CONCENTRATION AND CONSISTENCY RESULTS FOR CANONICAL AND CURVED EXPONENTIAL-FAMILY MODELS OF RANDOM GRAPHS

By Michael Schweinberger^{*} and Jonathan Stewart^{**}

Department of Statistics, Rice University, *m.s@rice.edu; ** jonathan.stewart@rice.edu

Statistical inference for exponential-family models of random graphs with dependent edges is challenging. We stress the importance of additional structure and show that additional structure facilitates statistical inference. A simple example of a random graph with additional structure is a random graph with neighborhoods and local dependence within neighborhoods. We develop the first concentration and consistency results for maximum likelihood and *M*-estimators of a wide range of canonical and curved exponential-family models of random graphs with local dependence. All results are nonasymptotic and applicable to random graphs with finite populations of nodes, although asymptotic consistency results can be obtained as well. In addition, we show that additional structure can facilitate subgraph-to-graph estimation, and present concentration results for subgraph-to-graph estimators. As an application, we consider popular curved exponential-family models of random graphs, with local dependence induced by transitivity and parameter vectors whose dimensions depend on the number of nodes.

1. Introduction. Models of network data have witnessed a surge of interest in statistics and related areas (e.g., [31]). Such data arise in the study of, for example, social networks, epidemics, insurgencies and terrorist networks.

Since the work of Holland and Leinhardt in the 1970s (e.g., [21]), it is known that network data exhibit a wide range of dependencies induced by transitivity and other interesting network phenomena (e.g., [39]). Transitivity is a form of triadic closure in the sense that, when a node k is connected to two distinct nodes i and j, then i and j are likely to be connected as well, which suggests that edges are dependent (e.g., [39]). A large statistical framework for modeling dependencies among edges is given by discrete exponential-family models of random graphs, called exponential-family random graphs (e.g., [15, 18, 20, 24, 36, 39, 53, 57]). Such models are popular among network scientists for the same reason Ising models are popular among physicists: Both classes of models enable scientists to model a wide range of dependencies of scientific interest (e.g., [39]).

Despite the appeal of the discrete exponential-family framework and its relationship to other discrete exponential-family models for dependent random variables (e.g., Ising models and discrete Markov random fields, [3, 9, 42]), statistical inference for exponential-family random graphs is challenging. One reason is that some exponential-family random graphs are ill-behaved (e.g., the so-called triangle model, [2, 6, 10, 18, 28, 47]), though well-behaved alternatives have been developed, among them curved exponential-family random graphs [24, 53]. A second reason is that in most applications of exponential-family random graphs statistical inference is based on a single observation of a random graph with dependent edges. Establishing desirable properties of estimators, such as consistency, is nontrivial when no more than one observation of a random graph with dependent edges is available. While some consistency results have been obtained under independence assumptions [11, 35, 40, 44, 51,

Received January 2018; revised December 2018.

MSC2010 subject classifications. Primary 05C80; secondary 62B05, 62F10, 91D30.

Key words and phrases. Curved exponential families, exponential families, exponential-family random graph models, *M*-estimators, multilevel networks, social networks.

58–60] and restrictive dependence assumptions [40, 51, 58]—as discussed in Section 5—the existing consistency results do not cover the models most widely used in practice (e.g., [39]): canonical and curved exponential-family random graphs with dependence among edges induced by transitivity and other interesting network phenomena [39].

We stress the importance of additional structure and show that additional structure facilitates statistical inference. We consider here a simple and common form of additional structure, called multilevel structure. Network data with multilevel structure are popular in network science, as the recent monograph of Lazega and Snijders [37] and a growing number of applications demonstrate (e.g., [22, 38, 52, 54, 56, 62]). A simple form of multilevel structure is given by a partition of a population of nodes into subsets of nodes, called neighborhoods. In applications, neighborhoods may correspond to school classes within schools, departments within companies and units of armed forces. It is worth noting that in multilevel networks the partition of the population of nodes is observed and models of multilevel networks attempt to capture dependencies within and between neighborhoods (e.g., [22, 38, 52, 54, 56, 62]), whereas the well-known class of stochastic block models [41] assumes that the partition is unobserved and that edges are independent conditional on the partition.

Additional structure in the form of multilevel structure offers opportunities in terms of statistical inference. We take advantage of these opportunities to develop the first statistical theory which shows that statistical inference for many canonical and curved exponential-family random graphs with dependent edges is possible. The main idea is based on a simple and general exponential-family argument that may be of independent interest. It helps establish nonasymptotic probability statements about estimators of canonical and curved exponential families for dependent random variables under weak conditions, as long as additional structure helps control the amount of dependence induced by the model and the sufficient statistics are sufficiently smooth functions of the random variables. We exploit the main idea to develop the first concentration and consistency results for maximum likelihood and M-estimators of canonical and curved exponential-family random graphs with dependent edges, under correct and incorrect model specifications. All results are nonasymptotic and applicable to random graphs with finite populations of nodes, although asymptotic consistency results can be obtained as well. In addition, we show that multilevel structure facilitates subgraph-to-graph estimation, and present concentration results for subgraph-to-graph estimators. As an application, we consider popular curved exponential-family random graphs [24, 53], with local dependence induced by transitivity and parameter vectors whose dimensions depend on the number of nodes.

These concentration and consistency results have important implications, both in statistical theory and practice:

- The most important implication is that statistical inference for transitivity and other network phenomena of great interest to network scientists is possible. To date, it has been widely believed that statistical inference for transitivity based on exponential-family random graphs is challenging (e.g., [10, 51]), but additional structure in the form of multilevel structure facilitates it.
- Network scientists can benefit from collecting network data with multilevel structure, because multilevel structure can facilitate statistical inference for exponential-family random graphs with dependent edges.

Last, but not least, it is worth noting that these concentration and consistency results cover two broad inference scenarios:

• Inference scenarios with finite populations of nodes. In many applications of exponentialfamily random graphs, there is a finite population of nodes and a population graph is assumed to have been generated by an exponential-family random graph model. A common goal of statistical inference, then, is to estimate the parameters of the data-generating exponential-family random graph model based on a complete or incomplete observation of the population graph. Our concentration results cover inference scenarios with finite populations of nodes, when the whole population graph is observed or when neighborhoods are sampled and the subgraphs induced by the sampled neighborhoods are observed.

• Inference scenarios with populations of nodes growing without bound. In addition, our concentration results can be used to obtain asymptotic consistency results by allowing the number of neighborhoods to grow without bound. The resulting asymptotic consistency results resemble asymptotic consistency results in other areas of statistics, albeit with two notable differences: first, the units of statistical analysis are subsets of nodes (neighborhoods) rather than nodes or edges; and, second, the sizes of the units need not be identical, but are similar in a well-defined sense.

Since the first application is more interesting than the second one, we state all results with the first application in mind, that is, all results focus on random graphs with finite populations of nodes, although we do mention some asymptotic consistency results along the way.

The remainder of our paper is structured as follows. Section 2 introduces models. Section 3 describes concentration and consistency results for maximum likelihood and M-estimators, under correct and incorrect model specifications. Section 4 shows that multilevel structure facilitates subgraph-to-graph estimation. A comparison with existing consistency results can be found in Section 5. Section 6 presents simulation results.

2. Exponential-family random graphs with multilevel structure. We introduce exponential-family random graphs with multilevel structure.

A simple and common form of multilevel structure is a partition of a population of nodes into $K \ge 2$ nonempty subsets of nodes $\mathcal{A}_1, \ldots, \mathcal{A}_K$, called neighborhoods. We note that in multilevel networks the partition of the population of nodes is observed (e.g., [22, 38, 52, 54, 56, 62]) and that some neighborhoods may be larger than others. We consider random graphs with undirected edges that may be either absent or present or may have weights, where the weights are elements of a countable set. Extensions to random graphs with directed edges are straightforward. Let $X = (X_k)_{k=1}^K$ and $Y = (Y_{k,l})_{k<l}^K$ be sequences of within- and between-neighborhood edge variables based on a sequence of neighborhoods $\mathcal{A}_1, \ldots, \mathcal{A}_K$, where $X_k = (X_{i,j})_{i \in \mathcal{A}_k < j \in \mathcal{A}_k}$ and $Y_{k,l} = (Y_{i,j})_{i \in \mathcal{A}_k, j \in \mathcal{A}_l}$ (k < l) correspond to within- and between-neighborhood edge variables $X_{i,j} \in \mathbb{X}_{i,j}$ and $Y_{i,j} \in \mathbb{Y}_{i,j}$, taking on values in countable sets $\mathbb{X}_{i,j}$ and $\mathbb{Y}_{i,j}$, respectively. We exclude self-edges, assuming that $X_{i,i} = 0$ holds with probability 1 ($i \in \mathcal{A}_k, k = 1, \ldots, K$), and write $\mathbb{X}_k = \prod_{i \in \mathcal{A}_k < j \in \mathcal{A}_k} \mathbb{X}_{i,j}$, $\mathbb{X} = \prod_{k=1}^K \prod_{i \in \mathcal{A}_k < j \in \mathcal{A}_k} \mathbb{X}_{i,j}$ and $\mathbb{Y} = \prod_{k<l}^K \prod_{i \in \mathcal{A}_k, j \in \mathcal{A}_l} \mathbb{Y}_{i,j}$. We assume that within-neighborhood edges X are independent of between-neighborhood

We assume that within-neighborhood edges X are independent of between-neighborhood edges Y, that is,

$$\mathbb{P}(X \in \mathcal{X}, Y \in \mathcal{Y}) = \mathbb{P}(X \in \mathcal{X})\mathbb{P}(Y \in \mathcal{Y}) \quad \text{for all } \mathcal{X} \times \mathcal{Y} \subseteq \mathbb{X} \times \mathbb{Y},$$

where \mathbb{P} denotes a probability distribution with support $\mathbb{X} \times \mathbb{Y}$. We do not assume that edges are independent, neither within nor between neighborhoods.

While in principle both within-neighborhood edge variables X and between-neighborhood edge variables Y may be of interest, we focus on within-neighborhood edge variables, which are of primary interest in applications (e.g., [22, 37, 38, 52, 54, 56, 62]). We therefore restrict attention to the probability law of X and do not make assumptions about the probability law of Y. We assume that the parameter vectors of the probability laws of X and Y are variation-independent, that is, the parameter space is a product space, so that statistical inference concerning the parameter vector of the probability law of X can be based on X without requiring knowledge of Y.

The distribution of within-neighborhood edge variables X is presumed to belong to an exponential family with local dependence, defined as follows.

DEFINITION (Exponential family with local dependence). An exponential family with local dependence is an exponential family of distributions with countable support X, having densities with respect to counting measure of the form

(2.1)
$$p_{\eta}(\boldsymbol{x}) = \exp(\langle \boldsymbol{\eta}, \boldsymbol{s}(\boldsymbol{x}) \rangle - \psi(\boldsymbol{\eta})) \boldsymbol{\nu}(\boldsymbol{x})$$
$$= \exp\left(\sum_{k=1}^{K} \langle \boldsymbol{\eta}_{k}, \boldsymbol{s}_{k}(\boldsymbol{x}_{k}) \rangle - \psi(\boldsymbol{\eta})\right) \boldsymbol{\nu}(\boldsymbol{x}),$$

...

where

$$\psi(\boldsymbol{\eta}) = \log \sum_{\boldsymbol{x}_1 \in \mathbb{X}_1} \cdots \sum_{\boldsymbol{x}_K \in \mathbb{X}_K} \exp\left(\sum_{k=1}^K \langle \boldsymbol{\eta}_k, s_k(\boldsymbol{x}_k) \rangle\right) \nu(\boldsymbol{x})$$

and $\nu(\mathbf{x}) = \prod_{k=1}^{K} \nu_k(\mathbf{x}_k)$.

In other words, edges may depend on other edges in the same neighborhood, but do not depend on edges in other neighborhoods [48]. Here, $\langle \eta, s(\boldsymbol{x}) \rangle = \sum_{k=1}^{K} \langle \eta_k, s_k(\boldsymbol{x}_k) \rangle$ is the inner product of a natural parameter vector $\eta \in \mathbb{R}^m$ and a sufficient statistic vector $s : \mathbb{X} \mapsto \mathbb{R}^m$ while $\eta_k \in \mathbb{R}^{m_k}$ and $s_k : \mathbb{X}_k \mapsto \mathbb{R}^{m_k}$ denote the natural parameter vector and sufficient statistic vector of neighborhood \mathcal{A}_k , respectively (k = 1, ..., K). The functions $v : \mathbb{X} \mapsto \mathbb{R}^+ \cup \{0\}$ and $v_k : \mathbb{X}_k \mapsto \mathbb{R}^+ \cup \{0\}$ (k = 1, ..., K) along with the sample space \mathbb{X} determine the reference measure of the exponential family. A careful discussion of the choice of reference measure can be found in Krivitsky [33].

