Tchetgen et al. (2020). Note the analogous connections between node 3 and 1 are omitted for clarity. Block arrows denote multiple arrows into the variables described therein. The full DAG with all arrows becomes quickly unwieldy. In the notation of Section 2, the nodal treatments are $\{Z_1, Z_2, Z_3\}$, nodal outcome variables $\{Y_1, Y_2, Y_3\}$, and we denote the undirected edges between nodes i and j as $A_{i,j}$. Directed network edges as well as fixed nodal covariates can also be built-in, but are omitted here for clarity. Superscripts denote the status of a variable at that time step, e.g. Y_2^t is the outcome for node 2 at time t .

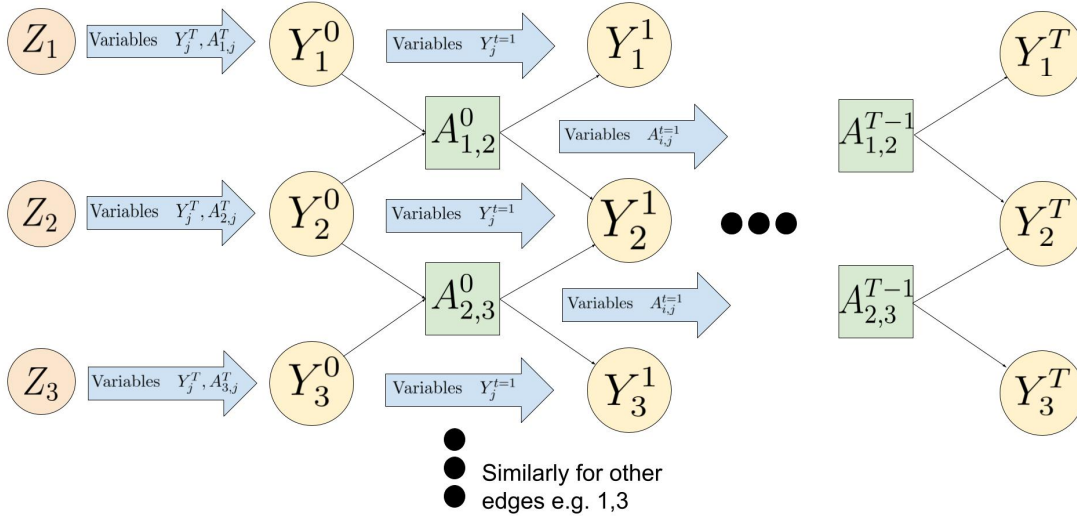


Fig. 1. Full DAG of temporal network formation feedback process, in the three node network case. Treatment variables Z_j are permitted to causally effect all of the other variables. Outcomes Y_i^t and Y_j^t are permitted to causally effect edges $A_{i,j}^t$, as well as outcomes Y_i^{t+1} and Y_j^{t+1} . Edges $A_{i,j}^t$ are permitted to causally effect edges $A_{k,l}^{t+1}$ with no restrictions on k and l . In practice, the functional form of models will usually restrict the causal impact of node i 's outcome and treatment to only the set of neighbouring nodes $\{j\}$, that is, where there is an edge present between node i and j .

Under this DAG, each of the treatment variables Z_j are permitted to causally affect all of the other outcome variables, as well as edges involving j . Outcomes Y_i^t and Y_j^t are permitted to causally effect edges $A_{i,j}^t$, as well as outcomes Y_i^{t+1} and Y_j^{t+1} . Edges $A_{i,j}^t$ are permitted to causally effect edges $A_{k,l}^{t+1}$ with no restrictions on k and l .

We note that there are key edges that are not present in this DAG. Edges and outcomes,

from time step t can only influence edges and outcomes at time step $t + 1$, not others. We argue that this is plausible under slow evolution of a network process over time. In general, it is exceedingly rare that enough data are available to identify all the relationships posited in Figure 1.

The DAG for our system is highly dense, and we usually observe little incremental time data on the social process. Thus, operations as introduced in Pearl (1995) to reduce $P(y|do(X = x))$ to expressions involving only distributions, from which we have realizations, e.g., $P(y|x)$, are not possible. That is the true causal effect is unidentifiable. We pursue “approximate causal inference” through approximating the DAG for our social process over time with a chain graph.

As noted in Lauritzen and Richardson (2002), the equilibrium distribution of a so called infinite DAG can be represented as a chain graph. Similarly to Ogburn et al. (2020), but modelling edges as random, the DAG in Figure 1 can be approximated by the chain graph shown in Figure 2. The chain component (undirected component) of this chain graph is close to being complete since we allow every nodal outcome to influence every other nodal outcome, however we only allow outcomes Y_i and Y_j to influence edge $A_{i,j}$. In the full DAG, backdoor paths through previous time steps result in this not being the case.

Under the assumption of independent treatment assignment, this suggests that the complex structure of such a DAG can be approximated with the chain graph factorization:

$$P(Y, A, Z) = \left(\prod_{i=1}^n P(z_i) \right) P(Y, A|Z) \quad (9)$$

As an approximation to the true temporal DAG a chain graph model of causality serves to render the causal effects tractable in practice. See Ogburn et al. (2020) and Lauritzen and Richardson (2002) for a fuller explanation.

As an illustration of the generality of this chain graph’s dependence structure, we define a conditional independence property. We denote independence between random variables with $\perp\!\!\!\perp$. We refer to the following as *local conditional independence*.

$$Y_i \perp\!\!\!\perp Y_j, Z_j \quad \left| \quad A_{i,j} = 0, \quad Z_i, \quad \{Y_l : A_{i,l} = 1\} \quad (10)$$

That is, nodal outcomes are conditionally independent given all neighbours, and that they are not connected. Equation (10) does not hold *a priori* due to the dependence induced by the random edge $A_{i,j}$ unlike in the fixed network case where it does. That is, the chain graph retains dependencies between outcomes of non-connected nodes, even when conditioning on neighbours. In fact, due to the close-to-complete nature of the chain component, there are few conditional independence assumptions that can be concretely made. Indeed the complexity of such systems is the reason why social network modelling has proven to be difficult.

Conditional independence properties have been considered for ERGMs (Snijders et al., 2006). Common properties induced by model choice in ERGMs are the so-called *Markov* property (Frank and Strauss, 1986) and the so-called *social circuit* property that requires

