The Annals of Statistics
2013, Vol. 41, No. 5, 2428–2461
DOI: 10.1214/13-AOS1155

ESTIMATING AND UNDERSTANDING EXPONENTIAL RANDOM GRAPH MODELS

BY SOURAV CHATTERJEE[1] AND PERSI DIACONIS[2]

Stanford University

We introduce a method for the theoretical analysis of exponential random graph models. The method is based on a large-deviations approximation to the normalizing constant shown to be consistent using theory developed by Chatterjee and Varadhan [*European J. Combin.* **32** (2011) 1000–1017]. The theory explains a host of difficulties encountered by applied workers: many distinct models have essentially the same MLE, rendering the problems "practically" ill-posed. We give the first rigorous proofs of "degeneracy" observed in these models. Here, almost all graphs have essentially no edges or are essentially complete. We supplement recent work of Bhamidi, Bresler and Sly [*2008 IEEE 49th Annual IEEE Symposium on Foundations of Computer Science* (*FOCS*) (2008) 803–812 IEEE] showing that for many models, the extra sufficient statistics are useless: most realizations look like the results of a simple Erdős–Rényi model. We also find classes of models where the limiting graphs differ from Erdős–Rényi graphs. A limitation of our approach, inherited from the limitation of graph limit theory, is that it works only for dense graphs.

1. Introduction. Graph and network data are increasingly common and a host of statistical methods have emerged in recent years. Entry to this large literature may be had from the research papers and surveys in Fienberg [21, 22]. One mainstay of the emerging theory are the exponential families

$$p_\beta(G) = \exp\left(\sum_{i=1}^{k} \beta_i T_i(G) - \psi(\beta)\right), \tag{1.1}$$

where $\beta = (\beta_1, \ldots, \beta_k)$ is a vector of real parameters, $T_1, T_2, \ldots, T_k$ are functions on the space of graphs (e.g., the number of edges, triangles, stars, cycles, …), and ψ is the normalizing constant. In this paper, T_1 is usually taken to be the number of edges (or a constant multiple of it).

We review the literature of these models in Section 2.1. Estimating the parameters in these models has proved to be a challenging task. First, the normalizing

Received March 2011; revised July 2013.
[1]Supported in part by NSF Grant DMS-10-05312 and a Sloan Research Fellowship.
[2]Supported in part by NSF Grant DMS-08-04324.
MSC2010 subject classifications. Primary 62F10, 05C80; secondary 62P25, 60F10.
Key words and phrases. Random graph, Erdős–Rényi, graph limit, exponential random graph models, parameter estimation.

constant $\psi(\beta)$ is unknown. Second, very different values of β can give rise to essentially the same distribution on graphs.

Here is an example: consider the model on simple graphs with n vertices,

$$p_{\beta_1,\beta_2}(G) = \exp\left(2\beta_1 E + \frac{6\beta_2}{n}\Delta - n^2\psi_n(\beta_1,\beta_2)\right), \tag{1.2}$$

where E, Δ denote the number of edges and triangles in the graph G. The normalization of the model ensures nontrivial large n limits. Without scaling, for large n, almost all graphs are empty or full. This model is studied by Strauss [54], Park and Newman [45, 46], Häggstrom and Jonasson [30], and many others.

Theorems 3.1 and 4.1 will show that for n large and nonnegative β_2,

$$\psi_n(\beta_1,\beta_2) \simeq \sup_{0\le u\le 1}\left(\beta_1 u + \beta_2 u^3 - \frac{1}{2}u\log u - \frac{1}{2}(1-u)\log(1-u)\right). \tag{1.3}$$

The maximizing value of the right-hand side is denoted $u^*(\beta_1,\beta_2)$. A plot of this function appears in Figure 1. Theorem 4.2 shows that for any $\beta_1 \in \mathbb{R}$ and $\beta_2 > 0$, with high probability, a pick from p_{β_1,β_2} is essentially the same as an Erdős–Rényi graph generated by including edges independently with probability $u^*(\beta_1,\beta_2)$. This phenomenon has previously been identified by Bhamidi et al. [6] and is discussed further in Section 2.1. Figure 2 shows the contour lines for Figure 1. All the (β_1,β_2) values on the same contour line lead to the same Erdős–Rényi model in the limit.

Our development uses the emerging tools of graph limits as developed by Lovász and coworkers. We give an overview in Section 2.2. Briefly, a sequence of graphs G_n converges to a limit if the proportion of edges, triangles, and other small subgraphs in G_n converges. There is a limiting object and the space of all

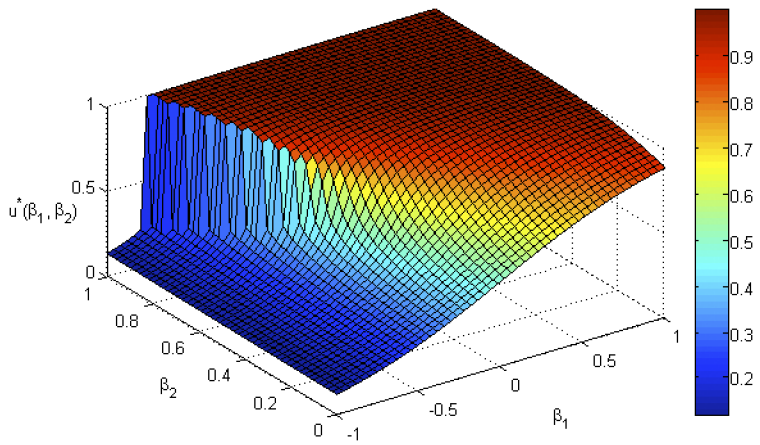

FIG. 1. *The plot of u^* against (β_1,β_2). There is a discontinuity on the left where u^* jumps from near* 0 *to near* 1; *this corresponds to a phase transition.* (*Picture by Sukhada Fadnavis.*)

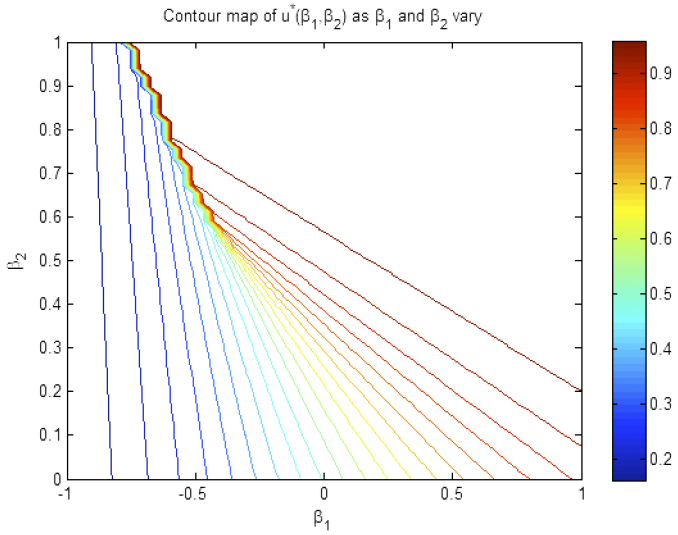

FIG. 2. *Contour lines for Figure 1. All pairs* (β_1, β_2) *on the same contour line correspond to the same value of* u^* *and hence those models will correspond to the same Erdős–Rényi model in the limit. The phase transition region is seen in the upper left-hand corner where all contour lines converge.* (*Picture by Sukhada Fadnavis*.)

these limiting objects serves as a useful compactification of the set of all graphs. Our theory works for functions $T_i(G)$ which are continuous in this topology. In their study of the large deviations of Erdős–Rényi random graphs, Chatterjee and Varadhan [15] derived the associated rate functions in the language of graph limit theory. Their work is crucial in the present development and is reviewed in Section 2.3.

Our main results are in Section 3 through Section 7. These sections contain only the statements of the theorems; all proofs are given in Section 8.

Working with general exponential models, Section 3 gives an extension of the approximation (1.3) for ψ_n (Theorem 3.1) and shows that, in the limit, almost all graphs from the model (1.1) are close to graphs where a certain functional is maximized. As will emerge, sometimes this maximum is taken on at a unique Erdős–Rényi model.

The main statistical motivation of this paper comes from the formula for the limit of the normalizing constant given in Theorem 3.1, since the normalizing constant is crucial for the computation of maximum likelihood estimates. At present, the computational tools used by practitioners to compute the normalizing constants in exponential random graph models become prohibitively time-consuming even for moderately large n. The theory initiated in this paper hopes to circumvent this problem by providing analytical formulas. As mentioned in the abstract, the limitation of our approach is that as of now, it applies only to dense graphs.

Incidentally, in a recent meeting at the American Institute of Mathematics, computer-intensive calculations carried out by Mark Handcock and David Hunter indicated that the formula given in Theorem 3.1 is actually a pretty good approximation to the exact value of the normalizing constant even for n as small as 20.

Section 4 studies the problem for the model (1.1) when $\beta_2, \ldots, \beta_k$ are positive (β_1 may have any sign). When the T_i's are subgraph counts, positive β_i's were originally intentioned (e.g., in [24]) to "encourage" the presence of the corresponding subgraphs. It is shown that the large-deviations approximation for ψ_n can be easily calculated as a one-dimensional maximization (Theorem 4.1). Further, amplifying the results of Bhamidi et al. [6], it is shown that in these cases, almost all realizations of the model (1.1) are close to an Erdős–Rényi graph (or perhaps a finite mixture of Erdős–Rényi graphs) (Theorem 4.2). These mixture cases actually occur for natural parameter values. Section 5 also gives a careful account of the phase transitions and near-degeneracies observed in the edge-triangle model (1.3).

Sections 6 and 7 investigate cases where β_i is allowed to be negative. While the general case remains open (and appears complicated), in Section 6 it is shown that Theorems 4.1 and 4.2 hold as stated if $(\beta_i)_{2\le i\le k}$ are sufficiently small in magnitude. This requires a careful study of associated Euler–Lagrange equations. Section 7 shows how the results change for the model containing edges and triangles when β_2 is negative. For sufficiently large negative β_2, typical realizations look like a random bipartite graph (where "random" means that the two parts, of equal size, are chosen uniformly at random from all possible choices). This is very different from the Erdős–Rényi model. The result generalizes to other models via an interesting analogy with the Erdős–Stone theorem from extremal graph theory.

A longer version of this paper with more pictures and additional results is available as "version 3" on arXiv (http://arxiv.org/pdf/1102.2650v3.pdf).

2. Background. This section gives needed background and notation in three areas. Exponential graph models (Section 2.1), graph limits (Section 2.2), and large deviations (Section 2.3).

2.1. *Exponential random graphs.* Let $\mathcal{G}_n$ be the space of all simple graphs on n labeled vertices ("simple" means undirected, with no loops or multiple edges). Thus, $\mathcal{G}_n$ contains $2^{\binom{n}{2}}$ elements. A variety of models in active use can be presented in exponential form

$$p_\beta(G) = \exp\left(\sum_{i=1}^{k} \beta_i T_i(G) - \psi(\beta)\right), \tag{2.1}$$

where $\beta = (\beta_1, \ldots, \beta_k)$ is a vector of real parameters, $T_1, T_2, \ldots, T_k$ are real-valued functions on $\mathcal{G}_n$, and $\psi(\beta)$ is a normalizing constant. Usually, T_i are taken to be counts of various subgraphs, for example, $T_1(G) = \#$ edges in G, $T_2(G) = \#$

triangles in $G, \ldots$. The main results of Section 3 work for more general "continuous functions" on graph space, such as the degree sequence or the eigenvalues of the adjacency matrix. This allows models with sufficient statistics of the form $\sum_{i=1}^{n} \beta_i d_i(G)$ with $d_i(G)$ the degree of vertex i. See, for example, [14].

These exponential models were used by Holland and Leinhardt [32] in the directed case. Frank and Strauss [24] developed them, showing that if T_i are chosen as edges, triangles, and stars of various sizes, the resulting random graph edges form a Markov random field. A general development is in Wasserman and Faust [57]. Newer developments, consisting mainly of new sufficient statistics and new ranges for parameters that give interesting and practically relevant structures, are summarized in Snijders et al. [53]. Finally, Rinaldo et al. [49] develop the geometric theory for this class of models with extensive further references.

A major problem in this field is the evaluation of the constant $\psi(\beta)$ which is crucial for carrying out maximum likelihood and Bayesian inference. As far as we know, there is no feasible analytic method for approximating ψ when n is large. Physicists have tried the technique of mean-field approximations; see Park and Newman [45, 46] for the case where T_1 is the number of edges and T_2 is the number of two-stars or the number of triangles. Mean-field approximations have no rigorous foundation, however, and are known to be unreliable in related models such as spin glasses [55]. For exponential graph models, Chatterjee and Dey [13] prove that they work for some restricted ranges of $\{\beta_i\}$: values where the graphs are shown to be essentially Erdős–Rényi graphs (see Theorem 4.2 below and [6]).

A host of techniques for approximating the normalizing constant using various Monte Carlo schemes have been proposed. These include the MCMLE procedure of Geyer and Thompson [29]. The bridge sampling approach of Gelman and Meng [28] also builds on techniques suggested by physicists to estimate free energy [$\psi(\beta)$ in our context]. The equi-energy sampler of Kou et al. [36] can also be harnessed to estimate ψ.