We consider here a wide range of exponential families with local dependence. A specific example of an exponential family with local dependence can be found in Section 3.3. In the following, we introduce selected exponential-family terms in order to distinguish exponential families from subfamilies of exponential families. Subfamilies of exponential families give rise to distinct theoretical challenges, and thus require a separate treatment. We therefore introduce the classic exponential-family notions of full and nonfull exponential families, canonical and curved exponential families, and minimal exponential families. These exponentialfamily terms are taken from the monographs on exponential families by Barndorff-Nielsen [1] and Brown [5] and have been used in other recent works as well: see, for example, Lauritzen, Rinaldo and Sadeghi [36], Rinaldo, Fienberg and Zhou [43] and Geyer [16]. To help ensure that parameters are identifiable, we assume that exponential families of the form (2.1)are minimal in the sense of Barndorff-Nielsen [1] and Brown [5], that is, the closure of the convex hull of the set $\{s(\mathbf{x}) : v(\mathbf{x}) > 0\}$ is not contained in a proper affine subspace of \mathbb{R}^m (e.g., [5], page 2). It is well known that all nonminimal exponential families can be reduced to minimal exponential families (e.g., [5], Theorem 1.9, page 13). We consider both full and nonfull exponential families of the form (2.1). An exponential family $\{\mathbb{P}_n, \eta \in \Xi\}$ is full if $\Xi = \mathfrak{N}$ and nonfull if $\Xi \subset \mathfrak{N}$, where $\mathfrak{N} = \{ \eta \in \mathbb{R}^m : \psi(\eta) < \infty \}$ is the natural parameter space, that is, the largest set of possible values the natural parameter vector η can take on. While full exponential families may be more convenient on mathematical grounds, nonfull exponential families—the most important example being curved exponential families (e.g., [14, 29])-offer parsimonious parameterizations of exponential families where the dimension m_k of neighborhood-dependent natural parameter vectors η_k is an increasing function of the number of nodes in neighborhoods A_k (k = 1, ..., K), and have turned out to be useful in practice [24, 53]. A simple approach to generating nonfull exponential families is to assume that η : int(Θ) \mapsto int(\mathfrak{N}) is a known function of a parameter vector $\theta \in \Theta$, where $\Theta \subseteq \{\theta \in \mathbb{R}^q : \psi(\eta(\theta)) < \infty\}$, int(Θ) and int(\mathfrak{N}) denote the interiors of Θ and \mathfrak{N} , respectively, and $q \leq m$. It is convenient to distinguish exponential families that can be reduced to canonical exponential families with natural parameter vectors of the form $\eta(\theta) = \theta$ from those that cannot. An exponential family can be reduced to a canonical exponential family with $\eta(\theta) = \theta$ when the map $\eta : int(\Theta) \mapsto int(\mathfrak{N})$ is affine. In other words, if $\eta(\theta) = A\theta + b$ with $A \in \mathbb{R}^{m \times q}$ and $b \in \mathbb{R}^m$, then the exponential family can be reduced to a canonical exponential family with $\eta(\theta) = \theta$ by absorbing A into the sufficient statistic vector and b into the reference measure. We therefore call all exponential families with affine maps $\eta : int(\Theta) \mapsto int(\mathfrak{N})$ canonical exponential families, and call all exponential families with nonaffine maps $\eta : int(\Theta) \mapsto int(\mathfrak{N})$ curved exponential families. We note that our definition of a curved exponential family is broader than the one used in Efron [13, 14], Brown ([5], pages 81–84), and Kass and Vos [29]. The main reason is that we do not restrict the map $\eta : int(\Theta) \mapsto int(\mathfrak{N})$ to be differentiable, because our concentration and consistency results in Sections 3 and 4 do not require differentiability.

Throughout, we assume that the neighborhoods are of the same order of magnitude and that the neighborhood-dependent natural parameters $\eta_{k,i}(\theta)$ are of the form $\eta_{k,i}(\theta) = \eta_i(\theta)$ $(i = 1, ..., m_k, k = 1, ..., K)$. We define neighborhoods of the same order of magnitude as follows.

DEFINITION (Neighborhoods of the same order of magnitude). A sequence of neighborhoods $\mathcal{A}_1, \ldots, \mathcal{A}_K$ is of the same order of magnitude if there exists a universal constant A > 1 such that $\max_{1 \le k \le K} |\mathcal{A}_k| \le A \min_{1 \le k \le K} |\mathcal{A}_k|$ ($K = 1, 2, \ldots$).

In other words, the largest neighborhood size is a constant multiple of the smallest neighborhood size, so that the sizes of neighborhoods may not be identical, but are similar in a well-defined sense. The definition is satisfied when the sizes of neighborhoods are bounded above. When the number of neighborhoods K grows and the sizes of neighborhoods grow with K, then the definition implies that the sizes of neighborhoods grow at the same rate.

We note that when the neighborhoods are not of the same order of magnitude, the natural parameters of neighborhoods may have to depend on the order of magnitude of neighborhoods (e.g., [7, 34, 35]), because there are good reasons to believe that small and large within-neighborhood subgraphs are not governed by the same natural parameters [12, 36, 51]. Size-dependent parameterizations have an important place in the exponential-family random graph framework, and some promising size-dependent parameterizations have been proposed. Most of them assume that natural parameters consist of size-invariant parameters and size-dependent deviations. The size-dependent deviations may be size-dependent offsets (e.g., [7, 34, 35]) or functions of size-dependent covariates [52]. Some of those sizedependent deviations can be absorbed into the sufficient statistic vector and reference measure, and are hence covered by our main concentration and consistency results in Sections 3 and 4. However, the topic of size-dependent parameterizations is an important topic in its own right, and deserves a separate treatment that is beyond the scope of our paper.

The assumption $\eta_{k,i}(\theta) = \eta_i(\theta)$ $(i = 1, ..., m_k, k = 1, ..., K)$ implies that the exponential families considered here can be reduced to exponential families with natural parameter vectors of the form

$$\boldsymbol{\eta}(\boldsymbol{\theta}) = (\eta_1(\boldsymbol{\theta}), \dots, \eta_m(\boldsymbol{\theta}))$$

and sufficient statistic vectors of the form

$$s(\boldsymbol{x}) = (s_1(\boldsymbol{x}), \ldots, s_m(\boldsymbol{x})),$$

where $s_i(\mathbf{x}) = \sum_{k=1}^{K} s_{k,i}(\mathbf{x}_k)$ (i = 1, ..., m) and $m = \max_{1 \le k \le K} m_k$. We assume that the dimensions m_k of the neighborhood-dependent natural parameter vectors $\boldsymbol{\eta}_k(\boldsymbol{\theta})$ are nondecreasing functions of the sizes $|\mathcal{A}_k|$ of neighborhoods \mathcal{A}_k (k = 1, ..., K), which implies that

 $m = \max_{1 \le k \le K} m_k$ is a nondecreasing function of $||\mathcal{A}||_{\infty} = \max_{1 \le k \le K} |\mathcal{A}_k|$. The dimensions m_k (k = 1, ..., K) and m do not depend on other quantities. The dimension q of parameter vector $\boldsymbol{\theta}$ satisfies $q \le m$, as mentioned above.

Notation. To prepare the ground for the concentration and consistency results in Sections 3 and 4, we introduce mean-value parameterizations of exponential families along with additional notation. Mean-value parameterizations facilitate concentration and consistency results, because concentration inequalities [4] bound probabilities of deviations from means and the mean-value parameters of an exponential family are the means of the sufficient statistics, defined by $\mu(\eta(\theta)) = \mathbb{E}_{\eta(\theta)}s(X) \in \operatorname{rint}(\mathbb{M})$ ([5], page 2 and page 75). Here, $\mathbb{E}_{\eta(\theta)}s(X)$ is the expectation of s(X) with respect to exponential-family distributions $\mathbb{P}_{\eta(\theta)}$ having densities of the form (2.1), \mathbb{M} is the convex hull of the set $\{s(x) : \nu(x) > 0\}$, and $\operatorname{rint}(\mathbb{M})$ is the relative interior of \mathbb{M} . We denote the data-generating parameter vector by $\theta^* \in \operatorname{int}(\Theta)$ and write $\mathbb{P} \equiv \mathbb{P}_{\eta(\theta^*)}$ and $\mathbb{E} \equiv \mathbb{E}_{\eta(\theta^*)}$. An open ball in \mathbb{R}^{ν} ($\nu \geq 1$) with center $c \in \mathbb{R}^{\nu}$ and radius $\rho > 0$ is denoted by $\mathcal{B}(c, \rho)$. We write $\|\cdot\|_1, \|\cdot\|_2$ and $\|\cdot\|_\infty$ to refer to the ℓ_1 -, ℓ_2 - and ℓ_∞ norm of vectors, respectively. Throughout, uppercase letters $A, B, C, C_1, C_2, \ldots > 0$ denote universal constants, which are independent of all other quantities of interest and may be recycled from line to line. The function $d : \mathbb{X} \times \mathbb{X} \mapsto \{0, 1, \ldots\}$ denotes the Hamming metric, defined by

$$d(\boldsymbol{x}_1, \boldsymbol{x}_2) = \sum_{k=1}^{K} \sum_{i \in \mathcal{A}_k < j \in \mathcal{A}_k} \mathbb{1}(x_{1,i,j} \neq x_{2,i,j}), \quad (\boldsymbol{x}_1, \boldsymbol{x}_2) \in \mathbb{X} \times \mathbb{X},$$

where $\mathbb{1}(x_{1,i,j} \neq x_{2,i,j})$ is an indicator function, which is 1 if $x_{1,i,j} \neq x_{2,i,j}$ and is 0 otherwise.

3. Concentration and consistency results: Maximum likelihood and *M*-estimators. In many applications of exponential-family random graphs, the parameter vector of primary interest is θ . To estimate the parameter vector θ of a wide range of full and nonfull, curved exponential families under weak assumptions on the map η : int(Θ) \mapsto int(\mathfrak{N}), we consider an estimating function [17]—a function $g : \Theta \times \mathbb{X} \mapsto \mathbb{R}$ of both θ and X—of the form

(3.1)
$$g(\theta; \mu(\eta(\theta^{\star}))) = \|\mu(\eta(\theta^{\star})) - \mu(\eta(\theta))\|_{2}, \quad \theta \in \Theta,$$

which is an approximation of

$$g(\boldsymbol{\theta}; \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))) = \|\boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star})) - \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}))\|_{2}, \quad \boldsymbol{\theta} \in \boldsymbol{\Theta},$$

where $\mu(\eta(\theta^*)) = s(X)$ is an estimator of the data-generating mean-value parameter vector $\mu(\eta(\theta^*)) = \mathbb{E}_{\eta(\theta^*)} s(X) \in \operatorname{rint}(\mathbb{M})$. The fact that $g(\theta; \mu(\eta(\theta^*)))$ is an approximation of $g(\theta; \mu(\eta(\theta^*)))$ follows from the triangle inequality, which shows that

$$\begin{split} \|g(\theta; \widehat{\mu(\eta(\theta^{\star}))}) - g(\theta; \mu(\eta(\theta^{\star})))\| \\ &\leq \|\widehat{\mu(\eta(\theta^{\star}))} - \mu(\eta(\theta^{\star}))\|_{2}, \quad \theta \in \Theta, \end{split}$$

along with the fact that, under suitable conditions, $\|\mu(\eta(\theta^*)) - \mu(\eta(\theta^*))\|_2$ is small with high probability, as we will show in Proposition 2 in Section 3.2.

Estimating function (3.1) has at least three advantages. First, estimating function (3.1) admits concentration and consistency results under weak assumptions on the map $\eta : int(\Theta) \mapsto$ int(\mathfrak{N}). Indeed, the map $\eta : int(\Theta) \mapsto int(\mathfrak{N})$ satisfies the main assumptions of Section 3 as long as the map is one-to-one and continuous, but it need not be differentiable. The weakness of these assumptions implies that the main results of Section 3 cover a vast range of full and nonfull exponential families—including, but not limited to curved exponential families—and it is possible to verify these assumptions in some of the most challenging applications,

as demonstrated in Section 3.3. Second, estimating function (3.1) is natural, because maximum likelihood estimation of the data-generating natural parameter vector η^{\star} of an exponential family with natural parameter vector η and sufficient statistic vector $s(\mathbf{x})$ can be based on the gradient of the loglikelihood function $\nabla_n \log p_n(\mathbf{x}) = \widehat{\mu(\eta^{\star})} - \mu(\eta)$ provided $\eta \in int(\mathfrak{N})$, where $\mu(\eta^{\star}) = s(x)$ ([5], Lemma 5.3, page 146). Therefore, maximum likelihood estimation of η^* can be based on estimating functions of the form $\|\hat{\mu}(\eta^*) - \mu(\eta)\|_2$. By the parameterization-invariance of maximum likelihood estimators, maximum likelihood estimation of functions of η^* , such as θ^* , can be based on $\|\mu(\eta(\theta^*)) - \mu(\eta(\theta))\|_2$ provided the map η : int(Θ) \mapsto int(\mathfrak{N}) is one-to-one. We note that estimating function (3.1) is chosen for mathematical convenience, facilitating concentration and consistency results for maximum likelihood estimators of many full and nonfull, curved exponential families under weak assumptions on the map η : int(Θ) \mapsto int(\mathfrak{N}), and is not chosen for computational convenience. Last, but not least, the simple form of estimating function (3.1) helps determine when minimizers of (3.1) exist and are unique, and how the minimizers are related to each other when there is more than one minimizer. These advantages are most useful in nonfull exponential families, in particular curved exponential families.