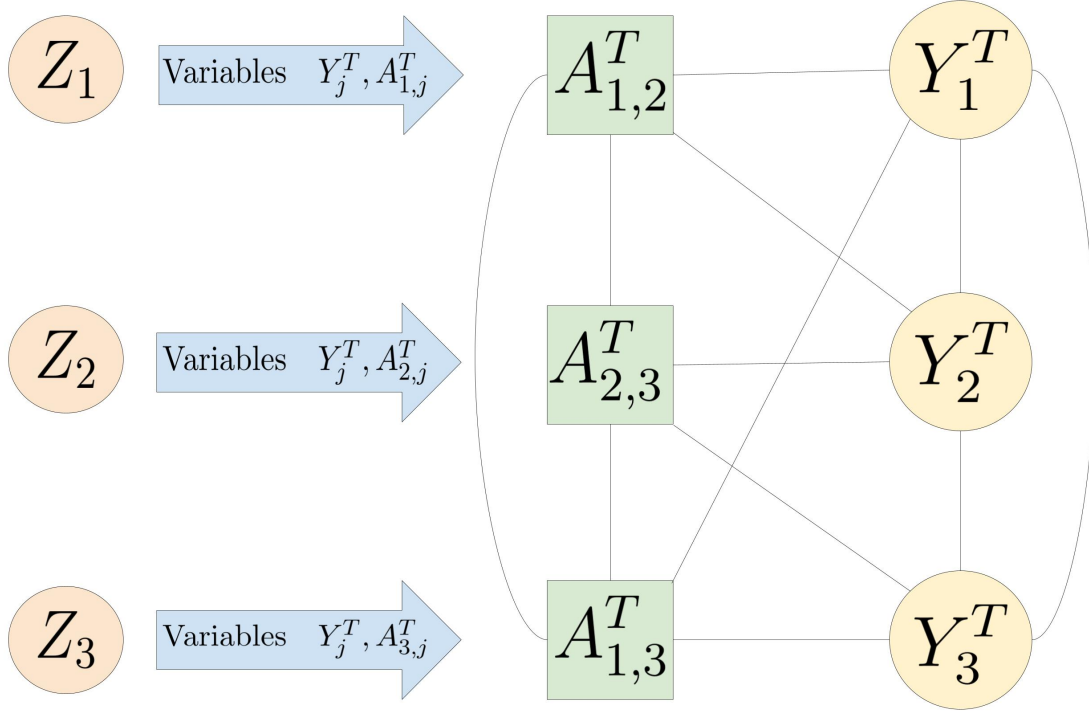


Fig. 2. Chain graph approximation of a three node temporal social network process with treatment

the Markov property in addition to a cycle condition (Snijders et al., 2006). Both of these commonly imposed assumptions severely limit possible dependence structures. In practice, the social network modelling community has found the social circuit assumption to yield well fitting models in a wide variety of situations (Goldenberg et al., 2010).

We do not comment on the validity of the approximation of the DAG by the chain graph. Ogburn et al. (2020) gave simulations supporting their version of this approximation. We will only consider recovering causal estimands based on the assumption that the chain graph approximation holds.

The remainder of this paper concerns modelling $P(Y, A|Z)$ in order to estimate causal quantities. We specify our causal estimands using equilibrium distribution potential outcomes, which we now explicate. Let $(A, Z, Y_i^c)^*$ be some intervention on the network. Under our formulation, we have:

$$P(Y_i = y | (A, Z, Y_i^c) = (A, Z, Y_i^c)^*) = P(Y_i = y | \text{do}((A, Z, Y_i^c) = (A, Z, Y_i^c)^*)) \quad (11)$$

Note that intervening directly on this equilibrium distribution is typically not possible in practice. For example, we can not manipulate the equilibrium number of treated friends due to the complex social processes at play. Our causal effects are of hypothetical interventions that results in k -treated neighbors in the equilibrium distribution of the

network. For example, in our case-study, imagine for the first 6 months of the school year there is highly active socialising as friendships form and behaviours change, and that after this the network reaches an equilibrium. The interpretation of the two-treatment effect is the difference in probability of smoking between having two smoker neighbors verses no smoker neighbors after the period of highly active socialising that establishes the network. Thus these causal effects are of explanatory interest as they make clear the role neighbors play in the complex social processes. A similar argument is made in Ogburn et al. (2020) and it is implicit in other current potential outcome based approaches (Aronow and Samii, 2017; Toulis and Kao, 2013). We interpret our causal estimands as an average, of the possible treatments, and social processes that led to a given treatment in equilibrium.

The lack of conditional independence assumptions that we are able to make in Section 5, is the direct consequence of the nearly complete chain graph model specified in this section.

5. Estimation and Identifying Assumptions

Section 3 introduced network potential outcomes and our estimands of interest. Section 4 justified their casual interpretation. In this section, we introduce models that can be practically used to infer the missing potential outcomes so as to produce concrete estimates of causal effects.

Our strategy is to pursue model-based Bayesian imputation of potential outcomes (Imbens and Rubin, 2015).

5.1. Structural Assumptions

We next state possible assumptions restricting the dependence structure of the process that will enable us to feasibly impute missing potential outcomes from network simulations. In our assumptions, we explicitly include fixed nodal covariates X .

ASSUMPTION 2. *Unconfounded Treatment Assignment Assumption Under Network Assignment*

$$P(Z|X, A, \mathbb{Y}) = P(Z|X, A, \mathbb{Y}') \quad \forall Z, X, A, \mathbb{Y} \quad \text{and} \quad \mathbb{Y}'$$

where A is a network, X is a set of nodal covariates on the nodes, and \mathbb{Y} and \mathbb{Y}' are the potential outcomes.

To account for uncertainty in the network it is possible to make the assumption, as in Kao (2017), that the causal link between the network and outcomes can be broken by the inclusion of nodal covariates. Specifically:

ASSUMPTION 3. *Unconfounded Network Assumption under Network Interference*

$$P(A|X, \mathbb{Y}) = P(A|X, \mathbb{Y}') \quad \forall A, X, \mathbb{Y} \quad \text{and} \quad \mathbb{Y}'$$

where A is a network, X is a set of nodal covariates on the nodes, and \mathbb{Y} and \mathbb{Y}' are the potential outcomes.

We note that Assumption 2 also follows as a consequence of the chain graph approximation, as in the three node example in Figure 2 (See Lauritzen and Richardson, 2002, for discussion of this).

Assumption 3 essentially assumes away the main problem of network causal inference, that the network structure and the potential outcomes are related. The idea is that after including enough nodal covariates, the association between the network and the nodal potential outcomes breaks down. Note that this assumption trivially holds if the network is considered fixed.

Denote the missing potential outcomes as \mathbb{Y}_{miss} and the observed as \mathbb{Y}_{obs} . Using Assumptions 1, 2 and 3, Kao (2017) show that :

$$P(\mathbb{Y}_{miss}|X, Z, A, \mathbb{Y}_{obs}) = P(\mathbb{Y}_{miss}|X, \mathbb{Y}_{obs}) \quad (12)$$

Kao (2017) propose choosing covariates X to include so that Assumption 3 is met, and suggest that the addition of covariates derived from modelling the network A may be sufficient to achieve this.

We require only Assumptions 1 and 2, and modelling the potential outcomes jointly with the network. We argue that this is more realistic in most situations, and forces the researchers into more realistic modelling choices to arrive at causal inference for nodal outcomes.

The relaxation of Assumption 3 invalidates the proof of Equation (12), with the corresponding result under the relaxed assumption being:

$$P(\mathbb{Y}_{miss}|X, Z, A, \mathbb{Y}_{obs}) = P(\mathbb{Y}_{miss}|X, A, \mathbb{Y}_{obs}) \quad (13)$$

The retention of the network A in the conditioning variables, requires that we model its posterior distribution in the imputation of the missing potential outcomes.

5.2. Modelling

To facilitate modelling when the network is not fixed, we use the parametric ERNM class (Fellows and Handcock, 2012). The ERNM class can be viewed as a generalization of ERGM to allow for random nodal covariates. Alternatively, and equivalently, it can be viewed as a generalization of Markov Random Fields (MRF) that allows for random edges. The basic formulation is an exponential-family statistical distribution for network a with nodal covariates y :

$$P(A = a, Y = y|\theta, X = x) = \frac{1}{c(\theta, \mathcal{N})} \exp(\theta \cdot g(a, y, x)) \quad (14)$$

where the model parameter is $\theta \in H \subseteq \mathbb{R}^q$ and $g : \mathcal{A} \times \mathcal{Y} \times \mathcal{X} \rightarrow \mathbb{R}^q$ is a vector of network statistics, that are sufficient statistics for the network generative stochastic process. The choice of $g(a, y, x)$ specifies a different family of distributions parameterised by θ . Examples are the network density ($\frac{2}{N(N-1)} \sum_{i,j}^{N,N} a_{i,j}$) and homophily on y ($\sum_{i,j=1:N}^{N,N} \mathbb{I}(y_i = y_j)$). Both of these statistics are defined over dyads (sub-graphs of pairs of nodes). Most models

include statistics defined over triads (sub-graphs of triplets of nodes), especially measures of transitivity: the propensity for ties between pairs of nodes who are tied to a common third node. The sample space, \mathcal{N} , is the space of all possible binary edge realizations together with all of the possible random nodes, e.g., $\mathcal{N} \subset 2^{\mathbb{A}} \times \mathcal{Y}^n \times \mathcal{X}^n$, where $2^{\mathbb{A}}$ is the power set of all the dyads and $\mathcal{Y}^n \times \mathcal{X}^n$ is the joint sample space of the nodal covariates (stochastic and fixed). There is an extensive literature on exponential-family models for networks, including network statistic specification, inference and computation (see the citations in Section 1). For specific development of ERNM, see Fellows and Handcock (2012).