Alas, at present writing these procedures seem useful only for relatively small graphs. For bigger graphs, the run-time of the Monte Carlo algorithms become unpleasantly long. Snijders [52] and Handcock [31] demonstrate this empirically with further discussion in [53]. One theoretical explanation for the poor performance of these techniques comes from the work of Bhamidi et al. [6]. Most of the algorithms above require a sample from the model (2.1). This is most often done by using a local Markov chain based on adding or deleting edges via Metropolis or Glauber dynamics (Gibbs sampling). These authors show that if the parameters are nonnegative, then for large n,

- either the p_β model is essentially the same as an Erdős–Rényi model (in which case the Markov chain mixes in $n^2 \log n$ steps);
- or the Markov chain takes exponential time to mix.

Thus, in cases where the model is not essentially trivial, the Markov chains required to carry MCMLE procedures cannot be usefully run to stationarity.

Two other approaches to estimation are worth mentioning. The pseudo-likelihood approach of Besag [5] is widely used because of its ease of implementation. Its properties are at best poorly understood: it does not directly maximize the likelihood and in empirical comparisions (see, e.g., [17]), has appreciably larger variability than the MLE. Comets and Janžura [16] prove consistency and asymptotic normality of the maximum pseudo-likelihood estimator in certain Markov random field models. Chatterjee [12] shows that it is consistent for estimating the temperature parameter of the Sherrington–Kirkpatrick model of spin glasses. The second approach is Snijders' [52] suggestion to use the Robbins–Monro optimization procedure [50] to compute solutions to the moment equations $E_\beta(T(G)) = T(G^*)$ where G^* is the observed graph. While promising, the approach requires generating points from p_β for arbitrary β. The only way to do this at present is by MCMC and the results of [6] suggest this may be impractical.

2.2. *Graph limits*. In a sequence of papers [9–11, 25, 37–43], Laszlo Lovász and coauthors V. T. Sós, B. Szegedy, C. Borgs, J. Chayes, K. Vesztergombi, A. Schrijver, and M. Freedman have developed a beautiful, unifying theory of graph limits. (See also the related work of Austin [3] and Diaconis and Janson [18] which traces this back to work of Aldous [1], Hoover [33] and Kallenberg [35].) This body of work sheds light on various graph-theoretic topics such as graph homomorphisms, Szemerédi's regularity lemma, quasi-random graphs, graph testing and extremal graph theory, and has even found applications in statistics and related areas (see, e.g., [14]). Their theory has been developed for dense graphs (number of edges comparable to the square of number of vertices) but parallel theories for sparse graphs are beginning to emerge [7].

Lovász and coauthors define the limit of a sequence of dense graphs as follows. We quote the definition verbatim from [40] (see also [10, 11, 18]). Let G_n be a sequence of simple graphs whose number of nodes tends to infinity. For every fixed simple graph H, let $|\hom(H, G)|$ denote the number of homomorphisms of H into G [i.e., edge-preserving maps $V(H) \to V(G)$, where $V(H)$ and $V(G)$ are the vertex sets]. This number is normalized to get the homomorphism density

$$t(H, G) := \frac{|\hom(H, G)|}{|V(G)|^{|V(H)|}}. \tag{2.2}$$

This gives the probability that a random mapping $V(H) \to V(G)$ is a homomorphism.

Note that $|\hom(H, G)|$ is not the count of the number of copies of H in G, but is a constant multiple of that if H is a complete graph. For example, if H is a triangle, $|\hom(H, G)|$ is the number of triangles in G multiplied by six. On the other hand if H is, say, a 2-star (i.e., a triangle with one edge missing) and G is a triangle, then the number of copies of H in G is zero, while $|\hom(H, G)| = 3^3 = 27$.

Suppose that the graphs G_n become more and more similar in the sense that $t(H, G_n)$ tends to a limit $t(H)$ for every H. One way to define a limit of the

sequence $\{G_n\}$ is to define an appropriate limit object from which the values $t(H)$ can be read off.

The main result of [40] (following the earlier equivalent work of Aldous [1] and Hoover [33]) is that indeed there is a natural "limit object" in the form of a function $h \in \mathcal{W}$, where $\mathcal{W}$ is the space of all measurable functions from $[0,1]^2$ into $[0,1]$ that satisfy $h(x,y) = h(y,x)$ for all x, y.

Conversely, every such function arises as the limit of an appropriate graph sequence. This limit object determines all the limits of subgraph densities: if H is a simple graph with $V(H) = [k] = \{1, \ldots, k\}$, let

$$t(H,h) = \int_{[0,1]^k} \prod_{\{i,j\} \in E(H)} h(x_i, x_j)\, dx_1 \cdots dx_k. \tag{2.3}$$

Here $E(H)$ denotes the edge set of H. A sequence of graphs $\{G_n\}_{n \geq 1}$ is said to converge to h if for every finite simple graph H,

$$\lim_{n \to \infty} t(H, G_n) = t(H,h). \tag{2.4}$$

Intuitively, the interval $[0,1]$ represents a "continuum" of vertices, and $h(x,y)$ denotes the probability of putting an edge between x and y. For example, for the Erdős–Rényi graph $G(n,p)$, if p is fixed and $n \to \infty$, then the limit graph is represented by the function that is identically equal to p on $[0,1]^2$. Clearly, this framework is therefore useful only when p does not tend to zero when $n \to \infty$, that is, in the case of dense Erdős–Rényi graphs.

These limit objects, that is, elements of $\mathcal{W}$, are called "graph limits" or "graphons" in [10, 11, 40]. A finite simple graph G on $\{1, \ldots, n\}$ can also be represented as a graph limit f^G is a natural way, by defining

$$f^G(x,y) = \begin{cases} 1, & \text{if } (\lceil nx \rceil, \lceil ny \rceil) \text{ is an edge in } G, \\ 0, & \text{otherwise.} \end{cases} \tag{2.5}$$

The definition makes sense because $t(H, f^G) = t(H, G)$ for every simple graph H and therefore the constant sequence $\{G, G, \ldots\}$ converges to the graph limit f^G. Note that this allows *all* simple graphs, irrespective of the number of vertices, to be represented as elements of a single abstract space, namely $\mathcal{W}$.

With the above representation, it turns out that the notion of convergence in terms of subgraph densities outlined above can be captured by an explicit metric on $\mathcal{W}$, the so-called *cut distance* (originally defined for finite graphs by Frieze and Kannan [26]). Start with the space $\mathcal{W}$ of measurable functions $f(x,y)$ on $[0,1]^2$ that satisfy $0 \leq f(x,y) \leq 1$ and $f(x,y) = f(y,x)$. Define the cut distance

$$d_\square(f,g) := \sup_{S,T \subseteq [0,1]} \left| \int_{S \times T} \big[f(x,y) - g(x,y) \big]\, dx\, dy \right|. \tag{2.6}$$

Introduce in $\mathcal{W}$ an equivalence relation: let Σ be the space of measure preserving bijections $\sigma : [0,1] \to [0,1]$. Say that $f(x,y) \sim g(x,y)$ if $f(x,y) = g_\sigma(x,y) :=$

$g(\sigma x, \sigma y)$ for some $\sigma \in \Sigma$. Denote by $\widetilde{g}$ the closure in $(\mathcal{W}, d_\square)$ of the orbit $\{g_\sigma\}$. The quotient space is denoted by $\widetilde{\mathcal{W}}$ and τ denotes the natural map $g \to \widetilde{g}$. Since $d_\square$ is invariant under σ one can define on $\widetilde{\mathcal{W}}$, the natural distance $\delta_\square$ by

$$\delta_\square(\widetilde{f}, \widetilde{g}) := \inf_\sigma d_\square(f, g_\sigma) = \inf_\sigma d_\square(f_\sigma, g) = \inf_{\sigma_1, \sigma_2} d_\square(f_{\sigma_1}, g_{\sigma_2})$$

making $(\widetilde{\mathcal{W}}, \delta_\square)$ into a metric space. To any finite graph G, we associate f^G as in (2.5) and its orbit $\widetilde{G} = \tau f^G = \widetilde{f}^G \in \widetilde{\mathcal{W}}$.

The papers by Lovász and coauthors establish many important properties of the metric space $\widetilde{\mathcal{W}}$ and the associated notion of graph limits. For example, $\widetilde{\mathcal{W}}$ is compact. A pressing objective is to understand what functions from $\widetilde{\mathcal{W}}$ into $\mathbb{R}$ are continuous. Fortunately, it is an easy fact that the homomorphism density $t(H, \cdot)$ is continuous for any finite simple graph H [10, 11]. There are other, more complicated functions that are continuous. For example, the degree distribution is continuous with respect to this topology, as is the distribution of eigenvalues. See [3, 4] for further discussions.

2.3. *Large deviations for random graphs.* Let $G(n, p)$ be the random graph on n vertices where each edge is added independently with probability p. This model has been the subject of extensive investigations since the pioneering work of Erdős and Rényi [19], yielding a large body of literature (see [8, 34] for partial surveys).

Recently, Chatterjee and Varadhan [15] formulated a large deviation principle for the Erdős–Rényi graph, in the same way as Sanov's theorem [51] gives a large deviation principle for an i.i.d. sample. The formulation and proof of this result makes extensive use of the properties of the topology described in Section 2.2.

Let $I_p : [0, 1] \to \mathbb{R}$ be the function

$$I_p(u) := \frac{1}{2} u \log \frac{u}{p} + \frac{1}{2}(1-u) \log \frac{1-u}{1-p}. \tag{2.7}$$

The domain of the function I_p can be extended to $\mathcal{W}$ as

$$I_p(h) := \int_0^1 \int_0^1 I_p\big(h(x, y)\big)\, dx\, dy. \tag{2.8}$$

The function I_p can be defined on $\widetilde{\mathcal{W}}$ by declaring $I_p(\widetilde{h}) := I_p(h)$ where h is any representative element of the equivalence class $\widetilde{h}$. Of course, this raises the question whether I_p is well defined on $\widetilde{\mathcal{W}}$. It was proved in [15] that the function I_p is indeed well defined on $\widetilde{\mathcal{W}}$ and is lower semicontinuous under the cut metric $\delta_\square$.

The random graph $G(n, p)$ induces probability distributions $\mathbb{P}_{n,p}$ on the space $\mathcal{W}$ through the map $G \to f^G$ and $\widetilde{\mathbb{P}}_{n,p}$ on $\widetilde{\mathcal{W}}$ through the map $G \to f^G \to \widetilde{f}^G = \widetilde{G}$. The large deviation principle for $\widetilde{\mathbb{P}}_{n,p}$ on $(\widetilde{\mathcal{W}}, \delta_\square)$ is the main result of [15].

THEOREM 2.1 (Chatterjee and Varadhan [15]). *For each fixed $p \in (0,1)$, the sequence $\widetilde{\mathbb{P}}_{n,p}$ obeys a large deviation principle in the space $\widetilde{\mathcal{W}}$ (equipped with the cut metric) with rate function I_p defined by (2.8). Explicitly, this means that for any closed set $\widetilde{F} \subseteq \widetilde{\mathcal{W}}$,*

$$\limsup_{n\to\infty} \frac{1}{n^2} \log \widetilde{\mathbb{P}}_{n,p}(\widetilde{F}) \le -\inf_{\widetilde{h}\in\widetilde{F}} I_p(\widetilde{h}) \tag{2.9}$$

and for any open set $\widetilde{U} \subseteq \widetilde{\mathcal{W}}$,

$$\liminf_{n\to\infty} \frac{1}{n^2} \log \widetilde{\mathbb{P}}_{n,p}(\widetilde{U}) \ge -\inf_{\widetilde{h}\in\widetilde{U}} I_p(\widetilde{h}). \tag{2.10}$$

3. The main result. Let $T : \widetilde{\mathcal{W}} \to \mathbb{R}$ be a bounded continuous function on the metric space $(\widetilde{\mathcal{W}}, \delta_\square)$. Fix n and let $\mathcal{G}_n$ denote the set of simple graphs on n vertices. Then T induces a probability mass function p_n on $\mathcal{G}_n$ defined as

$$p_n(G) := e^{n^2(T(\widetilde{G})-\psi_n)}.$$

Here $\widetilde{G}$ is the image of G in the quotient space $\widetilde{\mathcal{W}}$ as defined in Section 2.2 and ψ_n is a constant such that the total mass of p_n is 1. Explicitly,

$$\psi_n = \frac{1}{n^2} \log \sum_{G\in\mathcal{G}_n} e^{n^2 T(\widetilde{G})}. \tag{3.1}$$

The coefficient n^2 is meant to ensure that ψ_n tends to a nontrivial limit as $n \to \infty$. (Note that T does not vary with n.) To describe this limit, define a function $I : [0,1] \to \mathbb{R}$ as

$$I(u) := \tfrac{1}{2} u \log u + \tfrac{1}{2}(1-u)\log(1-u)$$

and extend I to $\widetilde{\mathcal{W}}$ in the usual manner:

$$I(\widetilde{h}) = \int\!\!\int_{[0,1]^2} I\big(h(x,y)\big)\,dx\,dy, \tag{3.2}$$

where h is a representative element of the equivalence class $\widetilde{h}$. As mentioned before, it follows from a result of [15] that I is well defined and lower semicontinuous on $\widetilde{\mathcal{W}}$. The following theorem is the first main result of this paper.

THEOREM 3.1. *If $T : \widetilde{\mathcal{W}} \to \mathbb{R}$ is a bounded continuous function and ψ_n and I are defined as above, then*

$$\psi := \lim_{n\to\infty} \psi_n = \sup_{\widetilde{h}\in\widetilde{\mathcal{W}}} \big(T(\widetilde{h}) - I(\widetilde{h})\big).$$

We will see later that the supremum on the right-hand side is actually a maximum, that is, there is some $\widetilde{h}$ where the supremum is attained. This is significant because such maximizing $\widetilde{h}$'s describe the structure of the random graph in the large n limit.

