ON REPLICA SYMMETRY OF LARGE DEVIATIONS IN RANDOM GRAPHS

EYAL LUBETZKY AND YUFEI ZHAO

Abstract. The following question is due to Chatterjee and Varadhan (2011). Fix $0 < p < r < 1$ and take $G \sim \mathcal{G}(n,p)$, the Erdős-Rényi random graph with edge density p, conditioned to have at least as many triangles as the typical $\mathcal{G}(n,r)$. Is G close in cut-distance to a typical $\mathcal{G}(n,r)$? Via a beautiful new framework for large deviation principles in $\mathcal{G}(n,p)$, Chatterjee and Varadhan gave bounds on the *replica symmetric* phase, the region of (p,r) where the answer is positive. They further showed that for any small enough p there are at least *two* phase transitions as r varies.

We settle this question by identifying the replica symmetric phase for triangles and more generally for any fixed d-regular graph. By analyzing the variational problem arising from the framework of Chatterjee and Varadhan we show that the replica symmetry phase consists of all (p,r) such that $(r^d, h_p(r))$ lies on the convex minorant of $x \mapsto h_p(x^{1/d})$ where h_p is the rate function of a binomial with parameter p. In particular, the answer for triangles involves $h_p(\sqrt{x})$ rather than the natural guess of $h_p(x^{1/3})$ where symmetry was previously known. Analogous results are obtained for linear hypergraphs as well as the setting where the largest eigenvalue of $G \sim \mathcal{G}(n,p)$ is conditioned to exceed the typical value of the largest eigenvalue of $\mathcal{G}(n,r)$. Building on the work of Chatterjee and Diaconis (2012) we obtain additional results on a class of exponential random graphs including a new range of parameters where symmetry breaking occurs. En route we give a short alternative proof of a graph homomorphism inequality due to Kahn (2001) and Galvin and Tetali (2004).

1. Introduction

The following question was raised by Chatterjee and Varadhan [8] concerning large deviations in $\mathcal{G}(n,p)$, the Erdős-Rényi random graph on n vertices with edge density p.

> Fix $0 < p < r < 1$ and let G_n be an instance of $\mathcal{G}(n,p)$ conditioned on the rare event of having at least as many triangles as a typical instance of $\mathcal{G}(n,r)$. Is it the case that as $n \to \infty$ the graph G_n is close in cut-distance to a typical $\mathcal{G}(n,r)$ graph?

(A more formal statement, including the definition of the graph cut-metric, is postponed to §1.1.) This amounts to asking whether the likely reason for too many triangles is an overwhelming number of edges, uniformly distributed, or some fewer edges arranged in a special structure, e.g., a clique. Dubbed *replica symmetry* and *symmetry breaking*, resp., the dichotomy between these scenarios turns out to depend on p and r. Intriguingly, it was known that for small enough p there are at least *two* phase transitions as r increases from p to 1, with replica symmetry near the two endpoints.

In this work we analyze the variational problems arising from the framework of Chatterjee and Varadhan and obtain a full answer for the question above, as depicted in Fig. 1. More generally, we identify the phase diagram for upper tails of any fixed regular subgraph and derive related results in other random graph settings, e.g., exponential random graphs, random hypergraphs, etc.

1.1. Subgraph densities and spectral radii.

Large deviations for subgraph densities in random graphs have been extensively studied (see, e.g., [29, 46, 32, 28, 30, 5, 13, 14] as well as [3, 27] and the references therein). A representing example which drew significant attention is upper tails of triangle counts, i.e., estimating the probability that $\mathcal{G}(n,p)$ has at least $\binom{n}{3}r^3$ triangles where $r = (1+\eta)p$ for fixed $\eta > 0$ (allowing p to vary with n), a problem whose understanding is still incomplete. The order of the rate function (the normalized logarithm of this probability) when $p \to 0$ was only very recently settled: Chatterjee [5] and DeMarco and Kahn [14] independently established it to be $n^2p^2\log(1/p)$ when $p \gtrsim \log n/n$, and yet the exact rate function remains unknown in this range of p. We now turn to what was known for fixed p, our focus in this paper.

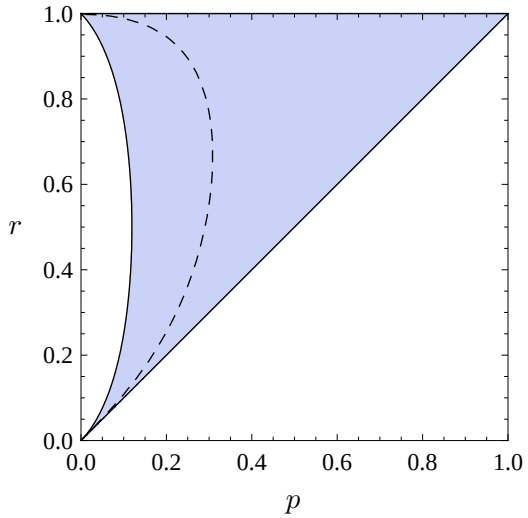

FIGURE 1. Phase diagram for the upper tail of triangle counts. Shaded region is the replica symmetric phase; the region to its left is the symmetry breaking phase. Previous results [6, 8] established replica symmetry to the right of the dashed curve.

Clearly, if the total number of edges in $\mathcal{G}(n,p)$ deviates to $m \sim \binom{n}{2} r$ then one will arrive at a random graph with m uniformly distributed edges featuring the desired number of triangles. Thus, the large deviation rate function for encountering $\binom{n}{3} r^3$ triangles in $\mathcal{G}(n,p)$ is at most $h_p(r)$, where

$$h_p(x) := x \log \frac{x}{p} + (1-x) \log \frac{1-x}{1-p} \quad \text{for } p \in (0,1) \text{ and } x \in [0,1] \tag{1.1}$$

is the rate function associated to the binomial distribution with probability p. However, it is possible that other configurations with broken symmetry would give rise to lower rate functions.

As an application of Stein's method for concentration inequalities, Chatterjee and Dey [6] found a range of (p,r) where the large deviation rate function for triangles is equal to $h_p(r)$, namely when $p \geq 2/(2+e^{3/2}) \approx 0.31$ or when r is suitably close either to p or to 1. This symmetry region was explicitly stated in [8, Theorem 4.3] as all pairs (p,r) where $(r^3, h_p(r))$ lies on the convex minorant of $x \mapsto h_p(x^{1/3})$. The breakthrough work of Chatterjee and Varadhan [8] introduced a remarkable general framework for large deviation principles in $\mathcal{G}(n,p)$ via Szemerédi's regularity lemma [45] and the theory of graph limits by Lovász et al. [34, 35, 4]. It expressed the large deviation rate function, and moreover the *structure* of the random graph conditioned on the large deviation, in terms of a variational problem on *graphons*, the infinite-dimensional limit objects for graph sequences.

Although often this variational problem is intractable, for triangles in the mentioned range of (p,r) it was shown in [8] to have a unique and symmetric solution. To formalize this symmetry, we say a graph G is close in cut-distance to a typical $\mathcal{G}(n,r)$ graph if all induced subgraphs on a linear number of vertices have edge density close to r. More precisely, for a graph G and $r \in [0,1]$ let

$$\delta_\square(G,r) := \sup_{A,B \subset V(G)} \frac{1}{|V(G)|^2} \big| e_G(A,B) - r\,|A|\,|B| \big|,$$

where $e_G(A,B)$ is the number of pairs $(a,b) \in A \times B$ with $ab \in E(G)$. Chatterjee and Varadhan showed that, in the above range of (p,r), if $G_n \sim \mathcal{G}(n,p)$ is conditioned to have at least $\binom{n}{3} r^3$ triangles then $\delta_\square(G_n, r) \to 0$ in probability as $n \to \infty$. The function $x \mapsto h_p(x^{1/3})$ governing that region is the natural candidate for the phase boundary as the cube-root accounts for the 3 edges of the triangle (see, e.g., [15, §4.5.2] for related literature), and indeed Chatterjee and Varadhan asked whether this precisely characterizes the full replica symmetric phase. As it turns out, however, the replica symmetric phase is strictly larger, being governed instead by $x \mapsto h_p(\sqrt{x})$. (See Fig. 1.)

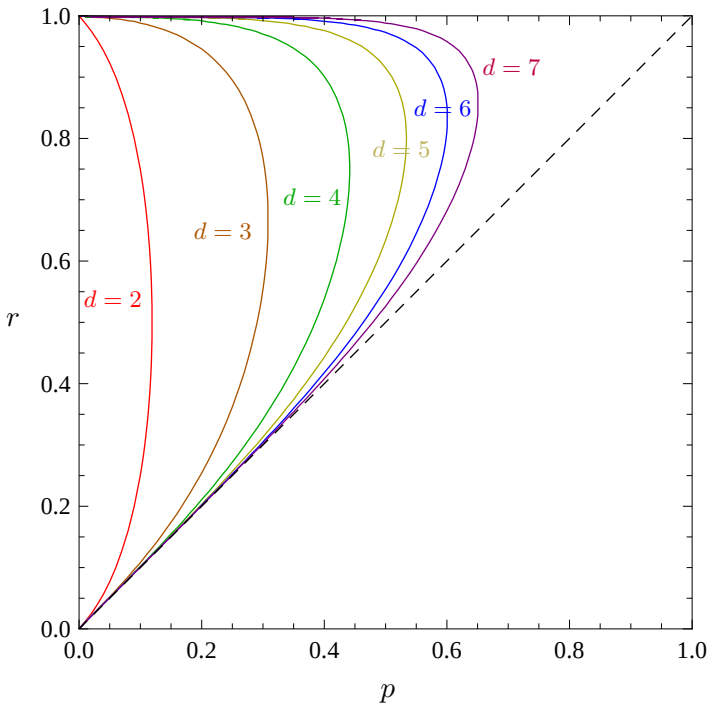

FIGURE 2. The phase boundary for counts of d-regular fixed subgraphs in $\mathcal{G}(n,p)$.

For any graph G let $e(G) = |E(G)|$, and for any two graphs G and H let $\hom(H,G)$ denote the number of homomorphisms from H to G (i.e., maps $V(H) \to V(G)$ that carry edges to edges). Let

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

be the probability that a random map $V(H) \to V(G)$ is a graph homomorphism. We now state our main result on the phase diagram for large deviations in densities of d-regular subgraphs.

Theorem 1.1. *Fix $0 < p \le r < 1$ and let H be a fixed d-regular graph for some $d \ge 2$. Let $G_n \sim \mathcal{G}(n,p)$ be the Erdős-Rényi random graph on n vertices with edge probability p.*

(i) If the point $(r^d, h_p(r))$ lies on the convex minorant of the function $x \mapsto h_p(x^{1/d})$ then

$$\lim_{n\to\infty} \frac{1}{\binom{n}{2}} \log \mathbb{P}\left(t(H,G_n) \ge r^{e(H)}\right) = -h_p(r)$$

and furthermore, for every $\varepsilon > 0$ there exists some $C = C(H,\varepsilon,p,r) > 0$ such that for all n,

$$\mathbb{P}\left(\delta_\square(G_n,r) < \varepsilon \;\middle|\; t(H,G_n) \ge r^{e(H)}\right) \ge 1 - e^{-Cn^2}.$$

(ii) If the point $(r^d, h_p(r))$ does not lie on the convex minorant of the function $x \mapsto h_p(x^{1/d})$ then

$$\lim_{n\to\infty} \frac{1}{\binom{n}{2}} \log \mathbb{P}\left(t(H,G_n) \ge r^{e(H)}\right) > -h_p(r)$$

and furthermore, there exist $\varepsilon, C > 0$ such that for all n,

$$\mathbb{P}\left(\inf\left\{\delta_\square(G_n,s) : 0 \le s \le 1\right\} > \varepsilon \;\middle|\; t(H,G_n) \ge r^{e(H)}\right) \ge 1 - e^{-Cn^2}.$$

In particular, when $d = 2$, case (ii) occurs if and only if $p < \left[1 + (r^{-1}-1)^{1/(1-2r)}\right]^{-1}$.

The boundary curves for various values of d are plotted in Fig. 2. It is easy to verify (Lemma A.1) that the rightmost point in the curve for d-regular subgraphs is $(p,r) = \left(\frac{d-1}{d-1+e^{d/(d-1)}}, \frac{d-1}{d}\right)$.

We give an analogous result for large deviations of the spectral radius of an Erdős-Rényi random graph. The phase boundary in this case coincides with that of triangles.

Theorem 1.2. *Fix $0 < p \le r < 1$. Let $G_n \sim \mathcal{G}(n,p)$ be an Erdős-Rényi random graph on n vertices with edge probability p, and let $\lambda_1(G_n)$ denote the largest eigenvalue of its adjacency matrix.*

(i) If $p \ge \left[1 + (r^{-1}-1)^{1/(1-2r)}\right]^{-1}$ then

$$\lim_{n\to\infty} \frac{1}{\binom{n}{2}} \log \mathbb{P}\left(\lambda_1(G_n) \ge nr\right) = -h_p(r)$$

and furthermore, for every $\varepsilon > 0$ there exists some $C = C(\varepsilon, p, r) > 0$ such that for all n,

$$\mathbb{P}\left(\delta_\square(G_n, r) < \varepsilon \mid \lambda_1(G_n) \ge nr\right) \ge 1 - e^{-Cn^2}.$$

(ii) If $p < \left[1 + (r^{-1}-1)^{1/(1-2r)}\right]^{-1}$ then

$$\lim_{n\to\infty} \frac{1}{\binom{n}{2}} \log \mathbb{P}\left(\lambda_1(G_n) \ge nr\right) > -h_p(r)$$

and furthermore, there exist $\varepsilon, C > 0$ such that for all n,

$$\mathbb{P}\left(\inf\left\{\delta_\square(G_n, s) : 0 \le s \le 1\right\} > \varepsilon \;\middle|\; \lambda_1(G_n) \ge nr\right) \ge 1 - e^{-Cn^2}.$$

Both theorems are proved through an analysis of the graphon variational problems rising from the framework of Chatterjee and Varadhan [8]. We show that throughout the replica symmetric region its unique solution is the symmetric one (a consequence of a generalized form of Hölder's inequality), whereas elsewhere one can construct graphons that outperform the symmetric candidate. Note that Theorem 1.2 addresses spectral large deviations, whereas the framework of [8] was tailored for subgraph densities (the recent work [9] broadens it to general random matrix properties with respect to an appropriately defined spectral distance. Here we consider concretely large deviations in the spectral norm of random graphs). Fortunately, the results of [8] easily extend to a wide family of graph parameters with respect to the cut-metric, including the operator norm, the extension of the (normalized) spectral norm to the space of graphons. This generalization is detailed in §2.

1.2. **Exponential random graphs.** We now turn our attention to a different random graph model, the basic setting of which assigns a probability $p_\beta(G)$ to every graph G on n labeled vertices as a function of its edge density $t(K_2, G)$, its triangles density $t(K_3, G)$ and a weight vector $\beta = (\beta_1, \beta_2)$ for these two quantities.[1] Namely, the graph G appears with the following probability:

$$p_\beta(G) = \frac{1}{Z_n} \exp\left(\binom{n}{2}\left(\beta_1 t(K_2, G) + \beta_2 t(K_3, G)\right)\right), \tag{1.2}$$

where Z_n is a normalizing factor (the partition function). When $\beta_2 > 0$ the model favors graphs with more triangles whereas triangles are discouraged for $\beta_2 < 0$. There is a rich literature on both flavors of the model, motivated in part by applications in social networking: the reader is referred to [25, 37, 38, 44] as well as [2, 7] and the references therein.

As shown by Bhamidi, Bresler and Sly [2] and Chatterjee and Diaconis [7], when $\beta_2 \ge 0$ and n is large, a typical random graph drawn from the distribution has a trivial structure — essentially the same one as an Erdős-Rényi random graph with a suitable edge density. This somewhat disappointing conclusion accounts for some of the practical difficulties with statistical parameter estimation for such models. It was further shown in [7] that if we allow β_2 to be sufficiently negative, then the model does behave appreciably differentially from an Erdős-Rényi model. In this part of our work we focus on the case $\beta_2 > 0$, and propose a natural generalization that will enable the model to exhibit a nontrivial structure instead of the previously observed Erdős-Rényi behavior.

[1] The general model allows for an arbitrary (fixed) collection of subgraphs. While the majority of our arguments can be extended to the general setting, we focus on the two-term model for simplicity and clarity.