MRF models can be seen as the ERNM conditional on the network in Equation (15), and ERGMs the ERNM conditional on the nodal attributes. The normalising constant for the conditional distribution for the MRF is written $c(\theta, \mathcal{N}(a), a)$ to reflect summing over the restricted space (Fellows and Handcock, 2012).

$$P(Y = y|A = a, \theta, X = x) = \frac{1}{c(\theta, \mathcal{N}(a), a)} \exp(\theta \cdot g(a, y, x)) \quad (15)$$

We will suppress the dependence on the fixed covariates $X = x$ for notational simplicity. For example a typical MRF model on a network with treatment effect, outcome homophily as well as outcome and treatment neighbour effects can be written as:

$$\begin{aligned} P(Y = y|A = a, Z = z) = \frac{1}{c(\theta, \mathcal{N}(a), a)} \exp & \left(\theta_1 \cdot \sum_{i=1}^N y_i + \theta_2 \cdot \sum_{i=1}^N y_i z_i + \right. \\ & \theta_3 \cdot \sum_{i,j=1, a_{i,j}=1}^N \mathbb{I}(y_i = y_j) + \\ & \left. \theta_4 \cdot \sum_{i,j=1, a_{i,j}=1}^N y_i y_j + \theta_5 \cdot \sum_{i,j=1, a_{i,j}=1}^N y_i z_j \right) \end{aligned} \quad (16)$$

The full ERNM, however, is permitted to include terms to account for common social phenomenon, e.g., density, the overall propensity for connection, social transitivity, the tendency for edges to complete triangles of edges within triads, or social popularity - the tendency for some nodes to have many more connections than others. A typical ERNM model which in addition to the MRF terms in Equation (16) accounts for transitivity, and centralisation can be written as:

$$\begin{aligned} P(Y = y, A = a|Z = z) = \frac{1}{c(\theta, \mathcal{N}(a), a)} \exp & \left(\theta_1 \cdot \frac{2}{N(N-1)} \sum_{i=1, j=1}^N a_{i,j} + \right. \\ & \theta_2 \cdot \text{GWDEG}(a) + \theta_3 \cdot \text{GWESP}(a) + \theta_4 \cdot \sum_{i=1}^N y_i + \\ & \theta_5 \cdot \sum_{i=1}^N y_i z_i + \theta_6 \cdot \sum_{i,j=1, a_{i,j}=1}^N \mathbb{I}(y_i = y_j) + \\ & \left. \theta_7 \cdot \sum_{i,j=1, a_{i,j}=1}^N y_i y_j + \theta_8 \cdot \sum_{i,j=1, a_{i,j}=1}^N y_i z_j \right) \end{aligned}$$

where the first term measures the density of edges, and $\text{GWDEG}(a)$ and $\text{GWESP}(a)$ are the so called geometrically weighted degree (GWDEG) and geometrically weighted edgewise shared partner (GWESP) terms, respectively (Snijders et al., 2006). These statistics can include an additional so-called decay parameter that can in, principle, be estimated and can also be simply specified (Hunter and Handcock, 2006; Handcock et al., 2021). $\text{GWDEG}(a)$ measures the propensity for the network a to be centralized and $\text{GWESP}(a)$ is a measure of transitivity of the network a . These terms are motivated by consideration of dependency graphs for social network models (Snijders et al., 2006) to have weaker conditional independence restrictions than earlier models (Frank and Strauss, 1986). They are more able to avoid the model degeneracy resulting from the use of apparently simpler terms (Handcock, 2003; Snijders et al., 2006). There is an extensive literature on the choice and characteristics of such terms, (Robins et al., 2007; Morris et al., 2008; Handcock et al., 2021) where these and other useful terms are discussed in detail.

Tchetgen Tchetgen et al. (2020) utilized Markov random field (MRF) models with coding or pseudo-likelihood estimators to power a Gibbs sampling procedure, from which they estimated causal effects of treatment in situations where a single network was observed, and the observed data treatment was permitted to be affected by covariates. Ogburn et al. (2020) gave an example with the underlying network estimated by maximizing “penalized node-conditional likelihoods”. On top of the inferred network, they estimated MRF parameters from multiple observations of the nodes’ statuses.

Fellows and Handcock (2012) extended extensive work on MCMC MLE estimation for ERGM (Snijders, 2002; Hunter and Handcock, 2006) to ERNM. However, following the Bayesian paradigm, we simulate from the posterior distribution of the missing potential outcomes conditional on the observed data. This accounts for uncertainty in a theoretically consistent manner (O’Hagan and Kendall, 1993), and removes the need for asymptotic assumptions on the node set which are unrealistic, or bootstrapping (Efron, 1979) as in Ogburn et al. (2020).

Toulis and Kao (2013) also followed this paradigm, allowing for edge uncertainty with a Poisson edge model, and a linear outcome model. We note that their model does not account for the dyad dependent nature of real social processes, which is perhaps the most important feature of social network data.

We note that sampling from the posterior distribution of an ERNM is non-trivial as it requires the use of the exchange algorithm for so called doubly intractable distributions (Murray et al., 2006). Details are contained in the supplement.

With a suitably simple MRF model with a fixed network or a separable network and outcome model, the causal effects can usually be computed directly from the realized parameter values. For an ERNM this is not the case. Noting that, in full generality, each of the nodal potential outcomes depends on the whole network and all other nodal potential outcomes. The equilibrium distribution of a missing binary potential outcome for node i can be written:

$$P(\mathbb{Y}_{miss,i}(a, y^{-i})|X, A_{obs}, Y_{obs}) = \int_{\Theta} P(\mathbb{Y}_{miss,i}(a, y^{-i})|\theta, X, A_{obs}, Y_{obs})P(\theta|X, A_{obs}, Y_{obs})d\theta \quad (17)$$

We can then approximate this by simulating a large number of networks from Equation 14, $\{(A, Y)_j\}_{j=1}^M$ where $\{\theta_j\}_{j=1}^M$ are posterior draws of θ . Set $M(a, y^{-i}) = \{(a', y') \in \{(A, Y)_j\}_{j=1}^M : a' = a, y'^{-i} = y^{-i}\}$ be the number of these networks with the required network a and other nodal covariates y^{-i} . This yields

$$P(\mathbb{Y}_{miss,i}(a, y^{-i}) = 1 | X, A_{obs}, Y_{obs}) \approx \frac{1}{|M(a, y^{-i})|} \sum_{(a,y) \in M(a, y^{-i})} Y_{miss,i}(a, y) \quad (18)$$

Simulating enough networks that have the required network and nodal covariates is infeasible for networks of realistic size. Although, arbitrary k -step are possible, here we consider one-step neighbourhoods, that is, we allow only nodes connected to the ego to effect the nodal outcome. Thus, we dramatically reduce the number of unique potential outcomes for any given node, by requiring that only the treatment assignment, edges involving a node and outcomes of the neighbours of the nodes matter, not the whole network. We note that this is, in fact, a highly restrictive assumption, though it is useful to feasibly simulate the missing potential outcomes.