As mentioned in the Introduction, evaluation of the normalizing constant is one of the key problems in statistical applications of exponential random graphs. Incidentally, even the existence of the limit in Theorem 3.1 has an important consequence. Suppose that a computer program can evaluate the exact value of the normalizing constant for moderate sized n. Then if n is large, one can choose a "scaled down" model with a smaller number of nodes, and use the exact value of the normalizing constant in the scaled down model as an approximation to the normalizing constant in the larger model.

Theorem 3.1 gives an asymptotic formula for ψ_n. However, it says nothing about the behavior of a random graph drawn from the exponential random graph model. Some aspects of this behavior can be described as follows. Let $\widetilde{F}^*$ be the subset of $\widetilde{\mathcal{W}}$ where $T(\widetilde{h}) - I(\widetilde{h})$ is maximized. By the compactness of $\widetilde{\mathcal{W}}$, the continuity of T and the lower semi-continuity of I, $\widetilde{F}^*$ is a nonempty compact set. Let G_n be a random graph on n vertices drawn from the exponential random graph model defined by T. The following theorem shows that for n large, $\widetilde{G}_n$ must lie close to $\widetilde{F}^*$ with high probability. In particular, if $\widetilde{F}^*$ is a singleton set, then the theorem gives a weak law of large numbers for G_n.

THEOREM 3.2. *Let $\widetilde{F}^*$ and G_n be defined as in the above paragraph. Let $\mathbb{P}$ denote the probability measure on the underlying probability space on which G_n is defined. Then for any $\eta > 0$ there exist $C, \gamma > 0$ such that for any n,*

$$\mathbb{P}\big(\delta_\square\big(\widetilde{G}_n, \widetilde{F}^*\big) > \eta\big) \le Ce^{-n^2\gamma}.$$

4. An application. Let $H_1, \ldots, H_k$ be finite simple graphs, where H_1 is the complete graph on two vertices (i.e., just a single edge), and each H_i contains at least one edge. Let $\beta_1, \ldots, \beta_k$ be k real numbers. For any $h \in \mathcal{W}$, let

$$T(h) := \sum_{i=1}^{k} \beta_i t(H_i, h), \tag{4.1}$$

where $t(H_i, h)$ is the homomorphism density of H_i in h, defined in (2.3). Note that there is nothing special about taking H_1 to be a single edge; if we do not want H_1 in our sufficient statistic, we just take $\beta_1 = 0$; all theorems would remain valid.

As remarked in Section 2.2, T is continuous with respect to the cut distance on $\mathcal{W}$, and hence admits a natural definition on $\widetilde{\mathcal{W}}$. Note that for any finite simple graph G that has at least as many nodes as the largest of the H_i's,

$$T(\widetilde{G}) = \sum_{i=1}^{k} \beta_i t(H_i, G).$$

For example, if $k = 2$, and H_2 is a triangle, and G has at least 3 nodes, then

$$T(\widetilde{G}) = \frac{2\beta_1(\#\text{edges in } G)}{n^2} + \frac{6\beta_2(\#\text{triangles in } G)}{n^3}.$$

Let ψ_n be as in (3.1), and let G_n be the n-vertex exponential random graph with sufficient statistic T. Theorem 3.1 gives a formula for $\lim_{n\to\infty} \psi_n$ as the solution of a variational problem. Surprisingly the variational problem is explicitly solvable if $\beta_2, \ldots, \beta_k$ are nonnegative.

THEOREM 4.1. *Let T, ψ_n and $H_1, \ldots, H_k$ be as above. Suppose $\beta_2, \ldots, \beta_k$ are nonnegative. Then*

$$\lim_{n\to\infty} \psi_n = \sup_{0\le u\le 1} \left(\sum_{i=1}^{k} \beta_i u^{e(H_i)} - I(u) \right), \tag{4.2}$$

where $I(u) = \frac{1}{2}u\log u + \frac{1}{2}(1-u)\log(1-u)$ and $e(H_i)$ is the number of edges in H_i. Moreover, each solution of the variational problem of Theorem 3.1 for this T is a constant function, where the constant solves the scalar maximization problem (4.2).

Theorem 4.1 gives the limiting value of ψ_n if $\beta_2, \ldots, \beta_k$ are nonnegative. The next theorem describes the behavior of the exponential random graph G_n under this condition if n is large.

THEOREM 4.2. *For each n, let G_n be an n-vertex exponential random graph with sufficient statistic T defined in* (4.1). *Assume that $\beta_2, \ldots, \beta_k$ are nonnegative. Then*:

(a) *If the maximization problem in* (4.2) *is solved at a unique value u^*, then G_n is indistinguishable from the Erdős–Rényi graph $G(n, u^*)$ in the large n limit, in the sense that $\widetilde{G}_n$ converges to the constant function u^* in probability as $n \to \infty$.*

(b) *Even if the maximizer is not unique, the set U of maximizers is a finite subset of* $[0, 1]$ *and*

$$\min_{u\in U} \delta_\square(\widetilde{G}_n, \widetilde{u}) \to 0 \qquad \text{in probability as } n \to \infty,$$

where $\widetilde{u}$ denotes the image of the constant function u in $\widetilde{\mathcal{W}}$. In other words, G_n behaves like an Erdős–Rényi graph $G(n, u)$ where u is picked randomly from some probability distribution on U.

It may be noted here that the conclusion of Theorem 4.2 was proved earlier by Bhamidi et al. [6] under certain restrictions on the parameters that they called a "high temperature condition." This is in analogy with spin systems, since random graphs may be interpreted as systems of particles (corresponding to edges) each

having spin 0 or 1 (i.e., closed or open). With this interpretation, it is straightforward to check that when $\beta_2, \ldots, \beta_k$ are nonnegative, the model defined above satisfies the so-called FKG property [23]. Stated simply, the FKG property means that if f and g are monotone functions of the random graph (i.e., functions whose values cannot decrease if more edges are added to the graph), then f and g are positively correlated random variables. The FKG property has important consequences; for instance, it implies that the expected value of $t(H_i, G)$ is an increasing function of β_j for any i and j. We will see some further consequences of the FKG property in our proof of Theorem 5.1 in the next section.

5. Phase transitions and near-degeneracy. To illustrate the results of the previous section, recall the exponential random graph model (1.2) with edges and triangles as sufficient statistics:

$$\begin{aligned} T(\widetilde{G}) &= 2\beta_1 \frac{\#\text{edges in } G}{n^2} + 6\beta_2 \frac{\#\text{triangles in } G}{n^3} \\ &= \beta_1 t(H_1, G) + \beta_2 t(H_2, G), \end{aligned} \tag{5.1}$$

where H_1 is a single edge and H_2 is a triangle. Fix β_1 and β_2 and let

$$\ell(u) := \beta_1 u + \beta_2 u^3 - I(u), \tag{5.2}$$

where $I(u) = \frac{1}{2}u \log u + \frac{1}{2}(1-u)\log(1-u)$, as usual. Let U be the set of maximizers of $\ell(u)$ in $[0, 1]$. Theorem 4.2 describes the limiting behavior of G_n in terms of the set U. In particular, if U consists of a single point $u^* = u^*(\beta_1, \beta_2)$, then G_n behaves like the Erdős–Rényi graph $G(n, u^*)$ when n is large.

It is likely that $u^*(\beta_1, \beta_2)$ does not have a closed form expression, other than when $\beta_2 = 0$, in which case

$$u^*(\beta_1, 0) = \frac{e^{2\beta_1}}{1 + e^{2\beta_1}}.$$

It is, however, quite easy to numerically approximate $u^*(\beta_1, \beta_2)$. Figure 3 plots $u^*(\beta_1, \beta_2)$ versus β_2 for four different fixed values of β_1, namely, $\beta_1 = 0.2$, -0.35, -0.45, and -0.8. The figures show that u^* is a continuous function of β_2 as long as β_1 is not too far down the negative axis.

But for β_1 below a threshold (e.g., when $\beta_1 = -0.45$), u^* shows a single jump discontinuity in β_2, signifying a phase transition. In physical terms, this is a first order phase transition, by the following logic. By Theorem 4.2, our random graph behaves like $G(n, u^*)$ when n is large. On the other hand, by a standard computation the expect number of triangles is the first derivative of the free energy ψ_n with respect to β_2. Therefore in the large n limit, a discontinuity in u^* as a function of β_2 signifies a discontinuity in the derivative of the limiting free energy, which is the physical definition of a first order phase transition.

At the point of discontinuity, $\ell(u)$ is maximized at two values of u, that is, the set U consists of two points. Lastly, as β_1 goes down the negative axis, the model starts

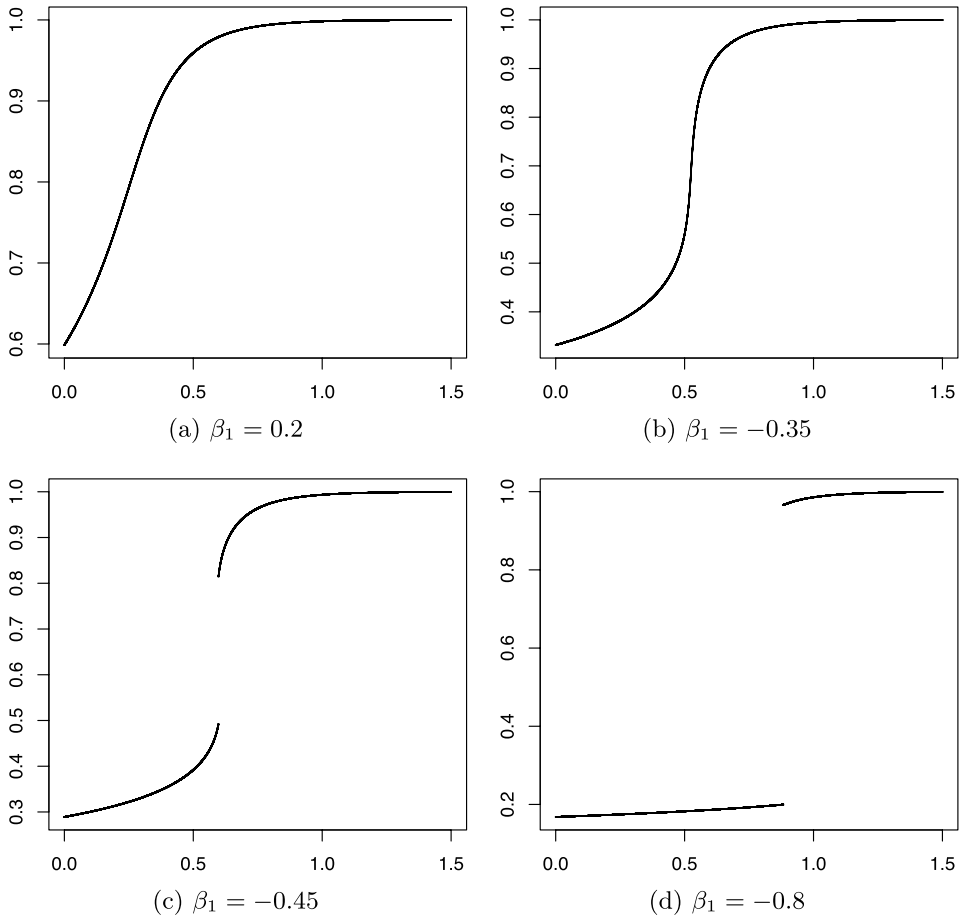

FIG. 3. *Plot of $u^*(\beta_1, \beta_2)$ on y-axis vs β_2 on x-axis for different fixed values of β_1. Part* (c) *demonstrates a phase transition. Part* (d) *demonstrates near-degeneracy.*

to exhibit "near-degeneracy" in the sense of Handcock [31] (see also [45]) as seen in the last frame of Figure 3. This means that when β_1 is a large negative number, then as β_2 varies, the model transitions from being a very sparse graph for low values of β_2, to a very dense graph for large values of β_2, completely skipping all intermediate structures. If this sentence is confusing, please see Theorem 5.1 below for a precise statement. This theorem gives a simple mathematical description of this phenomenon and hence the first rigorous proof of the degeneracy observed in exponential graph models. Related results are in Häggstrom and Jonasson [30].

THEOREM 5.1. *Let G_n be an exponential random graph with sufficient statistic T defined in* (5.1) *and let $\mathbb{P}$ be the probability measure on the underlying prob-*

ability space on which G_n is defined. Fix any $\beta_1 < 0$. Let

$$c_1 := \frac{e^{\beta_1}}{1+e^{\beta_1}}, \qquad c_2 := 1 + \frac{1}{2\beta_1}.$$

Suppose $|\beta_1|$ is so large that $c_1 < c_2$. Let $e(G_n)$ be the number of edges in G_n and let $f(G_n) := e(G_n)/\binom{n}{2}$ be the edge density. [*Note that $f(G_n) = \frac{n}{n-1} t(H_1, G_n)$, where H_1 is a single edge.*]

Then there exists $q = q(\beta_1) \in [0, \infty)$ such that if $-\infty < \beta_2 < q$, then

$$\lim_{n\to\infty} \mathbb{P}\big(f(G_n) > c_1\big) = 0,$$

and if $\beta_2 > q$, then

$$\lim_{n\to\infty} \mathbb{P}\big(f(G_n) < c_2\big) = 0.$$

In other words, if β_1 is a large negative number, then G_n is either sparse (if $\beta_2 < q$) or nearly complete (if $\beta_2 > q$).

The difference in the values of c_1 and c_2 can be quite striking even for relatively small values of β_1. For example, $\beta_1 = -5$ gives $c_1 \simeq 0.007$ and $c_2 = 0.9$. Significant extensions of Theorem 5.1 have been made in the recent manuscripts [2, 47, 48, 58].