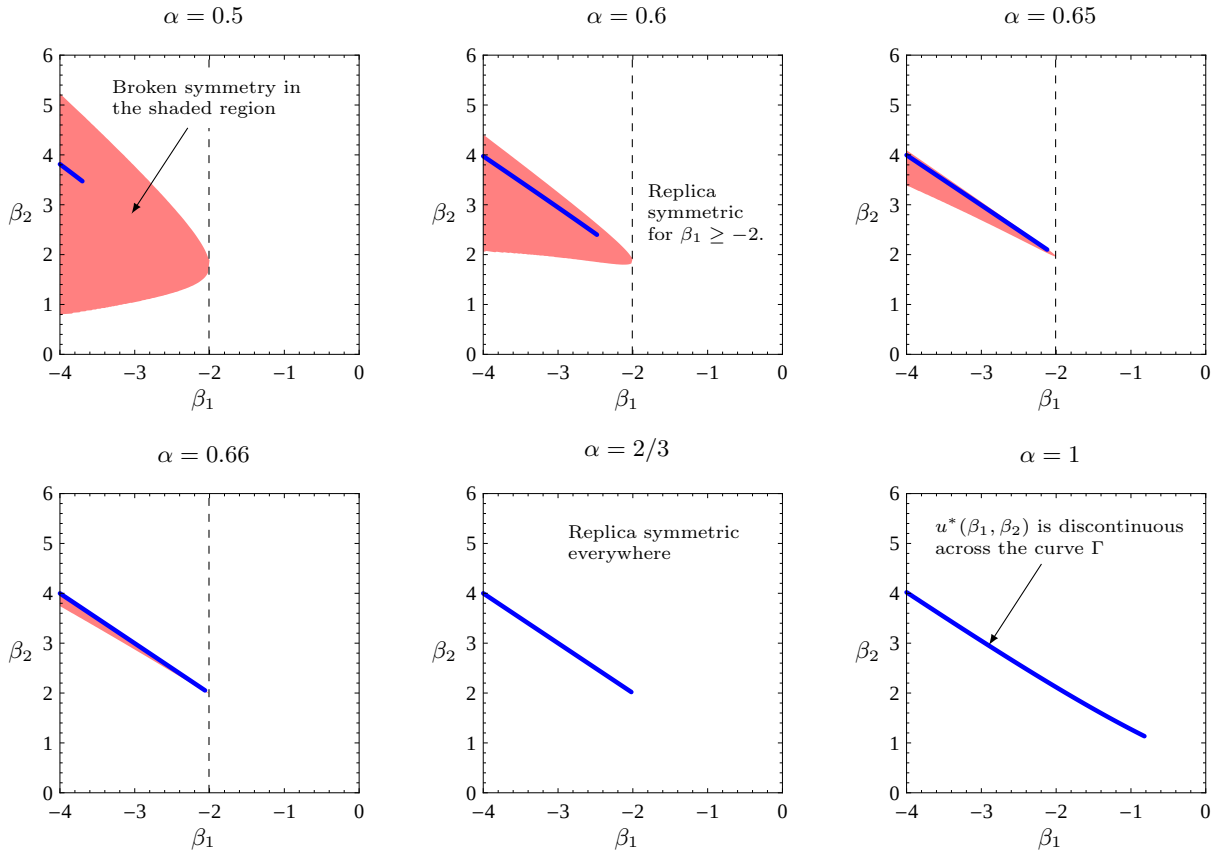

FIGURE 3. The (β_1, β_2)-phase diagrams for the exponential random graph model in (1.3) with $\beta_2 \geq 0$ and various values of $\alpha > 0$, as a special case of Theorem 1.3. When $\alpha < 2/3$, symmetry breaking occurs in the shaded region (at least) and replica symmetry occurs for $\beta_1 \geq -2$. When $\alpha \geq 2/3$, replica symmetry always occurs.

Consider the exponential random graph model which includes an additional exponent $\alpha > 0$ in the exponent of the triangle density term:

$$p_{\alpha,\beta}(G) = \frac{1}{Z_n} \exp\left(\binom{n}{2}(\beta_1 t(K_2, G) + \beta_2 t(K_3, G)^\alpha)\right). \tag{1.3}$$

We will show that this model exhibits a symmetry breaking phase transition even when $\beta_2 > 0$. When $\alpha \geq 2/3$, the generalized model features the Erdős-Rényi behavior, similar to the previously observed case of $\alpha = 1$. However, for $0 < \alpha < 2/3$, there exist regions of values of (β_1, β_2) for which a typical random graph drawn from this distribution has symmetry breaking. As was the case for Theorems 1.1 and 1.2, rather than just triangles we prove this result for any d-regular graph H.

Theorem 1.3. *Let H be a d-regular graph for some fixed $d \geq 2$ and fix $\beta_1 \in \mathbb{R}$ and $\beta_2, \alpha > 0$. Let G_n be an exponential random graph on n labeled vertices with law*

$$p_{\alpha,\beta}(G_n) = \frac{1}{Z_n} \exp\left(\binom{n}{2}(\beta_1 t(K_2, G_n) + \beta_2 t(H, G_n)^\alpha)\right). \tag{1.4}$$

(a) Suppose $\alpha \geq d/e(H)$. There exists a subset $\Gamma = \{(\beta_1, \varphi(\beta_1)) : \beta_1 < \log(e(H)\alpha - 1) - \frac{e(H)\alpha}{e(H)\alpha - 1}\}$ of $\mathbb{R}^2$ for some function $\varphi : \mathbb{R} \to \mathbb{R}$ such that for every $(\beta_1, \beta_2) \in \mathbb{R} \times (0, \infty) \setminus \Gamma$ there exists $0 < u^* < 1$ so that $\delta_\square(G_n, u^*) \to 0$ almost surely as $n \to \infty$, and for every $(\beta_1, \beta_2) \in \Gamma$ there exist $0 < u_1^* < u_2^* < 1$ such that $\min\{\delta_\square(G_n, u_1^*)\ ,\ \delta_\square(G_n, u_2^*)\} \to 0$ almost surely as $n \to \infty$.

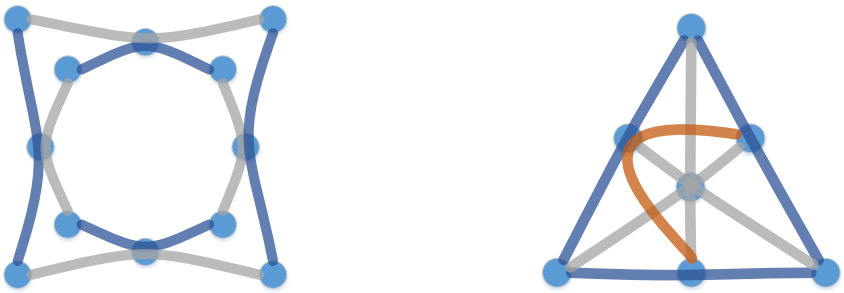

FIGURE 4. Linear d-regular hypergraphs: a cycle ($d = 2$) and the Fano plane ($d = 3$).

(b) Suppose $0 < \alpha < d/e(H)$ *and* $\beta_1 \geq \log(d-1) - d/(d-1)$. *Then there exists* $0 < u^* < 1$ *such that for every* $\varepsilon > 0$ *there is a* $C > 0$ *such that* $\delta_\square(G_n, u^*) \to 0$ *almost surely as* $n \to \infty$.

(c) Suppose $0 < \alpha < d/e(H)$ *and* $\beta_1 < \log(d-1) - d/(d-1)$. *Then there exists an open interval of values* $\beta_2 > 0$ *with the property that there exist* $\varepsilon, C > 0$ *such that for all* n,

$$\mathbb{P}\left(\inf\left\{\delta_\square(G_n, s) : 0 \leq s \leq 1\right\} > \varepsilon\right) \geq 1 - e^{-Cn^2}.$$

Note that, as in the previous theorems, for the replica symmetric phase one can quantify the rate of convergence, e.g., when $\delta_\square(G_n, u^*) \to 0$ almost surely we in fact have that for any $\varepsilon > 0$ there exists some $C > 0$ so that $\mathbb{P}(\delta_\square(G_n, u^*) \leq \varepsilon) \geq 1 - e^{-Cn^2}$ for every n.

1.3. **Linear hypergraphs.** Theorem 5.1 (see §5) extends Theorem 1.1 to the setting of random hypergraphs. A k-uniform hypergraph G consists of a set $V(G)$ of vertices and a set $E(G)$ of hyperedges, where each hyperedges is a k-element subsets of $V(G)$. It is said to be d-regular if every vertex is incident to exactly d edges, and *linear* if every two vertices are incident to at most one common hyperedge (see Fig. 4 for examples of d-regular 3-uniform linear hypergraphs). The random hypergraph $\mathcal{G}^{(k)}(n,p)$ is formed by starting with n vertices and adding k-element subset of the vertices as a hyperedge independently with probability p.

In order to generalize our arguments to large deviations in the density of H, an arbitrary d-regular linear hypergraph, one must first extend the theory developed by Chatterjee and Varadhan [8] to k-uniform hypergraphs. Thanks to the linearity of the hypergraph H there is a simple extension of Szemerédi's regularity lemma to hypergraphs that behaves well with respect to the density of H.

1.4. **Graph homomorphisms.** Alon [1] conjectured in 1991 that the number of independent sets in a d-regular graph H, denoted $i(H)$, satisfies $i(H) \leq i(K_{d,d})^{|V(H)|/(2d)}$, i.e., it is maximized when H is a union of complete bipartite graphs $K_{d,d}$. Kahn [31] verified this when H is bipartite using an ingenious application of the entropy method (specifically, Shearer's inequality). This result was thereafter extended by the second author [47] to all d-regular graphs via an elementary bijection. Using the entropy method of Kahn, Galvin and Tetali [22] extended [31] to graph homomorphisms:

Theorem 1.4 (Galvin and Tetali [22], following Kahn [31])**.** *Let* H *be a simple* d*-regular bipartite graph, and let* G *be a graph, possibly containing loops. We have*

$$\hom(H, G) \leq \hom(K_{d,d}, G)^{|V(H)|/(2d)}. \tag{1.5}$$

Observe that this inequality generalizes the independent set result, since $i(H) = \hom(H, \circlearrowright\!\!\bullet\!\!-\!\!\bullet)$. Previously, all the known proofs of these inequalities relied on entropy techniques. Regarding a more elementary proof, Kahn [31] wrote that "one would think that this simple and natural conjecture... would have a simple and natural proof." As a related aside, in §6 we give a short new entropy-free proof for (1.5) as an immediate consequence of the generalized form of Hölder's inequality.

1.5. **Organization.** In §2 we review graph limits as well as the large deviation principle for random graphs developed by Chatterjee and Varadhan. In §3 we apply the machinery of Chatterjee and Varadhan to prove Theorems 1.1 and 1.2, determining the phase diagram for large deviations of subgraph densities and the largest eigenvalue in $\mathcal{G}(n,p)$, resp. Section 4 focuses on exponential random graphs and gives the proof of Theorem 1.3. In §5 we extend Theorem 1.1 to densities of linear hypergraphs in random hypergraphs. Section 6 contains the short new proof of the inequalities of Kahn and Galvin-Tetali (Theorem 1.4). Finally, in §7 we discuss some open problems.

2. Graph limits and the framework of Chatterjee-Varadhan

The theory of Chatterjee-Varadhan reduces the problem of determining the rate function for large deviations in dense random graphs to solving a prescribed variational problem in graph limits. We will review the required definitions from graph limit theory and then describe the results of [8] in the broader context of "nice" graph parameters, generalizing subgraph counts.

2.1. **Graph limits.** Let $\mathcal{W}$ be the space of all bounded measurable functions $[0,1]^2 \to \mathbb{R}$ that are symmetric (i.e., $f(x,y) = f(y,x)$ for all $x, y \in [0,1]$). Further let $\mathcal{W}_0$ denote all symmetric measurable functions $[0,1]^2 \to [0,1]$, referred to as *graphons* or *kernels* (occasionally these are called *labeled graphons* since later we consider equivalence classes of $\mathcal{W}_0$ modulo measure-preserving bijections on $[0,1]$). Lovász and Szegedy [34] showed that the elements of $\mathcal{W}_0$ are limit objects for sequences of graphs w.r.t. all subgraph densities. Specifically, for any $f \in \mathcal{W}$ and any simple graph H with $V(H) = [m] = \{1, 2, \ldots, m\}$, define

$$t(H,f) = \int_{[0,1]^m} \prod_{(i,j)\in E(H)} f(x_i, x_j)\, dx_1 \cdots dx_m\,.$$

(We shall omit the domain of integration when there is no ambiguity.) Any simple graph G on vertices $\{1, 2, \ldots, n\}$ can be represented as a graphon f^G by

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

In particular, $t(H,G) = t(H, f^G)$ for any two graphs H and G.

A sequence of graphs $\{G_n\}_{n\geq 1}$ is said to converge if the sequence of subgraph densities $t(H, G_n)$ converges for every fixed finite simple graph H. It was shown in [34] that for any such convergent graph sequence there is a limit object $f \in \mathcal{W}_0$ such that $t(H, G_n) \to t(H, f)$ for every fixed H. Conversely, any $f \in \mathcal{W}_0$ arises as a limit of a convergent graph sequence.

We will consider several norms on $\mathcal{W}$, beginning with the standard L^p norm

$$\|f\|_p := \left(\int |f(x,y)|^p \; dxdy \right)^{1/p}.$$

Each $f \in \mathcal{W}$ can be viewed as a Hilbert-Schmidt kernel operator T_f on $L^2([0,1])$ by

$$(T_f u)(x) = \int_0^1 f(x,y) u(y)\, dy \quad \text{for any } u \in L^2([0,1]),$$

and the operator norm for f is then given by

$$\|f\|_{\mathrm{op}} := \min\left\{ c \geq 0 : \|T_f u\|_2 \leq c \|u\|_2 \text{ for all } u \in L^2([0,1]) \right\}.$$

As T_f is self-adjoint, its operator norm is equal to its spectral radius (see, e.g., [41, Thm. 12.25]).

The cut-norm on $\mathcal{W}$ is given by

$$\|f\|_\square := \sup_{S,T\subset[0,1]} \left|\int_{S\times T} f(x,y)\,dxdy\right|$$
$$= \sup_{u,v\colon [0,1]\to[0,1]} \left|\int_{[0,1]^2} f(x,y)u(x)v(y)\,dxdy\right|,$$

where the two suprema are equal since one only needs to consider $\{0,1\}$-valued u and v by the linearity of the integral. The second definition is useful for giving upper bounds using the cut-norm.

For any measure-preserving map $\sigma\colon [0,1]\to[0,1]$ and $f\in\mathcal{W}$, define $f^\sigma\in\mathcal{W}$ to be given by $f^\sigma(x,y)=f(\sigma(x),\sigma(y))$. We define the cut-distance on $\mathcal{W}$ by

$$\delta_\square(f,g)=\inf_\sigma \|f-g^\sigma\|_\square$$

where σ ranges over all measure-preserving bijections on $[0,1]$. For the case of graphons this gives a pseudometric space $(\mathcal{W}_0,\delta_\square)$, which can be turned into a genuine metric space $\widetilde{\mathcal{W}}_0$, equipped with the same cut-metric, by taking a quotient w.r.t. the equivalence relation $f\sim g$ iff $\delta_\square(f,g)=0$. The following theorem can be viewed as a topological interpretation of Szemerédi's regularity lemma.

Theorem 2.1 ([34])**.** *The metric space $(\widetilde{\mathcal{W}}_0,\delta_\square)$ is compact.*

It was shown in [4] that a sequence of graphs $\{G_n\}_{n\ge1}$ converges in the sense of subgraph densities if and only if the sequence of graphons $f^{G_n}\in\mathcal{W}_0$ converge in $\mathcal{W}_0$ w.r.t. the cut-distance. Equivalently, the topology on $\mathcal{W}_0$ induced by $\delta_\square$ is the weakest topology that is continuous w.r.t. the subgraph densities $t(H,\cdot)$ for every H. This underlines one of the reasons making the cut-metric topology a natural choice for the space of graphons.

2.2. **Graph parameters in the cut-metric topology.** We shall focus on graph parameters whose extensions to the space of graphons behave well under the cut-metric topology. One example of such a graph parameter is the subgraph density $t(H,\cdot)$ for an arbitrary finite simple graph H, which was defined in §2.1 directly on the full space of graphons such that $t(H,G)=t(H,f^G)$ for any graph G. A crucial feature of $t(H,\cdot)$ is being continuous w.r.t. the cut-metric ([4]), related to the existence of a "counting lemma" in the regularity lemma literature. This will be a prerequisite for applying the large deviations machinery of Chatterjee and Varadhan.

Definition 2.2. A *nice graph parameter* is a function $\tau\colon\widetilde{\mathcal{W}}_0\to\mathbb{R}$ that is continuous w.r.t. $\delta_\square$ and such that every local extrema of τ w.r.t. $L^\infty(\mathcal{W}_0)$ is necessarily a global extrema. We extend such a function τ to $\mathcal{W}_0$ in the obvious manner and further write $\tau(G)=\tau(f^G)$ for any graph G.

Another way to state the local extrema condition is that if $f\in\mathcal{W}_0$ is not a global maximum (resp. minimum) of τ then for every $\varepsilon>0$ there exists $g\in\mathcal{W}_0$ with $\|f-g\|_\infty<\varepsilon$ and $\tau(g)>\tau(f)$ (resp. $\tau(g)<\tau(f)$). This technical condition will later imply the continuity of the rate function.

Since the metric space $(\widetilde{\mathcal{W}}_0,\delta_\square)$ is compact and path-connected, the image of τ as above is a finite closed interval. In particular, its maximum is attained by a non-empty closed subset of $\widetilde{\mathcal{W}}_0$.