In the following, we assume that the map η : int(Θ) \mapsto int(\mathfrak{N}) is one-to-one. A natural class of estimators is hence given by

$$\widehat{\boldsymbol{\theta}} = \left\{ \boldsymbol{\theta} \in \boldsymbol{\Theta} : g(\boldsymbol{\theta}; \mu(\widehat{\boldsymbol{\eta}(\boldsymbol{\theta}^{\star})})) = \inf_{\boldsymbol{\dot{\theta}} \in \boldsymbol{\Theta}} g(\boldsymbol{\dot{\theta}}; \mu(\widehat{\boldsymbol{\eta}(\boldsymbol{\theta}^{\star})})) \right\}.$$

If the set $\hat{\theta}$ is nonempty, it may contain one element (e.g., in full exponential families) or more than one element (e.g., in nonfull exponential families). If the set $\hat{\theta}$ contains more than one element, then all elements of the set $\hat{\theta}$ map to mean-value parameter vectors $\mu(\eta(\hat{\theta}))$ that have the same ℓ_2 -distance from $\mu(\eta(\theta^*))$ by construction of estimating function (3.1). In addition, if the set $\hat{\theta}$ is nonempty, then all minimizers $\hat{\theta}$ of $g(\theta; \mu(\eta(\theta^*)))$ are approximations of the minimizer of $g(\theta; \mu(\eta(\theta^*)))$. The minimizer of $g(\theta; \mu(\eta(\theta^*)))$ is unique and is given by the data-generating parameter vector θ^* provided $\mu(\eta(\theta^*))) \in \operatorname{rint}(\mathbb{M})$:

$$\boldsymbol{\theta}^{\star} = \Big\{ \boldsymbol{\theta} \in \boldsymbol{\Theta} : g(\boldsymbol{\theta}; \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))) = \inf_{\boldsymbol{\dot{\theta}} \in \boldsymbol{\Theta}} g(\boldsymbol{\dot{\theta}}; \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))) \Big\}.$$

The data-generating parameter vector θ^* is the unique minimizer of $g(\theta; \mu(\eta(\theta^*))) = \|\mu(\eta(\theta^*)) - \mu(\eta(\theta))\|_2$, because $\theta^* \in int(\Theta)$ and the map $\eta : int(\Theta) \mapsto int(\mathfrak{N})$ is one-to-one by assumption, while the map $\mu : int(\mathfrak{N}) \mapsto rint(\mathbb{M})$ is one-to-one by classic exponential-family theory ([5], Theorem 3.6, page 74). Therefore, $\|\mu(\eta(\theta^*)) - \mu(\eta(\theta))\|_2 = 0$ holds if and only if $\theta = \theta^*$.

In the remainder of the section, we show that the estimator $\hat{\theta}$ is close to the data-generating parameter vector θ^* with high probability under weak conditions. We first sketch the main idea in Section 3.1 and then discuss concentration and consistency results for maximum like-lihood and *M*-estimators in Sections 3.3 and 3.4, respectively. An application to popular curved exponential-family random graphs is presented in Section 3.3.

3.1. Main idea: A nonasymptotic approach to full and nonfull, curved exponential families for dependent random variables. Establishing concentration and consistency results for estimators of full and nonfull, curved exponential-family random graphs with dependent edges is nontrivial, for at least three reasons. First, the map η : int(Θ) \mapsto int(\mathfrak{N}) may not be affine and may not be differentiable. Second, in many nonfull exponential families, the mean-value parameter vector $\mu(\eta(\theta))$ is not available in closed form and there is no simple and known relationship between the mean-value parameter vector $\mu(\eta(\hat{\theta}))$ evaluated at $\hat{\theta}$ and the sufficient statistic vector s(X). Third, concentration results for functions of edges, such as s(X),



FIG. 1. The figure demonstrates the main assumption: for all $\epsilon > 0$ small enough so that $\mathbb{B}(\theta^{\star}, \epsilon) \subseteq \operatorname{int}(\Theta)$, there exist $\gamma > 0$ and $\delta > 0$ such that $\theta \notin \mathbb{B}(\theta^{\star}, \epsilon)$ implies $\eta(\theta) \notin \mathbb{B}(\eta(\theta^{\star}), \gamma)$ and $\mu(\eta(\theta)) \notin \mathbb{B}(\mu(\eta(\theta^{\star})), \delta)$.

are more challenging when edges are dependent rather than independent. As a result, studying the behavior of the estimating function $||s(X) - \mu(\eta(\hat{\theta}))||_2$ and its minimizer $\hat{\theta}$ is not straightforward.

Our main idea is based on a simple and general exponential-family argument that may be of independent interest. It helps establish nonasymptotic probability statements about estimators of full and nonfull, curved exponential families for dependent random variables under weak conditions.

We make a single weak assumption: for all $\epsilon > 0$ small enough so that $\mathcal{B}(\theta^*, \epsilon) \subseteq int(\Theta)$, there exist $\gamma > 0$ and $\delta > 0$ such that

$$\boldsymbol{\theta} \notin \mathbb{B}(\boldsymbol{\theta}^{\star}, \boldsymbol{\epsilon}) \quad \Longrightarrow \quad \boldsymbol{\eta}(\boldsymbol{\theta}) \notin \mathbb{B}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}), \boldsymbol{\gamma}) \quad \Longrightarrow \quad \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta})) \notin \mathbb{B}(\boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star})), \boldsymbol{\delta}).$$

A graphical representation of the main assumption can be seen in Figure 1. A formal statement of the assumption can be found in Theorem 1 in Section 3.3, where γ and δ depend on ϵ and δ depends on the sizes of neighborhoods $\mathcal{A}_1, \ldots, \mathcal{A}_K$. The main assumption is satisfied when the map η : int(Θ) \mapsto int(\mathfrak{N}) is one-to-one and continuous, but it need not be differentiable. Note that the map μ : int(\mathfrak{N}) \mapsto rint(\mathbb{M}) is one-to-one and continuous by classic exponential-family theory ([5], Theorem 3.6, page 74).

The main assumption has a simple, but important implication. As no element of \mathfrak{N} outside of the ball $\mathfrak{B}(\eta(\theta^*), \gamma)$ maps to an element of the ball $\mathfrak{B}(\mu(\eta(\theta^*)), \delta)$, any element $\mu(\eta(\theta))$ of $\mathfrak{B}(\mu(\eta(\theta^*)), \delta)$ must correspond to an element $\eta(\theta)$ of $\mathfrak{B}(\eta(\theta^*), \gamma)$, which in turn must correspond to an element θ of $\mathfrak{B}(\theta^*, \epsilon)$, so that

$$\boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta})) \in \mathcal{B}(\boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star})), \delta) \implies \boldsymbol{\eta}(\boldsymbol{\theta}) \in \mathcal{B}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}), \gamma) \implies \boldsymbol{\theta} \in \mathcal{B}(\boldsymbol{\theta}^{\star}, \epsilon).$$

As a result, the probability of event $\hat{\theta} \in \mathcal{B}(\theta^*, \epsilon)$ can be bounded by bounding the probability of event $\mu(\eta(\hat{\theta})) \in \mathcal{B}(\mu(\eta(\theta^*)), \delta)$.

A challenge, which complicates probability statements about the event $\mu(\eta(\hat{\theta})) \in \mathbb{B}(\mu(\eta(\theta^*)), \delta)$, is that in many nonfull exponential families $\mu(\eta(\hat{\theta}))$ is not available in closed form and there is no simple and known relationship between $\mu(\eta(\hat{\theta}))$ and the sufficient statistic vector s(X). To appreciate the difficulty of the problem, suppose that $s(x) \in \operatorname{rint}(\mathbb{M})$ is observed, so that $\mu(\eta(\theta^*)) = s(x) \in \operatorname{rint}(\mathbb{M})$. The subset $\mathbb{M}(\Theta)$ of $\operatorname{rint}(\mathbb{M})$ induced by Θ is defined by

$$\mathbb{M}(\Theta) = \{ \mu' \in \operatorname{rint}(\mathbb{M}) : \text{ there exists } \theta \in \Theta \text{ such that } \mu(\eta(\theta)) = \mu' \}.$$

In full exponential families, $\mathbb{M}(\Theta) = \operatorname{rint}(\mathbb{M})$. Thus, there exists a minimizer $\widehat{\theta} \in \operatorname{int}(\Theta)$ of the estimating function $||s(\mathbf{x}) - \mu(\eta(\theta))||_2$ satisfying $\mu(\eta(\widehat{\theta})) = s(\mathbf{x})$. In fact, the minimizer is unique, because the maps $\eta : \operatorname{int}(\Theta) \mapsto \operatorname{int}(\mathfrak{N})$ and $\mu : \operatorname{int}(\mathfrak{N}) \mapsto \operatorname{rint}(\mathbb{M})$ are one-to-one ([5], Theorem 3.6, page 74). As a result, there is a simple and known relationship between $\mu(\eta(\widehat{\theta}))$ and $s(\mathbf{x})$. By contrast, in nonfull exponential families, $\mathbb{M}(\Theta)$ is a proper subset of rint(\mathbb{M}), because nonfull exponential families are subfamilies of exponential families that exclude some natural parameter vectors along with the corresponding mean-value parameter vectors are excluded, because the mean-value parameter vectors are not available in closed form. As a consequence, there is no simple and known relationship between $\mu(\eta(\widehat{\theta}))$ and $s(\mathbf{x})$, and it is not straightforward to determine where $\mu(\eta(\widehat{\theta}))$ is located in rint(\mathbb{M}), provided $\mu(\eta(\widehat{\theta}))$ is nonempty. So bounding the probability of event $\mu(\eta(\widehat{\theta})) \in \mathbb{B}(\mu(\eta(\theta^*)), \delta)$ in nonfull exponential families, in particular curved exponential families, is nontrivial.

But not all is lost. Despite the challenge of characterizing $\mathbb{M}(\Theta) \subset \operatorname{rint}(\mathbb{M})$ and hence $\mu(\eta(\widehat{\theta})) \subset \mathbb{M}(\Theta)$, it can be shown that, under suitable conditions, the set $\mu(\eta(\widehat{\theta})) \subset \mathbb{M}(\Theta)$ is nonempty and all elements of $\mu(\eta(\widehat{\theta}))$ are close to $\mu(\eta(\theta^*))$ with high probability. Indeed, when the set $\widehat{\theta}$ is nonempty, each element of the set $\widehat{\theta}$ satisfies

$$\|\boldsymbol{\mu}(\boldsymbol{\eta}(\widehat{\boldsymbol{\theta}})) - \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))\|_{2} \leq \|\boldsymbol{s}(\boldsymbol{x}) - \boldsymbol{\mu}(\boldsymbol{\eta}(\widehat{\boldsymbol{\theta}}))\|_{2} + \|\boldsymbol{s}(\boldsymbol{x}) - \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))\|_{2}$$

While characterizing $\mathbb{M}(\Theta)$ is difficult, we do know one fact about $\mathbb{M}(\Theta)$: $\mathbb{M}(\Theta)$ contains the data-generating mean-value parameter vector $\mu(\eta(\theta^*))$, which implies that the ℓ_2 -distance of $\mu(\eta(\widehat{\theta}))$ from s(x) cannot exceed the ℓ_2 -distance of s(x) from $\mu(\eta(\theta^*)) \in \mathbb{M}(\Theta)$. Thus, we obtain the upper bound

(3.2)
$$\|\boldsymbol{\mu}(\boldsymbol{\eta}(\widehat{\boldsymbol{\theta}})) - \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))\|_{2} \leq 2\|\boldsymbol{s}(\boldsymbol{x}) - \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))\|_{2}.$$

If the right-hand side of (3.2) can be shown to be small with high probability, then the problem of bounding the probability of event $\hat{\theta} \in \mathcal{B}(\theta^*, \epsilon)$ can be converted into the problem of bounding the probability of the event that s(X) is close to $\mu(\eta(\theta^*)) = \mathbb{E}_{\eta(\theta^*)}s(X) \in \operatorname{rint}(\mathbb{M})$ in a well-defined sense. All we need to bound the probabilities of those events are concentration results for the sufficient statistic vector s(X), which can be established as long as (a) the dependence among edges is sufficiently weak; and (b) the sufficient statistics are sufficiently smooth functions of edges. Concentration results are facilitated by additional structure that helps control the amount of dependence induced by the model and the smoothness of sufficient statistics. We focus here on a simple form of additional structure in the form of multilevel structure, which controls the dependence among edges by constraining it to neighborhoods. But there are many other forms of additional structure that could help address (a) and (b), e.g., other forms of multilevel structure or spatial structure.