6. The symmetric phase, symmetry breaking, and the Euler–Lagrange equations. The purpose of this section is to extend the analysis of the model from Section 4 beyond the case of nonnegative parameters. We begin with a standard approach to solving variational problems.

6.1. *Euler–Lagrange equations.* We return to the exponential random graph model with sufficient statistic T defined in (4.1) in terms of the densities of k fixed graphs $H_1, \dots, H_k$, where H_1 is a single edge. Theorems 4.1 and 4.2 analyze this model when $\beta_2, \dots, \beta_k$ are nonnegative. What if they are not? One can still try to derive the Euler–Lagrange equations (or Euler's equation; see [27]) for the related variational problem of maximizing $T(\widetilde{h}) - I(\widetilde{h})$. The following theorem presents the outcome of this effort.

For a finite simple graph H, let $V(H)$ and $E(H)$ denote the sets of vertices and edges of H. Given a symmetric measurable function $h : [0,1]^2 \to \mathbb{R}$, for each $(r,s) \in E(H)$ and each pair of points $x_r, x_s \in [0,1]$, define

$$\Delta_{H,r,s} h(x_r, x_s) := \int_{[0,1]^{|V(H)\setminus\{r,s\}|}} \prod_{\substack{(r',s')\in E(H)\\(r',s')\neq(r,s)}} h(x_{r'}, x_{s'}) \prod_{\substack{v\in V(H)\\v\neq r,s}} dx_v.$$

For $x, y \in [0,1]$ define

$$\Delta_H h(x,y) := \sum_{(r,s)\in E(H)} \Delta_{H,r,s} h(x,y). \tag{6.1}$$

For example, when H is a triangle, then $V(H) = \{1, 2, 3\}$ and

$$\Delta_{H,1,2}h(x, y) = \Delta_{H,1,3}h(x, y) = \Delta_{H,2,3}h(x, y) = \int_0^1 h(x, z)h(y, z)\,dz$$

and therefore $\Delta_H h(x, y) = 3\int_0^1 h(x, z)h(y, z)\,dz$. When H contains exactly one edge, define $\Delta_H h \equiv 1$ for any h, by the usual convention that the empty product is 1. The following theorem gives the Euler–Lagrange equations for the optimizer h of Theorem 3.1 in terms of these $\Delta_H h$'s.

THEOREM 6.1. *Let $T : \widetilde{\mathcal{W}} \to \mathbb{R}$ be defined as in (4.1) and the operator Δ_H be defined as in (6.1). If $\widetilde{h} \in \widetilde{\mathcal{W}}$ maximizes $T(\widetilde{h}) - I(\widetilde{h})$, then any representative element $h \in \widetilde{h}$ must satisfy for almost all $(x, y) \in [0, 1]^2$,*

$$h(x, y) = \frac{e^{2\sum_{i=1}^k \beta_i \Delta_{H_i} h(x,y)}}{1 + e^{2\sum_{i=1}^k \beta_i \Delta_{H_i} h(x,y)}}.$$

Moreover, any maximizing function must be bounded away from 0 and 1.

Unfortunately, these equations may have many solutions and therefore do not uniquely identify the optimizer. The next subsection gives a sufficient condition under which the solution in unique.

6.2. *The replica symmetric phase.* Borrowing terminology from spin glasses, we define the *replica symmetric phase* or simply the *symmetric phase* of a variational problem like maximizing $T(h) - I(h)$ as the set of parameter values for which all the maximizers are constant functions. When the parameters are such that all maximizers are nonconstant functions we say that the parameter vector is in the region of broken replica symmetry, or simply broken symmetry. There may be another situation, where some optimizers are constant while others are nonconstant, although we do not know of such examples. (This third region may be called a region of partial symmetry.)

Statistically, the exponential random graph behaves like an Erdős–Rényi graph in the symmetric region of the parameter space, while such behavior breaks down in the region of broken symmetry. This follows easily from Theorem 3.2.

Theorem 4.2 shows that for the sufficient statistic T defined in (4.1), each $(\beta_1, \beta_2, \ldots, \beta_k)$ in $\mathbb{R} \times \mathbb{R}_+^{k-1}$ falls in the replica symmetric region. Does symmetry hold only when $\beta_2, \ldots, \beta_k$ are nonnegative? The following theorem (proven with the aid of the Euler–Lagrange equations of Theorem 6.1), shows that this is not the case; $(\beta_1, \ldots, \beta_k)$ is in the replica symmetric region whenever $|\beta_2|, \ldots, |\beta_k|$ are small enough. Of course, this does not supersede Theorem 4.2 since it does not cover large positive values of $\beta_2, \ldots, \beta_k$. However, it proves replica symmetry for small negative values of $\beta_2, \ldots, \beta_k$, which is not covered by Theorem 4.2.

THEOREM 6.2. *Consider the exponential random graph with sufficient statistic T defined in (4.1). Suppose $\beta_1, \dots, \beta_k$ are such that*

$$\sum_{i=2}^{k} |\beta_i| e(H_i)\big(e(H_i) - 1\big) < 2,$$

where $e(H_i)$ is the number of edges in H_i. Then the conclusions of Theorems 4.1 and 4.2 hold true for this value of the parameter vector $(\beta_1, \dots, \beta_k)$.

6.3. *Symmetry breaking.* Theorems 4.2 and 6.2 establish various regions of symmetry in the exponential random graph model with sufficient statistic T defined in (4.1). That leaves the question: is there a region where symmetry breaks? We specialize to the simple case where $k = 2$ and H_2 is a triangle, that is, the example of Section 5. In this case, it turns out that replica symmetry breaks whenever β_2 is less than a sufficiently large negative number depending on β_1.

THEOREM 6.3. *Consider the exponential random graph with sufficient statistic T defined in (5.1). Then for any given value of β_1, there is a positive constant $C(\beta_1)$ sufficiently large so that whenever $\beta_2 < -C(\beta_1)$, $T(h) - I(h)$ is not maximized at any constant function. Consequently, if G_n is an n-vertex exponential random graph with this sufficient statistic, then there exists $\varepsilon > 0$ such that*

$$\lim_{n\to\infty} \mathbb{P}\big(\delta_\square(\widetilde{G}_n, \widetilde{C}) > \varepsilon\big) = 1,$$

where $\widetilde{C}$ is the set of constant functions. In other words, G_n does not look like an Erdős–Rényi graph in the large n limit.

For interesting recent developments about symmetry breaking in exponential random graph models, see Lubetzky and Zhao [44].

6.4. *A completely solvable case.* A j-star is an undirected graph with one "root" vertex and j other vertices connected to the root vertex, with no edges between any of these j vertices. Let H_j be a j-star for $j = 1, \dots, k$. Let T be the sufficient statistic

$$T(G) = \sum_{j=1}^{k} \beta_j t(H_j, G). \tag{6.2}$$

Theorems 4.1 and 4.2 describe the behavior of this model when $\beta_2, \dots, \beta_k$ are all nonnegative. The following theorem completely solves this model for all values of $\beta_2, \dots, \beta_k$. The proof of this theorem was suggested by the anonymous referee, improving upon the version of the result given in an earlier draft.

THEOREM 6.4. *For the sufficient statistic T defined in (6.2), the conclusions of Theorems 4.1 and 4.2 hold for any $\beta_1, \dots, \beta_k \in \mathbb{R}$.*

7. Extremal behavior. In the sections above, we have been assuming that $\beta_2, \ldots, \beta_k$ are positive or barely negative. In this section, we investigate what happens when $k = 2$ and β_2 is large and negative. The limits are describable but far from Erdős–Rényi. Our work here is inspired by related results of Sukhada Fadvanis who has a different argument (using Turán's theorem [56]) for the case of triangles.

Suppose H is any finite simple graph containing at least one edge. Let T be the sufficient statistic

$$T(\widetilde{G}) = 2\beta_1 \frac{\#\text{edges in } G}{n^2} + \beta_2 t(H, G).$$

Let G_n be the exponential random graph on n vertices with this sufficient statistic and let ψ_n be the associated normalizing constant as defined in (3.1). Then Theorem 3.1 gives

$$\lim_{n\to\infty} \psi_n = \sup_{h\in\mathcal{W}} \big(T(h) - I(h)\big) =: \psi,$$

where I is defined in (3.2). We also know (by Theorem 3.2) that

$$\delta_\square(\widetilde{G}_n, \widetilde{F}^*) \to 0 \qquad \text{in probability as } n \to \infty,$$

where $\widetilde{F}^*$ is the subset of $\widetilde{\mathcal{W}}$ where $T - I$ is maximized. (Note that $\widetilde{F}^*$ is a closed set since $T - I$ is an upper semicontinuous map.)

We can compute $\widetilde{F}^*$ and ψ when β_2 is positive, or negative with small magnitude. We are unable to carry out the explicit computation in the case of large negative β_2, unless H is a convenient object like a j-star. However, a qualitative description can still be given by analyzing the behavior of $\widetilde{F}^*$ and ψ as $\beta_2 \to -\infty$. Fixing β_1, we consider these objects as functions of β_2 and write $\widetilde{F}^*(\beta_2)$, $\psi(\beta_2)$ and T_{β_2} instead of $\widetilde{F}^*$, ψ and T. Recall that the chromatic number of a graph is the minimum number of colors required to color the edges so that no two neighbors get the same color.

THEOREM 7.1. *Fixing H and β_1, let $\widetilde{F}^*(\beta_2)$ and $\psi(\beta_2)$ be as above. Let $\chi(H)$ be the chromatic number of H, and define*

$$g(x, y) := \begin{cases} 1, & \text{if } \big[(\chi(H) - 1)x\big] \neq \big[(\chi(H) - 1)y\big], \\ 0, & \text{otherwise}, \end{cases} \tag{7.1}$$

where $[x]$ denotes the integer part of a real number x. Let $p = e^{2\beta_1}/(1 + e^{2\beta_1})$. Then

$$\lim_{\beta_2\to-\infty} \sup_{\widetilde{f}\in\widetilde{F}^*(\beta_2)} \delta_\square(\widetilde{f}, p\widetilde{g}) = 0$$

and

$$\lim_{\beta_2\to-\infty} \psi(\beta_2) = \frac{(\chi(H) - 2)}{2(\chi(H) - 1)} \log \frac{1}{1 - p}.$$

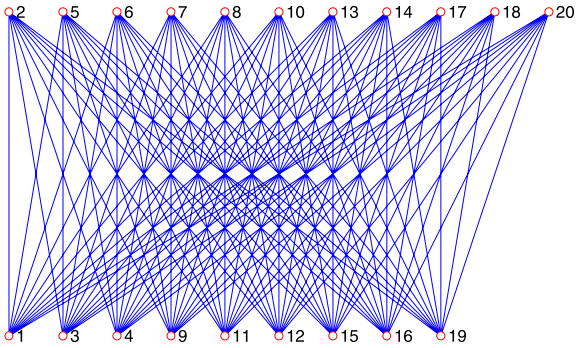

FIG. 4. *A simulated realization of the exponential random graph model on* 20 *nodes with edges and triangles as sufficient statistics, where* $\beta_1 = 120$ *and* $\beta_2 = -400$. (*Picture by Sukhada Fadnavis. Gibbs sampling used in simulations. We are unsure about the convergence of the MCMC algorithm used to generate these grapgs, but the structure of the simulated graphs match the predictions of Theorem* 7.1.)

Intuitively, the above result means that if β_2 is a large negative number and n is large, then an exponential random graph G_n with sufficient statistic T looks roughly like a complete $(\chi(H)-1)$-equipartite graph with $1-p$ fraction of edges randomly deleted, where $p = e^{2\beta_1}/(1+e^{2\beta_1})$. In particular, if H is bipartite, then G_n must be very sparse, since a 1-equipartite graph has no edges. Figure 4 gives a simulation result for the triangle model with large negative β_2.

Theorem 7.1 is closely related to the Erdős–Stone theorem [20] from extremal graph theory (or equivalently, Turán's theorem in the case of triangles as in the work of Fadnavis). Indeed, it may be possible to prove some parts of our theorem using the Erdős–Stone theorem, but we prefer a bare-hands argument given in Section 8. Due to this connection with extremal graph theory, we refer to behavior of the graph in the "large negative β_2" domain as *extremal behavior*.