Example 2.3 (Subgraph density)**.** For any fixed finite simple graph H, the subgraph density $t(H,\cdot)$ is a nice graph parameter. As mentioned above, $t(H,\cdot)$ is continuous w.r.t. $\delta_\square$ and in fact the map $f\mapsto t(H,f)$ is Lipschitz-continuous in the metric $\delta_\square$ ([4, Theorem 3.7]). The local extrema condition is fulfilled since the function $g^+=\min\{f+\varepsilon,1\}$ satisfies $t(H,g^+)>t(H,f)$ unless $t(H,f)=1$, and similarly $g^-=(1-\varepsilon)f$ has $t(H,g^-)<t(H,f)$ unless $t(H,f)=0$.

The next two examples are of graph parameters that do not meet the criteria of Definition 2.2.

Example 2.4 (Frobenius norm)**.** Let τ be the function that maps a weighted graph G on n vertices with adjacency matrix A_G to the normalized Frobenius norm $\|A_G\|_{\mathrm{F}}/n$. Then τ is discontinuous in $(\widetilde{\mathcal{W}}_0, \delta_\square)$ and therefore is not nice. Indeed, fix $0 < p < 1$ and let $G_n \sim \mathcal{G}(n,p)$. The sequence G_n is known to converge in $\widetilde{\mathcal{W}}_0$ almost surely to the constant graphon p (see [4, Theorem 4.5]), whose Frobenius norm is p. In contrast to this limiting value, $\|A_{G_n}\|_{\mathrm{F}}/n = \left\|f^{G_n}\right\|_2 \to \sqrt{p}$ almost surely.

Example 2.5 (Max-cut)**.** The function $\tau(G) = \mathrm{maxcut}(G)/|V(G)|^2$ extends to $\widetilde{\mathcal{W}}_0$ via

$$\tau(f) = \sup_{U \subset [0,1]} \int_{U \times ([0,1] \setminus U)} f(x,y)\, dxdy\,.$$

Despite being continuous w.r.t. $\delta_\square$ as well as monotone, the max-cut density is not nice. The continuity of τ follows from the fact $|\tau(f) - \tau(g)| \le \delta_\square(f,g)$, as any cut for either f or g translates into a cut for the other with value differing by at most $\delta_\square(f,g)$. To see that τ does not satisfy the local maxima condition, let f be the graphon defined to be 1 on $[0,\frac{1}{3}] \times [\frac{1}{3},1] \cup [\frac{1}{3},1] \times [0,\frac{1}{3}]$ and 0 elsewhere. We have $\tau(f) = \frac{2}{9}$, induced by $U = [0,\frac{1}{3}]$. This is not the global maximum for τ, which is $\frac{1}{4}$ for the constant function 1. However, we claim that f is a local maximizer of τ with respect to the L^∞-topology. By monotonicity, showing that $\tau(g) = \frac{2}{9}$ for the function $g = \min\left\{f + \frac{1}{2}, 1\right\}$ will imply that $\tau(f_0) \le \tau(g) = \tau(f)$ for any f_0 with $\|f_0 - f\|_\infty \le \frac{1}{2}$. Indeed, if $\mu(U \cap [0,\frac{1}{3}]) = a$ and $\mu(U \cap [\frac{1}{3},1]) = b$, where μ is Lebesgue measure, then the cut density induced by U for g is equal to $\frac{1}{2}a(\frac{1}{3}-a) + \frac{1}{2}b(\frac{2}{3}-b) + a(\frac{2}{3}-b) + b(\frac{1}{3}-a)$, which is maximized at $(a,b) = (0,\frac{2}{3})$ and $(a,b) = (\frac{1}{3},0)$.

We will see in §3.2 (see Lemma 3.6) that, in contrast to the Frobenius norm, the spectral norm (the focus of Theorem 1.2) does behave well under the cut-metric topology, thus qualifying for an application of the large deviation theory of Chatterjee and Varadhan.

2.3. **Large deviations for random graphs.** A random graph $G_n \sim \mathcal{G}(n,p)$ corresponds to the random point $f^{G_n} \in \widetilde{\mathcal{W}}_0$, thus $\mathcal{G}(n,p)$ induces a probability distribution $\mathbb{P}_{n,p}$ on $\widetilde{\mathcal{W}}_0$ supported on a finite set of points (graphs on n vertices). Recalling (1.1), we extend $h_p\colon [0,1] \to \mathbb{R}$ to $\mathcal{W}_0$ by

$$h_p(f) := \int_{[0,1]^2} h_p(f(x,y))\, dxdy \qquad \text{for any } f \in \mathcal{W}_0.$$

An important feature of h_p is that it is a convex function on $[0,1]$ and hence lower-semicontinuous on $\widetilde{\mathcal{W}}_0$ with respect to the cut-metric topology ([8, Lem. 2.1]).

Using Szemerédi's regularity lemma as well as tools from graph limits, Chatterjee and Varadhan [8] proved the following large deviation principle for random graphs.

Theorem 2.6 ([8])**.** *For each fixed $p \in (0,1)$, the sequence $\mathbb{P}_{n,p}$ obeys a large deviation principle in the space $(\widetilde{\mathcal{W}}_0, \delta_\square)$ with rate function h_p. Explicitly, for any closed set $F \subseteq \widetilde{\mathcal{W}}_0$,*

$$\limsup_{n\to\infty} \frac{1}{\binom{n}{2}} \log \mathbb{P}_{n,p}(F) \le -\inf_{f \in F} h_p(f)\,,$$

and for any open $U \subseteq \widetilde{\mathcal{W}}_0$,

$$\liminf_{n\to\infty} \frac{1}{\binom{n}{2}} \log \mathbb{P}_{n,p}(U) \ge -\inf_{f \in U} h_p(f)\,.$$

The machinery developed by Chatterjee and Varadhan reduces the problem determining the large deviation rate function for dense random graphs to solving a variational problem on graphons. For any nice graph parameter $\tau\colon \mathcal{W}_0 \to \mathbb{R}$, any $p \in [0,1]$, and any $t \in \tau(\mathcal{W}_0)$, let

$$\phi_\tau(p,t) := \inf\left\{h_p(f)\colon f \in \mathcal{W}_0\ ,\ \tau(f) \ge t\right\}\,. \tag{2.1}$$

Since h_p is lower-semicontinuous on $\widetilde{\mathcal{W}}_0$, the infimum in (2.1) is always attained.

The following result was stated in [8, Thm 4.1 and Prop. 4.2] for the graph parameter $\tau = t(K_3, \cdot)$. We state its generalization to the class of nice graph parameters as per Definition 2.2.

Theorem 2.7 (Variational problem)**.** *Let $\tau\colon \mathcal{W}_0 \to \mathbb{R}$ be a nice graph parameter and $G_n \sim \mathcal{G}(n,p)$. Fix $p \in (0,1)$ and $t < \max(\tau)$. Let $\phi_\tau(p,t)$ denote the solution to (2.1). Then*

$$\lim_{n\to\infty} \frac{1}{\binom{n}{2}} \log \mathbb{P}\left(\tau(G_n) \ge t\right) = -\phi_\tau(p,t)\,. \tag{2.2}$$

Let F^* be the set of minimizers for (2.1) and let $\widetilde{F}^*$ be its image in $\widetilde{\mathcal{W}}_0$. Then $\widetilde{F}^*$ is a non-empty compact subset of $\widetilde{\mathcal{W}}_0$. Moreover, for each $\varepsilon > 0$ there exists $C = C(\tau,\varepsilon,p,t) > 0$ so that for all n,

$$\mathbb{P}\left(\delta_\square(G_n, \widetilde{F}^*) < \varepsilon \;\middle|\; \tau(G_n) \ge t\right) \ge 1 - e^{-Cn^2}\,.$$

In particular, if $\widetilde{F}^* = \{f^*\}$ for some $f^* \in \widetilde{\mathcal{W}}_0$ then the conditional distribution of G_n given the event $\tau(G_n) \ge t$ converges to the point mass at f^* as $n \to \infty$.

Observe that by considering $-\tau$ (also a nice graph parameter) one obtains the same result for lower tail deviations. The intuition behind the second part of Theorem 2.7 is that the probability that $\delta_\square(G_n, \widetilde{F}^*) \ge \varepsilon$ conditioned on $\tau(G_n) \ge t$ can be again computed using Theorem 2.6 and shown to be exponentially smaller than that of the probability of the event $\tau(G_n) \ge t$.

The proof of Theorem 2.7 is a straightforward extension of the arguments of [8, Thm. 4.1] to nice graph parameters. A technical condition needed to complete the proof of Theorem 2.7 is given by the following lemma, key to which are the attributes of a nice graph parameter.

Lemma 2.8. *Let $\tau\colon \mathcal{W}_0 \to \mathbb{R}$ be a nice graph parameter. For any $p \in (0,1)$, the map $t \mapsto \phi_\tau(p,t)$ is continuous on $\tau(\mathcal{W}_0)$.*

Proof. Fix $0 < p < 1$. By definition, $t \mapsto \phi_\tau(p,t)$ is non-decreasing. To prove left-continuity, consider $a \in \mathbb{R}$. Since h_p is lower-semicontinuous on $\mathcal{W}_0$, the set $\{f : h_p(f) \le a\}$ is closed in $\mathcal{W}_0$, thus also compact by the compactness of $\mathcal{W}_0$. Since τ is continuous on $\mathcal{W}_0$, the set $\{\tau(f) : h_p(f) \le a\}$ is compact and in particular closed. Note that the latter is precisely the pre-image of $(-\infty, a]$ under the inverse of $t \mapsto \phi_\tau(p,t)$. As this pre-image is closed for any $a \in \mathbb{R}$, left-continuity follows.

In order to prove right-continuity it suffices to show that for every $t_0 \in \tau(\mathcal{W}_0)$ with $t_0 < \max(\tau)$ and every $\varepsilon > 0$ there exists some $f \in \mathcal{W}_0$ such that $h_p(f) < \phi_\tau(p,t_0) + \varepsilon$ and $\tau(f) > t_0$. Indeed, since $h_p\colon [0,1] \to \mathbb{R}$ is uniformly continuous (for any fixed p), there exists $\varepsilon' > 0$ so that $|h_p(x) - h_p(y)| < \varepsilon$ whenever $|x-y| < \varepsilon'$. Let f_0 be the minimizer of the variational problem (2.1) for $t = t_0$. The local extrema condition in Definition 2.2 now implies that there is some $f \in \mathcal{W}_0$ with $\tau(f) > t_0$ and $\|f - f_0\|_\infty < \varepsilon'$. Hence, $h_p(f) < h_p(f_0) + \varepsilon = \phi_\tau(p,t_0) + \varepsilon$, as desired. ∎

3. The phase diagram for subgraph densities and the spectral radius

3.1. Subgraph density.

In this section we prove Theorem 1.1, characterizing the phase diagram of upper tails (replica symmetry vs. symmetry breaking) of the density of a fixed d-regular subgraph.

Establishing the replica symmetric phase will hinge on a generalized form of Hölder's inequality which appeared in [18] (see also [19, Lemma 5.1]). We include its short proof for completeness.

Theorem 3.1 (Generalized Hölder's inequality)**.** *Let $\mu_1, \dots, \mu_n$ be measures on $\Omega_1, \dots, \Omega_n$, resp., and let $\mu = \prod_{i=1}^n \mu_i$ be the product measure on $\Omega = \prod_{i=1}^n \Omega_i$. Let $A_1, \dots, A_m$ be nonempty subsets of $[n] = \{1, \dots, n\}$ and write $\Omega_A = \prod_{\ell\in A} \Omega_\ell$ and $\mu_A = \prod_{\ell\in A} \mu_\ell$. Let $f_i \in L^{p_i}(\Omega_{A_i}, \mu_{A_i})$ with $p_i \ge 1$ for each $i \in [m]$ and suppose in addition that $\sum_{i:\ell\in A_i} (1/p_i) \le 1$ for each $\ell \in [n]$. Then*

$$\int \prod_{i=1}^m |f_i| \; d\mu \le \prod_{i=1}^m \left(\int |f_i|^{p_i} \; d\mu_{A_i} \right)^{1/p_i}.$$

In particular, when $p_i = d$ for every $i \in [m]$ we have $\int \prod_{i=1}^m |f_i| \; d\mu \le \prod \left(\int |f_i|^d \; d\mu_{A_i} \right)^{1/d}$.

Proof. The proof carries by induction on n with the trivial base case of $n = 0$. By Fubini's theorem,

$$\int_\Omega \prod_{i=1}^m |f_i| \; d\mu = \int_\Omega \prod_{i:n\in A_i} |f_i| \prod_{i:n\notin A_i} |f_i| \; d\mu = \int_{\Omega_{[n-1]}} \left(\int_{\Omega_n} \prod_{i:n\in A_i} |f_i| \; d\mu_n \right) \prod_{i:n\notin A_i} |f_i| \; d\mu_{[n-1]} .$$

where the argument of each f_i is the restriction of $x \in \Omega$ to the coordinates of A_i, denoted by x_{A_i}. Hölder's inequality (along with Jensen's inequality if $\sum_{i:n\in A_i}(1/p_i)$ is less than 1) implies that

$$\int_{\Omega_n} \prod_{i:n\in A_i} |f_i| \; d\mu_n \le \prod_{i:n\in A_i} \left(\int_{\Omega_n} |f_i|^{p_i} \; d\mu_n \right)^{1/p_i} ,$$

thus for each i with $n \in A_i$ we can let $f_i^* : \Omega_{[n-1]} \to \mathbb{R}$ denote the averaging map $\left(\int_{\Omega_n} |f_i|^{p_i} \; d\mu_n\right)^{1/p_i}$ and obtain that

$$\int_\Omega \prod_{i=1}^m |f_i| \; d\mu \le \int_{\Omega_{[n-1]}} \prod_{i:n\in A_i} f_i^* \prod_{i:n\notin A_i} |f_i| \; d\mu_{[n-1]} .$$

Now, the functions f_i^* correspond to $A_i^* = A_i \setminus \{n\}$ and therefore, thanks to the assumption that $\sum_{i:\ell\in A_i}(1/p_i) \le 1$ for each $\ell \in [n-1]$ we can apply the induction hypothesis and infer that

$$\begin{aligned} \int_\Omega \prod_{i=1}^m |f_i| \; d\mu &\le \prod_{i:n\in A_i} \left(\int_{\Omega_{[n-1]}} (f_i^*)^{p_i} \; d\mu_{[n-1]} \right)^{1/p_i} \prod_{i:n\notin A_i} \left(\int_{\Omega_{[n-1]}} |f_i|^{p_i} \; d\mu_{[n-1]} \right)^{1/p_i} \\ &= \prod_{i=1}^m \left(\int_\Omega |f_i|^{p_i} \; d\mu \right)^{1/p_i} , \end{aligned}$$

as required. ∎

It is helpful to compare Theorem 3.1 with the standard Hölder's inequality for the case where $p_i = d$ for all i. A direct application of Hölder's inequality produces the inequality $\|\prod f_i\|_1 \le \prod_i \|f_i\|_m$, whereas Theorem 3.1 exploits the extra assumption that $\#\{i : \ell \in A_i\} \le d$ for all $\ell \in [n]$ and gives the stronger inequality $\|\prod f_i\|_1 \le \prod_i \|f_i\|_d$. For instance,

$$\left(\int f_1(x,y) f_2(y,z) f_3(x,z) \; dxdydz \right)^2 \le \prod_i \left(\int f_i(x,y)^2 \; dxdy \right) .$$

Placing this in the context of subgraph densities, as an immediate corollary of Theorem 3.1 (in the special case that $p_i = d$ for all i) we have the following inequality.

Corollary 3.2. *Let H be a graph whose maximum degree is at most d, and let $f \in \mathcal{W}$. Then*

$$t(H, f) \le \|f\|_d^{e(H)} .$$

Recall that Theorem 2.7 reduces the problem of finding the phase boundary to determining whether the constant function r is a solution for the variational problem of minimizing $h_p(f)$ over $f \in \mathcal{W}_0$ subject to $t(H, f) \ge r^{e(H)}$, where H is some fixed graph. In light of the above corollary, it is important to estimate $h_p(f)$ for functions $f \in \mathcal{W}_0$ with $\|f\|_d = r$, as addressed next.

Lemma 3.3. *Let $0 < p < 1$ and let $f \in \mathcal{W}_0$. Suppose that $d \ge 1$ and $0 < r < 1$ are such that the point $(r^d, h_p(r))$ lies on the convex minorant of $x \mapsto h_p(x^{1/d})$. If in addition either*

(a) $p < r < 1$ and $\|f\|_d \ge r$, or
(b) $0 < r < p$ and $\|f\|_d \le r$,

then $h_p(f) \ge h_p(r)$, with equality occurring if and only if $f \equiv r$.

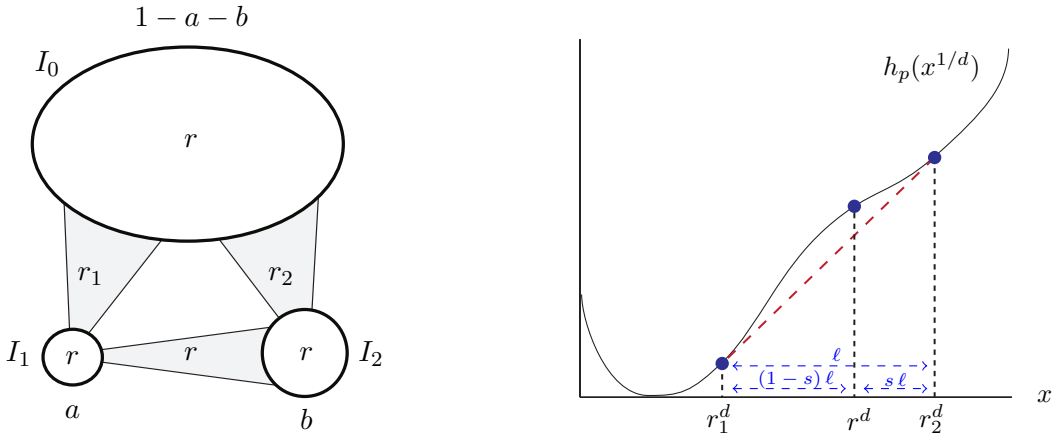

FIGURE 5. The construction in Lemma 3.4.