Concretely, to estimate the potential outcome for the k -peer effect estimand, for each i , instead of restricting to $M(a, y^{-i})$ we restrict to $M_{\mathbb{N}_i}^k(a, y^{-i}) = \{(a', y') \in \{(A, Y)_j\}_{j=1}^M : (a', y') \in (A, Y)_{\mathbb{N}_i}^k\}$. Where $(A, Y)_{\mathbb{N}_i}^k$ is defined in Definition 4. That is, we restrict to simulations where the correct neighbourhood is achieved to estimate the expected value of the missing potential outcomes.

$$P(\mathbb{Y}_{miss,i}(a, y^{-i}) = 1 | X, A_{obs}, Y_{obs}) \approx \frac{1}{|M_{\mathbb{N}_i}^k(a, y^{-i})|} \sum_{(a,y) \in M_{\mathbb{N}_i}^k(a, y^{-i})} Y_{miss,i}(a, y) \quad (19)$$

We can then use these expected potential outcomes, to estimate (Bayesian) expected versions of the causal estimands conditional on the observed data.

6. Example : Simulation Study

In this section, we conduct a simulation study to compare four different models. The primary goal is to assess how well each method can recover causal effects when the data generating process (DGP) is complex and unknown. The ERNM DGP allows stochastic tie and nodal variables and is capable of generating networks with complex transitivity, homophily and contagion. In addition, we can compute the posterior distribution of the causal estimands in each case. As such, we will use it as the primary generating process here. Other generating processes could have been selected, but would limit the scope of understanding the performance of our method as the true posterior would not be available.

6.1. A DAG compatible data generating process

In this section, we consider simulating from a 100 node network, with a procedure that follows the true DAG for a social process. We propose the following simulation procedure.

Algorithm 1: Figure 2 DAG simulation procedure**Result:** (A, Y) sampled from $2^{\mathbb{A}} \times \mathcal{Y}^n$, X fixed.Assign treatments Z to nodesLet $Y_i^0 = y^0 \quad \forall \quad i$ Let $A_{i,j}^0 = a_{i,j}^0 \quad \forall \quad i, j$ **for** $t = 1, 2, \dots, T$ **do** Simulate Y^t from $P(Y^t|Y^{t-1}, A^{t-1}, Z, X)$ Simulate A^t from $P(A^t|Y^t, A^{t-1}, Z, X)$ **end**

where (a^0, y^0) are initial values, typically a pre-treatment network.

The algorithm specified in Algorithm 1 is deliberately abstract. We do not specify the probability functions $P(Y^t|Y^{t-1}, A^{t-1}, Z, X)$ or $P(A^t|Y^t, A^{t-1}, Z, X)$ yet.

For our simulations, we chose $P(Y^t|Y^{t-1}, A^{t-1}, Z, X)$ and $P(A^t|Y^t, A^{t-1}, Z, X)$ as a logistic regression, using change statistics as predictors. Change statistics are simply the change in the network statistics due to a change in the network. That is, we allow for a proposed tie or node change to be more or less probable based on the corresponding change to some specified network statistics. We also suggest choosing change statistics in line with our intuition of the social processes. For example, edges that complete triangles of edges within triads are, all else equal, more likely to form than other edges.

We also make a slight simplification: We only allow a single edge or node to toggle at each time step. This results in the probability of a step being exactly the acceptance probability that would be used if we were using a Markov chain to sample from an ERNM, with a simple one edge or one vertex toggle as the proposal step. Thus, sampling from the DAG with this kind of model for a large enough T is equivalent to sampling from a Markov chain for the corresponding ERNM.

6.2. Model Specification

We consider four possible DGPs for $N = 100$ node networks where 50% for the nodes are treated before the social process evolves. The first is the ERNM itself. The second is the MRF model in the style of Tchetgen Tchetgen et al. (2020). The MRF formulation assumes a fixed network with parameters for outcome GWESP, outcome homophily, number of treated neighbours, positive outcome neighbours, main effect and intercept. This may represent a model that a researcher assuming a fixed network, with the simplistic chain graph approximation may adopt. Note that the MRF model can include terms that are functions of both edges and nodes, e.g., outcome GWESP and outcome homophily, but the calculation of these statistics only changes due to the nodes changing, not the edges changing. The third is an ERGM augmented with logistic regression. This accounts for stochastic variation in the network with the ERGM, and the stochastic variation in nodal outcomes with the logistic regression. Finally, we consider a pure logistic regression model where the network is only allowed to vary through the neighbour covariates in the logistic regression.

Table 1. Data Generating Process Summary. The first block are parameter values for the edge model, the second block are the parameter values for the node model, which in the case of the ERNMs are not separable. The third block gives a basic summary of the model classes. We consider networks generated by the ERNM and fit with the other DGPs, the parameter values for the other DGPs are shown to demonstrate the terms included in those models, not for model fitting. The Density parameters are $\times 10^4$.

	ERNM	MRF	ERGM+Logistic	Logistic
Density	-9.09	NA	-9.09	NA
GWESP	1	1	1	NA
GWDEG	1	NA	1	NA
Outcome Homophily	1	NA	1	NA
Intercept	1	1	1	1
Treatment	1	1	1	1
Neighbors Treated	0.1	0.1	0.1	0.1
Neighbors Outcomes	0.1	0.1	0.1	0.1
Stochastic Edges	Yes	No	Yes	No
Separable Likelihood	No	NA	Yes	NA

Table 1 shows the proposed parameter values and key properties of the model classes for the four DGPs. These ERNM parameters were chosen for simplicity and to achieve a mean degree of close to 3, which might be reasonable in, for example, a friendship network. We only present the results of simulation from the ERNM model, as it is the only model that represents the DAG.

The ERNM includes a mild spillover through a peer treatment effect, and contagion through a peer outcome effect, where in addition to homophily, peer outcomes and treatments also increase the chance of a positive nodal outcome. We also include a homophilous GWESP term on outcome, that is, a GWESP term measuring the transitivity among those with the same value of the outcome - outcome transitivity. We suggest that in many cases, a researcher would often believe that these effects are present in a social network formation process, and would fit such an ERNM to observed data.

We simulate networks from the ERNM DGP, and then fit the posterior distributions under the DGP which generated the ground truth posterior distribution. We then fit the posteriors of the remaining three mis-specified DGPs to the simulated networks, and compare the resulting posteriors to the ground truth posteriors.

6.3. Results

The causal estimands we consider are the treatment main effect, 1- to 5-peer outcome effects, 1- to 5-peer treatment effect. We simulate 100 network realizations from the ERNM, initialized with the empty network. For each of these realizations, we generated samples from the posterior parameter distribution for each of the 4 models. For each of the 400 posterior parameter distributions, we sampled 100 parameter realizations and estimated the causal estimands for those realizations. Thus, the output of the simulation

was 100 simulated networks with 4 posterior distributions for each network, for each of the causal estimands.

We note that this was a computationally demanding simulation. For each of 4 DGPs, we fit 100 posterior distributions to simulated networks from the ERNM. For each of these 100 distributions, we then drew 100 samples from the posterior, for each of which we simulated 100 networks to infer the missing potential outcomes. The fitting of each of the posterior distributions typically required of the order of 10^4 burn-in simulations with each step requiring a new ERNM MCMC, which required a toggle burn-in of order 10^4 . So each posterior fitting procedure required, the 10^8 ERNM toggles with associated change statistic calculation. As there were 4×10^2 posterior distribution required to be fit, the posterior fitting step required 4×10^{10} ERNM network toggles. Simulating and inferring the missing potential outcomes also requires MCMC burn-ins, though as multiple steps were not required it is a lower order component of the computation time.