8. Proofs.

PROOF OF THEOREM 3.1. For each Borel set $\widetilde{A} \subseteq \widetilde{\mathcal{W}}$ and each n, define

$$\widetilde{A}_n := \{\widetilde{h} \in \widetilde{A} : \widetilde{h} = \widetilde{G} \text{ for some } G \in \mathcal{G}_n\}.$$

Let $\mathbb{P}_{n,p}$ be the Erdős–Rényi measure defined in Section 2.3. Note that $\widetilde{A}_n$ is a finite set and

$$|\widetilde{A}_n| = 2^{n(n-1)/2}\mathbb{P}_{n,1/2}(\widetilde{A}_n) = 2^{n(n-1)/2}\mathbb{P}_{n,1/2}(\widetilde{A}).$$

Thus, if $\widetilde{F}$ is a closed subset of $\widetilde{\mathcal{W}}$ then by Theorem 2.1

$$\begin{aligned}\limsup_{n\to\infty} \frac{\log|\widetilde{F}_n|}{n^2} &\le \frac{\log 2}{2} - \inf_{\widetilde{h}\in\widetilde{F}} I_{1/2}(\widetilde{h}) \\ &= -\inf_{\widetilde{h}\in\widetilde{F}} I(\widetilde{h}).\end{aligned} \tag{8.1}$$

Similarly if $\widetilde{U}$ is an open subset of $\widetilde{\mathcal{W}}$,

$$\liminf_{n\to\infty} \frac{\log|\widetilde{U}_n|}{n^2} \geq -\inf_{\widetilde{h}\in\widetilde{U}} I(\widetilde{h}). \tag{8.2}$$

Fix $\varepsilon > 0$. Since T is a bounded function, there is a finite set R such that the intervals $\{(a, a+\varepsilon) : a \in R\}$ cover the range of T. For each $a \in R$, let $\widetilde{F}^a := T^{-1}([a, a+\varepsilon])$. By the continuity of T, each $\widetilde{F}^a$ is closed. Now,

$$e^{n^2\psi_n} \leq \sum_{a\in R} e^{n^2(a+\varepsilon)}|\widetilde{F}_n^a| \leq |R| \sup_{a\in R} e^{n^2(a+\varepsilon)}|\widetilde{F}_n^a|.$$

By (8.1), this shows that

$$\limsup_{n\to\infty} \psi_n \leq \sup_{a\in R}\Big(a+\varepsilon - \inf_{\widetilde{h}\in\widetilde{F}^a} I(\widetilde{h})\Big).$$

Each $\widetilde{h} \in \widetilde{F}^a$ satisfies $T(\widetilde{h}) \geq a$. Consequently,

$$\sup_{\widetilde{h}\in\widetilde{F}^a}\big(T(\widetilde{h}) - I(\widetilde{h})\big) \geq \sup_{\widetilde{h}\in\widetilde{F}^a}\big(a - I(\widetilde{h})\big) = a - \inf_{\widetilde{h}\in\widetilde{F}^a} I(\widetilde{h}).$$

Substituting this in the earlier display gives

$$\begin{aligned}\limsup_{n\to\infty} \psi_n &\leq \varepsilon + \sup_{a\in R}\sup_{\widetilde{h}\in\widetilde{F}^a}\big(T(\widetilde{h}) - I(\widetilde{h})\big)\\ &= \varepsilon + \sup_{\widetilde{h}\in\widetilde{\mathcal{W}}}\big(T(\widetilde{h}) - I(\widetilde{h})\big).\end{aligned} \tag{8.3}$$

For each $a \in R$, let $\widetilde{U}^a := T^{-1}((a, a+\varepsilon))$. By the continuity of T, $\widetilde{U}^a$ is an open set. Note that

$$e^{n^2\psi_n} \geq \sup_{a\in R} e^{n^2 a}|\widetilde{U}_n^a|.$$

Therefore by (8.2), for each $a \in R$

$$\liminf_{n\to\infty} \psi_n \geq a - \inf_{\widetilde{h}\in\widetilde{U}^a} I(\widetilde{h}).$$

Each $\widetilde{h} \in \widetilde{U}^a$ satisfies $T(\widetilde{h}) < a + \varepsilon$. Therefore,

$$\sup_{\widetilde{h}\in\widetilde{U}^a}\big(T(\widetilde{h}) - I(\widetilde{h})\big) \leq \sup_{\widetilde{h}\in\widetilde{U}^a}\big(a+\varepsilon - I(\widetilde{h})\big) = a+\varepsilon - \inf_{\widetilde{h}\in\widetilde{U}^a} I(\widetilde{h}).$$

Together with the previous display, this shows that

$$\begin{aligned}\liminf_{n\to\infty} \psi_n &\geq -\varepsilon + \sup_{a\in R}\sup_{\widetilde{h}\in\widetilde{U}^a}\big(T(\widetilde{h}) - I(\widetilde{h})\big)\\ &= -\varepsilon + \sup_{\widetilde{h}\in\widetilde{\mathcal{W}}}\big(T(\widetilde{h}) - I(\widetilde{h})\big).\end{aligned} \tag{8.4}$$

Since ε is arbitrary in (8.3) and (8.4), this completes the proof. $\square$

PROOF OF THEOREM 3.2. Take any $\eta > 0$. Let

$$\widetilde{A} := \{\widetilde{h} : \delta_\square(\widetilde{h}, \widetilde{F}^*) \geq \eta\}.$$

It is easy to see that $\widetilde{A}$ is a closed set. By compactness of $\widetilde{\mathcal{W}}$ and $\widetilde{F}^*$, and upper semi-continuity of $T - I$, it follows that

$$2\gamma := \sup_{\widetilde{h} \in \widetilde{\mathcal{W}}} \big(T(\widetilde{h}) - I(\widetilde{h})\big) - \sup_{\widetilde{h} \in \widetilde{A}} \big(T(\widetilde{h}) - I(\widetilde{h})\big) > 0.$$

Choose $\varepsilon = \gamma$ and define $\widetilde{F}^a$ and R as in the proof of Theorem 3.1. Let $\widetilde{A}^a := \widetilde{A} \cap \widetilde{F}^a$. Then

$$\mathbb{P}(G_n \in \widetilde{A}) \leq e^{-n^2\psi_n} \sum_{a \in R} e^{n^2(a+\varepsilon)} |\widetilde{A}_n^a| \leq e^{-n^2\psi_n} |R| \sup_{a \in R} e^{n^2(a+\varepsilon)} |\widetilde{A}_n^a|.$$

While bounding the last term above, it can be assumed without loss of generality that $\widetilde{A}^a$ is nonempty for each $a \in R$, for the other a's can be dropped without upsetting the bound. By (8.1) and Theorem 3.1 (noting that $\widetilde{A}^a$ is compact), the above display gives

$$\limsup_{n \to \infty} \frac{\log \mathbb{P}(G_n \in \widetilde{A})}{n^2} \leq \sup_{a \in R} \Big(a + \varepsilon - \inf_{\widetilde{h} \in \widetilde{A}^a} I(\widetilde{h})\Big) - \sup_{\widetilde{h} \in \widetilde{\mathcal{W}}} \big(T(\widetilde{h}) - I(\widetilde{h})\big).$$

Each $\widetilde{h} \in \widetilde{A}^a$ satisfies $T(\widetilde{h}) \geq a$. Consequently,

$$\sup_{\widetilde{h} \in \widetilde{A}^a} \big(T(\widetilde{h}) - I(\widetilde{h})\big) \geq \sup_{\widetilde{h} \in \widetilde{A}^a} \big(a - I(\widetilde{h})\big) = a - \inf_{\widetilde{h} \in \widetilde{A}^a} I(\widetilde{h}).$$

Substituting this in the earlier display gives

$$\begin{aligned}
\limsup_{n \to \infty} \frac{\log \mathbb{P}(G_n \in \widetilde{A})}{n^2} &\leq \varepsilon + \sup_{a \in R} \sup_{\widetilde{h} \in \widetilde{A}^a} \big(T(\widetilde{h}) - I(\widetilde{h})\big) - \sup_{\widetilde{h} \in \widetilde{\mathcal{W}}} \big(T(\widetilde{h}) - I(\widetilde{h})\big) \\
&= \varepsilon + \sup_{\widetilde{h} \in \widetilde{A}} \big(T(\widetilde{h}) - I(\widetilde{h})\big) - \sup_{\widetilde{h} \in \widetilde{\mathcal{W}}} \big(T(\widetilde{h}) - I(\widetilde{h})\big). \\
&= \varepsilon - 2\gamma = -\gamma.
\end{aligned}$$

This completes the proof. □

PROOF OF THEOREM 4.1. By Theorem 3.1,

$$\lim_{n \to \infty} \psi_n = \sup_{h \in \mathcal{W}} \big(T(h) - I(h)\big). \tag{8.5}$$

By Hölder's inequality,

$$t(H_i, h) \leq \int\!\!\int_{[0,1]^2} h(x, y)^{e(H_i)}\, dx\, dy.$$

Thus, by the nonnegativity of $\beta_2, \dots, \beta_k$,

$$T(h) \le \beta_1 t(H_1, h) + \sum_{i=2}^{k} \beta_i \iint_{[0,1]^2} h(x, y)^{e(H_i)}\, dx\, dy$$

$$= \iint_{[0,1]^2} \sum_{i=1}^{k} \beta_i h(x, y)^{e(H_i)}\, dx\, dy.$$

On the other hand, the inequality in the above display becomes an equality if h is a constant function. Therefore, if u is a point in $[0, 1]$ that maximizes

$$\sum_{i=1}^{k} \beta_i u^{e(H_i)} - I(u),$$

then the constant function $h(x, y) \equiv u$ solves the variational problem (8.5). To see that constant functions are the only solutions, assume that there is at least one i such that the graph H_i has at least one vertex with two or more neighbors. The above steps show that if h is a maximizer, then for each i,

$$t(H_i, h) = \iint_{[0,1]^2} h(x, y)^{e(H_i)}\, dx\, dy. \tag{8.6}$$

In other words, equality holds in Hölder's inequality. Suppose that H_i has vertex set $\{1, 2, \dots, k\}$ and vertices 2 and 3 are both neighbors of 1 in H_i. Recall that

$$t(H_i, h) = \int_{[0,1]^k} \prod_{\{j,l\} \in E(H_i)} h(x_j, x_l)\, dx_1 \cdots dx_k.$$

In particular, the integrand contains the product $h(x_1, x_2)h(x_1, x_3)$. From this and the criterion for equality in Hölder's inequality, it follows that $h(x_1, x_2) = h(x_1, x_3)$ for almost every (x_1, x_2, x_3). Using the symmetry of h one can now easily conclude that h is almost everywhere a constant function.

If the condition does not hold, then each H_i is a union of vertex-disjoint edges. Assume that some H_i has more than one edge. Then again by (8.6) it follows that h must be a constant function.

Finally, if each H_i is just a single edge, then the maximization problem (8.5) can be explicitly solved and the solutions are all constant functions. □

PROOF OF THEOREM 4.2. The assertions about graph limits in this theorem are direct consequences of Theorems 3.2 and 4.1. Since $\sum_{i=1}^{k} \beta_i u^{e(H_i)}$ is a polynomial function of u and $I(u)$ is sufficiently well-behaved, showing that U is a finite set is a simple analytical exercise. □

PROOF OF THEOREM 5.1. Fix $\beta_1 < 0$ such that $c_1 < c_2$. As a preliminary step, let us prove that for any $\beta_2 > 0$,

$$\lim_{n \to \infty} \mathbb{P}\big(f(G_n) \in (c_1, c_2)\big) = 0. \tag{8.7}$$

Fix $\beta_2 > 0$. Let u be any maximizer of ℓ. Then by Theorem 4.2, it suffices to prove that either $u < e^{\beta_1}/(1+e^{\beta_1})$ or $u > 1 + 1/2\beta_1$. This is proved as follows. Define a function $g : [0,1] \to \mathbb{R}$ as

$$g(v) := \ell\big(v^{1/3}\big).$$

Then ℓ is maximized at u if and only if g is maximized at u^3. Since ℓ is a bounded continuous function and $\ell'(0) = \infty$, $\ell'(1) = -\infty$, ℓ cannot be maximized at 0 or 1. Therefore, the same is true for g. Let v be a point in $(0,1)$ at which g is maximized. Then $g''(v) \le 0$. A simple computation shows that

$$g''(v) = \frac{1}{9v^{5/3}}\biggl(-2\beta_1 + \log\frac{v^{1/3}}{1-v^{1/3}} - \frac{1}{2(1-v^{1/3})}\biggr).$$

Thus, $g''(v) \le 0$ only if

$$\log\frac{v^{1/3}}{1-v^{1/3}} \le \beta_1 \quad \text{or} \quad -\frac{1}{2(1-v^{1/3})} \le \beta_1.$$

This shows that a maximizer u of ℓ must satisfy $u \le c_1$ or $u \ge c_2$. Now, if $u = c_1$, then $u < c_2$, and therefore the above computations show that $g''(v) > 0$, where $v = u^3$. Similarly, if $u = c_2$ then $u > c_1$ and again $g''(v) > 0$. Thus, we have proved that $u < c_1$ or $u > c_2$. By Theorem 3.2, this completes the proof of (8.7) when $\beta_2 > 0$.

Now notice that as $\beta_2 \to \infty$, $\sup_{u \le a} \ell(u) \sim \beta_2 a^3$ for any fixed $a \le 1$. This shows that as $\beta_2 \to \infty$, any maximizer of ℓ must eventually be larger than $1 + 1/2\beta_1$. Therefore, for sufficiently large β_2,

$$\lim_{n\to\infty} \mathbb{P}\big(f(G_n) < c_2\big) = 0. \tag{8.8}$$

Next consider the case $\beta_2 \le 0$. Let $\widetilde{F}^*$ be the set of maximizers of $T(\widetilde{h}) - I(\widetilde{h})$. Take any $\widetilde{h} \in \widetilde{F}^*$ and let h be a representative element of $\widetilde{h}$. Let $p = e^{2\beta_1}/(1+e^{2\beta_1})$. An easy verification shows that

$$T(h) - I(h) = \beta_2 t(H_2, h) - I_p(h) - \tfrac{1}{2}\log(1-p),$$

where $I_p(h)$ is defined as in (2.8). Define a new function

$$h_1(x,y) := \min\big\{h(x,y), p\big\}.$$

Since the function I_p defined in (2.7) is minimized at p, it follows that for all $x, y \in [0,1]$, $I_p(h_1(x,y)) \le I_p(h(x,y))$. Consequently, $I_p(h_1) \le I_p(h)$. Again, since $\beta_2 \le 0$ and $h_1 \le h$ everywhere, $\beta_2 t(H_2, h_1) \ge \beta_2 t(H_2, h)$. Combining these observations, we see that $T(h_1) - I(h_1) \ge T(h) - I(h)$. Since h maximizes $T - I$ it follows that equality must hold at every step in the above deductions, from which it is easy to conclude that $h = h_1$ a.e. In other words, $h(x,y) \le p$ a.e. This is true

for every $\widetilde{h} \in \widetilde{F}^*$. Since $p < c_1$, the above deduction coupled with Theorem 3.2 proves that when $\beta_2 \le 0$,

$$\lim_{n\to\infty} \mathbb{P}\big(f(G_n) > c_1\big) = 0. \tag{8.9}$$

Recalling that β_1 is fixed, define

$$a_n(\beta_2) := \mathbb{P}\big(f(G_n) > c_1\big), \qquad b_n(\beta_2) := \mathbb{P}\big(f(G_n) < c_2\big).$$

Let A_n and B_n denote the events in brackets in the above display. A simple computation shows that

$$a_n'(\beta_2) = \frac{6}{n}\operatorname{Cov}\big(1_{A_n}, \Delta(G_n)\big) \quad \text{and} \quad b_n'(\beta_2) = \frac{6}{n}\operatorname{Cov}\big(1_{B_n}, \Delta(G_n)\big),$$

where $\Delta(G_n)$ is the number of triangles in G_n. As noted at the end of Section 4, the exponential random graph model with $\beta_2 \ge 0$ satisfies the FKG criterion [23]. Therefore, the above identities show that on the nonnegative axis, a_n is a nondecreasing function and b_n is a nonincreasing function.