Proof. Let $\psi(x) = h_p(x^{1/d})$ and let $\hat{\psi}$ be the convex minorant of ψ. By Jensen's inequality,

$$h_p(f) = \int \psi(f(x,y)^d)\, dxdy \geq \int \hat{\psi}(f(x,y)^d)\, dxdy \geq \hat{\psi}\left(\int f(x,y)^d\, dxdy\right) = \hat{\psi}\left(\|f\|_d^d\right).$$

Since $h_p(x)$ (and hence $\hat{\psi}$) is decreasing along $[0,p]$ and increasing along $[p,1]$ (see §A), under either of the assumptions in Parts (a),(b) we have $\hat{\psi}(\|f\|_d^d) \geq \hat{\psi}(r^d) = \psi(r^d) = h_p(r)$; thus, $h_p(f) \geq h_p(r)$. Moreover, as $\hat{\psi}$ is not linear in any neighborhood of r^d, equality can occur if and only if $f \equiv r$. ∎

The final element needed for the proof of Theorem 1.1 is a construction that outperforms the constant graphon in the symmetry breaking regime. This is achieved by the following lemma.

Lemma 3.4. *Let H be a d-regular graph. Fix $0 < p \leq r < 1$ so that $(r^d, h_p(r))$ is not on the convex minorant of $x \mapsto h_p(x^{1/d})$. Then there exists $f \in \mathcal{W}_0$ with $t(H,f) > r^{e(H)}$ and $h_p(f) < h_p(r)$.*

Proof. Since $(r^d, h_p(r))$ does not lie on the convex minorant of $x \mapsto h_p(x^{1/d})$, there necessarily exist $0 \leq r_1 < r < r_2 \leq 1$ such that the point $(r^d, h_p(r))$ lies strictly above the line segment joining $(r_1^d, h_p(r_1))$ and $(r_2^d, h_p(r_2))$. Letting s be such that

$$r^d = sr_1^d + (1-s)r_2^d,$$

we therefore have

$$sh_p(r_1) + (1-s)h_p(r_2) < h_p(r). \tag{3.1}$$

Let $\varepsilon > 0$ and define

$$a = s\varepsilon^2, \qquad b = (1-s)\varepsilon^2 + \varepsilon^3,$$
$$I_0 = [a, 1-b], \qquad I_1 = [0,a], \qquad I_2 = [1-b,1],$$

noting that for $a < 1-b$ for any sufficiently small ε. Define $f_\varepsilon \in \mathcal{W}_0$ by

$$f_\varepsilon(x,y) = \begin{cases} r_1 & \text{if } (x,y) \in (I_0 \times I_1) \cup (I_1 \times I_0), \\ r_2 & \text{if } (x,y) \in (I_0 \times I_2) \cup (I_2 \times I_0), \\ r & \text{otherwise}. \end{cases}$$

(See Fig. 5 for an illustration of this construction.) We claim that

$$t(H, f_\varepsilon) - r^{e(H)} = v(H)\left(a(r_1^d - r^d) + b(r_2^d - r^d)\right) r^{e(H)-d} + O(\varepsilon^4). \tag{3.2}$$

Indeed, the only embeddings of H that have values different from $r^{e(H)}$ are those where at least one vertex of H is mapped to $I_1 \cup I_2$. Since a and b are both $O(\varepsilon^2)$, in order to compute $t(H,f_\varepsilon) - r^{e(H)}$

up to an $O(\varepsilon^4)$ error we need only consider embeddings of H where precisely one vertex gets mapped to $I_1 \cup I_2$. Denote this vertex of H by u, and observe that if u is mapped to I_1 then the contribution to $t(H, f_\varepsilon) - r^{e(H)}$ is $(r_1^d - r^d)r^{e(H)-d}$ since H is d-regular. Similarly, if u is mapped to I_2 then the contribution is $(r_2^d - r^d)r^{e(H)-d}$. Putting everything together yields (3.2).

By definition of a, b and s we have

$$a(r_1^d - r^d) + b(r_2^d - r^d) = s\varepsilon^2(r_1^d - r^d) + ((1-s)\varepsilon^2 + \varepsilon^3)(r_2^d - r^d) = \varepsilon^3(r_2^d - r^d)\,.$$

Recalling that $r_2 > r$ and plugging the last equation in (3.2) it now follows that $t(H, f_\varepsilon) > r^{e(H)}$ for any sufficiently small $\varepsilon > 0$. At the same time, we also have

$$\begin{aligned} h_p(f_\varepsilon) - h_p(r) &= 2a(1-a-b)(h_p(r_1) - h_p(r)) + 2b(1-a-b)(h_p(r_2) - h_p(r)) \\ &= 2(1-a-b)\left(ah_p(r_1) + bh_p(r_2) - (a+b)h_p(r)\right) \\ &= 2(1-a-b)\varepsilon^2\left(sh_p(r_1) + (1-s)h_p(r_2) - h_p(r) + (h_p(r_2) - h_p(r))\varepsilon\right)\,. \end{aligned}$$

Revisiting (3.1) we conclude that $h_p(f_\varepsilon) < h_p(r)$ for any sufficiently small $\varepsilon > 0$. ∎

We now have all the ingredients needed for establishing the phase diagram of upper tail deviations for subgraph densities.

Proof of Theorem 1.1. For Part (i), by applying Theorem 2.7 to the graph parameter $t(H, \cdot)$, it suffices to show that the constant function r is the unique element $f \in \mathcal{W}_0$ minimizing $h_p(f)$ subject to $t(H, F) \geq r^{e(H)}$. Indeed, by Corollary 3.2, $t(H, F) \geq r^{e(H)}$ implies that $\|f\|_d \geq r$, and by Lemma 3.3 Part (a), $h_p(f) \geq h_p(r)$ with equality if and only if f is the constant function r.

To prove Part (ii), let $F^* \subset \widetilde{\mathcal{W}}_0$ be the set of minimizers for the variational problem (2.1) with the graph parameter $t(H, \cdot)$. Then F^* does not contain the constant function r by Lemma 3.4, nor does it contain any constant function of value $r' \neq r$ (when $r' > r$ one has $h_p(r') > h_p(r)$, whereas if $r' < r$ then $t(H, f) < r^{e(H)}$). Let $\widetilde{\mathcal{C}} \subset \widetilde{\mathcal{W}}_0$ be the set of constant graphons. Since F^* and $\widetilde{\mathcal{C}}$ are disjoint and both are compact, $\delta_\square(F^*, \widetilde{\mathcal{C}}) > 0$. The desired result follows from applying Theorem 2.7 with $\varepsilon = \delta_\square(F^*, \widetilde{\mathcal{C}})/2$.

When $d = 2$, the phase boundary is explicitly given by Lemma A.2, thus concluding the proof. ∎

One can also ask what the phase diagram is for lower tail deviations of subgraph densities. We next show that for certain bipartite graphs there is replica symmetry everywhere for lower tails.

A beautiful conjecture of Erdős and Simonovits [43] and Sidorenko [42] (from here on referred to as *Sidorenko's conjecture*) states that every bipartite graph H satisfies $t(H, G) \geq t(K_2, G)^{e(H)}$ for every graph G. The conjecture was verified for various graphs H (e.g., trees, even cycles [42], hypercubes [26], bipartite graphs with one vertex complete to the other part [11]). As it turns out, for any such graph H the lower tail deviations are always replica symmetric (no phase transition).

Proposition 3.5. *Fix $0 < r \leq p < 1$, let $G_n \sim \mathcal{G}(n, p)$ be the Erdős-Rényi random graph and let H be a fixed bipartite graph for which Sidorenko's conjecture holds. Then*

$$\lim_{n\to\infty} \frac{1}{\binom{n}{2}} \log \mathbb{P}\left(t(H, G_n) \leq r^{e(H)}\right) = -h_p(r)$$

and furthermore, for every $\varepsilon > 0$ there exists some constant $C = C(H, \varepsilon, p, r) > 0$ such that

$$\mathbb{P}\left(\delta_\square(G_n, r) < \varepsilon \;\middle|\; t(H, G_n) \leq r^{e(H)}\right) \geq 1 - e^{-Cn^2}\,.$$

Proof. Applying Theorem 2.7 with $-t(H, \cdot)$, it suffices to show that the constant function r is the unique element $f \in \mathcal{W}_0$ minimizing $h_p(f)$ subject to $t(H, f) \leq r^{e(H)}$. Since H satisfies Sidorenko's conjecture, $t(H, f) \geq \|f\|_1^{e(H)}$ for all $f \in \mathcal{W}_0$. Thus, if $t(H, f) \leq r^{e(H)}$ then $\|f\|_1 \leq r$, and so by Lemma 3.3 Part (b) (applied to the case $d = 1$, noting that then $h_p(x^{1/d})$ is itself convex) we have $h_p(f) \geq h_p(r)$ with equality if and only if f is the constant function r. ∎

3.2. **Largest eigenvalue.** In this section we prove Theorem 1.2, addressing the phase boundary for large deviations in the spectral norm. The proof will follow a similar route as the previous section, yet first we must show that the spectral norm is a nice graph parameter.

For a graph G on n vertices, let $\lambda_1(G)$ be the largest eigenvalue of its adjacency matrix A_G. Since A_G is symmetric, $\lambda_1(G) = \|A_G\|_{\mathrm{op}}$, and therefore over $\mathcal{W}$ we have $\|f^G\|_{\mathrm{op}} \geq \lambda_1(G)/n$. It is easy to verify (see Lemma 3.6 below) that in fact $\|f^G\|_{\mathrm{op}} = \lambda_1(G)/n$, thus the operator norm on $\mathcal{W}$ is the graphon extension of the (normalized) largest eigenvalue. Furthermore, as we show below, $\|\cdot\|_{\mathrm{op}}$ is (uniformly) continuous w.r.t. the cut-metric, and the local extrema condition in Definition 2.2 is satisfied as well, thus $\|\cdot\|_{\mathrm{op}}$ is a nice graph parameter.

Lemma 3.6. *The function $\|\cdot\|_{\mathrm{op}}$ is a continuous extension of the normalized graph spectral norm, i.e., $\lambda_1(G)/n$ for a graph G on n vertices, to $(\widetilde{\mathcal{W}}_0, \delta_\square)$. Moreover, $\|\cdot\|_{\mathrm{op}}$ is a nice graph parameter.*

Proof. We first show that $\|f\|_{\mathrm{op}} = \lambda_1(G)/n$ for any graph G on n vertices. Clearly, the largest eigenvector of A_G, the adjacency matrix of G, can be turned into a step function $u\colon [0,1] \to \mathbb{R}$ such that $T_{f^G} u = (\lambda_1(G)/n)u$, and so $\|f^G\|_{\mathrm{op}} \geq \lambda_1(G)/n$. Conversely, for any $u\colon [0,1] \to \mathbb{R}$ we consider the step function $u_n\colon [0,1] \to \mathbb{R}$ such that for any $1 \leq i \leq n$, on the interval $(\frac{i-1}{n}, \frac{i}{n}]$ it is equal to the average of u over that interval. Let $v \in \mathbb{R}^n$ be the vector of values of u_n. Since f^G is constant in every box $(\frac{i-1}{n}, \frac{i}{n}] \times (\frac{j-1}{n}, \frac{j}{n}]$, we have

$$\|T_f u\|_2 = \|T_f u_n\|_2 = \|A_G v\|_2 / n^2 \leq \lambda_1(G) \|v\|_2 / n^2 = \lambda_1(G) \|u_n\|_2 / n \leq \lambda_1(G) \|u\|_2 / n,$$

where the last inequality is due to convexity. It follows that $\|f\|_{\mathrm{op}} = \lambda_1(G)/n$.

Next, we will argue that

$$\|f\|_{\mathrm{op}}^4 \leq 4 \|f\|_\square \qquad \text{for any symmetric measurable } f\colon [0,1]^2 \to [-1,1]. \tag{3.3}$$

Let $u \in L^2([0,1])$ with $\|u\|_2 = 1$. By Cauchy-Schwarz we can infer that

$$\begin{aligned}
\|T_f u\|_2^4 &= \left(\int \left(\int f(x,y) u(y) \, dy \right)^2 dx \right)^2 = \left(\int f(x,y) f(x,y') u(y) u(y') \, dx dy dy' \right)^2 \\
&\leq \left(\int \left(\int f(x,y) f(x,y') \, dx \right)^2 dy dy' \right) \left(\int u(y)^2 u(y')^2 \, dy dy' \right) . \\
&= \int f(x,y) f(x,y') f(x',y) f(x',y') \, dx dx' dy dy' .
\end{aligned}$$

For any fixed x', y' we can let $v_{y'}(x) = f(x,y')$ and $w_{x'}(y) = f(x',y)$, thus rewriting the above as

$$\int \left(\int f(x,y) v_{y'}(x) w_{x'}(y) \, dx dy \right) f(x',y') \, dx' dy' \leq 4 \|f\|_\square ,$$

with the last inequality justified by the fact that for any $g \in \mathcal{W}$ and $v, w\colon [0,1] \to [-1,1]$ we have $\left| \int g(x,y) v(x) w(y) \, dx dy \right| \leq 4 \|g\|_\square$ by the definition of the cut-norm, thereby establishing Eq. (3.3). (The factor of 4 above was due to splitting v, w into positive and negative parts. Indeed, for $f\colon [0,1]^2 \to [0,1]$ the bound in (3.3) remains valid without the factor of 4 in the right-hand side.)

Consider $f, g \in \widetilde{\mathcal{W}}_0$ and let σ vary over all measure-preserving bijections on $[0,1]$. We then have

$$\left| \|f\|_{\mathrm{op}} - \|g\|_{\mathrm{op}} \right| \leq \inf_\sigma \|f - g^\sigma\|_{\mathrm{op}} \leq \sqrt{2} \inf_\sigma \|f - g^\sigma\|_\square^{1/4} = \sqrt{2} \delta_\square(f,g)^{1/4} ,$$

thus implying that $\|\cdot\|_{\mathrm{op}}$ is uniformly continuous in $(\widetilde{\mathcal{W}}_0, \delta_\square)$.

Finally, we need to verify the local extrema condition in Definition 2.2. Let $f \in \mathcal{W}_0$. There are no local non-global minima since $\|(1-\varepsilon)f\|_{\mathrm{op}} = (1-\varepsilon)\|f\|_{\mathrm{op}} < \|f\|_{\mathrm{op}}$ unless $\|f\|_{\mathrm{op}} = 0$ already. In addition, we claim that $g = \min\{f + \varepsilon, 1\}$ satisfies $\|g\|_{\mathrm{op}} > \|f\|_{\mathrm{op}}$ unless $\|f\|_{\mathrm{op}} = 1$ already. Indeed,

take some $u \in L^2([0,1])$ which is nonzero (a.e.) and satisfies $u \geq 0$ and $T_f u = \|f\|_{\mathrm{op}} u$. It suffices to show that $T_g u > T_f u$ on some subset of $[0,1]$ with positive measure. Let $A = \{x \in [0,1] : u(x) > 0\}$ be the support of u, which by our choice has positive Lebesgue measure $\mu(A) > 0$. If $\mu(A) < 1$ then $T_f u = u = 0$ (a.e.) on $A^c := [0,1] \setminus A$, so that $f = 0$ on $A^c \times A$. Hence, $g = \varepsilon$ on $A^c \times A$ and $T_g u = \varepsilon \|u\|_1 > 0 = T_f u$ on A^c, as desired. Suppose therefore that $\mu(A) = 1$. If $\|f\|_{\mathrm{op}} < 1$ we must have $g > f$ on a subset of positive measure, and hence also $T_g u > T_f u$ on a subset of positive measure, as u has full support. This shows that f cannot be a local maximum unless $\|f\|_{\mathrm{op}} = 1$. ∎

In light of the above lemma, the variational problem under consideration in Theorem 1.2 becomes

$$\inf\{h_p(f) : f \in \mathcal{W}_0, \|f\|_{\mathrm{op}} \geq r\}\,. \tag{3.4}$$

We will need the following straightforward inequality relating the operator norm, $\|\cdot\|_1$ and $\|\cdot\|_2$.