The most important implication, then, is that statistical inference for many exponentialfamily random graphs is possible and can be justified by statistical theory, provided a suitable form of additional structure is available. Indeed, our main idea helps establish concentration and consistency results for estimators of:

- many full and nonfull, curved exponential families;
- many models with dependent edges;
- finite populations of nodes;

as long as there is additional structure that helps address (a) and (b).

It is worth noting that verifying the main assumption is by no means trivial. Its verification is easiest when the sufficient statistics are monotone functions of edges, that is, functions of a graph that do not decrease (increase) when edges are added to the graph. It is less straightforward when the sufficient statistics are not monotone functions of edges, as is the case with many popular curved exponential-family random graphs with geometrically weighted model terms [24, 53]. But we demonstrate in Section 3.3 that the main assumption can be verified even when the sufficient statistics are not monotone functions of a graph, using curved exponential-family random graphs with geometrically weighted model terms as an example.

We make these ideas rigorous in Theorem 1 in Section 3.3. An application to curved exponential-family random graphs is presented in Corollary 1 in Section 3.3. More general results for M-estimators, under correct and incorrect model specifications, are mentioned in Section 3.4. To prepare the ground for these results, we first introduce concentration results for sufficient statistics in Section 3.2.

3.2. *Concentration results for sufficient statistics*. To obtain concentration results for sufficient statistics, we need concentration results for functions of random graphs with dependent edges.

Such concentration results are nontrivial for at least two reasons. First, exponential families of the form (2.1) may induce strong dependence within neighborhoods and the sizes of neighborhoods need not be identical. Second, exponential families of the form (2.1) can induce a wide range of dependencies within neighborhoods. Therefore, we need generalpurpose concentration inequalities that cover a wide range of dependencies.

The following general-purpose concentration inequality addresses the challenges discussed above. It shows that the dependence induced by exponential families of the form (2.1) may be strong within neighborhoods but is sufficiently weak overall to obtain concentration results as long as the neighborhoods are not too large.

PROPOSITION 1. Consider an exponential family with countable support X and local dependence. Let $f : X \mapsto \mathbb{R}$ be Lipschitz with respect to the Hamming metric $d : X \times X \mapsto \{0, 1, 2, ...\}$ with Lipschitz coefficient $||f||_{\text{Lip}} > 0$ and assume that $\mathbb{E} f(X)$ exists. Then there exists C > 0 such that, for all t > 0,

$$\mathbb{P}\left(\left|f(\boldsymbol{X}) - \mathbb{E}f(\boldsymbol{X})\right| \ge t\right) \le 2\exp\left(-\frac{t^2}{C\sum_{k=1}^{K} \binom{|\mathcal{A}_k|}{2} \|\mathcal{A}\|_{\infty}^4 \|f\|_{\mathrm{Lip}}^2}\right)$$

Proposition 1 covers a wide range of exponential-family random graphs with local dependence. The assumption that the function of interest is smooth, in the sense that it is Lipschitz with respect to the Hamming metric, is common in the concentration-of-measure literature: see, for example, the concentration results for dependent random variables by Samson [46], Chatterjee ([8], Theorem 4.3, page 75), and Kontorovich and Ramanan [32]. The smoothness assumption can be weakened by using divide-and-conquer strategies: for example, one may divide the domain of a function of interest into high- and low-probability regions and require the function to be smooth on high-probability regions, but not on low-probability regions. Such divide-and-conquer strategies were used by, for example, Vu [55], Kim and Vu [30] and Yang et al. ([61], Lemma 9). While exploring divide-and-conquer strategies for exponentialfamily random graphs with local dependence would be interesting, Proposition 1 suffices for the purpose of demonstrating that statistical inference for many exponential-family random graphs with local dependence is possible.

Proposition 1 paves the way for concentration results for sufficient statistics. Proposition 2 shows that the sufficient statistic vector $\mu(\eta(\theta^*)) = s(X)$ is close to the data-generating mean-value parameter vector $\mu(\eta(\theta^*)) \in \operatorname{rint}(\mathbb{M})$ with high probability provided the number of neighborhoods K is large relative to the size of the largest neighborhood $||\mathcal{A}||_{\infty}$ and the dimension m of $\mu(\eta(\theta^*))$.

PROPOSITION 2. Consider a full or nonfull, curved exponential family with countable support X and local dependence. Assume that there exists A > 0 such that

$$(3.3) \|s(\boldsymbol{x}_1) - s(\boldsymbol{x}_2)\|_{\infty} \le A \, d(\boldsymbol{x}_1, \boldsymbol{x}_2) \|\mathcal{A}\|_{\infty} \quad \text{for all } (\boldsymbol{x}_1, \boldsymbol{x}_2) \in \mathbb{X} \times \mathbb{X}.$$

Then there exists C > 0 such that, for all deviations of the form $t = \delta \sum_{k=1}^{K} {\binom{|\mathcal{A}_k|}{2}}^{\alpha}$ with $\delta > 0$ and $0 \le \alpha \le 1$,

$$\mathbb{P}(\|\widehat{\boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))} - \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))\|_{2} \ge t) \le 2 \exp\left(-\frac{\delta^{2} C K}{m \|\mathcal{A}\|_{\infty}^{4(2-\alpha)}} + \log m\right).$$

The smoothness condition (3.3) of Proposition 2 is satisfied as long as changing an edge cannot change the within-neighborhood sufficient statistics by more than a constant multiple of $||A||_{\infty}$. It is verified in Corollary 1 in Section 3.3.

REMARK 1 (The relationship between α and sparsity). Proposition 2 shows how the concentration of sufficient statistics depends on the power $\alpha \in [0, 1]$ of deviations of size $t = \delta \sum_{k=1}^{K} {\binom{|\mathcal{A}_k|}{2}}^{\alpha}$. The power α can be interpreted as the level of sparsity of a random graph, with lower values of α corresponding to higher levels of sparsity. The conventional definition of a sparse random graph is based on the scaling of the expected number of edges, that is, the sufficient statistic of Bernoulli random graphs with independent edges. We use the term sparse random graph to refer to the scaling of the expectations of all sufficient statistics of exponential-family random graphs. If $\alpha = 1$, the within-neighborhood subgraphs may be called dense in the sense that the expectations of within-neighborhood sufficient statistics are nonnegligible fractions of the number of edge variables $\binom{|\mathcal{A}_k|}{2}$ in neighborhood \mathcal{A}_k ($k = 1, \ldots, K$). Otherwise, the within-neighborhood subgraphs may be called sparse. Note that the interpretation of α in terms of sparsity makes more sense when the neighborhoods grow than when the neighborhoods are bounded above. But, regardless of whether the neighborhoods on the power α .

REMARK 2 (Sharpness). The concentration results discussed above are not, and cannot be sharp, because these results cover many models and many dependencies. It goes without saying that in special cases sharper results can be obtained. We are here not interested in sharp bounds in special cases, because the main appeal of the exponential-family framework is that it can capture many dependencies.

3.3. *Maximum likelihood estimators*. The main idea of Section 3.1 is made rigorous in Theorem 1, which establishes concentration results for maximum likelihood estimators of full and nonfull, curved exponential-family random graphs with local dependence.

THEOREM 1. Consider a full or nonfull, curved exponential-family random graph with countable support X and local dependence. Let

$$\boldsymbol{\Theta} \subseteq \{\boldsymbol{\theta} \in \mathbb{R}^q : \psi(\boldsymbol{\eta}(\boldsymbol{\theta})) < \infty\}.$$

Assume that $\theta^* \in int(\Theta)$. Let $\eta : int(\Theta) \mapsto int(\mathfrak{N})$ be one-to-one and assume that, for all $\epsilon > 0$ small enough so that $\mathbb{B}(\theta^*, \epsilon) \subseteq int(\Theta)$, there exists $\gamma(\epsilon) > 0$ such that, for all $\theta \in \Theta \setminus \mathbb{B}(\theta^*, \epsilon)$, we have $\eta(\theta) \in \mathfrak{N} \setminus \mathbb{B}(\eta(\theta^*), \gamma(\epsilon))$. In addition, assume that there exist $\delta(\epsilon) > 0$ and A > 0 such that, for all $\eta(\theta) \in \mathfrak{N} \setminus \mathbb{B}(\eta(\theta^*), \gamma(\epsilon))$,

(3.4)
$$\|\boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star})) - \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}))\|_{2} \ge \delta(\epsilon) \sum_{k=1}^{K} {|\mathcal{A}_{k}| \choose 2}^{\alpha} \text{ for some } 0 \le \alpha \le 1$$

and, for all $(\mathbf{x}_1, \mathbf{x}_2) \in \mathbb{X} \times \mathbb{X}$,

(3.5)
$$\|s(\mathbf{x}_1) - s(\mathbf{x}_2)\|_{\infty} \le Ad(\mathbf{x}_1, \mathbf{x}_2) \|\mathcal{A}\|_{\infty}$$

Then, for all $\epsilon > 0$ small enough so that $\mathbb{B}(\theta^*, \epsilon) \subseteq \operatorname{int}(\Theta)$, there exist $\kappa(\epsilon) > 0$ and C > 0 such that

$$\mathbb{P}(\widehat{\boldsymbol{\theta}} \in \mathcal{B}(\boldsymbol{\theta}^{\star}, \epsilon)) \geq 1 - 2\exp\left(-\frac{\kappa(\epsilon)^2 C K}{m \|\mathcal{A}\|_{\infty}^{4(2-\alpha)}} + \log m\right).$$

If the exponential family is full, then $\widehat{\theta}$ is unique in the event $\widehat{\theta} \in \mathbb{B}(\theta^*, \epsilon)$.

Theorem 1 shows that the estimator $\hat{\theta}$ is close to the data-generating parameter vector θ^* with high probability provided the number of neighborhoods *K* is large relative to $||\mathcal{A}||_{\infty}$ and *m*. An application of Theorem 1 to popular curved exponential-family random graphs can be found in Corollary 1. These concentration results cover inference scenarios with a finite population of nodes and a population graph generated by an exponential-family random graph model, and assume that the population graph can be observed. Inference scenarios where the population graph cannot be observed but subgraphs of the population graph can be observed are considered in Section 4. Asymptotic consistency results can be obtained by allowing the number of neighborhoods *K* to increase without bound. We discuss asymptotic consistency results in Remark 3 following Corollary 1.

Conditions (3.4) and (3.5) of Theorem 1 are verified in Corollary 1. As pointed out in Section 3.1, the assumptions are satisfied when the map $\eta : int(\Theta) \mapsto int(\mathfrak{N})$ is one-toone and continuous, but it need not be differentiable. Condition (3.4) is an identifiability assumption and covers both sparse ($0 \le \alpha < 1$) and dense ($\alpha = 1$) within-neighborhood subgraphs. The power α can be interpreted as the level of sparsity of a random graph, as explained in Section 3.2. Theorem 1 shows that sparsity comes at a cost, because the probability of event $\hat{\theta} \notin B(\theta^*, \epsilon)$ decays slower when the within-neighborhood subgraphs are sparse rather than dense. The fact that sparsity weakens concentration results is well known in the concentration-of-measure literature on random graphs with independent edges (e.g., [27, 30]). Condition (3.5) is a smoothness condition, which is satisfied as long as changing an edge cannot change the within-neighborhood sufficient statistics by more than a constant multiple of $\|\mathcal{A}\|_{\infty}$.

It is worth noting that in full exponential families the set $\hat{\theta}$ contains a single element when it is nonempty, whereas in nonfull exponential families it may contain more than one element. A pleasant feature of estimating function (3.1) is that, with high probability, the minimizers $\hat{\theta}$ of (3.1) do not give rise to global minima that are separated by large distances, under the assumptions made. The reason is that, if the set $\hat{\theta}$ contains more than one element, then all elements of the set $\hat{\theta}$ map to mean-value parameter vectors $\mu(\eta(\hat{\theta}))$ whose ℓ_2 -distance from $\mu(\eta(\theta^*))$ is identical and whose ℓ_2 -distance from $\mu(\eta(\theta^*))$ is bounded above by

$$\|\boldsymbol{\mu}(\boldsymbol{\eta}(\widehat{\boldsymbol{\theta}})) - \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))\|_{2} \leq 2\|\boldsymbol{\mu}(\widehat{\boldsymbol{\eta}(\boldsymbol{\theta}^{\star})}) - \boldsymbol{\mu}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))\|_{2}$$

as explained in Section 3.1. By Proposition 2, $\mu(\eta(\theta^*))$ is close to $\mu(\eta(\theta^*))$ with high probability provided the number of neighborhoods *K* is sufficiently large. Therefore, all elements of the set $\mu(\eta(\hat{\theta}))$ are close to $\mu(\eta(\theta^*))$ with high probability and hence, by the identifiability conditions of Theorem 1, all elements of the set $\hat{\theta}$ are close to θ^* with high probability.