Let $q_1 := \sup\{x \in \mathbb{R} : \lim_{n\to\infty} a_n(x) = 0\}$. By equation (8.8), $q_1 < \infty$ and by equation (8.9) $q_1 \ge 0$. Similarly, if $q_2 := \inf\{x \in \mathbb{R} : \lim_{n\to\infty} b_n(x) = 0\}$, then $0 \le q_2 < \infty$. Also, clearly, $q_1 \le q_2$ since $a_n + b_n \ge 1$ everywhere. We claim that $q_1 = q_2$. This would complete the proof by the monotonicity of a_n and b_n.

To prove that $q_1 = q_2$, suppose not. Then $q_1 < q_2$. Then for any $\beta_2 \in (q_1, q_2)$, $\limsup a_n(\beta_2) > 0$ and $\limsup b_n(\beta_2) > 0$. Now,

$$0 \le a_n(\beta_2) + b_n(\beta_2) - 1 = \mathbb{P}\big(f(G_n) \in (c_1, c_2)\big).$$

Therefore by (8.7),

$$\lim_{n\to\infty} \big(a_n(\beta_2) + b_n(\beta_2) - 1\big) = 0.$$

Thus, for any $\beta_2 \in (q_1, q_2)$, $\limsup(1 - b_n(\beta_2)) > 0$. By Theorem 4.2, this implies that the function ℓ has a maximum in $[c_2, 1]$. Similarly, for any $\beta_2 \in (q_1, q_2)$, $\limsup(1 - a_n(\beta_2)) > 0$ and therefore the function ℓ has a maximum in $[0, c_1]$. Now fix $q_1 < \beta_2 < \widetilde{\beta}_2 < q_2$, and let ℓ and $\widetilde{\ell}$ denote the two ℓ-functions corresponding to β_2 and $\widetilde{\beta}_2$, respectively. That is,

$$\ell(u) = \beta_1 u + \beta_2 u^3 - I(u), \qquad \widetilde{\ell}(u) = \beta_1 u + \widetilde{\beta}_2 u^3 - I(u).$$

By the above argument, ℓ attains its maximum at some point $u_1 \in [0, c_1]$ and at some point $u_2 \in [c_2, 1]$. (There may be other maxima, but that is irrelevant for us.) Note that

$$\max_{u \le c_1} \widetilde{\ell}(u) = \max_{u \le c_1} \big(\ell(u) + (\widetilde{\beta}_2 - \beta_2)u^3\big) \le \ell(u_1) + (\widetilde{\beta}_2 - \beta_2)c_1^3.$$

On the other hand

$$\max_{u \ge c_2} \widetilde{\ell}(u) \ge \widetilde{\ell}(u_2) = \ell(u_2) + (\widetilde{\beta}_2 - \beta_2)u_2^3 \ge \ell(u_2) + (\widetilde{\beta}_2 - \beta_2)c_2^3.$$

Since $\ell(u_1) = \ell(u_2)$, $\widetilde{\beta}_2 > \beta_2$ and $c_2 > c_1$, this shows that

$$\max_{u \le c_1} \widetilde{\ell}(u) < \max_{u \ge c_2} \widetilde{\ell}(u),$$

contradicting our previous deduction that $\widetilde{\ell}$ has maxima in both $[0, c_1]$ and $[c_2, 1]$. This proves that $q_1 = q_2$. $\square$

PROOF OF THEOREM 6.1. Let g be a symmetric bounded measurable function from $[0, 1]$ into $\mathbb{R}$. For each $u \in \mathbb{R}$, let

$$h_u(x, y) := h(x, y) + ug(x, y).$$

Then h_u is a symmetric bounded measurable function from $[0, 1]$ into $\mathbb{R}$. First, suppose that h is bounded away from 0 and 1. Then $h_u \in \mathcal{W}$ for every u sufficiently small in magnitude. Since h maximizes $T(h) - I(h)$ among all elements of $\mathcal{W}$, therefore under the above assumption, for all u sufficiently close to zero,

$$T(h_u) - I(h_u) \le T(h) - I(h).$$

In particular,

$$\frac{d}{du}\big(T(h_u) - I(h_u)\big)\Big|_{u=0} = 0. \tag{8.10}$$

It is easy to check that $T(h_u) - I(h_u)$ is differentiable in u for any h and g. In particular, the derivative is given by

$$\frac{d}{du}\big(T(h_u) - I(h_u)\big) = \sum_{i=1}^{k} \beta_i \frac{d}{du} t(H_i, h_u) - \frac{d}{du} I(h_u).$$

Now,

$$\begin{aligned}\frac{d}{du} I(h_u) &= \int\int \frac{d}{du} I\big(h(x, y) + ug(x, y)\big)\, dy\, dx \\ &= \frac{1}{2}\int\int g(x, y) \log \frac{h_u(x, y)}{1 - h_u(x, y)}\, dy\, dx.\end{aligned}$$

Consequently,

$$\frac{d}{du} I(h_u)\Big|_{u=0} = \frac{1}{2}\int\int g(x, y) \log \frac{h(x, y)}{1 - h(x, y)}\, dy\, dx.$$

Next, note that

$$\begin{aligned}&\frac{d}{du} t(H_i, h_u) \\ &\quad = \int_{[0,1]^{V(H)}} \sum_{(r,s) \in E(H_i)} g(x_r, x_s) \prod_{\substack{(r',s') \in E(H_i) \\ (r',s') \ne (r,s)}} h_u(x_{r'}, x_{s'}) \prod_{v \in V(H)} dx_v \\ &\quad = \int\int g(x, y) \Delta_{H_i} h_u(x, y)\, dy\, dx.\end{aligned}$$

Combining the above computations and (8.10), we see that for any symmetric bounded measurable $g : [0, 1] \to \mathbb{R}$,

$$\int\int g(x, y)\left(\sum_{i=1}^{k} \beta_i \Delta_{H_i} h(x, y) - \frac{1}{2}\log\frac{h(x, y)}{1 - h(x, y)}\right) dy\, dx = 0.$$

Taking $g(x, y)$ equal to the function within the brackets (which is bounded since h is assumed to be bounded away from 0 and 1), the conclusion of the theorem follows.

Now note that the theorem was proved under the assumption that h is bounded away from 0 and 1. We claim that this is true for any h that maximizes $T(h) - I(h)$. To prove this claim, take any such h. Fix $p \in (0, 1)$. For each $u \in [0, 1]$, let

$$h_{p,u}(x, y) := h(x, y) + u\big(p - h(x, y)\big)_+.$$

In other words, $h_{p,u}$ is simply h_u with $g = (p - h)_+$. Then certainly, $h_{p,u}$ is a symmetric bounded measurable function from $[0, 1]^2$ into $[0, 1]$. Note that

$$\frac{d}{du} h_{p,u}(x, y) = \big(p - h(x, y)\big)_+.$$

Using this, an easy computation as above shows that

$$\begin{aligned}
&\frac{d}{du}\big(T(h_{p,u}) - I(h_{p,u})\big)\Big|_{u=0} \\
&\quad = \int\int \left(\sum_{i=1}^{k} \beta_i \Delta_{H_i} h(x, y) - \frac{1}{2}\log\frac{h(x, y)}{1 - h(x, y)}\right)\big(p - h(x, y)\big)_+ dy\, dx \\
&\quad \geq \int\int \left(-C - \frac{1}{2}\log\frac{h(x, y)}{1 - h(x, y)}\right)\big(p - h(x, y)\big)_+ dy\, dx,
\end{aligned}$$

where C is a positive constant depending only on $\beta_1, \ldots, \beta_k$ and $H_1, \ldots, H_k$ (and not on p or h). When $h(x, y) = 0$, the integrand is interpreted as ∞, and when $h(x, y) = 1$, the integrand is interpreted as 0.

Now, if p is so small that

$$-C - \frac{1}{2}\log\frac{p}{1 - p} > 0,$$

then the previous display proves that the derivative of $T(h_{p,u}) - I(h_{p,u})$ with respect to u is strictly positive at $u = 0$ if $h < p$ on a set of positive Lebesgue measure. Hence, h cannot be a maximizer of $T - I$ unless $h \geq p$ almost everywhere. This proves that any maximizer of $T - I$ must be bounded away from zero. A similar argument with $g = -(h - p)_+$ shows that it must be bounded away from 1, and hence completes the proof of the theorem. □

PROOF OF THEOREM 6.2. It suffices to prove that the maximizer of $T(h) - I(h)$ as h varies over $\mathcal{W}$ is unique. This is because: if h is a maximizer, then so

is $h_\sigma(x, y) := h(\sigma x, \sigma y)$ for any measure preserving bijection $\sigma : [0, 1] \to [0, 1]$. The only functions that are invariant under such transforms are constant functions.

Let Δ_H be the operator defined in Section 6.1. Let $\|\cdot\|_\infty$ denote the L^∞ norm on $\mathcal{W}$ (i.e., the essential supremum of the absolute value). Let h and g be two maximizers of $T - I$. For any finite simple graph H, a simple computation shows that

$$\begin{aligned}\|\Delta_H h - \Delta_H g\|_\infty &\le \sum_{(r,s)\in E(H)} \|\Delta_{H,r,s} h - \Delta_{H,r,s} g\|_\infty \\ &\le e(H)\big(e(H) - 1\big)\|h - g\|_\infty.\end{aligned}$$

Using the above inequality, Theorem 6.1 and the inequality

$$\left|\frac{e^x}{1 + e^x} - \frac{e^y}{1 + e^y}\right| \le \frac{|x - y|}{4}$$

(easily proved by the mean value theorem) it follows that for almost all x, y,

$$\begin{aligned}\big|h(x, y) - g(x, y)\big| &= \left|\frac{e^{2\sum_{i=1}^k \beta_i \Delta_{H_i} h(x,y)}}{1 + e^{2\sum_{i=1}^k \beta_i \Delta_{H_i} h(x,y)}} - \frac{e^{2\sum_{i=1}^k \beta_i \Delta_{H_i} g(x,y)}}{1 + e^{2\sum_{i=1}^k \beta_i \Delta_{H_i} g(x,y)}}\right| \\ &\le \frac{1}{2}\sum_{i=1}^k |\beta_i| \|\Delta_{H_i} h - \Delta_{H_i} g\|_\infty \\ &\le \frac{1}{2}\|h - g\|_\infty \sum_{i=1}^k |\beta_i| e(H_i)\big(e(H_i) - 1\big).\end{aligned}$$

If the coefficient of $\|h - g\|_\infty$ in the last expression is strictly less than 1, it follows that h must be equal to g a.e. $\square$

PROOF OF THEOREM 6.3. Fix β_1. Let $p = e^{2\beta_1}/(1 + e^{2\beta_1})$ and $\gamma := -\beta_2$, so that for any $h \in \mathcal{W}$,

$$T(h) - I(h) = -\gamma t(H_2, h) - I_p(h) - \tfrac{1}{2}\log(1 - p).$$

Assume without loss of generality that $\beta_2 < 0$. Suppose u is a constant such that the function $h(x, y) \equiv u$ maximizes $T(h) - I(h)$, that is, minimizes $\gamma t(H_2, h) + I_p(h)$. Note that

$$\gamma t(H_2, h) + I_p(h) = \gamma u^3 + I_p(u).$$

Clearly, the definition of u implies that $\gamma u^3 + I_p(u) \le \gamma x^3 + I_p(x)$ for all $x \in [0, 1]$. This implies that u must be in $(0, 1)$, because the derivative of $x \mapsto \gamma x^3 + I_p(x)$ is $-\infty$ at 0 and ∞ at 1. Thus,

$$0 = \frac{d}{dx}\big(\gamma x^3 + I_p(x)\big)\Big|_{x=u} = 3\gamma u^2 + \frac{1}{2}\log\frac{u}{1 - u} - \frac{1}{2}\log\frac{p}{1 - p},$$

which shows that $u \le c(\gamma)$, where $c(\gamma)$ is a function of γ such that

$$\lim_{\gamma\to\infty} c(\gamma) = 0.$$

This shows that

$$\lim_{\gamma\to\infty} \min_{0\le x\le 1} \bigl(\gamma x^3 + I_p(x)\bigr) = I_p(0) = \frac{1}{2}\log\frac{1}{1-p}. \tag{8.11}$$