Lemma 3.7. *For every $f \in \mathcal{W}_0$ we have $\|f\|_1 \leq \|f\|_{\mathrm{op}} \leq \|f\|_2$.*

Proof. For the left inequality, observe that since $f \geq 0$ we have that

$$\|f\|_1 = \|T_f \mathbf{1}\|_1 \leq \|T_f \mathbf{1}\|_2 \leq \|f\|_{\mathrm{op}}\,.$$

For the right inequality in the statement of the lemma, let $u\colon [0,1] \to \mathbb{R}$. By Cauchy-Schwarz,

$$\|T_f u\|_2^2 = \int \left(\int f(x,y)u(y)\,dy\right)^2 dx \leq \left(\int f(x,y)^2\,dydx\right)\left(\int u(y)^2\,dy\right) = \|f\|_2^2\,\|u\|_2^2\,,$$

and therefore $\|f\|_{\mathrm{op}} \leq \|f\|_2$, as claimed. ∎

The following lemma is the operator norm analogue of Lemma 3.4, providing a construction that beats the constant graphon in the symmetry breaking regime.

Lemma 3.8. *Let $0 < p \leq r < 1$ be such that $(r^2, h_p(r))$ does not lie on the convex minorant of $x \mapsto h_p(\sqrt{x})$. Then there exists some $f \in \mathcal{W}_0$ with $\|f\|_{\mathrm{op}} > r$ and $h_p(f) < h_p(r)$.*

Proof. Let $\varepsilon > 0$. Let f_ε be the construction from the proof of Lemma 3.4 with $d = 2$, and define the parameters of that construction $r_1, r_2, s, a, b, I_0, I_1, I_2$ as given there. Having already demonstrated that $h_p(f_\varepsilon) < h_p(r)$ for any small enough $\varepsilon > 0$, it remains to show that $\|f_\varepsilon\|_{\mathrm{op}} > r$. To this end, it suffices to exhibit a function $u \in L^2([0,1])$ such that $(T_{f_\varepsilon} u)(x) > ru(x)$ for all $x \in [0,1]$. Let

$$u(x) = \begin{cases} (1-a-b)r_1 & \text{if } x \in I_1\,, \\ (1-a-b)r_2 & \text{if } x \in I_2\,, \\ r & \text{if } x \in I_0\,. \end{cases}$$

Recall that $f_\varepsilon(x,y)$ is r except when $(x,y) \in (I_0 \times I_i) \cup (I_i \times I_0)$ where it is r_i for $i = 1, 2$. It now follows that for every $x \in I_1 = [0,a]$

$$(T_{f_\varepsilon} u)(x) = a(1-a-b)r_1 r + b(1-a-b)r_2 r + (1-a-b)r_1 r > (1-a-b)r_1 r = ru(x)\,.$$

Similarly, for any $x \in I_2 = [1-b,1]$ we have

$$(T_{f_\varepsilon} u)(x) > (1-a-b)r_2 r = ru(x)\,.$$

Finally, if $x \in I_0 = [a, 1-b]$ then

$$(T_{f_\varepsilon} u)(x) = a(1-a-b)r_1^2 + b(1-a-b)r_2^2 + (1-a-b)r^2 = (1-a-b)(r^2 + ar_1^2 + br_2^2)\,.$$

Plugging in the facts that $r^2 = sr_1^2 + (1-s)r_2^2$ while $a = s\varepsilon^2$ and $b = (1-s)\varepsilon^2 + \varepsilon^3$, we get that

$$(T_{f_\varepsilon} u)(x) = (1-\varepsilon^2-\varepsilon^3)(r^2 + r^2\varepsilon^2 + r_2^2\varepsilon^3) = r^2 + (r_2^2 - r^2)\varepsilon^3 + O(\varepsilon^4) > r^2 = ru(x)\,,$$

where the strict inequality is valid for any sufficiently small $\varepsilon > 0$ since $r_2 > r$.

Altogether, $(T_{f_\varepsilon} u)(x) > ru(x)$ for all $x \in [0,1]$ and so $\|f\|_{\mathrm{op}} > r$, as required. ∎

Proof of Theorem 1.2. To prove Part (i), by Theorem 2.7 applied to the graph parameter $\|\cdot\|_{\mathrm{op}}$ it suffices to show that the constant function r is the unique element $f \in \mathcal{W}_0$ minimizing $h_p(f)$ subject to $\|f\|_{\mathrm{op}} \geq r$. Indeed, by Lemma 3.7, $\|f\|_2 \geq \|f\|_{\mathrm{op}} \geq r$. By Lemma 3.3 Part (a) and Lemma A.2 we know that $h_p(f) \geq h_p(r)$, with equality if and only if f is the constant function r.

For Part (ii), similar to the proof of Part (ii) of Theorem 1.1 in §3.1, Lemma 3.8 implies that the set of minimizers of the variational problem (2.1) is disjoint from the set of constant graphons. We then apply Theorem 2.7 to conclude the proof, with the phase boundary given by Lemma A.2. ∎

The behavior of the lower tails deviations in the spectral norm is similar to that of the subgraph densities in Proposition 3.5, where replica symmetry is exhibited everywhere (no phase transition).

Proposition 3.9. *Let $0 < r \leq p < 1$. Let $G_n \sim \mathcal{G}(n,p)$ be the Erdős-Rényi random graph and let $\lambda_1(G_n)$ denote the largest eigenvalue of its adjacency matrix. Then*

$$\lim_{n\to\infty} \frac{1}{\binom{n}{2}} \log \mathbb{P}\left(\lambda_1(G_n) \leq r\right) = -h_p(r)$$

and furthermore, for every $\varepsilon > 0$ there is some $C = C(\varepsilon, p, r) > 0$ such that for all n,

$$\mathbb{P}\left(\delta_\square(G_n, r) < \varepsilon \mid \lambda_1(G_n) \leq nr\right) \geq 1 - e^{-Cn^2}.$$

Proof. Applying Theorem 2.7 with $\tau = -\|\cdot\|_{\mathrm{op}}$, it suffices to show that the constant function r is the unique element $f \in \mathcal{W}_0$ minimizing $h_p(f)$ subject to $\|f\|_{\mathrm{op}} \leq r$. By Lemma 3.7, if $f \in \mathcal{W}_0$ with $\|f\|_{\mathrm{op}} \leq r$ then $\|f\|_1 \leq \|f\|_{\mathrm{op}} \leq r$. It now follows from Lemma 3.3 Part (b) (used with $d = 1$ bearing in mind that $h_p(x)$ is convex) that $h_p(f) \geq h_p(r)$ with equality if and only if $f \equiv r$. ∎

4. Exponential random graph models

Let us review the tools developed by Chatterjee and Diaconis [7] to analyze exponential random graphs. Define

$$h(x) := x \log x + (1-x)\log(1-x) \qquad \text{for } x \in [0,1],$$

and for any graphon $f \in \mathcal{W}_0$ let

$$h(f) := \int_{[0,1]^2} h(f(x,y))\, dxdy.$$

The following result from [7, Thm. 3.1 and Thm. 3.2] reduces the analysis of the exponential random graph model in the large n limit to a variational problem. It was proven with the help of the theory developed by Chatterjee and Varadhan [8] for large deviations in random graphs.

Theorem 4.1 (Chatterjee and Diaconis [7])**.** *Let $\tau\colon \widetilde{\mathcal{W}}_0 \to \mathbb{R}$ be a bounded continuous function. Let $Z_n = \sum_G \exp\left(\binom{n}{2}\tau(G)\right)$ where the sum is taken over all $2^{\binom{n}{2}}$ simple graphs G on n labeled vertices. Let $\psi_n = \binom{n}{2}^{-1} \log Z_n$. Then*

$$\psi := \lim_{n\to\infty} \psi_n = \sup_{f \in \widetilde{\mathcal{W}}_0} \left(\tau(f) - h(f)\right), \tag{4.1}$$

and the set $F^* \subset \widetilde{\mathcal{W}}_0$ of maximizers of this variational problem is nonempty and compact.

Let G_n be a random graph on n vertices drawn from the exponential random graph model defined by τ, i.e., with distribution $Z_n^{-1} \exp\left(\binom{n}{2}\tau(\cdot)\right)$. Then for every $\eta > 0$ there exists $C = C(\tau, \eta) > 0$ such that for all n,

$$\mathbb{P}\left(\delta_\square(G_n, F^*) > \eta)\right) \leq e^{-Cn^2}.$$

We say that the exponential random graph model has *replica symmetry* if the set of maximizers F^* for the variational problem $\tau(f) - h(f)$ contains only constant functions, and we say that it has *replica symmetry breaking* if no constant function is a maximizer. Intuitively, Theorem 4.1 implies that when there is replica symmetry, for large n, the random graph behaves like an Erdős-Rényi random graph (or a mixture of Erdős-Rényi random graphs), while this is not the case when there is broken symmetry. More precisely we have the following result (see [7, Thm. 6.2]).

Corollary 4.2. *Continuing with Theorem 4.1. Let $\widetilde{\mathcal{C}} \subset \widetilde{\mathcal{W}}_0$ be the set of constant functions. If $F^* \cap \widetilde{\mathcal{C}} = \emptyset$ then there exist $C, \varepsilon > 0$ such that for all n,*

$$\mathbb{P}(\delta_\square(G_n, \widetilde{\mathcal{C}}) > \varepsilon) \geq 1 - e^{-Cn^2}.$$

To prove Theorem 1.3 using the above tools, we need to analyze the following variational problem:

$$\sup_{f \in \widetilde{\mathcal{W}}_0} \left(\beta_1 t(K_2, f) + \beta_2 t(H, f)^\alpha - h(f)\right). \tag{4.2}$$

Here is the main result of this section, from which Theorem 1.3 follows by the results above.

Theorem 4.3. *Let H be a d-regular graph ($d \geq 2$) and fix $\beta_1 \in \mathbb{R}$ and $\alpha, \beta_2 > 0$. Let $\mathcal{E}$ denote the corresponding exponential random graph model on n labeled vertices as specified in (1.4).*

(a) If $\alpha \geq d/e(H)$, then $\mathcal{E}$ has replica symmetry. Moreover, there exists a set $\Gamma \subset \mathbb{R}^2$ of the form

$$\Gamma = \{(\beta_1, \varphi(\beta_1)) : \beta_1 < \log(e(H)\alpha - 1) - \tfrac{e(H)\alpha}{e(H)\alpha - 1}\} \subset \mathbb{R}^2 \quad \text{for some function } \varphi : \mathbb{R} \to \mathbb{R}$$

such that when $(\beta_1, \beta_2) \in \mathbb{R} \times (0, \infty) \setminus \Gamma$ the set of maximizers of the variational problem (4.2) is a single constant function, and when $(\beta_1, \beta_2) \in \Gamma$ the set of maximizers consists of exactly two distinct constant functions.

(b) If $0 < \alpha < d/e(H)$ and $\beta_1 \geq \log(d-1) - d/(d-1)$ then $\mathcal{E}$ has replica symmetry. Moreover, the variational problem (4.2) is maximized by a unique constant function.

(c) If $0 < \alpha < d/e(H)$ and $\beta_1 < \log(d-1) - d/(d-1)$ then there exists an open interval of values $\beta_2 > 0$ for which $\mathcal{E}$ has broken symmetry, i.e., the set of maximizers of the variational problem (4.2) does not contain any constant function. Furthermore, this open interval can be taken to be $(\underline{\beta}_2, \overline{\beta}_2)$ with $\underline{\beta}_2 = \underline{u}^{1-e(H)\alpha} h_p'(\underline{u})/(e(H)\alpha)$ and $\overline{\beta}_2 = \overline{u}^{1-e(H)\alpha} h_p'(\overline{u})/(e(H)\alpha)$, where $(\underline{u}^d, h_p(\underline{u}))$ and $(\overline{u}^d, h_p(\overline{u}))$ are the two points where the lower common tangent of $x \mapsto h_p(x^{1/d})$ touches the curve for $p = 1/(1 + e^{-\beta_1})$.

When restricted only to constant functions in $\widetilde{\mathcal{W}}_0$, the variational problem (4.2) becomes the one-dimensional optimization problem

$$\sup_{0 \leq u \leq 1} \left(\beta_1 u + \beta_2 u^{e(H)\alpha} - h(u)\right). \tag{4.3}$$

Note that for both (4.2) and (4.3) the supremum is in fact a maximum due to compactness. Let u^* be the maximizer for (4.3). When there is replica symmetry, the maximum values attained in (4.2) and (4.3) are equal, and the exponential random graph behaves like an Erdős-Rényi random graph with edge density u^*. It is possible that there are two distinct maximizers u^*, in which case the model behaves like a (possibly trivial) distribution over two separate Erdős-Rényi models. In the work of Chatterjee and Diaconis [7], where the $\alpha = 1$ case was considered, it was shown that u^* as a function of (β_1, β_2) experiences a discontinuity across a curve in the parameter space. Radin and Yin [39] later showed that (when $\alpha = 1$) the limiting free energy density ψ from (4.1), as a function in the parameter space (β_1, β_2), is analytic except on a first order phase transition curve ending in a critical point with second order phase transition. (See Fig. 3 in §1 for a plot of the location of the discontinuity in the (β_1, β_2)-phase diagram.)

Here we focus less on the discontinuity of u^* and more on symmetry breaking. Nevertheless we shall start our analysis by giving a simple geometric interpretation of the discontinuity of u^*.

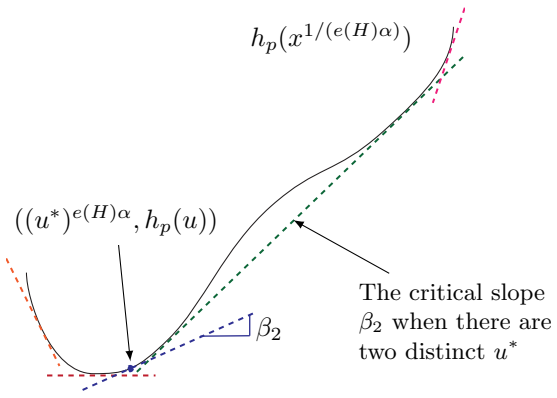

FIGURE 6. Discontinuity in the symmetric solution u^* due to the geometry of $x \mapsto h_p(x^{1/(e(H)\alpha)})$.

By definition,

$$h_p(x) = h(x) - x \log \frac{p}{1-p} - \log(1-p)\,.$$

Setting

$$p = \frac{1}{1+e^{-\beta_1}}\,,$$

so that $\beta_1 = \log \frac{p}{1-p}$, we absorb the linear term in (4.2) and (4.3) into the entropy term, at which point these two optimization problems respectively become

$$\sup_{f \in \widetilde{\mathcal{W}}_0} \left(\beta_2 t(H,f)^\alpha - h_p(f) - \log(1-p)\right) \tag{4.4}$$

and

$$\sup_{0 \le u \le 1} \left(\beta_2 u^{e(H)\alpha} - h_p(u) - \log(1-p)\right). \tag{4.5}$$

By a change of variables $u = x^{1/(e(H)\alpha)}$ in (4.5) we get the equivalent optimization problem

$$\sup_{0 \le x \le 1} \left(\beta_2 x - h_p(x^{1/(e(H)\alpha)}) - \log(1-p)\right). \tag{4.6}$$

Observe that $x = x^*$ maximizes (4.6) iff the tangent to the curve defined by $x \mapsto h_p(x^{1/(e(H)\alpha)})$ at $x = x^*$ has slope β_2 and lies below the curve. Thanks to Lemma A.1 from the appendix, we know that $x \mapsto h_p(x^{1/\gamma})$ is convex if $0 < \gamma \le 1$ or if

$$\gamma > 1 \quad \text{and} \quad p \ge p_0(\gamma) := \frac{\gamma - 1}{\gamma - 1 + e^{\gamma/(\gamma-1)}}\,. \tag{4.7}$$

Otherwise, $h_p(x^{1/\gamma})$ has exactly two inflection points to the right of $x = p^\gamma$, so that the curve starts convex, becomes concave, and finally turns convex again. In addition, $h_p(x^{1/\gamma})$ has an infinite slope at both endpoints. For any $\beta_2 \in \mathbb{R}$, there is a unique lower tangent of slope β_2 to the curve $h_p(x^{1/(e(H)\alpha)})$, touching the curve at $x = x^* = (u^*)^{e(H)\alpha}$. As β_2 varies, u^* increases continuously with β_2, except in the situation where the curve of $h_p(x^{1/(e(H)\alpha)})$ is not convex and β_2 is the slope of the unique lower tangent that touches the curve at two points. In that case, (4.5) is optimized at two distinct values of u, denoted by $0 < \underline{u} < \overline{u} < 1$, and as β_2 increases through this critical point, u^* jumps over the interval $(\underline{u}, \overline{u})$ corresponding to the part of the curve lying above the convex minorant, then increases continuously afterwards. When $h_p(x^{1/(e(H)\alpha)})$ is convex, this jump does not occur. See Fig. 6 for an illustration of this process (the function $h_p(x^{1/(e(H)\alpha)})$ is plotted not-to-scale in order to highlight its features). The uniqueness of u^* is stated below as a lemma.