Ordinarily the researcher would observe one network, fit one posterior and simulate networks to infer the causal effect, which is feasible for networks of the order of hundreds of nodes.

Table 2 shows the mean posterior-mean and the Frequentist coverage rates of the 95% Bayesian credible intervals, together with the true causal estimands of the DGP. The coverage rates are included to enable calibration of the credible intervals (Little, 2011). We also show the mean mean-a-posteriori to justify that, on average, the posteriors are centred around the true value.

We note that the ERGM logistic and pure logistic models recover some outcome effects on average, but perform very poorly on treatment effects. The ERNM and MRF posteriors seem to broadly be centred close to the true effects, though the MRF posteriors have much lower Frequentist coverage than the ERNM model, perhaps suggesting optimistically low variance in the posterior distribution. This is as expected as the MRF model does not account for randomness in the edges of the network. In addition, the MRF model was unable to identify higher order peer treatment effects, denoted as NA in Table 2. This is because 4- and 5- peer treatments were not observed in any of the simulated networks.

However, if we work consistently in the Bayesian framework, as the networks were generated from an ERNM, the posterior causal estimand distribution derived from the ERNM posterior, is the “ground truth” in the Bayesian sense. Thus, the correct assessment of the performance of any given method should be done by comparing its posterior causal estimand distribution to the ground truth distribution. The comparison for a given method is to compare its posterior fit based on each of the 100 simulated networks to the corresponding posterior derived from the true DGP. Therefore, understanding the performance of each method reduces to comparing distributions. We use the relative distribution (Handcock and Morris, 1999) to this end. We consider the relative rank distribution of each of the pairs of models, using boundary adjusted kernel density estimation using the `reldist` package (Handcock, 2015).

Table 3 shows the estimated Kullback-Leibler (KL) divergences between the relative rank distribution and the uniform distribution. To calibrate the size of the divergences, the

Table 2. Mean mean a-posteriori causal estimands fitted to 100 network simulations from ERNM. The coverage of the true mean by the 100 estimated 95% credible intervals is shown in brackets.

	True	ERNM	MRF	ERG+Logistic	Logistic
main	0.28	0.27 (65%)	0.28 (67%)	-0.03 (0%)	-0.17 (0%)
1-peer-out	0.28	0.27 (69%)	0.36 (31%)	0.18 (59%)	0.16 (8%)
2-peer-out	0.50	0.50 (68%)	0.64 (21%)	0.45 (98%)	0.39 (53%)
3-peer-out	0.66	0.65 (63%)	0.77 (34%)	0.66 (97%)	0.58 (72%)
4-peer-out	0.77	0.74 (57%)	0.82 (52%)	0.76 (95%)	0.70 (77%)
5-peer-out	0.82	0.80 (58%)	0.83 (62%)	0.80 (95%)	0.76 (81%)
1-peer-treat	0.13	0.14 (80%)	0.16 (69%)	-0.06 (0%)	0.00 (11%)
2-peer-treat	0.27	0.27 (77%)	0.28 (70%)	-0.12 (0%)	0.00 (13%)
3-peer-treat	0.39	0.39 (71%)	0.40 (70%)	-0.16 (0%)	0.00 (14%)
4-peer-treat	0.50	0.48 (72%)	NA	-0.20 (0%)	0.01 (15%)
5-peer-treat	0.56	0.54 (70%)	NA	-0.23 (1%)	0.01 (14%)

KL divergence between two unit variance Gaussian distributions is equal to one-half the squared difference between their means. On this scale, a KL divergence of d corresponds to a $\sqrt{2d}$ mean difference. As the ERNM model is being compared against itself, we expect the divergence to be 0. The others have large KL divergences from the ERNM posterior, suggesting that they are not able to recreate the true posterior distribution of important causal estimands, when misspecified.

7. Case-Study of Smoking Behavior within a High School

In this section, we give a real data case-study utilizing ERNM to relax the fixed network assumption as well as conditional unconfoundedness of the edges and nodal potential outcomes. In this case we do not know the ground truth, so the purpose of this section is to demonstrate that our method produces plausible posterior estimand distributions and to highlight the differences between methods for these data.

We consider one of the school social networks from the National Longitudinal Study of Adolescent Health (Harris et al., 2007). The network we used represents all students in the school. The social network is constructed by data collected from an in-school survey of all students in grades 7-12 of the school ($n = 869$). The instrument asked them to nominate up to five of their best friends of each sex within the school. Two students were friends if they each nominated each other. There were 462 were male and 407 female students, with 344 having reported trying a cigarette at least once (“smokers”). We note that the outside world (e.g., friends outside of school and family members) will likely have an effect, even though within school and grades have strong direct effects. Our modeling, in common with almost all modeling of these networks, assumes the social network processes are unconfounded by the outside world.

We consider the following estimands and estimate them under different frameworks:

- (a) k -peer effect of the gender of peers on smoker status of the ego

Table 3. Mean KL divergence of relative rank distributions of posteriors for causal estimands across 100 network simulations from the ERNM

	ERNM	MRF	ERGM+Logistic	Logistic
main	0	1.09	3.77	4.44
1-peer-out	0	2.45	2.18	3.33
2-peer-out	0	2.59	1.18	1.88
3-peer-out	0	2.20	1.08	1.26
4-peer-out	0	1.56	1.03	1.13
5-peer-out	0	1.22	0.93	1.08
1-peer-treat	0	1.24	4.23	3.52
2-peer-treat	0	1.25	4.24	3.52
3-peer-treat	0	1.34	4.23	3.48
4-peer-treat	0	NA	4.21	3.44
5-peer-treat	0	NA	4.19	3.41

(b) k -peer outcome effect of peer smoker status on smoker status of the ego.

Within our Bayesian framework, we consider the following models for imputing the required potential outcomes to claim causal inference:

- (a) Full ERNM with potential outcome as a random nodal covariates.
- (b) Markov random field model with fixed network
- (c) Logistic regression.

ERNMs allow a much wider choice of terms than the other models and we choose them to represent density (Density), homophily on grade, sex, and smoking (Grade Homophily, Sex Homophily, Smoke Homophily), transitivity within grade (Grade GWESP), centrality (GWDEG), as well as terms able to be represented in the MRF and logistic regression models. Here we use the ERNM-style homophily term developed in Fellows and Handcock (2012).

In our framing, the outputs are posterior distributions of causal estimands. For information, we also show the results of the MRF model, with parameters estimated through maximum pseudo-likelihood estimation and potential outcomes derived from these as in Tchetgen Tchetgen et al. (2020). In line with the known bias of pseudo-likelihood estimates for ERGM (Duijn et al., 2009) we believe this method can perform poorly.

We used a version of the exchange algorithm (Murray et al., 2006), with an extension which allows for efficient sampling from the ERNM posterior. The development is given in the supplement. While informative priors are compatible with the computational framework, here we report based on a uniform prior over all parameters. The results do not appear to be sensitive to the choice of prior.

Table 4 gives a summary of the posterior distributions for each of the models, showing the posterior means with the posterior standard errors in parentheses. We note that

Table 4. Summary of posterior distributions of network models. Posterior means are shown with posterior standard errors in parentheses. The Density parameters are $\times 10^4$.