Next, let g be the function

$$g(x,y) := \begin{cases} 0, & \text{if } x, y \text{ on same side of } 1/2,\\ p, & \text{if not.}\end{cases}$$

Clearly, for almost all (x,y,z), $g(x,y)g(y,z)g(z,x) = 0$. Thus, $t(H_2,g) = 0$. A simple computation shows that

$$I_p(g) = \frac{1}{4}\log\frac{1}{1-p}.$$

Thus, $\gamma t(H_2,g) + I_p(g) = \frac{1}{4}\log\frac{1}{1-p}$. This shows that if γ is large enough (depending on p and hence β_1), then $T - I$ cannot be maximized at a constant function. The rest of the conclusion follows easily from Theorem 3.2 and the compactness of $\widetilde{\mathcal{W}}$. $\square$

PROOF OF THEOREM 6.4. Take any $h \in \mathcal{W}$. Note that

$$\begin{aligned} t(H_j,h) &= \int_{[0,1]^j} h(x_1,x_2)h(x_1,x_3)\cdots h(x_1,x_j)\,dx_1\cdots dx_j\\ &= \int_0^1 M(x)^j\,dx, \end{aligned}$$

where

$$M(x) = \int_0^1 h(x,y)\,dy.$$

Since I is a convex function,

$$\int_0^1 I\bigl(h(x,y)\bigr)\,dy \ge I\bigl(M(x)\bigr)$$

with equality if and only if $h(x,y)$ is the same for almost all y. Thus, putting

$$P(u) := \sum_{j=1}^k \beta_j u^j,$$

we get

$$T(h) - I(h) = \int_0^1 P\bigl(M(x)\bigr)\,dx - I(h) \le \int_0^1 \bigl(P\bigl(M(x)\bigr) - I\bigl(M(x)\bigr)\bigr)\,dx$$

with equality if and only if for almost all x, (a) $h(x, y)$ is constant as a function of y, and (b) $M(x)$ equals a value u^* that maximizes $P(u) - I(u)$. By the symmetry of h, the condition (a) implies that h is constant almost everywhere. The condition (b) gives the set of possible values of this constant. The rest follows as in the proofs of Theorems 4.1 and 4.2. □

LEMMA 8.1. *Let r be any integer $\geq \chi(H)$. Let K_r be the complete graph on r vertices. Then for any symmetric measurable $h : [0, 1]^2 \to \{0, 1\}$, if $t(K_r, h) > 0$ then $t(H, h) > 0$.*

PROOF. Let $h_n(x, y)$ be the average value of h in the dyadic square of width 2^{-n} containing the point (x, y). A standard martingale argument implies that the sequence of functions $\{h_n\}_{n\geq 1}$ converges to h almost everywhere. For any positive integer u, let K_r^u denote the complete r-partite graph on ru vertices, where each partition consists of u vertices (so that $K_r^1 = K_r$). Since $r \geq \chi(H)$, it is easy to see that there exists u so large that H is a subgraph of K_r^u [i.e., $V(H) \subseteq V(K_r^u)$ and $E(H) \subseteq E(K_r^u)$]. Fix such a u.

By the almost everywhere convergence of h_n to h and the assumption that $t(K_r, h) > 0$, there is a set of r distinct points $x_1, \ldots, x_r \in [0, 1]$ that do not lie on the boundary of any dyadic interval, such that $h(x_i, x_j) > 0$ and $\lim_{n\to\infty} h_n(x_i, x_j) = h(x_i, x_j)$ for each $1 \leq i \neq j \leq r$. Since h is $\{0, 1\}$-valued, $h(x_i, x_j) = 1$ for each $i \neq j$. Choose n so large that for each $i \neq j$,

$$h_n(x_i, x_j) \geq 1 - \varepsilon,$$

where $\varepsilon = 1/2r^2u^2$. Let $(X_i^s)_{1\leq i\leq r, 1\leq s\leq u}$ be independent random variables, where X_i^s is uniformly distributed in the dyadic interval of width 2^{-n} containing x_i. Then for each $1 \leq i \neq j \leq r$, $1 \leq q, s \leq u$,

$$\mathbb{P}\big(h\big(X_i^q, X_j^s\big) = 1\big) = h_n(x_i, x_j) \geq 1 - \varepsilon.$$

Therefore,

$$\mathbb{P}\big(h\big(X_i^q, X_j^s\big) = 1 \text{ for all } 1 \leq i \neq j \leq r, 1 \leq q, s \leq u\big) \geq 1 - r^2u^2\varepsilon = 1/2.$$

Let $(Y_i^s)_{1\leq i\leq r, 1\leq s\leq u}$ be independent random variables uniformly distributed in $[0, 1]$. Conditional on the event that Y_i^s belongs to the dyadic interval of width 2^{-n} containing x_i, Y_i^s has the same distribution as X_i^s. As a consequence of the last display, this shows that

$$\begin{aligned} t\big(K_r^u, h\big) &= \mathbb{P}\big(h\big(Y_i^q, Y_j^s\big) = 1 \text{ for all } 1 \leq i \neq j \leq r, 1 \leq q, s \leq u\big) \\ &\geq 2^{-nru}\mathbb{P}\big(h\big(X_i^q, X_j^s\big) = 1 \text{ for all } 1 \leq i \neq j \leq r, 1 \leq q, s \leq u\big) > 0. \end{aligned}$$

Since H is a subgraph of K_r^u, therefore $t(H, h) > 0$. □

THEOREM 8.2. *Let g be the function defined in (7.1). Take any $p \in (0,1)$. If f is any element of $\mathcal{W}$ that minimizes $I_p(f)$ among all f satisfying $t(H,f)=0$, then $\widetilde{f} = p\widetilde{g}$.*

PROOF. Take any minimizer f. (Minimizers exist due to the Lovász–Szegedy compactness theorem [41], Theorem 5.1, and the lower semicontinuity of I_p.) First, note that $f \le p$ almost everywhere: if not, then $I_p(f)$ can be decreased by replacing f with $\min\{f,p\}$, which retains the condition $t(H,f)=0$.

Next, note that for almost all x, y, $f(x,y)=0$ or p. If not, then redefine f to be equal to p wherever f was positive. This decreases the entropy while retaining the condition $t(H,f)=0$.

Let $h = f/p$. Then h takes value 0 or 1 almost everywhere and h minimizes $I_p(ph)$ among all symmetric measurable $h:[0,1]^2 \to \{0,1\}$ satisfying $t(H,h)=0$. Equivalently, h maximizes $\iint h(x,y)\,dx\,dy$ among all symmetric measurable $h:[0,1]^2 \to \{0,1\}$ satisfying $t(H,h)=0$. Our goal is to show that $\widetilde{h} = \widetilde{g}$.

Let $r := \chi(H)$. Let $X_0, X_1, X_2, \ldots$ be a sequence of i.i.d. random variables uniformly distributed in $[0,1]$. Let

$$\mathcal{R} := \bigl\{i : h(X_i, X_j) = 1 \text{ for all } 1 \le j < i\bigr\},$$

and let $R := |\mathcal{R}|$. Let $\lambda(x) := \int h(x,y)\,dy$, so that for any given i,

$$\mathbb{P}\bigl(h(X_i,X_j)=1 \text{ for all } 1\le j<i\bigr) = \mathbb{E}\bigl(\lambda(X_i)^{i-1}\bigr) = \mathbb{E}\bigl(\lambda(X_0)^{i-1}\bigr).$$

Thus,

$$\begin{aligned}
\mathbb{E}(R) &= \sum_{i=1}^{\infty} \mathbb{P}\bigl(h(X_i,X_j)=1 \text{ for all } 1\le j<i\bigr) \\
&= \sum_{i=1}^{\infty} \mathbb{E}\bigl(\lambda(X_0)^{i-1}\bigr) \qquad (8.12)\\
&\ge \sum_{i=1}^{\infty} \bigl(\mathbb{E}\lambda(X_0)\bigr)^{i-1} = \frac{1}{1-\mathbb{E}\lambda(X_0)} = \frac{1}{1-\iint h(x,y)\,dx\,dy}.
\end{aligned}$$

Let g be the function defined in (7.1). Suppose the vertex set of H is $\{1,\ldots,k\}$ for some integer k. If $t(H,g)>0$, then there exist $x_1,\ldots,x_k$ such that $g(x_i,x_j)=1$ whenever $\{i,j\}$ is an edge in H. By the definition of g, this implies that H can be colored by $r-1$ colors so that no two adjacent vertices receive the same color; since this is false, therefore $t(H,g)$ must be zero. By the optimality property of h, this gives

$$\iint h(x,y)\,dx\,dy \ge \iint g(x,y)\,dx\,dy = 1 - \frac{1}{r-1}.$$

Therefore by (8.12),

$$\mathbb{E}(R) \geq r - 1.$$

Again by Lemma 8.1, $t(K_r, h) = 0$. Therefore, $R \leq r - 1$ almost surely. Combined with the above display, this shows that equality must hold in (8.12) and $R = r - 1$ almost surely. In particular, $\mathbb{E}(\lambda(X_0)^2) = (\mathbb{E}\lambda(X_0))^2$ and $\mathbb{E}\lambda(X_0) = 1 - 1/(r-1)$, which shows that

$$\lambda(x) = 1 - \frac{1}{r-1} \qquad \text{a.e.}$$

For each x, let $A(x) := \{y : h(x, y) = 0\}$. Then $|A(x)| = 1/(r-1)$ a.e., where $|A(x)|$ denotes the Lebesgue measure of $A(x)$.

Define a random graph G on $\{0, 1, 2, \ldots\}$ by including the edge (i, j) if and only if $h(X_i, X_j) = 1$. Since $t(K_r, h) = 0$, G cannot contain any copy of K_r. Thus, with probability 1, $h(X_0, X_i) = 0$ for some $i \in \mathcal{R}$. In other words, $\bigcup_{i\in\mathcal{R}} A(X_i)$ cover almost all of $[0, 1]$. Again, $|A(X_i)| = 1/(r-1)$ for all $i \in \mathcal{R}$ and $|\mathcal{R}| = r - 1$ almost surely. All this together imply that with probability 1, $A(X_i) \cap A(X_j)$ has Lebesgue measure zero for all $i \neq j \in \mathcal{R}$, since

$$\sum_{i,j\in\mathcal{R}, i<j} |A(X_i) \cap A(X_j)| \leq \sum_{i\in\mathcal{R}} |A(X_i)| - \left|\bigcup_{i\in\mathcal{R}} A(X_i)\right| = 0.$$

Let $Y_1, Y_2, \ldots$ and $Z_1, Z_2, \ldots$ be i.i.d. random variables uniformly distributed in $[0, 1]$, that are independent of the sequence $X_1, X_2, \ldots$. Since $t(K_r, h) = 0$, with probability 1 there cannot exist l and a set B of integers of size $r - 2$ such that $h(Y_l, X_i) = h(Z_l, X_i) = 1$ for all $i \in B$, $h(X_i, X_j) = 1$ for all $i \neq j \in B$, and $h(Y_l, Z_l) = 1$.

Now fix a realization of $X_1, X_2, \ldots$. This fixes the set $\mathcal{R}$. Take any $i \in \mathcal{R}$. Let I be the smallest integer such that both Y_I and Z_I are in $A(X_i)$. Clearly Y_I and Z_I are independent and uniformly distributed in $A(X_i)$, conditional on the sequence $X_1, X_2, \ldots$ and our choice of $i \in \mathcal{R}$. By the observation from the preceding paragraph, $h(Y_I, Z_I) = 0$ with probability 1, since the set $\mathcal{R} \setminus \{i\}$ serves the role of B.

This shows that given $X_1, X_2, \ldots$, the sets $A(X_i)$ have the property that for almost all $y, z \in A(X_i)$, $h(y, z) = 0$. Since $\lambda(x) = 1 - 1/(r-1)$ a.e. and $|A(X_i)| = 1/(r-1)$, this shows that for almost all $y \in A(X_i)$ and almost all $z \notin A(X_i)$, $h(y, z) = 1$.

The properties of $(A(X_i))_{i\in\mathcal{R}}$ that we established can be summarized as follows: the sets $A(X_i)$ are disjoint up to errors of measure zero; each $A(X_i)$ has Lebesgue measure $1/(r-1)$ and together they cover the whole of $[0, 1]$; for almost all $y, z \in [0, 1]$, $h(y, z) = 0$ if they belong to the same $A(X_i)$, and $h(y, z) = 1$ if $y \in A(X_i)$ and $z \in A(X_j)$ for some $i \neq j$. These properties immediately show that h is the same as the function g up to a rearrangement; the formal argument can be completed as follows.