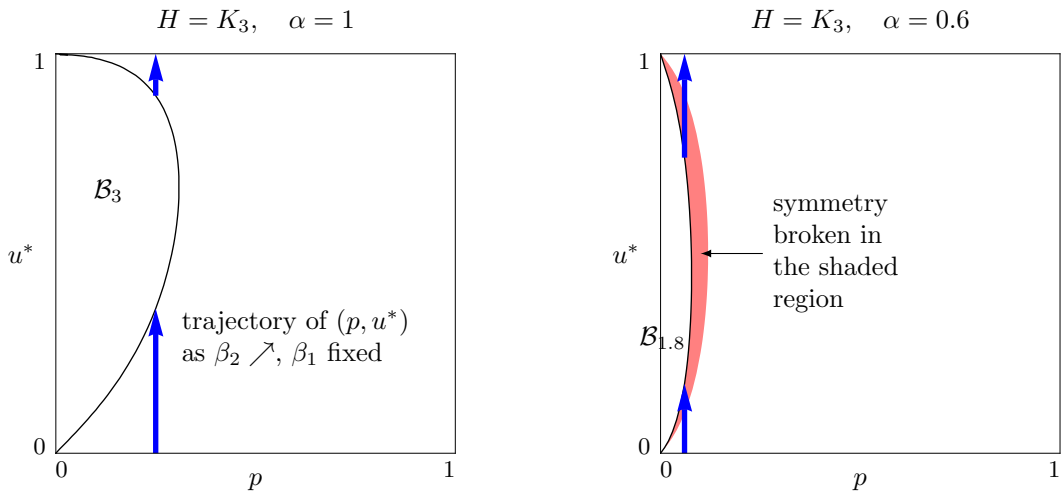

FIGURE 7. Discontinuity in the symmetric solution u^* as reflected in the (p, u)-phase diagram.

Lemma 4.4. *If $0 < e(H)\alpha \le 1$ or if $e(H)\alpha > 1$ and $p \ge p_0(e(H)\alpha)$ as defined in (4.7), then the optimization problem (4.3) is maximized at a unique value of u. Otherwise, (4.3) is maximized at a unique u except for a single value of β_2, where the maximum is attained at two distinct u's.*

We can also represent this jump of u^* in the (p, u)-phase diagram as follows. For each $\gamma > 0$, consider the region $\mathcal{B}_\gamma \subset [0,1]^2$ containing all points (p, u) such that $(u^\gamma, h_p(u))$ does not lie on the convex minorant of $x \mapsto h_p(x^{1/\gamma})$. When $\gamma < 1$ the region $\mathcal{B}_\gamma$ is empty, but otherwise it is nonempty. The geometric argument in the previous paragraph shows that u can never appear as a maximizer to (4.5) if $(p, u) \in \mathcal{B}_{e(H)\alpha}$, but all other values of u can. Thus, for a fixed $p = 1/(1 + e^{-\beta_1})$, as β_2 increases from $-\infty$ to ∞, the point (p, u^*) moves up continuously from 0 in the (p, u)-phase diagram, and jumps over $\mathcal{B}_{e(H)\alpha}$ as it reaches it. Thereafter it resumes moving up until hitting 1. This process is illustrated on the left of Fig. 7 when $e(H)\alpha = 3$, e.g., when $H = K_3$ and $\alpha = 1$.

Turning to large deviations, we know from Lemma 3.4 that there is broken symmetry for the density of copies of a d-regular graph H whenever we are in $\mathcal{B}_d$. It turns out that the same is true for the corresponding exponential random graph model given in (1.4). As we just saw though, one must remove $\mathcal{B}_{e(H)\alpha}$ from the possible solution space for (p, u^*). Whenever $\gamma < \gamma'$ we have $\mathcal{B}_\gamma \subset \mathcal{B}_{\gamma'}$ (see Lemma A.5), and consequently, if $e(H)\alpha \ge d$ then $\mathcal{B}_{e(H)\alpha}$ covers $\mathcal{B}_d$ and so the entire symmetry breaking phase is removed, leaving replica symmetry everywhere. This agrees with the results of Chatterjee and Diaconis [7] for the case $\alpha = 1$. However, when $e(H)\alpha < d$ it is possible to have $(p, u^*) \in \mathcal{B}_d$, in which case the construction from Lemma 3.4 breaks the symmetry. This is shown on the right of Fig. 7 for the case $d = 2$ and $e(H)\alpha = 1.8$, e.g., for $H = K_3$ and $\alpha = 0.6$.

Proof of Theorem 4.3. The proof of Part (a) is essentially the same as the proof of Theorem 4.1 in the work of Chatterjee and Diaconis [7], except now we use the generalized Hölder's inequality (Theorem 3.1) instead of the usual Hölder's inequality.

Suppose that $\alpha \ge d/e(H)$. We first need to show that in this case the only maximizers for the variational problem (4.4) are constant functions. Applying Corollary 3.2, for any $f \in \widetilde{\mathcal{W}}_0$ we have

$$\beta_2 t(H, f)^\alpha - h_p(f) \le \beta_2 \|f\|_d^{e(H)\alpha} - h_p(f) \le \beta_2 \|f\|_{e(H)\alpha}^{e(H)\alpha} - h_p(f)\,,$$

where the last inequality used the assumption on α. This in turn is equal to

$$\int (\beta_2 f(x,y)^{e(H)\alpha} - h_p(f(x,y)))\, dxdy \le \sup_{0 \le u \le 1} \beta_2 u^{e(H)\alpha} - h_p(u)\,,$$

showing that the $\beta_2 t(H,f) - h_p(f)$ is indeed maximized at constant functions. Furthermore, when $\beta_2 u^{e(H)\alpha} - h_p(u)$ is maximized at a unique u^*, then equality holds in place of the above inequalities only for the constant function $f = u^*$. An additional argument is needed to treat the case when $\beta_2 u^{e(H)\alpha} - h_p(u)$ is maximized at two distinct values. Either by checking the equality conditions in the proof of Theorem 3.1, or by referring to [18], we know that equality in Corollary 3.2 occurs if and only if $f(x,y) = g(x)g(y)$ for some function $g\colon [0,1] \to [0,\infty)$. It can then be easily checked that equality in the above sequence of inequalities can only occur when f is a constant function. The value of this constant u^* is given by the optimization problem (4.5) and its the uniqueness is addressed in Lemma 4.4.

We now turn to prove Part (b). Since $\beta_1 \ge \log(d-1) - d/(d-1)$,

$$p = \frac{1}{1+e^{-\beta_1}} \ge \frac{1}{1+e^{-\log(d-1)+d/(d-1)}} = \frac{d-1}{d-1+e^{d/(d-1)}} = p_0(d)\,.$$

By Lemma A.1, $x \mapsto h_p(x^{1/d})$ is convex for this value of p. Hence, $h_p(f) \ge h_p(\|f\|_d)$ by Jensen's inequality with equality if and only if f is a constant function. Since $t(H,f) \le \|f\|_d$ by Corollary 3.2,

$$\beta_2 t(H,f)^\alpha - h_p(f) \le \beta_2 \|f\|_d^{e(H)\alpha} - h_p(\|f\|_d) \le \sup_{0\le u\le 1} (\beta_2 u^{e(H)\alpha} - h_p(u))\,,$$

with equality iff f is the constant function equal to u^*, the unique maximizer of $\beta_2 u^{e(H)\alpha} - h_p(u)$. The uniqueness of u^* follows from Lemma 4.4 together with noting that when $e(H)\alpha > 1$ we have $p \ge p_0(d) > p_0(e(H)\alpha)$ as $d > e(H)\alpha$ and $p_0(\cdot)$ is increasing in $[1,\infty)$.

It remains to prove Part (c). We have $0 < p < p_0(d)$. Let $0 < \underline{u} < \overline{u} < 1$ be such that the lower common tangent to $x \mapsto h_p(x^{1/d})$ touches the curve at $x = \underline{u}^d$ and $\overline{u}^d$. Since $e(H)\alpha < d$, Lemma A.5 implies that the points $(\underline{u}^{e(H)\alpha}, h_p(\underline{u}))$ and $(\overline{u}^{e(H)\alpha}, h_p(\overline{u}))$ both lie on the convex minorant of $x \mapsto h_p(x^{1/(e(H)\alpha)})$) and do not lie on the common lower tangent (if there is one). Let $\underline{\beta}_2$ and $\overline{\beta}_2$ be as in the theorem statement, observing that these are the values of the derivative of $h_p(x^{1/(e(H)\alpha)})$ at $x = \underline{u}^{e(H)\alpha}$ and $\overline{u}^{e(H)\alpha}$, respectively. Then for any $\beta \in (\underline{\beta}_2, \overline{\beta}_2)$, using the slope interpretation of β_2 given in the discussion proceeding this proof, we see that the optimization problem (4.5) is maximized for some $u^* \in (\underline{u}, \overline{u})$. At the same time, by Lemma 3.4 there exists some $f \in \mathcal{W}_0$ such that $t(H,f) > (u^*)^{e(H)}$ and $h_p(f) < h_p(u^*)$. It follows that

$$\sup_{f\in\widetilde{\mathcal{W}}_0} (\beta_2 t(H,f)^\alpha - h_p(f)) > \beta_2 (u^*)^{e(H)\alpha} - h_p(u^*) = \sup_{0\le u\le 1} (\beta_2 u^{e(H)\alpha} - h_p(u))\,,$$

and hence $\beta_2 t(H,f)^\alpha - h_p(f)$ is not maximized at any constant function. ∎

5. Densities of linear hypergraphs in random hypergraphs

In this section, we extend our results to densities of linear hypergraphs in random hypergraphs. Homomorphisms and densities are defined as in graphs. Similarly, for any $r \in [0,1]$ let

$$\delta_\square(G,r) := \sup_{A_1,\dots,A_k\subset V(G)} \frac{1}{|V(G)|^k} \big| e_G(A_1,\dots,A_k) - r\,|A_1|\cdots|A_k|\,\big|$$

where $e_G(A_1,\dots,A_k)$ is the number of (ordered) hyperedges of the form $(a_1,\dots,a_k) \in A_1\times\cdots\times A_k$. The main result in this section is the following:

Theorem 5.1. *Fix $d,k \ge 2$ and $0 < p \le r < 1$. Let H be a d-regular k-uniform linear hypergraph. Let $G_n \sim \mathcal{G}^{(k)}(n,p)$ be the random k-uniform hypergraph on n vertices with hyperedge probability p.*

(i) If the point $(r^d, h_p(r))$ lies on the convex minorant of the function $x \mapsto h_p(x^{1/d})$ then

$$\lim_{n\to\infty} \frac{1}{\binom{n}{k}} \log \mathbb{P}\left(t(H,G_n) \ge r^{e(H)}\right) = -h_p(r)$$

and furthermore, for every $\varepsilon > 0$ there exists some constant $C = C(H, \varepsilon, p, r) > 0$ such that

$$\mathbb{P}\left(\delta_\square(G_n, r) < \varepsilon \;\middle|\; t(H, G_n) \geq r^{e(H)}\right) \geq 1 - e^{-Cn^k}.$$

(ii) If the point $(r^d, h_p(r))$ does not lie on the convex minorant of the function $x \mapsto h_p(x^{1/d})$ then

$$\lim_{n\to\infty} \frac{1}{\binom{n}{k}} \log \mathbb{P}\left(t(H, G_n) \geq r^{e(H)}\right) > -h_p(r)$$

and furthermore, there exist $\varepsilon, C > 0$ such that

$$\mathbb{P}\left(\inf\left\{\delta_\square(G_n, s) : 0 \leq s \leq 1\right\} > \varepsilon \;\middle|\; t(H, G_n) \geq r^{e(H)}\right) \geq 1 - e^{-Cn^k}.$$

In particular, when $d = 2$, case (ii) occurs if and only if $p < \left[1 + (r^{-1} - 1)^{1/(1-2r)}\right]^{-1}$.

Let $k \geq 2$ be an integer and let $\mathcal{W}^{(k)}$ be the space of all bounded measurable functions $[0,1]^k \to \mathbb{R}$ that are symmetric (i.e., $f(x_1, x_2, \ldots, x_k) = f(x_{\pi(1)}, x_{\pi(2)}, \ldots, x_{\pi(k)})$ for any permutation π of $[k] = \{1, 2, \ldots, k\}$). Let $\mathcal{W}_0^{(k)}$ denote all symmetric measurable functions $[0,1]^k \to [0,1]$. Every k-uniform hypergraph G corresponds to a point $f^G \in \mathcal{W}_0^{(k)}$ similar to the case for graphs. As before, we can endow $\mathcal{W}$ with usual L^p-norm and in addition have the following cut norm:

$$\|f\|_\square := \sup_{S_1,\ldots,S_k \subset [0,1]} \int_{S_1 \times \cdots \times S_k} f(x_1, \ldots, x_k)\, dx_1 \cdots dx_k.$$

This gives rise to the cut distance: for any $f, g \in \mathcal{W}_0^{(k)}$,

$$\delta_\square(f, g) := \inf_\sigma \|f - g^\sigma\|_\square$$

where σ ranges over all measure-preserving bijections on $[0,1]$, and $g^\sigma \in \mathcal{W}_0^{(k)}$ is defined by $g^\sigma(x_1, \ldots, x_k) = g(\sigma(x_1), \ldots, \sigma(x_k))$. Let $\widetilde{\mathcal{W}}_0^{(k)}$ be the metric space formed by taking equivalences of points in $\mathcal{W}_0^{(k)}$ with zero cut-distance.

The space $\mathcal{W}_0^{(k)}$ is a straightforward generalization of the space $\mathcal{W}_0$ of graphons. Unfortunately, it does not fully capture the richness of the structure of hypergraphs. This notion is closely related to some initial attempts at generalizing Szemerédi's regularity lemma to hypergraphs (e.g., [10]). The main issue is that while the regularity lemma generalizes easily to this setting, there is no corresponding counting lemma for embedding a fixed hypergraph H unless H is linear (recall that a hypergraph is linear if every pair of vertices is contained in at most one hyperedge). The difficulty in extending the results to general H is related to the intricacies of hypergraph regularity (see, e.g., Gowers [24] and Nagle, Rödl, Schacht, and Skokan [36, 40], as well as the recent progress in this direction by Elek and Szegedy [16, 17]). Here we restrict ourselves to the basic setting above which suffices for controlling densities of linear hypergraphs.

For any $f \in \mathcal{W}^{(k)}$ and any k-uniform hypergraph H, write $V(H) = [m]$ and define

$$t(H, f) = \int_{[0,1]^k} \prod_{\{i_1,\ldots,i_k\} \in E(H)} f(x_{i_1}, \ldots, x_{i_k})\, dx_1 \cdots dx_m.$$

The Chatterjee-Varadhan theory can be generalized to derive rate functions for large deviations of H-counts, where H is a fixed linear hypergraph. We outline the modifications and omit the complete details, as the changes required in the original proofs are mostly straightforward.

We start with a statement generalizing the weak regularity lemma of Frieze and Kannan [20]. The analytic form of this statement for graphs can be found in Lovász and Szegedy [35, Lem. 3.1].

Theorem 5.2. *For every $\varepsilon > 0$ there exists some $M(\varepsilon) > 0$ such that for every $f \in \mathcal{W}_0^{(k)}$ there exist some $m \le M(\varepsilon)$ and some $g \in \mathcal{W}_0^{(k)}$ with $\delta_\square(f,g) \le \varepsilon$, and such that g is constant in each box $(\frac{i_1-1}{m}, \frac{i_1}{m}] \times \cdots \times (\frac{i_k-1}{m}, \frac{i_k}{m}]$.*

Using Theorem 5.2, the proof of Lovász and Szegedy [35, Thm. 5.1] can be modified to give the following topological interpretation of this result.

Theorem 5.3. *For any integer $k \ge 2$, the metric space $(\widetilde{\mathcal{W}}_0^{(k)}, \delta_\square)$ is compact.*

Theorems 5.2 and 5.3 allow us to generalize the framework of Chatterjee and Varadhan to $(\widetilde{\mathcal{W}}_0^{(k)}, \delta_\square)$. The random hypergraph graph $\mathcal{G}^{(k)}(n,p)$ corresponds to a random point $f^{\mathcal{G}^{(k)}(n,p)} \in \widetilde{\mathcal{W}}^{(k)}$, and therefore it induces a probability distribution $\mathbb{P}_{n,p}$ on $\widetilde{\mathcal{W}}^{(k)}$ supported on a finite set of points corresponding to hypergraphs on n vertices.

Theorem 5.4. *For each fixed $p \in (0,1)$, the sequence $\mathbb{P}_{n,p}$ obeys a large deviation principle in the space $(\widetilde{\mathcal{W}}_0^{(k)}, \delta_\square)$ with rate function h_p. Explicitly, for any closed set $F \subseteq \widetilde{\mathcal{W}}_0^{(k)}$,*

$$\limsup_{n\to\infty} \frac{1}{\binom{n}{k}} \log \mathbb{P}_{n,p}(F) \le -\inf_{f\in F} h_p(f)\,,$$

and for any open set $U \subseteq \widetilde{\mathcal{W}}_0^{(k)}$,

$$\liminf_{n\to\infty} \frac{1}{\binom{n}{k}} \log \mathbb{P}_{n,p}(U) \ge -\inf_{f\in U} h_p(f)\,.$$

To derive large deviation results for subgraph densities in random graphs, it was crucial that the subgraph densities $t(H,\cdot)$ behaved continuously with respect to the cut topology. The next result implies that the same is true when H is a linear hypergraph. The proof is a straightforward generalization of the proof for graphs (see [4, Thm. 3.7]).

Theorem 5.5. *Let H be a k-uniform linear hypergraph. Then for any $f, g \in \mathcal{W}_0^{(k)}$,*

$$|t(H,f) - t(H,g)| \le e(H)\delta_\square(f,g)\,.$$

The rate function for large deviations in H-counts is then determined by the following variational problem.