Applications. We present two applications of Theorem 1. An application to canonical exponential-family random graphs can be found in Appendix A (see the Supplementary Material, [50]). Here, we focus on curved exponential-family random graphs with geometrically

weighted model terms, which are popular in practice (e.g., [39]) but are challenging on theoretical grounds.

As a specific example, consider curved exponential-family random graphs with support $\mathbb{X} = \{0, 1\}^{\sum_{k=1}^{K} {\binom{|\mathcal{A}_k|}{2}}}$ and within-neighborhood edge and geometrically weighted edgewise shared partner terms [24, 53]. Such models are based on sufficient statistics of the form

$$s_{k,1}(\boldsymbol{x}_k) = \sum_{a \in \mathcal{A}_k < b \in \mathcal{A}_k} x_{a,b},$$

$$s_{k,i+1}(\boldsymbol{x}_k) = \sum_{a \in \mathcal{A}_k < b \in \mathcal{A}_k} x_{a,b} f_{a,b,i}(\boldsymbol{x}_k), \quad i = 1, \dots, |\mathcal{A}_k| - 2,$$

where $f_{a,b,i}(\mathbf{x}_k) = \mathbb{1}(\sum_{c \in \mathcal{A}_k, c \neq a, b} x_{a,c} x_{b,c} = i)$ is an indicator function, which is 1 if nodes *a* and *b* are both connected to *i* other nodes in neighborhood \mathcal{A}_k and is 0 otherwise (k = 1, ..., K). The natural parameters are of the form

$$\eta_{k,1}(\boldsymbol{\theta}) = \theta_1,$$

$$\eta_{k,i+1}(\boldsymbol{\theta}) = \exp(\vartheta) [1 - (1 - \exp(-\vartheta))^i], \quad i = 1, \dots, |\mathcal{A}_k| - 2,$$

where $\vartheta > 0$ controls the rate of decay of the geometric sequence $(1 - \exp(-\vartheta))^i$, i = 1, 2, ..., (k = 1, ..., K). For convenience, we consider here the parameterization $\theta_2 = \exp(-\vartheta) \in (0, 1)$, so that $\theta = (\theta_1, \theta_2) \in \mathbb{R} \times (0, 1)$. Such model terms are called geometrically weighted terms, because the natural parameters $\eta_{k,i+1}(\theta)$ are based on the geometric sequence $(1 - \exp(-\vartheta))^i$, i = 1, 2, ...

While complicated, such models are able to capture transitivity in neighborhoods. As explained in the Introduction, transitivity is one of the more interesting network phenomena, and induces dependence among edges. There are many models of transitivity, some of which are well-posed while others are ill-posed. An example of a model that is ill-posed in the large-graph limit is the so-called triangle model (e.g., [2, 10, 18, 28, 47]). The triangle model is a canonical exponential-family random graph model with the number of edges and triangles as sufficient statistics; note that a triangle in a random graph corresponds to three distinct nodes such that all three pairs of nodes are connected by edges. Compared with the triangle model, the curved exponential-family random graph model described above makes more reasonable assumptions:

- The curved exponential-family random graph model exploits multilevel structure to constrain the dependence among edges induced by transitivity to neighborhoods, that is, subsets of nodes. By contrast, the triangle model does not restrict transitivity to subsets of nodes, and allows each edge to depend on many other edges in the random graph.
- The curved exponential-family random graph model implies that within neighborhoods, for each pair of nodes, the value added by additional triangles to the log odds of the conditional probability of an edge decays at a geometric rate (e.g., [23, 54]). As a result, the model encourages triangles within neighborhoods, but discourages too many of them. By contrast, the added value of additional triangles under the triangle model is constant, so that the triangle model with a positive triangle parameter places more probability mass on graphs with more triangles (among graphs with the same number of edges).

While the problematic assumptions underlying the triangle model lead to undesirable behavior in large random graphs (e.g., [2, 10, 18, 28, 47]), curved exponential-family random graphs with geometrically weighted edgewise shared partner terms have turned out to be well behaved (e.g., [24, 47]) and have been widely used (see, e.g., [23–25, 39, 53, 54]). A fullfledged discussion of these complex models is beyond the scope of our paper. We therefore refer the interested reader to the above-cited literature and focus here on concentration and consistency results.

Curved exponential-family random graphs with within-neighborhood edge and geometrically weighted edgewise shared partner terms are popular in practice but are challenging on theoretical grounds, for several reasons. First, the dimension of the natural parameter vector $\eta(\theta) \in \mathbb{R}^{\|\mathcal{A}\|_{\infty}-1}$ is an increasing function of the number of nodes in the largest neighborhood(s), $\|\mathcal{A}\|_{\infty}$. Second, the natural parameter vector $\eta(\theta) \in \mathbb{R}^{\|\mathcal{A}\|_{\infty}-1}$ is a nonaffine function of a lower-dimensional parameter vector $\theta \in \mathbb{R} \times (0, 1)$. Third, the mean-value parameter vector $\mu(\eta(\theta))$ is not available in closed form. Finally, the sufficient statistics $s_2(x), \ldots, s_{\|\mathcal{A}\|_{\infty}-1}(x)$ are not monotone functions of graphs, which complicates the verification of the main assumption of Theorem 1, as mentioned in Section 3.1. Despite these challenges, it is possible to verify all conditions of Theorem 1 and obtain the following concentration result. It shows that the estimator $\hat{\theta}$ is close to the data-generating parameter vector θ^* with high probability provided K is large relative to $\|\mathcal{A}\|_{\infty}^6 \log \|\mathcal{A}\|_{\infty}$.

COROLLARY 1. Consider a curved exponential-family random graph with withinneighborhood edge and geometrically weighted edgewise shared partner terms. Let $\Theta = \mathbb{R} \times (0, 1)$ and assume that $\theta^* \in int(\Theta)$. Then all conditions of Theorem 1 are satisfied and hence, for all $\epsilon > 0$ small enough so that $\mathbb{B}(\theta^*, \epsilon) \subseteq int(\Theta)$, there exist $\kappa(\epsilon) > 0$ and C > 0such that

$$\mathbb{P}(\widehat{\boldsymbol{\theta}} \in \mathcal{B}(\boldsymbol{\theta}^{\star}, \epsilon)) \geq 1 - 2\exp\left(-\frac{\kappa(\epsilon)^2 CK}{\|\mathcal{A}\|_{\infty}^6} + \log\|\mathcal{A}\|_{\infty}\right),$$

provided $|\mathcal{A}_k| \ge 4$ $(k = 1, \ldots, K)$ and $K \ge 2$.

Corollary 1 is the first concentration result for estimators of exponential-family random graphs with dependence among edges induced by transitivity, one of the more interesting network phenomena. In addition, it is the first concentration result for curved exponential-family random graphs with geometrically weighted model terms, which are popular in practice (e.g., [24, 53]). The concentration result assumes that each neighborhood A_k consists of $|A_k| \ge 4$ nodes, because θ^* is not identifiable when $|A_k| \le 3$ (k = 1, ..., K). As mentioned above, the set $\hat{\theta}$ may contain more than one element, but all elements of the set $\hat{\theta}$ are close to θ^* with high probability provided K is large relative to $||A||_{\infty}^6 \log ||A||_{\infty}$. Concentration results for other curved exponential-family random graphs with geometrically weighted model terms (e.g., [24, 53]) can be established along the same lines.

REMARK 3 (Asymptotic consistency results). As pointed out in the Introduction, we state all theoretical results for finite populations of nodes, because in practice all populations are finite. Asymptotic consistency results can be obtained by allowing the number of neighborhoods *K* to grow without bound. If there exists a universal constant *C* > 0 such that $|\mathcal{A}_k| < C$ (k = 1, 2, ...), then the main idea described in Section 3.1 along with the concentration results in Section 3.2 imply that $\hat{\theta}$ is a consistent estimator of θ^* with rate of convergence $K^{1/2}$. As the units of statistical analysis are neighborhoods, the rate $K^{1/2}$ resembles the rate in classical statistical problems where the rate is the square root of the sample size, albeit with two notable differences: first, the units are subsets of nodes (neighborhoods) rather than nodes or edges; and, second, the sizes of units are not identical, but the size of the largest unit is a constant multiple of the size of the smallest. Last, but not least, it is possible to obtain asymptotic consistency results when *K* grows and the neighborhoods grow with *K*, which implies that $\|\mathcal{A}\|_{\infty} \log \|\mathcal{A}\|_{\infty}$) $\rightarrow \infty$, the probability of event $\hat{\theta} \in \mathbb{B}(\theta^*, \epsilon)$ tends to 1.

3.4. *M*-estimators, correct and incorrect model specifications. The concentration and consistency results for maximum likelihood estimators in Section 3.3 are special cases of more general results for *M*-estimators. To demonstrate, we introduce a natural class of *M*-estimators in Appendix B.1, which includes both likelihood- and moment-based estimators, and present concentration results in Appendix B.2 along with an application to misspecified models with omitted covariate terms. These results cover both correct and incorrect model specifications, as the example with omitted covariate terms demonstrates. Due to space restrictions, we provide details in Appendix B (see the Supplementary Material, [50]).

4. Extendability and subgraph-to-graph estimators. A question that has been asked about exponential-family random graphs is whether it is possible to extend, in a well-defined sense, an exponential-family random graph with a given set of nodes to an exponential-family random graph with more nodes [12, 36, 51]. We show that multilevel structure helps extend an exponential-family random graph with a given set of neighborhoods to an exponential-family random graph with more neighborhoods (Section 4.1) and hence facilitates subgraph-to-graph estimation (Section 4.2). The importance of these results lies in the fact that subgraph-to-graph estimation for exponential-family random graphs is believed to be difficult (e.g., [51]), but our results demonstrate that additional structure facilitates it.

4.1. *Extendability*. While many exponential-family random graphs with a given set of nodes cannot be extended to exponential-family random graphs with more nodes [12, 36, 51], an exponential-family random graph with a given set of neighborhoods can be extended to an exponential-family random graph with more neighborhoods.

To demonstrate, consider a population graph $(X_{\mathcal{L}}, Y_{\mathcal{L}})$ with a set of neighborhoods $\mathcal{L} = \{\mathcal{A}_1, \ldots, \mathcal{A}_L\}$, where $X_{\mathcal{L}} \in \mathbb{X}_{\mathcal{L}}$ and $Y_{\mathcal{L}} \in \mathbb{Y}_{\mathcal{L}}$ denote the sequences of within- and betweenneighborhood edge variables based on the set of neighborhoods \mathcal{L} , respectively. As before, assume that $X_{\mathcal{L}}$ is governed by an exponential family with countable support $\mathbb{X}_{\mathcal{L}}$ and local dependence, with neighborhood-dependent natural parameters $\eta_{\mathcal{A},i}(\theta) = \eta_i(\theta)$ and sufficient statistics $s_{\mathcal{A},i}(\mathbf{x}_{\mathcal{A}})$ ($i = 1, \ldots, m_{\mathcal{A}}, \mathcal{A} \in \mathcal{L}$). Therefore, the exponential family can be reduced to an exponential family with natural parameter vector

(4.1)
$$\boldsymbol{\eta}(\boldsymbol{\theta}) = \left(\eta_1(\boldsymbol{\theta}), \dots, \eta_m(\boldsymbol{\theta})\right)$$

and sufficient statistic vector

$$s(\boldsymbol{x}_{\mathcal{L}}) = (s_1(\boldsymbol{x}_{\mathcal{L}}), \ldots, s_m(\boldsymbol{x}_{\mathcal{L}})),$$

where $s_i(\mathbf{x}_{\mathcal{L}}) = \sum_{\mathcal{A} \in \mathcal{L}} s_{\mathcal{A},i}(\mathbf{x}_{\mathcal{A}})$ (i = 1, ..., m) and $m = \max_{\mathcal{A} \in \mathcal{L}} m_{\mathcal{A}}$.