Given $X_1, X_2, \ldots$, let $u : [0, 1] \to [0, 1]$ be the map defined as

$$u(x) := \text{minimum } i \in \mathcal{R} \text{ such that } x \in A(X_i).$$

Note that with probability 1, for almost all x there is a unique $i \in \mathcal{R}$ such that $x \in A(X_i)$. Let $\sigma : [0,1] \to [0,1]$ be a measure-preserving bijection such that $x \mapsto u(\sigma x)$ is a nonincreasing (we omit the construction). Then σ maps the intervals $[0, 1/(r-1)]$, $[1/(r-1), 2/(r-1)]$, $\ldots$, $[(r-2)/(r-1), 1]$ onto the sets $(A(X_i))_{i\in\mathcal{R}}$ up to errors of measure zero. By the properties of $A(X_i)$ established above, this shows that $h(\sigma x, \sigma y)$ is the same as $g(x, y)$ up to an error of measure zero. $\square$

PROOF OF THEOREM 7.1. First, note that

$$T_{\beta_2}(h) - I(h) = \beta_2 t(H,h) - I_p(h) - \tfrac{1}{2}\log(1-p),$$

where $p = e^{2\beta_1}/(1+e^{2\beta_1})$. Take a sequence $\beta_2^{(n)} \to -\infty$, and for each n, let $\widetilde{h}_n$ be an element of $\widetilde{F}^*(\beta_2^{(n)})$. Let $\widetilde{h}$ be a limit point of $\widetilde{h}_n$ in $\widetilde{\mathcal{W}}$. If $t(H,h) > 0$, then by the continuity of the map $t(H,\cdot)$ and the boundedness of I_p,

$$\lim_{n\to\infty} \psi(\beta_2^{(n)}) = -\infty.$$

But this is impossible since $\psi(\beta_2^{(n)})$ is uniformly bounded below, as can be easily seen by considering the function g defined in (7.1) as a test function in the variational problem. Thus, $t(H,h) = 0$. If f is a function such that $t(H,f) = 0$ and $I_p(f) < I_p(h)$, then for all sufficiently large n,

$$T_{\beta_2^{(n)}}(h_n) - I(h_n) < T_{\beta_2^{(n)}}(f) - I(f)$$

contradicting the definition of $\widetilde{F}^*(\beta_2)$. Thus, if f is a function such that $t(H,f) = 0$, then $I_p(f) \ge I_p(h)$. By Theorem 8.2, this shows that $\widetilde{h} = p\widetilde{g}$. The compactness of $\widetilde{\mathcal{W}}$ now proves the first part of the theorem.

For the second part, first note that

$$\begin{aligned}\liminf_{n\to\infty} \psi(\beta_2^{(n)}) &\ge \lim_{n\to\infty} \bigl(T_{\beta_2^{(n)}}(pg) - I(pg)\bigr) = -I_p(pg) - \frac{1}{2}\log(1-p)\\ &= \frac{(\chi(H)-2)}{2(\chi(H)-1)} \log\frac{1}{1-p}.\end{aligned}$$

Next, note that by the lower-semicontinuity of I_p and the fact that $\beta_2^{(n)}$ is eventually negative,

$$\begin{aligned}\limsup_{n\to\infty} \psi(\beta_2^{(n)}) &= \limsup_{n\to\infty} \bigl(\beta_2^{(n)} t(H,h_n) - I_p(h_n)\bigr) - \frac{1}{2}\log(1-p)\\ &\le \limsup_{n\to\infty} \bigl(-I_p(h_n)\bigr) - \frac{1}{2}\log(1-p)\\ &\le -I_p(pg) - \frac{1}{2}\log(1-p).\end{aligned}$$

The proof is complete. $\square$

Acknowledgements. We thank Hans Anderson, Charles Radin, Austen Head, Susan Holmes, Sumit Mukherjee, and especially Sukhada Fadnavis and Mei Yin for their substantial help with this paper. We are particularly grateful to the referee for an exceptionally thorough and useful report and the improvement to Theorem 6.4.

REFERENCES

[1] ALDOUS, D. J. (1981). Representations for partially exchangeable arrays of random variables. *J. Multivariate Anal.* **11** 581–598. MR0637937

[2] ARISTOFF, D. and RADIN, C. (2011). Emergent structures in large networks. Preprint. Available at http://arxiv.org/abs/1110.1912.

[3] AUSTIN, T. (2008). On exchangeable random variables and the statistics of large graphs and hypergraphs. *Probab. Surv.* **5** 80–145. MR2426176

[4] AUSTIN, T. and TAO, T. (2010). Testability and repair of hereditary hypergraph properties. *Random Structures Algorithms* **36** 373–463. MR2666763

[5] BESAG, J. (1975). Statistical analysis of non-lattice data. *Statistician* **24** 179–195.

[6] BHAMIDI, S., BRESLER, G. and SLY, A. (2008). Mixing time of exponential random graphs. In 2008 *IEEE* 49*th Annual IEEE Symposium on Foundations of Computer Science* (*FOCS*) 803–812. IEEE, Washington, DC.

[7] BOLLOBÁS, B. (2001). *Random Graphs*, 2nd ed. *Cambridge Studies in Advanced Mathematics* **73**. Cambridge Univ. Press, Cambridge. MR1864966

[8] BOLLOBÁS, B. and RIORDAN, O. (2009). Metrics for sparse graphs. In *Surveys in Combinatorics* 2009. *London Mathematical Society Lecture Note Series* **365** 211–287. Cambridge Univ. Press, Cambridge. MR2588543

[9] BORGS, C., CHAYES, J., LOVÁSZ, L., SÓS, V. T. and VESZTERGOMBI, K. (2006). Counting graph homomorphisms. In *Topics in Discrete Mathematics*. *Algorithms and Combinatorics* **26** 315–371. Springer, Berlin. MR2249277

[10] BORGS, C., CHAYES, J. T., LOVÁSZ, L., SÓS, V. T. and VESZTERGOMBI, K. (2008). Convergent sequences of dense graphs. I. Subgraph frequencies, metric properties and testing. *Adv. Math.* **219** 1801–1851. MR2455626

[11] BORGS, C., CHAYES, J. T., LOVÁSZ, L., SÓS, V. T. and VESZTERGOMBI, K. (2012). Convergent sequences of dense graphs II. Multiway cuts and statistical physics. *Ann. of Math.* (2) **176** 151–219. MR2925382

[12] CHATTERJEE, S. (2007). Estimation in spin glasses: A first step. *Ann. Statist.* **35** 1931–1946. MR2363958

[13] CHATTERJEE, S. and DEY, P. S. (2010). Applications of Stein's method for concentration inequalities. *Ann. Probab.* **38** 2443–2485. MR2683635

[14] CHATTERJEE, S., DIACONIS, P. and SLY, A. (2011). Random graphs with a given degree sequence. *Ann. Appl. Probab.* **21** 1400–1435. MR2857452

[15] CHATTERJEE, S. and VARADHAN, S. R. S. (2011). The large deviation principle for the Erdős–Rényi random graph. *European J. Combin.* **32** 1000–1017. MR2825532

[16] COMETS, F. and JANŽURA, M. (1998). A central limit theorem for conditionally centred random fields with an application to Markov fields. *J. Appl. Probab.* **35** 608–621. MR1659520

[17] CORANDER, J., DAHMSTRÖM, K. and DAHMSTRÖM, P. (2002). Maximum likelihood estimation for exponential random graph models. In *Contributions to Social Network Analysis, Information Theory and Other Topics in Statistics: A Festschrift in Honour of Ove Frank* (J. Hagberg, ed.) 1–17. Dept. Statistics, Univ. Stockholm.

[18] DIACONIS, P. and JANSON, S. (2008). Graph limits and exchangeable random graphs. *Rend. Mat. Appl.* (7) **28** 33–61. MR2463439

[19] ERDŐS, P. and RÉNYI, A. (1960). On the evolution of random graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl.* **5** 17–61. MR0125031

[20] ERDÖS, P. and STONE, A. H. (1946). On the structure of linear graphs. *Bull. Amer. Math. Soc.* (*N.S.*) **52** 1087–1091. MR0018807

[21] FIENBERG, S. E. (2010). Introduction to papers on the modeling and analysis of network data. *Ann. Appl. Stat.* **4** 1–4. MR2758081

[22] FIENBERG, S. E. (2010). Introduction to papers on the modeling and analysis of network data—II. *Ann. Appl. Stat.* **4** 533–534. MR2744531

[23] FORTUIN, C. M., KASTELEYN, P. W. and GINIBRE, J. (1971). Correlation inequalities on some partially ordered sets. *Comm. Math. Phys.* **22** 89–103. MR0309498

[24] FRANK, O. and STRAUSS, D. (1986). Markov graphs. *J. Amer. Statist. Assoc.* **81** 832–842. MR0860518

[25] FREEDMAN, M., LOVÁSZ, L. and SCHRIJVER, A. (2007). Reflection positivity, rank connectivity, and homomorphism of graphs. *J. Amer. Math. Soc.* **20** 37–51 (electronic). MR2257396

[26] FRIEZE, A. and KANNAN, R. (1999). Quick approximation to matrices and applications. *Combinatorica* **19** 175–220. MR1723039

[27] GELFAND, I. M. and FOMIN, S. V. (2000). *Calculus of Variations*. Dover, New York.

[28] GELMAN, A. and MENG, X.-L. (1998). Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statist. Sci.* **13** 163–185. MR1647507

[29] GEYER, C. J. and THOMPSON, E. A. (1992). Constrained Monte Carlo maximum likelihood for dependent data. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **54** 657–699. MR1185217

[30] HÄGGSTRÖM, O. and JONASSON, J. (1999). Phase transition in the random triangle model. *J. Appl. Probab.* **36** 1101–1115. MR1742153

[31] HANDCOCK, M. S. (2003). Assessing degeneracy in statistical models of social networks. Working Paper 39, Center for Statistics and the Social Sciences, Univ. Washington, Seattle, WA.

[32] HOLLAND, P. W. and LEINHARDT, S. (1981). An exponential family of probability distributions for directed graphs. *J. Amer. Statist. Assoc.* **76** 33–65. MR0608176

[33] HOOVER, D. N. (1982). Row-column exchangeability and a generalized model for probability. In *Exchangeability in Probability and Statistics* (*Rome*, 1981) 281–291. North-Holland, Amsterdam. MR0675982

[34] JANSON, S., ŁUCZAK, T. and RUCINSKI, A. (2000). *Random Graphs*. Wiley, New York. MR1782847

[35] KALLENBERG, O. (2005). *Probabilistic Symmetries and Invariance Principles*. Springer, New York. MR2161313

[36] KOU, S. C., ZHOU, Q. and WONG, W. H. (2006). Equi-energy sampler with applications in statistical inference and statistical mechanics. *Ann. Statist.* **34** 1581–1652. MR2283711

[37] LOVÁSZ, L. (2006). The rank of connection matrices and the dimension of graph algebras. *European J. Combin.* **27** 962–970. MR2226430

[38] LOVÁSZ, L. (2007). Connection matrices. In *Combinatorics, Complexity, and Chance. Oxford Lecture Series in Mathematics and its Applications* **34** 179–190. Oxford Univ. Press, Oxford. MR2314569

[39] LOVÁSZ, L. and SÓS, V. T. (2008). Generalized quasirandom graphs. *J. Combin. Theory Ser. B* **98** 146–163. MR2368030

[40] LOVÁSZ, L. and SZEGEDY, B. (2006). Limits of dense graph sequences. *J. Combin. Theory Ser. B* **96** 933–957. MR2274085

[41] LOVÁSZ, L. and SZEGEDY, B. (2007). Szemerédi's lemma for the analyst. *Geom. Funct. Anal.* **17** 252–270. MR2306658

[42] LOVÁSZ, L. and SZEGEDY, B. (2009). Contractors and connectors of graph algebras. *J. Graph Theory* **60** 11–30. MR2478356

[43] LOVÁSZ, L. and SZEGEDY, B. (2010). Testing properties of graphs and functions. *Israel J. Math.* **178** 113–156. MR2733066

[44] LUBETZKY, E. and ZHAO, Y. (2012). On replica symmetry of large deviations in random graphs. Preprint. Available at http://arxiv.org/abs/1210.7013.

[45] PARK, J. and NEWMAN, M. E. J. (2004). Solution of the two-star model of a network. *Phys. Rev. E* (3) **70** 066146, 5. MR2133810

[46] PARK, J. and NEWMAN, M. E. J. (2005). Solution for the properties of a clustered network. *Phys. Rev. E* (3) **72** 026136, 5.

[47] RADIN, C. and SADUN, L. (2013). Phase transitions in a complex network. *J. Phys. A* **46** 305002.

[48] RADIN, C. and YIN, M. (2011). Phase transitions in exponential random graphs. Preprint. Available at http://arxiv.org/abs/1108.0649.

[49] 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

[50] ROBBINS, H. and MONRO, S. (1951). A stochastic approximation method. *Ann. Math. Statist.* **22** 400–407. MR0042668

[51] SANOV, I. N. (1961). On the probability of large deviations of random variables. In *Select. Transl. Math. Statist. and Probability, Vol.* 1 213–244. Amer. Math. Soc., Providence, RI. MR0116378

[52] SNIJDERS, T. A. (2002). Markov chain Monte Carlo estimation of exponential random graph models. *J. Soc. Structure* **3**.

[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] STRAUSS, D. (1986). On a general class of models for interaction. *SIAM Rev.* **28** 513–527. MR0867682

[55] TALAGRAND, M. (2003). *Spin Glasses: A Challenge for Mathematicians. Ergebnisse der Mathematik und Ihrer Grenzgebiete*. 3. *Folge. A Series of Modern Surveys in Mathematics* [*Results in Mathematics and Related Areas*. 3*rd Series. A Series of Modern Surveys in Mathematics*] **46**. Springer, Berlin. MR1993891

[56] TURÁN, P. (1941). Eine Extremalaufgabe aus der Graphentheorie. *Mat. Fiz. Lapok* **48** 436–452. MR0018405

[57] WASSERMAN, S. and FAUST, K. (2010). *Social Network Analysis: Methods and Applications. Structural Analysis in the Social Sciences*, 2nd ed. Cambridge Univ. Press, Cambridge.

[58] YIN, M. (2012). Critical phenomena in exponential random graphs. Preprint. Available at http://arxiv.org/abs/1208.2992.

DEPARTMENT OF STATISTICS
STANFORD UNIVERSITY
SEQUOIA HALL
390 SERRA MALL
STANFORD, CALIFORNIA 94305
USA
E-MAIL: souravc@stanford.edu
diaconis@math.stanford.edu