Theorem 5.6. *Let H be a k-uniform linear hypergraph and let $G_n \sim \mathcal{G}^{(k)}(n,p)$ be the random k-uniform hypergraph on n vertices with hyperedge probability p. For any fixed $p, r \in (0,1)$,*

$$\lim_{n\to\infty} \frac{1}{\binom{n}{k}} \log \mathbb{P}\left(t(H,G_n) \ge r^{e(H)}\right) = -\inf\left\{h_p(f)\colon f \in \mathcal{W}_0^{(k)}, t(H,f) \ge r^{e(H)}\right\}. \tag{5.1}$$

Let F^* be the set of minimizers for (5.1) and let $\widetilde{F}^*$ be its image in $\widetilde{\mathcal{W}}_0^{(k)}$. Then $\widetilde{F}^*$ is a non-empty compact set. Moreover, for each $\varepsilon > 0$ there exists some $C(H,\varepsilon,p,r) > 0$ so that for any n

$$\mathbb{P}\left(\delta_\square(G_n, \widetilde{F}^*) \ge \varepsilon \;\middle|\; t(H,G_n) \ge r^{e(H)}\right) \le e^{-Cn^k}.$$

In particular, if $\widetilde{F}^* = \{f^*\}$ for some $f^* \in \widetilde{\mathcal{W}}_0^{(k)}$ then the conditional distribution of G_n given the event $t(H,G_n) \ge r^{e(H)}$ converges to the point mass at f^* as $n \to \infty$.

We now turn to study the variational problem (5.1) towards the proof of Theorem 5.1. The following inequality is an immediate consequence of Theorem 3.1.

Lemma 5.7. *Let H be a k-uniform hypergraph with maximum degree at most d, and let $f \in \mathcal{W}_0^{(k)}$. Then $t(H,f) \le \|f\|_d^{e(H)}$.*

The following lemma mirrors Lemma 3.4 for proving the symmetry breaking phase.

Lemma 5.8. *Let H be linear d-regular k-uniform hypergraph. Let $0 < p \le r < 1$ be such that $(r^d, h_p(r))$ does not lie on the convex minorant of $x \mapsto h_p(x^{1/d})$. Then there exists $f \in \mathcal{W}_0^{(k)}$ with $t(H,f) > r^{e(H)}$ and $h_p(f) < h_p(r)$.*

The proof of Lemma 5.8 is essentially the same as the proof of Lemma 3.4, with the following modification. One needs to adjust f_ε into $f_\varepsilon = r + (r_1 - r)\mathbb{1}_A + (r_2 - r)\mathbb{1}_B$, where $A \subset [0,1]^k$ is the union of the box $[0,a] \times [a,1-b]^{k-1}$ along with the $k-1$ other boxes formed by permuting the coordinates. Similarly B is the union of $[1-b,1] \times [a,1-b]^{k-1}$ and its coordinate permutations.

Proof of Theorem 5.1. Our starting point is an application of Theorem 5.6. Suppose $f \in \mathcal{W}_0^{(k)}$ satisfies $t(H,f) \ge r^{e(H)}$. By Lemma 5.7, $\|f\|_d \ge r$. For Part (i) of the theorem, Lemma 3.3 Part (a) (which is also valid for $\mathcal{W}_0^{(k)}$) implies that $h_p(f) \ge h_p(r)$ with equality if and only if f is the constant function r, so the variational problem on the right-hand side of (5.1) has the constant function r as the unique minimizer. For Part (ii) of the theorem, Lemma 5.8 implies that the constant function r is not in the set of minimizers of the variational problem (5.1), and the conclusion follows analogously to the proof of Part (ii) of Theorem 1.1. ∎

6. Graph homomorphism inequalities

Our goal in this section is to present a short new proof of Theorem 1.4, stating that for any graph G allowing loops and any d-regular bipartite graph H one has

$$\hom(H,G) \le \hom(K_{d,d},G)^{|V(H)|/(2d)}\,. \tag{6.1}$$

(This inequality is tight when H is a disjoint union of copies of the complete bipartite graph $K_{d,d}$.) Generalized to graphons, the inequality states that for any d-regular bipartite graph H and $f \in \mathcal{W}_0$

$$t(H,f) \le t(K_{d,d},f)^{|V(H)|/(2d)} \tag{6.2}$$

(the more general formulation follows from (6.1) via a standard limiting argument, e.g., see [22]). As mentioned in the introduction, all previously known proofs of this inequality involved entropy. In contrast, the following proof does not rely on entropy or limiting arguments and instead is an immediate consequence of the generalized Hölder's inequality (Theorem 3.1).

Proof of Theorem 1.4. Label the vertices on the left-bipartition of H by $[n] = \{1,\dots,n\}$, and let $A_i \subset [n]$ be the neighborhood of the i-th vertex on the right-bipartition of H. Define

$$g(x_1,\dots,x_d) := \int f(x_1,y)f(x_2,y)\cdots f(x_d,y)\,dy \qquad \text{for any } x_1,\dots,x_d \in [0,1]\,,$$

and write $g(x_A)$ for $x \in [0,1]^n$ and $A = \{i_1,\dots,i_d\}$ to denote $g(x_{i_1},\dots,x_{i_d})$. With this notation,

$$t(H,f) = \int \prod_{j=1}^{n}\prod_{i\in A_j} f(x_i,y_j)\,dx_1\cdots dx_n dy_1\cdots dy_n = \int_{[0,1]^n} g(x_{A_1})\cdots g(x_{A_n})\,dx \le \|g\|_d^n$$

by the generalized Hölder's inequality (Theorem 3.1). The result follows from noting that

$$\|g\|_d^d = \int \prod_{i=1}^{d}\prod_{j=1}^{d} f(x_i,y_j)\,dx_1\cdots dx_d dy_1\cdots dy_d = t(K_{d,d},f)\,. \qquad ∎$$

Remark. A natural question to ask is whether Theorem 1.4 can be extended to all d-regular graphs H, as in the case for independent sets [47]. Unfortunately, the answer to this question is negative. A simple counter-example is $H = K_3$ and f being the graphon corresponding to the 2×2 identity matrix. The second author [48] extended Theorem 1.4 to non-bipartite H for certain families of $f \in \mathcal{W}_0$, e.g., $\{0,1\}$-valued graphons that are non-decreasing in both coordinates. Galvin [21] conjectured that if G is a graph allowing loops and H is a simple d-regular graph then $\hom(H,G)$ is at most the maximum of $\hom(K_{d,d},G)^{|V(H)|/(2d)}$ and $\hom(K_{d+1},G)^{|V(H)|/(d+1)}$.

7. Open problems

It is natural to ask for extensions of the Chatterjee-Varadhan [8] large deviations theory to sparse Erdős-Rényi random graphs (i.e., $\mathcal{G}(n,p)$ where $p(n) \to 0$ as $n \to \infty$) or to densities of general (not necessarily linear) hypergraphs in random hypergraphs. These may require extensions of Szemerédi's regularity lemma to sparse graphs (see [33, 23, 12]) and hypergraphs [24, 36, 40]. Even for $\mathcal{G}(n,p)$ with fixed p, various problems remain open, several of which we highlight below.

Minimizers of the variational problem. It was pointed out by Chatterjee and Varadhan [8] that no solutions of the variational problem in Theorem 2.7 are known anywhere in the symmetry breaking phase. This remains the case. In fact, there is not a single point (p,r) in the symmetry breaking phase where we can even compute the large deviation rate. It would be interesting to see whether the minimizers are always 2-step graphons, i.e., graphons which are constant on each of $[0,w]^2$, $\big([0,w]\times(w,1]\big)\cup\big((w,1]\times[0,w]\big)$ and $[w,1]^2$ for some $w\in(0,1)$.

Phase boundary for non-regular graphs. In this paper, we identified the replica symmetric phase for upper tail deviations in the densities of d-regular graphs. The phase boundary for non-regular graphs remains unknown. We suspect that a modification of the construction in Lemma 3.4 can be used to establish the symmetry breaking phase for some (perhaps all) subgraph density deviations. However, at present we do not have matching boundaries for the replica symmetric phase.

Lower-tail phase transition. Proposition 3.5 shows that if a bipartite graph H satisfies Sidorenko's conjecture then there is replica symmetry everywhere in the lower tail deviation of H-densities. It is an open question whether all bipartite graphs satisfy Sidorenko's conjecture, although it is possible that Proposition 3.5 can be proved without the full resolution of Sidorenko's conjecture.

When H is not bipartite, the following argument shows that there exists symmetry breaking in the lower tail, at least for certain values of (p,r). Let f be the graphon taking the value 0 on $[0,\frac12]^2\cup[\frac12,1]^2$ and the value p elsewhere, so that $t(H,f)=0$ and $h_p(f)=h_p(0)/2$. Let $r_0\in(0,p)$ be such that $h_p(0)=2h_p(r_0)$. Then for any $r\in(0,r_0)$ we have $h_p(f)<h_p(r)$, and so (p,r) is in the symmetry breaking phase, resulting in a nontrivial phase diagram. We currently do not know the complete lower tail phase diagram for any non-bipartite H.

Symmetry breaking in exponential random graphs. Fig. 3 showed several (β_1,β_2)-phase plots for the symmetry breaking region given in Part (c) of Theorem 4.3. However, unlike the situation for large deviations, we do not know if that is the full region of symmetry breaking. It would be interesting to characterize the full set of triples (α,β_1,β_2) for which there is replica symmetry in Theorem 1.3.

Appendix A. The convex minorant of $h_p(x^{1/\gamma})$

This appendix contains some technical lemmas about the convex minorant of $h_p(x^{1/\gamma})$, which appears throughout the paper. Let us first informally summarize the claims. It is well known that

$$h_p(x) = x\log\frac{x}{p} + (1-x)\log\frac{1-x}{1-p}$$

is a convex function of x. When $\gamma>1$ (here we allow any real γ, not just integers), it turns out that there is some $p_0(\gamma)$ for which $x\mapsto h_p(x^{1/\gamma})$ is still a convex function when $p\ge p_0$. However, when $p<p_0$, the function $x\mapsto h_p(x^{1/\gamma})$ is no longer convex: it has exactly two inflection points, both to the right of its minimum at $x=p^\gamma$. The function is concave in the corresponding middle region, whereas it is convex in the two outer regions.

In the case when $x\mapsto h_p(x^{1/\gamma})$ is not convex, it has a unique lower tangent, touching the plot of the function at the points $(\underline{q}^\gamma, h_p(\underline{q}))$ and $(\overline{q}^\gamma, h_p(\overline{q}))$. The convex minorant of $x\mapsto h_p(x^{1/\gamma})$ is formed by replacing the middle segment $x\in(\underline{q}^\gamma,\overline{q}^\gamma)$ by the lower common tangent, as shown in

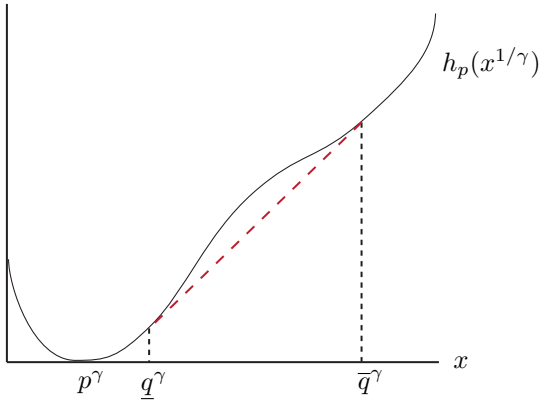

FIGURE 8. Illustration of the convex minorant of $x \mapsto h_p(x^{1\gamma})$.

Fig. 8. (The various not-to-scale plots of $h_p(x^{1/\gamma})$ are shown for illustrative purposes in order to highlight the features of the plots. In contrast, all the phase diagrams plots are drawn to scale.)

Let $\mathcal{B}_\gamma \subset [0,1]^2$ denote the set of all points (p,q) such that $(q^\gamma, h(q))$ does not lie on the convex minorant of $x \mapsto h_p(x^{1/\gamma})$. Each vertical section of $\mathcal{B}_\gamma$ is thus the interval $(\underline{q}^\gamma, \overline{q}^\gamma)$ described in the previous paragraph. For example, $\mathcal{B}_2$ is the shaded region in Fig. 9, and the boundaries of $\mathcal{B}_\gamma$ for additional values of γ were plotted in Figure 2 (see Section 1).

We shall show that each $\mathcal{B}_\gamma$ resembles a rotated V-shape. More importantly, we show that $\mathcal{B}_\gamma$ strictly contains $\mathcal{B}_{\gamma'}$ if $\gamma > \gamma'$.

The first lemma describes the shape of the function $x \mapsto h_p(x^{1/\gamma})$.

Lemma A.1. *The function $x \mapsto h_p(x^{1/\gamma})$ with domain $(0,1)$ is convex for $0 < \gamma \le 1$. If $\gamma > 1$, and*

$$p \ge p_0(\gamma) := \frac{\gamma - 1}{\gamma - 1 + e^{\gamma/(\gamma-1)}},$$

then the function is also convex. If $\gamma > 1$ and $0 < p < p_0$, then the function has exactly two inflection points (both to the right of $x = p^\gamma$), with a region of concavity in the middle. Finally the function has infinite derivatives at both endpoints of $(0,1)$.

Proof. When $0 < \gamma \le 1$, $x \mapsto x^{1/\gamma}$ is convex. Since the composition of two convex functions is convex we deduce that $x \mapsto h_p(x^{1/\gamma})$ is convex.

Next, assume $\gamma \ge 1$. We have

$$h_p'(x) = \log \frac{x}{1-x} - \log \frac{p}{1-p}, \qquad \text{and} \qquad h_p''(x) = \frac{1}{x(1-x)}.$$

Therefore,

$$\frac{d}{dx} h_p(x^{1/\gamma}) = \frac{1}{\gamma} x^{1/\gamma - 1} h_p'(x^{1/\gamma})$$

and

$$\begin{aligned} \frac{d^2}{dx^2} h_p(x^{1/\gamma}) &= \frac{1}{\gamma^2} x^{2/\gamma - 2} h''(x^{1/\gamma}) + \frac{1}{\gamma}\left(\frac{1}{\gamma} - 1\right) x^{1/\gamma - 2} h_p'(x^{1/\gamma}) \\ &= \frac{x^{1/\gamma - 2}}{\gamma^2}\left(x^{1/\gamma} h_p''(x^{1/\gamma}) - (\gamma - 1) h_p'(x^{1/\gamma})\right). \end{aligned}$$

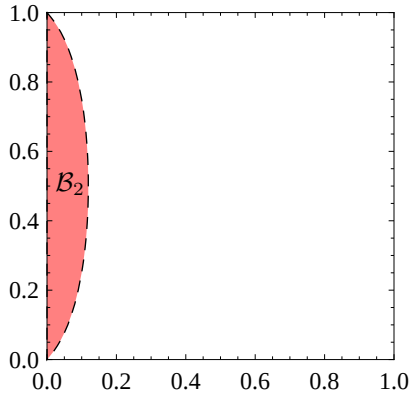

FIGURE 9. The region $\mathcal{B}_2$ consisting of all points (p, q) such that $(q^2, h_p(q))$ does not lie on the convex minorant of $x \mapsto h_p(\sqrt{x})$.

The claim on infinite derivatives easily follows from the above formulas. Setting $x = q^\gamma$, we have

$$\begin{aligned}\frac{d^2}{dx^2} h_p(x^{1/\gamma})\Big|_{x=q^\gamma} &= \frac{q^{1-2\gamma}}{\gamma^2}\left(q h_p''(q) - (\gamma-1)h_p'(q)\right)\\ &= \frac{q^{1-2\gamma}}{\gamma^2}\left(\frac{1}{1-q} - (\gamma-1)\log\frac{q}{1-q} + (\gamma-1)\log\frac{p}{1-p}\right). \qquad \text{(A.1)}\end{aligned}$$

Hence, $h_p(x^{1/\gamma})$ is convex at $x = q^\gamma$ whenever $\frac{1}{1-q} - (\gamma-1)\log\frac{q}{1-q} \geq -(\gamma-1)\log\frac{p}{1-p}$ and concave otherwise. The fact that

$$\frac{d}{dq}\left(\frac{1}{1-q} - (\gamma-1)\log\frac{q}{1-q}\right) = \frac{\gamma q - \gamma + 1}{q(1-q)^2}$$

implies that $\frac{1}{1-q} - (\gamma-1)\log\frac{q}{1-q}$ is decreasing until $q = (\gamma-1)/\gamma$ and then increasing afterwards. It diverges to $+\infty$ at both endpoints of $(0,1)$ and attains a minimum value of $\gamma - (\gamma-1)\log(\gamma-1)$ at $q = (\gamma-1)/\gamma$. The term $(\gamma-1)\log\frac{p}{1-p}$ of Eq. (A.1) is increasing for $p \in (0,1)$ and surjective onto the reals. Therefore, $h_p''(x^{1/d}) \geq 0$ for all x if $\gamma - (\gamma-1)\log(\gamma-1) + (\gamma-1)\log\frac{p}{1-p} \geq 0$, which is equivalent to having $p \geq \frac{\gamma-1}{\gamma-1+e^{\gamma/(\gamma-1)}}$. Additionally, if $p < \frac{\gamma-1}{\gamma-1+e^{\gamma/(\gamma-1)}}$, then $h_p''(x)$ starts as positive, becomes negative, then turns positive again. ∎

We next give an explicit description of the region $\mathcal{B}_2$.