Consider a subgraph $(X_{\mathcal{K}}, Y_{\mathcal{K}})$ induced by a subset of neighborhoods $\mathcal{K} \subset \mathcal{L}$. Then the subgraph $(X_{\mathcal{K}}, Y_{\mathcal{K}})$ with subset of neighborhoods \mathcal{K} is extendable to the population graph $(X_{\mathcal{L}}, Y_{\mathcal{L}})$ with set of neighborhoods $\mathcal{L} \supset \mathcal{K}$ as follows.

PROPOSITION 3. Consider a full or nonfull, curved exponential-family random graph with set of neighborhoods \mathcal{L} , countable support $\mathbb{X}_{\mathcal{L}}$, and local dependence. Assume that, for all $\mathbf{y}_{\mathcal{L}} \in \mathbb{Y}_{\mathcal{L}}$,

$$\mathbb{P}(Y_{\mathcal{L}} = y_{\mathcal{L}}) = \prod_{\mathbb{C} \in \mathcal{L}, \mathbb{D} \in \mathcal{L}, \mathbb{C} \neq \mathbb{D}} \mathbb{P}(Y_{\mathbb{C}, \mathbb{D}} = y_{\mathbb{C}, \mathbb{D}}),$$

where $\mathbf{Y}_{\mathcal{C},\mathcal{D}} = (Y_{i,j})_{i \in \mathcal{C}, j \in \mathcal{D}}$. Then, for all $\boldsymbol{\theta} \in \boldsymbol{\Theta} \subseteq \{\boldsymbol{\theta} \in \mathbb{R}^q : \psi_{\mathcal{L}}(\boldsymbol{\eta}(\boldsymbol{\theta})) < \infty\}$, all $\mathcal{K} \subset \mathcal{L}$, and all $\mathbf{x}_{\mathcal{K}} \in \mathbb{X}_{\mathcal{K}}$ and $\mathbf{y}_{\mathcal{K}} \in \mathbb{Y}_{\mathcal{K}}$,

$$\begin{split} \mathbb{P}_{\eta(\theta)}(X_{\mathcal{K}} = x_{\mathcal{K}}, Y_{\mathcal{K}} = y_{\mathcal{K}}, X_{\mathcal{L} \setminus \mathcal{K}} \in \mathbb{X}_{\mathcal{L} \setminus \mathcal{K}}, Y_{\mathcal{L} \setminus \mathcal{K}} \in \mathbb{Y}_{\mathcal{L} \setminus \mathcal{K}}) \\ = \mathbb{P}_{\eta(\theta)}(X_{\mathcal{K}} = x_{\mathcal{K}}, Y_{\mathcal{K}} = y_{\mathcal{K}}), \end{split}$$

where

$$\psi_{\mathcal{L}}(\boldsymbol{\eta}(\boldsymbol{\theta})) = \sum_{\mathcal{A} \in \mathcal{L}} \psi_{\mathcal{A}}(\boldsymbol{\eta}(\boldsymbol{\theta}))$$
$$= \sum_{\mathcal{A} \in \mathcal{L}} \log \sum_{\boldsymbol{x}_{\mathcal{A}} \in \mathbb{X}_{\mathcal{A}}} \exp(\langle \boldsymbol{\eta}(\boldsymbol{\theta}), s_{\mathcal{A}}(\boldsymbol{x}_{\mathcal{A}}) \rangle) \nu_{\mathcal{A}}(\boldsymbol{x}_{\mathcal{A}})$$

The marginal density of a subgraph $\mathbf{x}_{\mathcal{K}} \in \mathbb{X}_{\mathcal{K}}$ of $\mathbf{x}_{\mathcal{L}} \in \mathbb{X}_{\mathcal{L}}$ induced by $\mathcal{K} \subset \mathcal{L}$ is an exponential-family density with support $\mathbb{X}_{\mathcal{K}}$ and local dependence:

$$\sum_{\boldsymbol{x}_{\mathcal{L}\setminus\mathcal{K}}\in\mathbb{X}_{\mathcal{L}\setminus\mathcal{K}}}p_{\boldsymbol{\eta}(\boldsymbol{\theta})}(\boldsymbol{x}_{\mathcal{L}})=p_{\boldsymbol{\eta}(\boldsymbol{\theta})}(\boldsymbol{x}_{\mathcal{K}})=\exp(\langle \boldsymbol{\eta}(\boldsymbol{\theta}), \boldsymbol{s}(\boldsymbol{x}_{\mathcal{K}})\rangle-\psi_{\mathcal{K}}(\boldsymbol{\eta}(\boldsymbol{\theta})))\boldsymbol{v}_{\mathcal{K}}(\boldsymbol{x}_{\mathcal{K}}),$$

where $\psi_{\mathcal{K}}(\boldsymbol{\eta}(\boldsymbol{\theta})) = \sum_{\mathcal{A}\in\mathcal{K}} \psi_{\mathcal{A}}(\boldsymbol{\eta}(\boldsymbol{\theta})), \, \boldsymbol{\eta}(\boldsymbol{\theta}) = (\eta_1(\boldsymbol{\theta}), \dots, \eta_m(\boldsymbol{\theta})), \, s(\boldsymbol{x}_{\mathcal{K}}) = (\sum_{\mathcal{A}\in\mathcal{K}} s_{\mathcal{A},1}(\boldsymbol{x}_{\mathcal{A}}), \dots, \sum_{\mathcal{A}\in\mathcal{K}} s_{\mathcal{A},m}(\boldsymbol{x}_{\mathcal{A}})), \, and \, v_{\mathcal{K}}(\boldsymbol{x}_{\mathcal{K}}) = \prod_{\mathcal{A}\in\mathcal{K}} v_{\mathcal{A}}(\boldsymbol{x}_{\mathcal{A}}).$

Thus, in the above-mentioned sense, the exponential-family random graph induced by a subset of neighborhoods \mathcal{K} can be extended to the exponential-family random graph with set of neighborhoods $\mathcal{L} \supset \mathcal{K}$. A more restrictive result was proved by Schweinberger and Handcock ([48], Theorem 1).

4.2. Subgraph-to-graph estimators. The extendability of exponential-family random graphs with multilevel structure discussed in Section 4.1 facilitates subgraph-to-graph estimation.

To demonstrate, let \mathcal{L} be the set of neighborhoods of the population graph and assume that $\mathbf{x}_{\mathcal{L}} \in \mathbb{X}_{\mathcal{L}}$ was generated by an exponential family with countable support $\mathbb{X}_{\mathcal{L}}$ and local dependence. Suppose that it is infeasible to observe $\mathbf{x}_{\mathcal{L}} \in \mathbb{X}_{\mathcal{L}}$, but it is feasible to sample a subset of neighborhoods $\mathcal{K} \subset \mathcal{L}$ and collect data on the subgraphs induced by $\mathcal{K} \subset \mathcal{L}$. We assume henceforth that the sampling design is ignorable in the sense of Rubin [45] and Handcock and Gile [19], that is, the probability of observing subgraphs does not depend on the unobserved subgraphs. A simple example is a sampling design that samples neighborhoods at random and collects data on the subgraphs induced by the sampled neighborhoods.

By Proposition 3 and the ignorability of the sampling design [19, 45], the observed-data likelihood function based on the observed subgraph $x_{\mathcal{K}} \in \mathbb{X}_{\mathcal{K}}$ of $x_{\mathcal{L}} \in \mathbb{X}_{\mathcal{L}}$ is proportional to

$$\sum_{\boldsymbol{x}_{\mathcal{L}}\setminus\mathcal{K}\in\mathbb{X}_{\mathcal{L}\setminus\mathcal{K}}}p_{\boldsymbol{\eta}(\boldsymbol{\theta})}(\boldsymbol{x}_{\mathcal{L}})=p_{\boldsymbol{\eta}(\boldsymbol{\theta})}(\boldsymbol{x}_{\mathcal{K}}),\quad \boldsymbol{x}_{\mathcal{L}}\in\mathbb{X}_{\mathcal{L}},$$

where $\eta(\theta)$ is of the form (4.1). In other words, maximum likelihood estimation can be based on $p_{\eta(\theta)}(\mathbf{x}_{\mathcal{K}})$. Motivated by the same considerations we outlined in Section 3, we therefore consider an estimating function of the form

$$g_{\mathcal{K}}(\boldsymbol{\theta}; \boldsymbol{\mu}_{\mathcal{K}}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))) = \|\boldsymbol{\mu}_{\mathcal{K}}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star})) - \boldsymbol{\mu}_{\mathcal{K}}(\boldsymbol{\eta}(\boldsymbol{\theta}))\|_{2}$$

where $\mu_{\mathcal{K}}(\eta(\theta^{\star})) = s(\mathbf{x}_{\mathcal{K}}), \mu_{\mathcal{K}}(\eta(\theta)) = \mathbb{E}_{\eta(\theta)}s(\mathbf{X}_{\mathcal{K}})$, and $s(\mathbf{x}_{\mathcal{K}})$ is defined in Proposition 3. The data-generating parameter vector θ^{\star} of the population graph can hence be estimated by the estimator $\hat{\theta}_{\mathcal{K}}$ based on the observed subgraph $\mathbf{x}_{\mathcal{K}} \in \mathbb{X}_{\mathcal{K}}$ of $\mathbf{x}_{\mathcal{L}} \in \mathbb{X}_{\mathcal{L}}$:

$$\widehat{\boldsymbol{\theta}}_{\mathcal{K}} = \Big\{ \boldsymbol{\theta} \in \boldsymbol{\Theta} : g_{\mathcal{K}}(\boldsymbol{\theta}; \mu_{\mathcal{K}}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))) = \inf_{\boldsymbol{\dot{\theta}} \in \boldsymbol{\Theta}} g_{\mathcal{K}}(\boldsymbol{\dot{\theta}}; \mu_{\mathcal{K}}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star}))) \Big\}.$$

The following concentration result shows that, with high probability, the estimator $\hat{\theta}_{\mathcal{K}}$ based on the observed subgraph induced by $\mathcal{K} \subset \mathcal{L}$ is close to the data-generating parameter vector θ^* of the population graph as long as the number of sampled neighborhoods $|\mathcal{K}|$ is large relative to $\|\mathcal{L}\|_{\infty} = \max_{\mathcal{A} \in \mathcal{L}} |\mathcal{A}|$ and $m = \max_{\mathcal{A} \in \mathcal{L}} m_{\mathcal{A}}$. THEOREM 2. Consider a full or nonfull, curved exponential-family random graph with set of neighborhoods \mathcal{L} , countable support $\mathbb{X}_{\mathcal{L}}$, and local dependence. Let

$$\boldsymbol{\Theta} \subseteq \{ \boldsymbol{ heta} \in \mathbb{R}^q : \psi_{\mathcal{L}} (\boldsymbol{\eta}(\boldsymbol{ heta})) < \infty \}.$$

Assume that $\theta^* \in int(\Theta)$ and that, for all $\epsilon > 0$ small enough so that $\mathbb{B}(\theta^*, \epsilon) \subseteq int(\Theta)$, there exists $\gamma(\epsilon) > 0$ such that, for all $\theta \in \Theta \setminus \mathbb{B}(\theta^*, \epsilon)$, we have $\eta(\theta) \in \mathfrak{N} \setminus \mathbb{B}(\eta(\theta^*), \gamma(\epsilon))$. In addition, assume that there exist $\delta(\epsilon) > 0$ and A > 0 such that, for all $\mathcal{K} \subseteq \mathcal{L}$ and all $\eta(\theta) \in \mathfrak{N} \setminus \mathbb{B}(\eta(\theta^*), \gamma(\epsilon))$,

(4.2)
$$\|\boldsymbol{\mu}_{\mathcal{K}}(\boldsymbol{\eta}(\boldsymbol{\theta}^{\star})) - \boldsymbol{\mu}_{\mathcal{K}}(\boldsymbol{\eta}(\boldsymbol{\theta}))\|_{2} \ge \delta(\epsilon) \sum_{\mathcal{A}\in\mathcal{K}} {|\mathcal{A}| \choose 2}^{\alpha} \text{ for some } 0 \le \alpha \le 1$$

and, for all $\mathcal{K} \subseteq \mathcal{L}$ and all $(\mathbf{x}_1, \mathbf{x}_2) \in \mathbb{X}_{\mathcal{K}} \times \mathbb{X}_{\mathcal{K}}$,

(4.3)
$$\|s(\boldsymbol{x}_1) - s(\boldsymbol{x}_2)\|_{\infty} \le Ad(\boldsymbol{x}_1, \boldsymbol{x}_2) \|\mathcal{K}\|_{\infty},$$

where $\|\mathcal{K}\|_{\infty} = \max_{\mathcal{A}\in\mathcal{K}} |\mathcal{A}|$. Then, for all $\epsilon > 0$ small enough so that $\mathcal{B}(\boldsymbol{\theta}^{\star}, \epsilon) \subseteq \operatorname{int}(\boldsymbol{\Theta})$, there exist $\kappa(\epsilon) > 0$ and C > 0 such that

$$\mathbb{P}(\widehat{\boldsymbol{\theta}}_{\mathcal{K}} \in \mathcal{B}(\boldsymbol{\theta}^{\star}, \epsilon)) \geq 1 - 2\exp\left(-\frac{\kappa(\epsilon)^2 C|\mathcal{K}|}{m \|\mathcal{L}\|_{\infty}^{4(2-\alpha)}} + \log m\right).$$

If the exponential family is full, then $\widehat{\theta}_{\mathcal{K}}$ is unique in the event $\widehat{\theta}_{\mathcal{K}} \in \mathbb{B}(\theta^{\star}, \epsilon)$.