	ERNM	MRF	Logistic Regression
Density	-9.96 (0.06)	NA	NA
Grade GWESP	0.11 (0.00)	NA	NA
GWDEG	-1.67 (0.36)	NA	NA
Grade Homophily	4.13 (0.05)	NA	NA
Sex Homophily	0.67 (0.08)	NA	NA
Smoke Homophily	0.48 (0.06)	0.43 (0.07)	NA
Intercept	0.95 (0.11)	0.75 (0.12)	-1.07 (0.2)
Male	-0.07 (0.03)	0.02 (0.04)	0.47 (0.16)
Female neighbors	0.11 (0.02)	0.08 (0.02)	-0.17 (0.05)
Smoker neighbours	-0.64 (0.12)	-0.46 (0.12)	0.54 (0.06)
Stochastic Edges	Yes	No	No
Stochastic Covariates	Smoker Status	Smoker Status	Smoker Status
Separable Likelihood	No	NA	NA

parameters should not be compared across models, as the functional forms are different, we show this table to summarize the posteriors, but to also highlight the differences between the models. We do not interpret the posterior parameter distributions directly, rather we make comparison through the smoker peer effect.

Table 5 and Figure 3 show the k -smoker-peer effect estimates. These are estimated as the additional chance of smoking that having k smoker friends has over having no smoker friends. The ERNM and MRF model are in agreement for peers one to three, with some divergence after this. The logistic regression model is markedly different from the ERNM for one and two peer effects, while for higher effects the estimates are closer to the ERNM estimates. The MRF model estimates are quite different and the estimated model does not fit the data well, nor accurately estimate the causal estimands in this case.

We believe the ERNM to be most plausible from a theoretical perspective. Whilst in this example the effect size difference from the MRF model was not large, we believe it to be a more robust approach when estimating network causal effects. In particular, where there is strong transitivity interacting with nodal outcomes as well as for smaller networks, we expect the effect would be larger.


In the context of our problem, the particular advantage of ERNM is that for simulated networks the smoker nodes are observed within network sub-structures consistent with the observed network. Figures 4 and 5 compare the distributions of the proportion of smoker edges and triads in the networks simulated from each model. We consider the proportion of smoker triads, as all model simulations underestimate the absolute number of triads. We note that the ERNM model and the full MRF fit considerably better than the pseudo-MRF and the logistic regression model.

We show Bayesian posterior prediction goodness-of-fit graphics in the style of Hunter et al.

Table 5. Posterior means of the average k -peer smoker outcome effects, for various methods. The pseudo-MRF value is the mean simulated from the parameter estimate.

	ERNM	MRF	pseudo-MRF	Logistic Regression
$k=1$	0.12 (0.01)	0.13 (0.01)	0.17 (0.02)	0.08 (<0.01)
$k=2$	0.22 (0.02)	0.23 (0.02)	0.28 (0.02)	0.18 (0.01)
$k=3$	0.35 (0.02)	0.33 (0.02)	0.35 (0.02)	0.30 (0.02)
$k=4$	0.48 (0.04)	0.42 (0.03)	0.40 (0.02)	0.42 (0.03)
$k=5$	0.59 (0.04)	0.50 (0.03)	0.45 (0.02)	0.53 (0.04)
$k=6$	0.68 (0.05)	0.57 (0.04)	0.49 (0.02)	0.63 (0.04)
$k=7$	0.75 (0.04)	0.63 (0.04)	0.52 (0.02)	0.70 (0.04)

(2008) in the supplement. These demonstrate that networks simulated from the ERNM model posterior correspond closely to the observed data.

**Fig. 3.** Plot comparing the posteriors distribution of the ERNM, MRF and logistic regression estimated k -peer outcome effect. The MRF with parameters equal to the maximum pseudo-likelihood estimate is also shown

8. Discussion

In this paper, we model causality when the underlying population is networked and that network is endogenous to the social process. Our approach jointly stochastically models the links and nodal covariates in the network, better representing our state of knowledge and their codependency.

Considering a DAG, we suggest that estimating true causal effects in this setting, is almost always infeasible due to the lack of fine-grained temporal data, as well as the typically highly complex causal structure. As an alternative, we present a chain graph approximation to the DAG, which allows for a dependence structure that we believe to be compatible with such settings. We then frame approximate causal inference in terms of network equilibrium potential outcomes, that is, potential outcomes that are free to depend on nodes in the neighbourhood of the node in question. We propose the use of ERNMs to jointly model both random edges and nodal covariates. We also develop a simple modification




Fig. 4. Distribution of edge density when both incident nodes are smokers. ERNM and MRF posterior simulated distributions fit the observed data considerably better than logistic regression or the pseudo-likelihood estimated MRF.




Fig. 5. Distribution of triad density when all incident nodes are smokers. ERNM and MRF posterior simulated distributions fit the observed data considerably better than logistic regression or the pseudo-likelihood estimated MRF.

to the exchange algorithm, allowing for feasible sampling from the posterior distribution. We use the posterior distribution, through simulation, to impute the distribution of the missing potential outcomes, allowing the consideration of the distribution of the causal estimands. We show, using a school network from the National Longitudinal Study of Adolescent Health, that failing to account for the network structure of the problem could lead to erroneous qualitative conclusions, in particular when considering one- and two-peer outcome effects.

Our primary contributions lie in considering the consequences of relaxing the commonly made fixed network assumption and proposing the use of suitable social network models to estimate causal effects. In particular, we do not need to make Assumption 3 in Section 5.1. This relaxation complicates the causal structure considerably and necessitates the use of complex models to derive causal estimands. Clearly, the relaxation of this assumption also allows for greater generalizability.

In our simulations and case-study, we make direct comparisons to Tchetgen Tchetgen et al. (2020), which is the closest model in the literature to the setting of the paper. Aronow and Samii (2017) use a randomisation approach, while Hudgens and Halloran (2008); Ogburn et al. (2020, 2022); Kalisch et al. (2012) make assumptions that are less appropriate compared to the latest alternative, which is Tchetgen Tchetgen et al. (2020). Our results show our method is superior to Tchetgen Tchetgen et al. (2020). Aronow and Samii (2017) presumes a complete knowledge of the stochastic allocation of treatments to nodes (i.e., treatment assignment mechanism). This is not generally true for our situation and not true for our case-study, so we do not make a comparison here.

8.1. Limitations

There are a number of limitations of the approach, some based on the social process model and some the limit of causal inference. Our inferences hold for the given node set and social process, whereas assuming the fixed network narrows the scope to that observed network only. However, we suggest that the posteriors derived from the given network can serve as strong priors for “similar” networks. In general, qualitative features of the posteriors from the given network can usefully inform other analyses. It is not generally possible to make further statements than this for networks of different sizes.

Network generation processes of interest are typically complex and this complexity tends to increase with the network size (i.e., node set, edge set) (Schweinberger et al., 2020). Most network models do not match this complexity as it increases and so tend to poorly represent some aspects of this complexity. As such, our casual model, and indeed most statistical models we know of will be poor representations of large complex networks and should be used with caution in this case. We emphasize that this is an intrinsic mismatch between the complexity of network models and the complexity of large network generation processes and not a deficiency of our approach separately. Specifically, it is not that network density or sparsity are intrinsically an issue, but that likely complexity mismatch is the issue. In general, using this approach with a very large network would require strong prior belief in simple structure (e.g., a Erdos-Renyi model or a stochastic block model). Alternatively, the fixed network approach may be easier for very large networks,

as modeling of any kind may require overly strong assumptions.

Another limitation is the assumption of no outside influences. This is assumed, often implicitly, in almost all causal frameworks, including those that are not over networks. Separately, almost all social network models assume processes are unconfounded by the outside world given the observed world. For example, in the school friendship application of Section 7 there is an implicit assumption that the smoking behavior of the students in the school are unaffected by other friends outside of school, family or other community members. It is likely that these outside influences will have an effect, even though schools and classes have strong direct effects. Our framing brings these assumptions into focus, and they are almost always unquestioned assumptions of network and causal modeling.