Lemma A.2 ($\gamma = 2$ case)**.** *Let $p, q \in (0,1)$. The point $(q^2, h_p(q))$ lies strictly above the convex minorant of $x \mapsto h_p(\sqrt{x})$ if and only if $p < \left(1 + (q^{-1}-1)^{1/(1-2q)}\right)^{-1}$.*

Proof. We claim that the lower common tangent of $x \mapsto h_p(\sqrt{x})$ has slope $\log(\frac{1-p}{p})$. To show this, it suffices to check that $h_p(\sqrt{x}) - x\log(\frac{1-p}{p})$ has a horizontal common lower tangent, and it suffices to check the same thing for $h_p(x) - x^2\log(\frac{1-p}{p})$. Observe that

$$\begin{aligned}h_p(x) - x^2\log\frac{1-p}{p} &= x\log\frac{x}{p} + (1-x)\log\frac{1-x}{1-p} - x^2\log\frac{1-p}{p}\\ &= x\log x + (1-x)\log(1-x) - x(1-x)\log\frac{p}{1-p} - \log(1-p)\end{aligned}$$

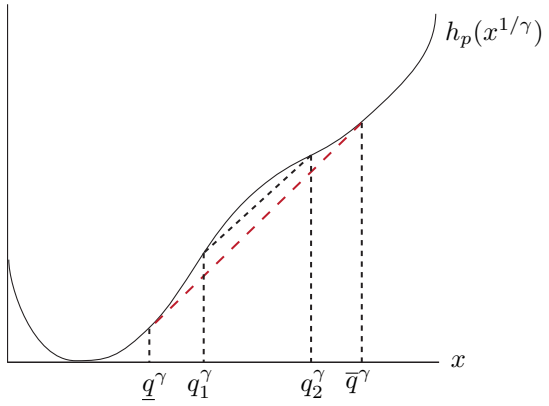

FIGURE 10. Illustration of the convex minorant of $x \mapsto h_p(x^{1/\gamma})$ in the setting of Lemma A.3.

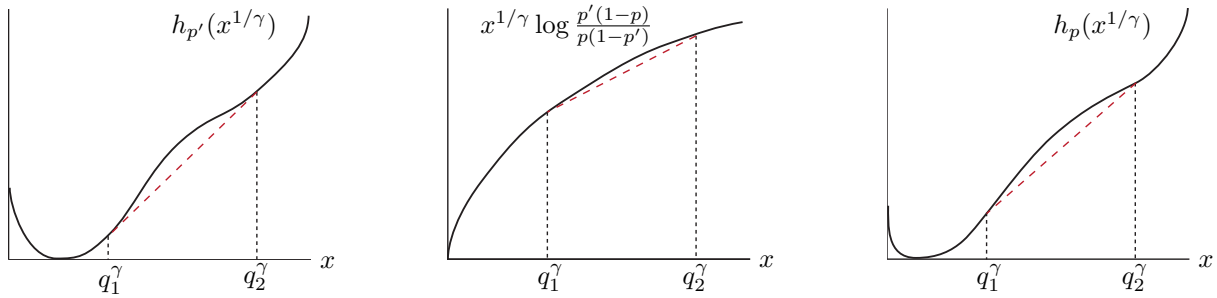

FIGURE 11. Illustration of the functions in Eq. (A.3).

is invariant under $x \mapsto 1-x$, so that its lower tangent must be horizontal by symmetry, and let it touch the curve at $x = \underline{q}, \overline{q}$, so that $0 < \underline{q} < \overline{q} < 1$ are the zeros of the derivative, namely

$$\log \frac{x}{1-x} - (1-2x) \log \frac{p}{1-p} \,. \tag{A.2}$$

It follows that $(q^2, h_p(q))$ lies strictly above the convex minorant if and only if $\underline{q} < q < \overline{q}$, which is equivalent to having $\frac{1}{1-2q} \log \frac{q}{1-q} \le \log \frac{p}{1-p}$. Rearranging the latter concludes the proof. ∎

For $\gamma > 1$ and $0 < p < p_0(\gamma)$, define $\underline{q} = \underline{q}(\gamma, p)$ and $\overline{q} = \overline{q}(\gamma, p)$ to be such that the lower common tangent to $x \mapsto h_p(x^{1/\gamma})$ touches the curve at points $(\underline{q}^\gamma, h_p(\underline{q})$ and $(\overline{q}^\gamma, h_p(\overline{q}))$. An examination of the geometry of the curve, as illustrated in Fig. 10, immediately leads to the following lemma:

Lemma A.3. *Let $\gamma > 1$ and $0 < p < p_0(\gamma)$. Let $0 < q_1 < q_2 < 1$. If the line segment joining points $(q_1^\gamma, h_p(q_1))$ and $(q_2^\gamma, h_p(q_2))$ lies below the curve $\{(q^\gamma, h_p(q)) : 0 \le q \le 1\}$ and is not tangent to the curve at one of the end points, then this segment lies strictly above the lower common tangent of the curve. Consequently, $\underline{q}(\gamma, p) < q_1 < q_2 < \overline{q}(\gamma, p)$.*

We now apply the above lemma to describe the shape of the regions $\mathcal{B}_\gamma$.

Lemma A.4. *If $\gamma > 1$ and $0 < p < p' < p_0(\gamma)$, then $\underline{q}(\gamma, p) < \underline{q}(\gamma, p') < \overline{q}(\gamma, p') < \overline{q}(\gamma, p)$. So $\mathcal{B}_\gamma$ is a rotated-V-shaped region.*

Proof. Let $q_1 = \underline{q}(\gamma, p')$ and $q_2 = \overline{q}(\gamma, p')$. As illustrated in Fig. 11, we have

$$h_p(x^{1/\gamma}) = h_{p'}(x^{1/\gamma}) + x^{1/\gamma} \log \frac{p'(1-p)}{p(1-p')} + \log \frac{1-p'}{1-p} \,. \tag{A.3}$$

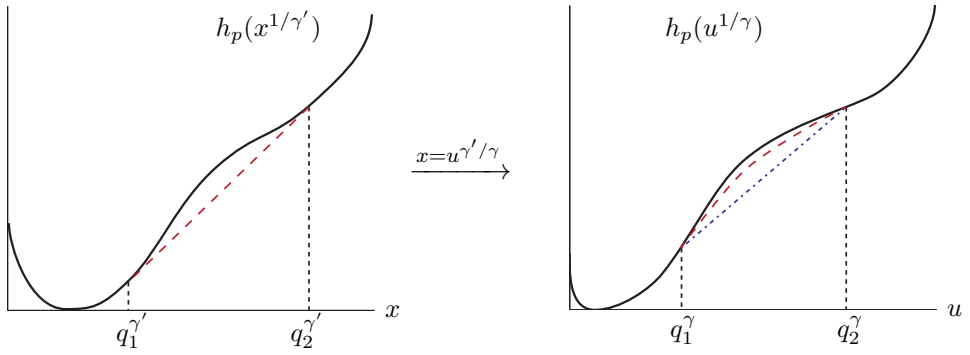

FIGURE 12. The plot of $x \mapsto h_p(x^{1\gamma})$ following a change of variable $x = u^{\gamma'/\gamma}$.

The segment joining $(q_1^\gamma, h_{p'}(q_1))$ and $(q_2^\gamma, h_{p'}(q_2))$ lies below $x \mapsto h_{p'}(x^{1/\gamma})$ by definition. Since $x \mapsto x^{1/\gamma} \log \frac{p'(1-p)}{p(1-p')}$ is concave, the segment joining $x = q_1^\gamma$ and $x = q_2^\gamma$ must also lie below the curve $x \mapsto h_p(x^{1/\gamma})$, and it is not tangent to either endpoint due to the $x^{1/\gamma}$ term. The conclusion now follows from Lemma A.3. ∎

Lemma A.5. *If $\gamma > \gamma' > 1$ and $0 < p < p_{\gamma'}$, then $\underline{q}(\gamma, p) < \underline{q}(\gamma', p) < \overline{q}(\gamma', p) < \overline{q}(\gamma, p)$. So $\mathcal{B}_\gamma$ strictly contains $\mathcal{B}_{\gamma'}$.*

Proof. Let $q_1 = \underline{q}(\gamma', p)$ and $q_2 = \underline{q}(\gamma', p)$. Let ℓ' denote the line segment joining points $(q_1^{\gamma'}, h_p(q_1))$ and $(q_2^{\gamma'}, h_p(q_2))$, so that ℓ' is tangent to the curve $x \mapsto h_p(x^{1/\gamma'})$.

Consider a transformation of the plots induced by the change of variable $x = u^{\gamma'/\gamma}$. The plot of $x \mapsto h_p(x^{1/\gamma'})$ becomes the plot of $u \mapsto h_p(u^{1/\gamma})$ (see Fig. 12). Originally ℓ' was a line segment of positive slope lying below the curve $x \mapsto h_p(x^{1/\gamma'})$ and tangent to it at both endpoints. Following the transformation $x = u^{\gamma'/\gamma}$ (recall that $\gamma'/\gamma < 1$) we see that ℓ' becomes a concave curve ℓ still lying below the new curve $u \mapsto h_p(u^{1/\gamma})$ and tangent to it at both endpoints. This implies that the line segment joining the two endpoints of ℓ in the new frame lies below both curves and is tangent to neither at the endpoints. The conclusion then follows from Lemma A.3. ∎

ACKNOWLEDGMENTS

We thank Amir Dembo and Ofer Zeitouni for fruitful discussions. This work was initiated while Y. Z. was an intern at the Theory Group of Microsoft Research, and he thanks the Theory Group for its hospitality.

REFERENCES

[1] N. Alon. Independent sets in regular graphs and sum-free subsets of finite groups. *Israel J. Math.*, 73(2):247–256, 1991.
[2] S. Bhamidi, G. Bresler, and A. Sly. Mixing time of exponential random graphs. *Ann. Appl. Probab.*, 21(6):2146–2170, 2011.
[3] B. Bollobás. *Random graphs*, volume 73 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2001.
[4] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, and K. Vesztergombi. Convergent sequences of dense graphs. I. Subgraph frequencies, metric properties and testing. *Adv. Math.*, 219(6):1801–1851, 2008.
[5] S. Chatterjee. The missing log in large deviations for triangle counts. *Random Structures Algorithms*, 40(4):437–451, 2012.

[6] S. Chatterjee and P. S. Dey. Applications of Stein's method for concentration inequalities. *Ann. Probab.*, 38(6):2443–2485, 2010.
[7] S. Chatterjee and P. Diaconis. Estimating and understanding exponential random graph models. *Ann. Statist.*, to appear.
[8] S. Chatterjee and S. R. S. Varadhan. The large deviation principle for the Erdős-Rényi random graph. *European J. Combin.*, 32(7):1000–1017, 2011.
[9] S. Chatterjee and S. R. S. Varadhan. Large deviations for random matrices. *Comm. Stoch. Analysis*, 6(1):1–13, 2012.
[10] F. R. K. Chung. Regularity lemmas for hypergraphs and quasi-randomness. *Random Structures Algorithms*, 2(2):241–252, 1991.
[11] D. Conlon, J. Fox, and B. Sudakov. An approximate version of Sidorenko's conjecture. *Geom. Funct. Anal.*, 20(6):1354–1366, 2010.
[12] D. Conlon, J. Fox, and Y. Zhao. Extremal results in sparse pseudorandom graphs. Preprint arXiv:1204.6645.
[13] B. DeMarco and J. Kahn. Tight upper tail bounds for cliques. *Random Structures Algorithms*.
[14] B. DeMarco and J. Kahn. Upper tails for triangles. *Random Structures Algorithms*, 40(4):452–459, 2012.
[15] A. Dembo and O. Zeitouni. *Large deviations techniques and applications*, volume 38 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin, 2010. Corrected reprint of the second (1998) edition.
[16] G. Elek and B. Szegedy. Limits of hypergraphs, removal and regularity lemmas. A non-standard approach. Preprint arxiv:0705.2179.
[17] G. Elek and B. Szegedy. A measure-theoretic approach to the theory of dense hypergraphs. Preprint arXiv:0810.4062.
[18] H. Finner. A generalization of Hölder's inequality and some probability inequalities. *Ann. Probab.*, 20(4):1893–1901, 1992.
[19] E. Friedgut. Hypergraphs, entropy, and inequalities. *Amer. Math. Monthly*, 111(9):749–760, 2004.
[20] A. Frieze and R. Kannan. Quick approximation to matrices and applications. *Combinatorica*, 19(2):175–220, 1999.
[21] D. Galvin. Maximizing H-colorings of regular graphs. *J. Graph Theory*, to appear.
[22] D. Galvin and P. Tetali. On weighted graph homomorphisms. In *Graphs, morphisms and statistical physics*, volume 63 of *DIMACS Ser. Disc. Math. Theoret. Comput. Sci.*, pages 97–104. Amer. Math. Soc., Providence, RI, 2004.
[23] S. Gerke and A. Steger. The sparse regularity lemma and its applications. In *Surveys in combinatorics 2005*, volume 327 of *London Math. Soc. Lecture Note Ser.*, pages 227–258. Cambridge Univ. Press, Cambridge, 2005.
[24] W. T. Gowers. Hypergraph regularity and the multidimensional Szemerédi theorem. *Ann. of Math. (2)*, 166(3):897–946, 2007.
[25] O. Häggström and J. Jonasson. Phase transition in the random triangle model. *J. Appl. Probab.*, 36(4):1101–1115, 1999.
[26] H. Hatami. Graph norms and Sidorenko's conjecture. *Israel J. Math.*, 175:125–150, 2010.
[27] S. Janson, T. Łuczak, and A. Rucinski. *Random graphs*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000.
[28] S. Janson, K. Oleszkiewicz, and A. Ruciński. Upper tails for subgraph counts in random graphs. *Israel J. Math.*, 142:61–92, 2004.
[29] S. Janson and A. Ruciński. The infamous upper tail. *Random Structures Algorithms*, 20(3):317–342, 2002. Probabilistic methods in combinatorial optimization.
[30] S. Janson and A. Ruciński. The deletion method for upper tail estimates. *Combinatorica*, 24(4):615–640, 2004.
[31] J. Kahn. An entropy approach to the hard-core model on bipartite graphs. *Combin. Probab. Comput.*, 10(3):219–237, 2001.
[32] J. H. Kim and V. H. Vu. Divide and conquer martingales and the number of triangles in a random graph. *Random Structures Algorithms*, 24(2):166–174, 2004.
[33] Y. Kohayakawa. Szemerédi's regularity lemma for sparse graphs. In *Foundations of computational mathematics (Rio de Janeiro, 1997)*, pages 216–230. Springer, Berlin, 1997.
[34] L. Lovász and B. Szegedy. Limits of dense graph sequences. *J. Combin. Theory Ser. B*, 96(6):933–957, 2006.
[35] L. Lovász and B. Szegedy. Szemerédi's lemma for the analyst. *Geom. Funct. Anal.*, 17(1):252–270, 2007.
[36] B. Nagle, V. Rödl, and M. Schacht. The counting lemma for regular k-uniform hypergraphs. *Random Structures Algorithms*, 28(2):113–179, 2006.
[37] J. Park and M. E. J. Newman. Solution of the two-star model of a network. *Phys. Rev. E (3)*, 70(6):066146, 5, 2004.
[38] J. Park and M. E. J. Newman. Solution for the properties of a clustered network. *Phys. Rev. E*, 72:026136, Aug 2005.
[39] C. Radin and M. Yin. Phase transitions in exponential random graphs. Preprint arXiv:1108.0649.

[40] V. Rödl and J. Skokan. Regularity lemma for k-uniform hypergraphs. *Random Structures Algorithms*, 25(1):1–42, 2004.
[41] W. Rudin. *Functional analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill Inc., New York, second edition, 1991.
[42] A. Sidorenko. A correlation inequality for bipartite graphs. *Graphs Combin.*, 9(2):201–204, 1993.
[43] M. Simonovits. Extremal graph problems, degenerate extremal problems, and supersaturated graphs. In *Progress in graph theory (Waterloo, Ont., 1982)*, pages 419–437. Academic Press, Toronto, ON, 1984.
[44] D. Strauss. On a general class of models for interaction. *SIAM Rev.*, 28(4):513–527, 1986.
[45] E. Szemerédi. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, volume 260 of *Colloq. Internat. CNRS*, pages 399–401. CNRS, Paris, 1978.
[46] V. H. Vu. A large deviation result on the number of small subgraphs of a random graph. *Combin. Probab. Comput.*, 10(1):79–94, 2001.
[47] Y. Zhao. The number of independent sets in a regular graph. *Combin. Probab. Comput.*, 19(2):315–320, 2010.
[48] Y. Zhao. The bipartite swapping trick on graph homomorphisms. *SIAM J. Discrete Math.*, 25(2):660–680, 2011.

E. Lubetzky
Microsoft Research, One Microsoft Way, Redmond, WA 98052-6399.
E-mail address: eyal@microsoft.com

Y. Zhao
Department of Mathematics, MIT, Cambridge, MA 02139-4307.
E-mail address: yufeiz@math.mit.edu