Theorem 2 shows that there are costs associated with observing a subset of neighborhoods $\mathcal{K} \subset \mathcal{L}$ rather than the whole set of neighborhoods \mathcal{L} of the population graph: The probability of event $\hat{\theta}_{\mathcal{K}} \notin \mathcal{B}(\theta^{\star}, \epsilon)$ decays with the number of sampled neighborhoods $|\mathcal{K}|$ and is hence lowest when the whole set of neighborhoods \mathcal{L} of the population graph is sampled.

As a specific example, consider the main example of Section 3.3: curved exponentialfamily random graphs with within-neighborhood edge and geometrically weighted edgewise shared partner terms.

COROLLARY 2. Consider a curved exponential-family random graph with set of neighborhoods \mathcal{L} , countable support $\mathbb{X}_{\mathcal{L}}$, and local dependence induced by within-neighborhood edge and geometrically weighted edgewise shared partner terms. Let $\Theta = \mathbb{R} \times (0, 1)$ and assume that $\theta^* \in int(\Theta)$. Then all conditions of Theorem 2 are satisfied and hence, for all $\epsilon > 0$ small enough so that $\mathcal{B}(\theta^*, \epsilon) \subseteq int(\Theta)$, there exist $\kappa(\epsilon) > 0$ and C > 0 such that

$$\mathbb{P}(\widehat{\boldsymbol{\theta}}_{\mathcal{K}} \in \mathcal{B}(\boldsymbol{\theta}^{\star}, \epsilon)) \geq 1 - 2\exp\left(-\frac{\kappa(\epsilon)^2 C |\mathcal{K}|}{\|\mathcal{L}\|_{\infty}^6} + \log\|\mathcal{L}\|_{\infty}\right),$$

provided $|\mathcal{A}| \ge 4$ ($\mathcal{A} \in \mathcal{L}$) and $|\mathcal{K}| \ge 2$.

REMARK 4 ("Bad" subsets of neighborhoods $\mathcal{K} \subseteq \mathcal{L}$). Since the neighborhoods need not have the same size, it is natural to ask whether it is possible to sample a "bad" subset of neighborhoods \mathcal{K} with too small or too large neighborhoods, which could make it challenging to estimate some of the parameters. However, the assumptions of Theorem 2 rule out "bad" subsets of neighborhoods \mathcal{K} , for two reasons. First, while some neighborhoods may be larger than others, Theorem 2 assumes that the neighborhoods are of the same order of magnitude, as defined in Section 2. In other words, the neighborhoods have similar sizes. Second, the conditions of Theorem 2 assume that the model satisfies identifiability and smoothness conditions for all possible subsets of neighborhoods $\mathcal{K} \subseteq \mathcal{L}$. Corollary 2 shows that, in the special case of curved exponential-family random graphs with within-neighborhood edge and geometrically weighted edgewise shared partner terms, the identifiability conditions require $|\mathcal{A}| \ge 4$ for all neighborhoods $\mathcal{A} \in \mathcal{L}$ of the population graph. Thus, no neighborhood can be too small, and no neighborhood can be too large, because all neighborhoods are of the same order of magnitude. As a consequence, under the stated assumptions, it is impossible to sample a "bad" subset of neighborhoods \mathcal{K} with too small or too large neighborhoods.

5. Comparison with existing consistency results. To compare our concentration and consistency results to existing consistency results, we focus on exponential-family random graphs with dependent edges. It is worth noting that there are consistency results for exponential-family random graphs with independence assumptions—see, for example, Diaconis, Chatterjee and Sly [11], Rinaldo, Petrovic and Fienberg [44], Krivitsky and Kolaczyk [35] and Yan, Leng and Zhu [60], Yan et al. [59]—but such independence assumptions may not be satisfied in applications, as discussed in the Introduction.

Concerning exponential-family random graphs with dependent edges, Shalizi and Rinaldo [51] showed that maximum likelihood estimators of natural parameters of fixed dimension are consistent provided exponential-family random graphs satisfy strong extendability or projectability assumptions. However, those projectability assumptions rule out dependencies induced by transitivity and many other interesting network phenomena. Xiang and Neville [58] reported consistency results under weak dependence assumptions, but did not give any example of an exponential-family random graph with dependent edges that satisfies those assumptions. Mukherjee [40] showed that consistent estimation of the so-called two-star model is possible, but those results have not been extended to other exponential-family random graphs. In addition, Shalizi and Rinaldo [51], Xiang and Neville [58] and Mukherjee [40] focus on consistency results for estimators of natural parameter vectors whose dimensions do not depend on the number of nodes. By contrast, we advance the statistical theory of exponential-family random graphs by providing the first concentration and consistency results that cover:

- a wide range of exponential-family random graphs with dependence among edges induced by transitivity and other interesting network phenomena;
- curved exponential-family random graphs with dependent edges and parameter vectors whose dimension depends on the number of nodes (Section 3);
- maximum likelihood and *M*-estimators (Section 3 and Appendix B);
- correct and incorrect model specifications (Section 3 and Appendix B);
- subgraph-to-graph estimators (Section 4).

These results underscore the importance of additional structure: It is the additional structure in the form of multilevel structure that facilitates these results.

6. Simulation results. To shed light on the finite-graph properties of maximum likelihood estimators, we generated data from the canonical and curved exponential-family random graphs mentioned in Section 3.3. We used R package hergm [49] to generate 1000 graphs from each model and estimated the data-generating parameter vector by Monte Carlo maximum likelihood estimators [24].

We first consider canonical exponential-family random graphs with support $\mathbb{X} = \{0, 1\}^{\sum_{k=1}^{K} \binom{|\mathcal{A}_k|}{2}}$ and local dependence induced by within-neighborhood edge and transitive edge terms [26]. Within-neighborhood edge and transitive edge terms correspond to neighborhood-dependent natural parameters $\eta_{k,1}(\boldsymbol{\theta}) = \theta_1$ and $\eta_{k,2}(\boldsymbol{\theta}) = \theta_2$ and sufficient



FIG. 2. 1000 estimates of the exponential-family random graph with within-neighborhood edge and transitive edge terms, where each graph consists of K = 100, 200 and 300 neighborhoods of size 50 with natural parameter vectors $\eta_k(\theta) = (\theta_1, \theta_2)$. The horizontal and vertical lines indicate the coordinates of the data-generating parameter vector $\theta^* = (\theta_1^*, \theta_2^*)$.

statistics $s_{k,1}(\mathbf{x}_k)$ and $s_{k,2}(\mathbf{x}_k)$ given by

$$s_{k,1}(\boldsymbol{x}_k) = \sum_{i \in \mathcal{A}_k < j \in \mathcal{A}_k} x_{i,j},$$

$$s_{k,2}(\boldsymbol{x}_k) = \sum_{i \in \mathcal{A}_k < j \in \mathcal{A}_k} x_{i,j} \max_{h \in \mathcal{A}_k, h \neq i,j} x_{i,h} x_{j,h},$$

where k = 1, ..., K. It is worth noting that the number of transitive edges is not the same as the number of triangles. A discussion of the model along with concentration results for maximum likelihood estimators can be found in Appendix A (see the Supplementary Material, [50]). Figure 2 shows 1000 estimates of the exponential-family random graph with withinneighborhood edge and transitive edge terms, where each graph consists of K = 100, 200 and 300 neighborhoods of size 50 with natural parameter vectors $\eta_k(\theta) = (\theta_1, \theta_2)$ (k = 1, ..., K). The figure suggests that the probability mass of estimators becomes more and more concentrated in a neighborhood of the data-generating parameters as the number of neighborhoods K increases from 100 to 300, demonstrating that the concentration results in Section 3 are manifest when K is in the low hundreds and $||\mathcal{A}||_{\infty} = 50$.

Figure 3 sheds light on the performance of a simple form of a size-dependent parameterization that allows small and large neighborhoods to have different parameters. We consider exponential-family random graphs with within-neighborhood edge and transitive edge terms, where each graph consists of 33 neighborhoods of size 25 with natural parameter vectors $\eta_k(\theta) = (\theta_1, \theta_2) \ (k = 1, ..., 33), 34$ neighborhoods of size 50 with natural parameter vectors $\eta_k(\theta) = (\theta_1 + \theta_3, \theta_2 + \theta_4) \ (k = 34, ..., 67),$ and 33 neighborhoods of size 75 with natural



FIG. 3. 1000 estimates of the exponential-family random graph with within-neighborhood edge and transitive edge terms, where each graph consists of 33 neighborhoods of size 25 with natural parameter vectors $\eta_k(\theta) = (\theta_1, \theta_2)$, 34 neighborhoods of size 50 with natural parameter vectors $\eta_k(\theta) = (\theta_1 + \theta_3, \theta_2 + \theta_4)$, and 33 neighborhoods of size 75 with natural parameter vectors $\eta_k(\theta) = (\theta_1 + \theta_5, \theta_2 + \theta_6)$. The horizontal and vertical lines indicate the coordinates of the data-generating parameter vector $\theta^* = (\theta_1^*, \theta_2^*, \theta_3^*, \theta_4^*, \theta_5^*, \theta_6^*)$.



FIG. 4. 1000 estimates of the curved exponential-family random graph with within-neighborhood edge and geometrically weighted edgewise shared partner terms, where each graph consists of K = 100 neighborhoods of size 50 with natural parameter vectors $\eta_k(\theta) = (\theta_1, \eta_{k,1}(\theta_2), \dots, \eta_{k,48}(\theta_2))$. The vertical lines indicate the coordinates of the data-generating parameter vector $\theta^* = (\theta_1^*, \theta_2^*)$.

parameter vectors $\eta_k(\theta) = (\theta_1 + \theta_5, \theta_2 + \theta_6)$ (k = 68, ..., 100). Figure 3 demonstrates that the estimates of the baseline edge and transitive edge parameters θ_1 and θ_2 tend to be closer to the data-generating parameters than the deviation parameters $\theta_3, \theta_4, \theta_5$ and θ_6 , because θ_1 and θ_2 are estimated from all neighborhoods whereas $\theta_3, \theta_4, \theta_5$ and θ_6 are estimated from a subset of neighborhoods.

Figure 4 shows 1000 estimates of the curved exponential-family random graph with within-neighborhood edge and geometrically weighted edgewise shared partner terms described in Section 3.3. Each graph consists of K = 100 neighborhoods of size 50 with natural parameter vectors $\eta_k(\theta) = (\theta_1, \eta_{k,1}(\theta_2), \dots, \eta_{k,48}(\theta_2))$, where θ_1 is the natural parameter of the edge term and $\eta_{k,1}(\theta_2), \dots, \eta_{k,48}(\theta_2)$ are the natural parameters of the geometrically weighted edgewise shared partner term $(k = 1, \dots, 100)$. The figure shows that the probability mass of the estimators is concentrated in a small neighborhood of the data-generating parameters.

7. Discussion. We have taken constructive steps to demonstrate that statistical inference for exponential-family random graphs with dependence among edges induced by transitivity and other interesting network phenomena is possible, provided additional structure in the form of multilevel structure is available. The theoretical results reported here underscore the importance of additional structure. In practice, many other forms of additional structure exist and could be used for the purpose of facilitating statistical inference for exponential-family random graphs (e.g., other forms of multilevel structure or spatial structure).

Last, but not least, while we have focused here on theoretical results showing that multilevel structure facilitates statistical inference, it is worth noting that multilevel structure has computational benefits as well: The contributions of neighborhoods to estimating functions for example, the expectations of within-neighborhood sufficient statistics—may be computed or approximated by exploiting parallel computing on computing clusters.

Acknowledgments. We are grateful to the Editor, the Associate Editor and three anonymous referees, whose constructive suggestions led to substantial improvements of the paper. This work was supported by NSE Cropts DMS 1512644 and DMS 1812110

This work was supported by NSF Grants DMS-1513644 and DMS-1812119.

SUPPLEMENTARY MATERIAL

Supplement to "Concentration and consistency results for canonical and curved exponential-family models of random graphs" (DOI: 10.1214/19-AOS1810SUPP; .pdf). All results are proved in the supplement [50]. In addition, the supplement contains concentration results for *M*-estimators. These results cover correct and incorrect model specifications.