Much more could be done if the network was observed at multiple time points. In this situation, dynamic models such as SAOM (Steglich et al., 2010; Niezink et al., 2019) could be used as a replacement to the ERNM. It is also possible to use SAOMs to represent the latent evolutionary processes, although the identification of these models would be challenging in most causal settings.

Notable by its absence is discussion of suitable prior distributions for ERNMs. We note there has been some work developing conjugate prior distributions for ERGMs (Wang, 2011) and strongly suspect that such an approach may also be applicable in our setting. In practice, flat Gaussian prior are often used for ERGMs (Caimo and Friel, 2011). For the purpose of demonstrating our approach, we used uniform priors, which make no account for the geometry of exponential families, but allow us to defer careful consideration of possible priors to future work, while demonstrating the utility of our method. We performed similar posterior fits with flat Gaussian priors, with very similar results.

The cost of our approach is the strong assumption that such a complex process can be adequately modelled by an ERNM. This is, in general, the main criticism of model based causal inference approaches: that models are mis-specified with unknown consequences. In a network setting this mis-specification is often acute e.g. constant marginal effect of additional smoker friends in the linear potential outcomes model. Our central argument is that a complex model is much less mis-specified than current approaches. We have sought to justify this with real data and simulations, though propose this as a future area of research. For example, how dependent do outcomes in networks need to be to invalidate conclusions made with unrealistic models?

The mis-specification may seem to be cause for pessimism. However we emphasize that network settings are indeed the extreme case of small data, as we usually only have one observation of a network on a fixed set of nodes. Thus, intuitively, we should expect strict functional form assumptions to be required to generate any meaningful statistical, and especially causal, inference. In fact, we argue that approaching network problems with simpler assumptions is problematic. While potentially less prone to mis-specification in the sense that simple models can be used, this glosses over the inherent difficulty of dealing with network data where nodes and edges are strongly dependent on other nodes and edges.

8.2. Conclusions

We believe that in the context of social network analysis, where individual attributes are heavily influenced by social context, such covariation of edges and nodal covariates is overwhelmingly more representative of many social processes. That is, in many social networks we believe it highly likely that individual characteristics and the connections that form between individuals are strongly dependent. This is especially true if the network evolves over time. Our approach is fully Bayesian on the fixed node set of the network and does not require assumptions that the number of nodes in the network is large for its validity.

We also note that specifying a model for the full network data generation process also allows inference in cases where the network is only partially observed (e.g., network sampling Handcock and Gile, 2010). Accounting for such sampling structure in ERNM is analogous to the method for ERGMs (Fellows and Handcock, 2023; Gile and Handcock, 2016). Accounting for this is not possible with the other methods considered in this work. In addition, the network process can be considered to evolve after treatment conditional on some pre-treatment network. Such an approach may lead to increased power for randomised control trials on a networked population, at the cost of our modelling assumptions.

We believe that meaningful steps can be made towards causal inference on networks, through careful consideration of the complex causal structure of such problems. Whilst we make strict assumption on the function form of this structure, if the researcher is unwilling to make such assumptions, we opine casual inference is out of reach. We suggest it is better to acknowledge the complexity of the situation, and therefore claim that causal inference is not possible, than employing highly restrictive assumptions on the dependence structure of the data generating process to allow simpler models to be employed.

9. Acknowledgements

This article is based upon work supported by the National Science Foundation(NSF, MMS-0851555, SES-1357619, IIS-1546259) and National Institute of Child Health and Human Development (NICHD, R21HD063000, R21HD075714 and R24-HD041022). The content is solely the responsibility of the authors and do not necessarily represent the official views of the National Institutes of Health or the National Science Foundation.

10. Supplementary Materials

Supplement The supplement contains an additional chain graph approximation diagram, a review of ERNMs and Bayesian computation for them, as well as an MCMC convergence analysis for the adolescent health network. (pdf)

References

Aronow, P. M. and C. Samii (2017, 12). Estimating average causal effects under general interference, with application to a social network experiment. *Ann. Appl. Stat.* 11(4), 1912–1947.

- Caimo, A. and N. Friel (2011). Bayesian inference for exponential random graph models. *Social Networks* 33(1), 41 – 55.
- Christakis, N. A. and J. H. Fowler (2007). The spread of obesity in a large social network over 32 years. *New England Journal of Medicine* 357(4), 370–379. PMID: 17652652.
- Christakis, N. A. and J. H. Fowler (2008). The collective dynamics of smoking in a large social network. *New England Journal of Medicine* 358(21), 2249–2258. PMID: 18499567.
- Christakis, N. A. and J. H. Fowler (2010, 09). Social network sensors for early detection of contagious outbreaks. *PLOS ONE* 5(9), 1–8.
- Daraganova, G. and G. Robins (2012). Autologistic actor attribute models. In D. Lusher, J. Koskinen, and G. Robins (Eds.), *Exponential Random Graph Models for Social Networks: Theory, Methods, and Applications*, Structural Analysis in the Social Sciences, pp. 102–114. Cambridge University Press.
- DeAmour, A. (2016). *Misspecification, Sparsity, and Superpopulation Inference for Sparse Social Networks*. Ph. D. thesis, Harvard University.
- Duijn, M., K. J. Gile, and M. S. Handcock (2009). A framework for the comparison of maximum pseudo likelihood and maximum likelihood estimation of exponential family random graph models. *Social networks* 31, 52–62.
- Durante, D., D. B. Dunson, and J. T. Vogelstein (2017). Nonparametric Bayes Modeling of Populations of Networks. *Journal of the American Statistical Association* 112(520), 1516–1530.
- Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* 7(1), 1 – 26.
- Fellows, I. and M. S. Handcock (2012). Exponential-family random network models. DOI: 10.48550/arXiv.1208.012, URL: <https://arxiv.org/abs/1208.0121>.
- Fellows, I. E. and M. S. Handcock (2023). Modeling of networked populations when data is sampled or missing. *Metron* 81(1), 21–35.
- Fosdick, B. K. and P. D. Hoff (2015). Testing and modeling dependencies between a network and nodal attributes. *Journal of the American Statistical Association* 110(511), 1047–1056.
- Frank, O. and D. Strauss (1986). Markov Graphs. *Journal of the American Statistical Association* 81(395), 832–842.
- Frydenberg, M. (1990). The Chain Graph Markov Property. *Scandinavian Journal of Statistics* 17(4), 333–353.
- Gile, K. J. and M. S. Handcock (2016). Analysis of Networks with Missing Data with Application to the National Longitudinal Study of Adolescent Health. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 66, 501–519.