REFERENCES

- BARNDORFF-NIELSEN, O. (1978). Information and Exponential Families in Statistical Theory. Wiley, Chichester. MR0489333
- [2] BHAMIDI, S., BRESLER, G. and SLY, A. (2011). Mixing time of exponential random graphs. Ann. Appl. Probab. 21 2146–2170. MR2895412 https://doi.org/10.1214/10-AAP740
- [3] BHATTACHARYA, B. B. and MUKHERJEE, S. (2018). Inference in Ising models. *Bernoulli* 24 493–525. MR3706767 https://doi.org/10.3150/16-BEJ886
- [4] BOUCHERON, S., LUGOSI, G. and MASSART, P. (2013). Concentration Inequalities: A Nonasymptotic Theory of Independence. Oxford Univ. Press, Oxford. MR3185193 https://doi.org/10.1093/acprof:oso/ 9780199535255.001.0001
- [5] BROWN, L. D. (1986). Fundamentals of Statistical Exponential Families with Applications in Statistical Decision Theory. Institute of Mathematical Statistics Lecture Notes—Monograph Series 9. IMS, Hayward, CA. MR0882001
- [6] BUTTS, C. T. (2011). Bernoulli graph bounds for general random graph models. Sociol. Method. 41 299– 345.
- [7] BUTTS, C. T. and ALMQUIST, Z. W. (2015). A flexible parameterization for baseline mean degree in multiple-network ERGMs. J. Math. Sociol. 39 163–167. MR3367715 https://doi.org/10.1080/ 0022250X.2014.967851
- [8] CHATTERJEE, S. (2005). Concentration inequalities with exchangeable pairs. Ph.D. thesis, Dept. Statistics, Stanford Univ., Satnford, CA.
- [9] CHATTERJEE, S. (2007). Estimation in spin glasses: A first step. Ann. Statist. 35 1931–1946. MR2363958 https://doi.org/10.1214/009053607000000109
- [10] CHATTERJEE, S. and DIACONIS, P. (2013). Estimating and understanding exponential random graph models. Ann. Statist. 41 2428–2461. MR3127871 https://doi.org/10.1214/13-AOS1155
- [11] CHATTERJEE, S., DIACONIS, P. and SLY, A. (2011). Random graphs with a given degree sequence. Ann. Appl. Probab. 21 1400–1435. MR2857452 https://doi.org/10.1214/10-AAP728
- [12] CRANE, H. and DEMPSEY, W. (2015). A framework for statistical network modeling. Available at https: //arxiv.org/abs/1509.08185.
- [13] EFRON, B. (1975). Defining the curvature of a statistical problem (with applications to second order efficiency). Ann. Statist. 3 1189–1242. MR0428531
- [14] EFRON, B. (1978). The geometry of exponential families. Ann. Statist. 6 362–376. MR0471152
- [15] FRANK, O. and STRAUSS, D. (1986). Markov graphs. J. Amer. Statist. Assoc. 81 832-842. MR0860518
- [16] GEYER, C. J. (2009). Likelihood inference in exponential families and directions of recession. *Electron. J. Stat.* 3 259–289. MR2495839 https://doi.org/10.1214/08-EJS349
- [17] GODAMBE, V. P. and KALE, B. K. (1991). Estimating Functions. Oxford Univ. Press, Oxford.
- [18] HANDCOCK, M. S. (2003). Statistical models for social networks: Inference and degeneracy. In *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers* (R. Breiger, K. Carley and P. Pattison, eds.) 1–12. National Academies Press, Washington, DC.
- [19] HANDCOCK, M. S. and GILE, K. J. (2010). Modeling social networks from sampled data. Ann. Appl. Stat. 4 5–25. MR2758082 https://doi.org/10.1214/08-AOAS221
- [20] HARRIS, J. K. (2013). An Introduction to Exponential Random Graph Modeling. Sage, Thousand Oaks, CA.
- [21] HOLLAND, P. W. and LEINHARDT, S. (1976). Local structure in social networks. Sociol. Method. 1-45.
- [22] HOLLWAY, J., LOMI, A., PALLOTTI, F. and STADTFELD, C. (2017). Multilevel social spaces: The network dynamics of organizational fields. *Network Science* 5 187–212.
- [23] HUNTER, D. R. (2007). Curved exponential family models for social networks. Soc. Netw. 29 216–230.
- [24] HUNTER, D. R., GOODREAU, S. M. and HANDCOCK, M. S. (2008). Goodness of fit of social network models. J. Amer. Statist. Assoc. 103 248–258. MR2394635 https://doi.org/10.1198/ 016214507000000446
- [25] HUNTER, D. R. and HANDCOCK, M. S. (2006). Inference in curved exponential family models for networks. J. Comput. Graph. Statist. 15 565–583. MR2291264 https://doi.org/10.1198/ 106186006X133069
- [26] HUNTER, D. R., KRIVITSKY, P. N. and SCHWEINBERGER, M. (2012). Computational statistical methods for social network models. J. Comput. Graph. Statist. 21 856–882. MR3005801 https://doi.org/10. 1080/10618600.2012.732921
- [27] JANSON, S. and RUCIŃSKI, A. (2002). The infamous upper tail. Random Structures Algorithms 20 317– 342. MR1900611 https://doi.org/10.1002/rsa.10031
- [28] JONASSON, J. (1999). The random triangle model. J. Appl. Probab. 36 852–867. MR1737058 https://doi.org/10.1239/jap/1032374639

- [29] KASS, R. E. and VOS, P. W. (1997). Geometrical Foundations of Asymptotic Inference. Wiley Series in Probability and Statistics: Probability and Statistics. Wiley, New York. MR1461540 https://doi.org/10. 1002/9781118165980
- [30] KIM, J. H. and VU, V. H. (2004). Divide and conquer martingales and the number of triangles in a random graph. *Random Structures Algorithms* 24 166–174. MR2035874 https://doi.org/10.1002/rsa.10113
- [31] KOLACZYK, E. D. (2009). Statistical Analysis of Network Data: Methods and Models. Springer Series in Statistics. Springer, New York. MR2724362 https://doi.org/10.1007/978-0-387-88146-1
- [32] KONTOROVICH, L. and RAMANAN, K. (2008). Concentration inequalities for dependent random variables via the martingale method. Ann. Probab. 36 2126–2158. MR2478678 https://doi.org/10.1214/ 07-AOP384
- [33] KRIVITSKY, P. N. (2012). Exponential-family random graph models for valued networks. *Electron. J. Stat.* 6 1100–1128. MR2988440 https://doi.org/10.1214/12-EJS696
- [34] KRIVITSKY, P. N., HANDCOCK, M. S. and MORRIS, M. (2011). Adjusting for network size and composition effects in exponential-family random graph models. *Stat. Methodol.* 8 319–339. MR2800354 https://doi.org/10.1016/j.stamet.2011.01.005
- [35] KRIVITSKY, P. N. and KOLACZYK, E. D. (2015). On the question of effective sample size in network modeling: An asymptotic inquiry. *Statist. Sci.* 30 184–198. MR3353102 https://doi.org/10.1214/ 14-STS502
- [36] LAURITZEN, S., RINALDO, A. and SADEGHI, K. (2018). Random networks, graphical models and exchangeability. J. R. Stat. Soc. Ser. B. Stat. Methodol. 80 481–508. MR3798875 https://doi.org/10.1111/ rssb.12266
- [37] LAZEGA, E. and SNIJDERS, T. A. B., eds. (2016). *Multilevel Network Analysis for the Social Sciences*. Springer, Cham.
- [38] LOMI, A., ROBINS, G. and TRANMER, M. (2016). Introduction to multilevel social networks. Soc. Netw. 44 266–268.
- [39] LUSHER, D., KOSKINEN, J. and ROBINS, G. (2013). *Exponential Random Graph Models for Social Networks*. Cambridge Univ. Press, Cambridge.
- [40] MUKHERJEE, S. (2013). Consistent estimation in the two star exponential random graph model. Technical report. Dept. Statistics, Columbia Univ. New York. Available at arXiv:1310.4526v1.
- [41] NOWICKI, K. and SNIJDERS, T. A. B. (2001). Estimation and prediction for stochastic blockstructures. J. Amer. Statist. Assoc. 96 1077–1087. MR1947255 https://doi.org/10.1198/016214501753208735
- [42] RAVIKUMAR, P., WAINWRIGHT, M. J. and LAFFERTY, J. D. (2010). High-dimensional Ising model selection using l₁-regularized logistic regression. Ann. Statist. 38 1287–1319. MR2662343 https://doi.org/10.1214/09-AOS691
- [43] RINALDO, A., FIENBERG, S. E. and ZHOU, Y. (2009). On the geometry of discrete exponential families with application to exponential random graph models. *Electron. J. Stat.* 3 446–484. MR2507456 https://doi.org/10.1214/08-EJS350
- [44] RINALDO, A., PETROVIĆ, S. and FIENBERG, S. E. (2013). Maximum likelihood estimation in the βmodel. Ann. Statist. 41 1085–1110. MR3113804 https://doi.org/10.1214/12-AOS1078
- [45] RUBIN, D. B. (1976). Inference and missing data. *Biometrika* 63 581–592. MR0455196 https://doi.org/10. 1093/biomet/63.3.581
- [46] SAMSON, P.-M. (2000). Concentration of measure inequalities for Markov chains and Φ-mixing processes. Ann. Probab. 28 416–461. MR1756011 https://doi.org/10.1214/aop/1019160125
- [47] SCHWEINBERGER, M. (2011). Instability, sensitivity, and degeneracy of discrete exponential families. J. Amer. Statist. Assoc. 106 1361–1370. MR2896841 https://doi.org/10.1198/jasa.2011.tm10747
- [48] SCHWEINBERGER, M. and HANDCOCK, M. S. (2015). Local dependence in random graph models: Characterization, properties and statistical inference. J. R. Stat. Soc. Ser. B. Stat. Methodol. 77 647–676. MR3351449 https://doi.org/10.1111/rssb.12081
- [49] SCHWEINBERGER, M. and LUNA, P. (2018). HERGM: Hierarchical exponential-family random graph models. J. Stat. Softw. 85 1–39.
- [50] SCHWEINBERGER, M. and STEWART, J. (2019). Supplement to "Concentration and consistency results for canonical and curved exponential-family models of random graphs." https://doi.org/10.1214/ 19-AOS1810SUPP.
- [51] SHALIZI, C. R. and RINALDO, A. (2013). Consistency under sampling of exponential random graph models. Ann. Statist. 41 508–535. MR3099112 https://doi.org/10.1214/12-AOS1044
- [52] SLAUGHTER, A. J. and KOEHLY, L. M. (2016). Multilevel models for social networks: Hierarchical Bayesian approaches to exponential random graph modeling. *Soc. Netw.* 44 334–345.
- [53] SNIJDERS, T. A. B., PATTISON, P. E., ROBINS, G. L. and HANDCOCK, M. S. (2006). New specifications for exponential random graph models. *Sociol. Method.* 36 99–153.

- [54] STEWART, J., SCHWEINBERGER, M., BOJANOWSKI, M. and MORRIS, M. (2019). Multilevel network data facilitate statistical inference for curved ERGMs with geometrically weighted terms. *Soc. Netw.* To appear.
- [55] VU, V. H. (2002). Concentration of non-Lipschitz functions and applications. *Random Structures Algo*rithms 20 262–316. MR1900610 https://doi.org/10.1002/rsa.10032
- [56] WANG, P., ROBINS, G., PATTISON, P. and LAZEGA, E. (2013). Exponential random graph models for multilevel networks. Soc. Netw. 35 96–115.
- [57] WASSERMAN, S. and PATTISON, P. (1996). Logit models and logistic regressions for social networks. I. An introduction to Markov graphs and p. Psychometrika 61 401–425. MR1424909 https://doi.org/10. 1007/BF02294547
- [58] XIANG, R. and NEVILLE, J. (2011). Relational learning with one network: An asymptotic analysis. In Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS) 1–10.
- [59] YAN, T., JIANG, B., FIENBERG, S. E. and LENG, C. (2019). Statistical inference in a directed network model with covariates. J. Amer. Statist. Assoc. 114 857–868. MR3963186
- [60] YAN, T., LENG, C. and ZHU, J. (2016). Asymptotics in directed exponential random graph models with an increasing bi-degree sequence. Ann. Statist. 44 31–57. MR3449761 https://doi.org/10.1214/ 15-AOS1343
- [61] YANG, E., RAVIKUMAR, P., ALLEN, G. I. and LIU, Z. (2015). Graphical models via univariate exponential family distributions. J. Mach. Learn. Res. 16 3813–3847. MR3450553
- [62] ZAPPA, P. and LOMI, A. (2015). The analysis of multilevel networks in organizations: Models and empirical tests. *Organizational Research Methods* **18** 542–569.