- Goldenberg, A., A. X. Zheng, S. E. Fienberg, and E. M. Airoldi (2010). A survey of statistical network models. *Foundations and Trends® in Machine Learning* 2(2), 129–233.
- Handcock, M. S. (2002). Degeneracy and inference for social network models. In *Paper presented at the Sunbelt XXII International Social Network Conference in New Orleans, LA*.
- Handcock, M. S. (2003). Assessing degeneracy in statistical models of social networks. Working paper #39, Center for Statistics and the Social Sciences, University of Washington.
- Handcock, M. S. (2015). *reldist: Relative Distribution Methods*. Los Angeles, CA: CRAN. Version 1.6-4. Project home page at <https://faculty.stat.ucla.edu/handcock/RelDist/>.
- Handcock, M. S. and K. J. Gile (2010). Modeling social networks from sampled data. *Annals of Applied Statistics* 4(1), 5–25.
- Handcock, M. S., D. R. Hunter, C. T. Butts, S. M. Goodreau, P. N. Krivitsky, and M. Morris (2021). *ergm: Fit, Simulate and Diagnose Exponential-Family Models for Networks*. The Statnet Project (<https://statnet.org>). R package version 4.0-6406.
- Handcock, M. S. and M. Morris (1999). *Relative Distribution Methods in the Social Sciences*. New York: Springer. ISBN 0-387-98778-9.
- Harris, K., C. Halpern, A. Smolen, and B. Haberstick (2007, 01). The national longitudinal study of adolescent health (add health) twin data. *Twin research and human genetics : the official journal of the International Society for Twin Studies* 9, 988–97.
- Hoff, P. D. (2005). Bilinear mixed-effects models for dyadic data. *Journal of the American Statistical Association* 100(469), 286–295.
- Hudgens, M. G. and M. E. Halloran (2008). Toward causal inference with interference. *Journal of the American Statistical Association* 103(482), 832–842. PMID: 19081744.
- Hunter, D. R., S. M. Goodreau, and M. S. Handcock (2008). Goodness of fit of social network models. *Journal of the American Statistical Association* 103(481), 248–258.
- Hunter, D. R. and M. S. Handcock (2006). Inference in curved exponential family models for networks. *Journal of Computational and Graphical Statistics* 15(3), 565–583.
- Hunter, D. R., M. S. Handcock, C. T. Butts, S. M. Goodreau, and M. Morris (2008). ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software* 24(3), 1–29.
- Imbens, G. W. and D. B. Rubin (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press.
- Kalisch, M., M. Mächler, D. Colombo, M. H. Maathuis, and P. Bühlmann (2012). Causal inference using graphical models with the R package pcalg. *Journal of Statistical Software* 47(11), 1–26.

- Kao, E. (2017). *Causal Inference Under Network Interference: A Framework for Experiments on Social Networks*. Ph. D. thesis, Harvard University.
- Koskinen, J. and G. Daraganova (2022, 04). Bayesian Analysis of Social Influence. *Journal of the Royal Statistical Society Series A: Statistics in Society* 185(4), 1855–1881.
- Krivitsky, P. N., M. S. Handcock, and M. Morris (2011). Adjusting for network size and composition effects in exponential-family random graph models. *Statistical Methodology* 8(4), 319–339.
- Lauritzen, S., A. Rinaldo, and K. Sadeghi (2018, 01). Random Networks, Graphical Models and Exchangeability. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 80(3), 481–508.
- Lauritzen, S. L. and T. S. Richardson (2002). Chain graph models and their causal interpretations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64(3), 321–348.
- Lee, Y. and E. L. Ogburn (2021). Network dependence can lead to spurious associations and invalid inference. *Journal of the American Statistical Association* 116(535), 1060–1074.
- Liang, F., I. H. Jin, Q. Song, and J. S. Liu (2016). An adaptive exchange algorithm for sampling from distributions with intractable normalizing constants. *Journal of the American Statistical Association* 111(513), 377–393.
- Little, R. (2011). Calibrated Bayes, for Statistics in General, and Missing Data in Particular. *Statistical Science* 26(2), 162 – 174.
- Morris, M., M. S. Handcock, and D. R. Hunter (2008). Specification of exponential-family random graph models: Terms and computational aspects. *Journal of Statistical Software* 24(4).
- Murray, I., Z. Ghahramani, and D. J. C. MacKay (2006). Mcmc for doubly-intractable distributions. In *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, UAI’06, Arlington, Virginia, USA, pp. 359–366. AUAI Press.
- Niezink, N. M. D., T. A. B. Snijders, and M. A. J. van Duijn (2019). No longer discrete: Modeling the dynamics of social networks and continuous behavior. *Sociological Methodology* 49(1), 295–340.
- Ogburn, E. L., I. Shpitser, and Y. Lee (2020). Causal inference, social networks and chain graphs. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 183(4), 1659–1676.
- Ogburn, E. L., O. Sofrygin, I. Díaz, and M. J. van der Laan (2022). Causal inference for social network data. *Journal of the American Statistical Association* 0(0), 1–15.
- Ogburn, E. L. and T. J. VanderWeele (2014, 11). Causal diagrams for interference. *Statist. Sci.* 29(4), 559–578.
- O’Hagan, A. and M. Kendall (1993). *Bayesian Inference*. Kendall’s advanced theory of statistics. Arnold.

- Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika* 82(4), 669–688.
- Pearl, J. (2009). *Causality: Models, Reasoning and Inference* (2nd ed.). USA: Cambridge University Press.
- Richardson, T. S. and J. M. Robins (2013). Single world intervention graphs (swigs): A unification of the counterfactual and graphical approaches to causality. *Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper 128*(30), 2013.
- Robins, G., P. Elliott, and P. Pattison (2001, 01). Network models for social selection processes. *Social Networks* 23, 1–30.
- Robins, G., P. Pattison, and P. Elliott (2001). Network models for social influence processes. *Psychometrika* 66(2), 161–189.
- Robins, G., T. Snijders, P. Wang, M. Handcock, and P. Pattison (2007). Recent developments in exponential random graph (p) models for social networks. *Social Networks* 29(2), 192–215.
- Sävje, F., P. M. Aronow, and M. G. Hudgens (2021). Average treatment effects in the presence of unknown interference. *The Annals of Statistics* 49(2), 673 – 701.
- Schweinberger, M. and M. S. Handcock (2015). Local dependence in random graph models: characterization, properties and statistical inference. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 77(3), 647–676.
- Schweinberger, M., P. N. Krivitsky, C. T. Butts, and J. R. Stewart (2020). Exponential-Family Models of Random Graphs: Inference in Finite, Super and Infinite Population Scenarios. *Statistical Science* 35(4), 627 – 662.
- Shalizi, C. and A. Thomas (2011, 05). Homophily and contagion are generically confounded in observational social network studies. *Sociological methods and research* 40, 211–239.
- Shpitser, I., E. T. Tchetgen, and R. Andrews (2021). Modeling interference via symmetric treatment decomposition.
- Snijders, T. A. B. (2002). Markov chain Monte Carlo estimation of exponential random graph models. *Journal of Social Structure* 3(2), 1–41.
- Snijders, T. A. B., P. E. Pattison, G. L. Robins, and M. S. Handcock (2006). New specifications for exponential random graph models. *Sociological Methodology* 36(1), 99–153.
- Sofrygin, O. and M. J. van der Laan (2017). Semi-parametric estimation and inference for the mean outcome of the single time-point intervention in a causally connected population. *Journal of Causal Inference* 5(1), 20160003.
- Spirtes, P., C. Glymour, and R. Scheines (2000). *Causation, Prediction, and Search* (2nd ed.). MIT press.

- Steglich, C., T. A. B. Snijders, and M. Pearson (2010). Dynamic networks and behavior: Separating selection from influence. *Sociological Methodology* 40(1), 329–393.
- Tchetgen Tchetgen, E. J., I. R. Fulcher, and I. Shpitser (2020). Auto-g-computation of causal effects on a network. *Journal of the American Statistical Association* 0(0), 1–12.
- Toulis, P. and E. Kao (2013, 17–19 Jun). Estimation of causal peer influence effects. In S. Dasgupta and D. McAllester (Eds.), *Proceedings of the 30th International Conference on Machine Learning*, Volume 28 of *Proceedings of Machine Learning Research*, Atlanta, Georgia, USA, pp. 1489–1497. PMLR.
- van der Laan, M. J. (2014). Causal inference for a population of causally connected units. *Journal of Causal Inference* 2(1), 13–74.
- Wang, R. (2011). *Likelihood-based inference of exponential-family random graph models for social networks*. Ph. D. thesis, University of Washington.