



BREAKING OF ENSEMBLE EQUIVALENCE FOR DENSE RANDOM GRAPHS UNDER A SINGLE CONSTRAINT

FRANK DEN HOLLANDER,^{*} AND
MAARTEN MARKERING ^{**}, *Leiden University*

Abstract

Two ensembles are frequently used to model random graphs subject to constraints: the *microcanonical ensemble* (= hard constraint) and the *canonical ensemble* (= soft constraint). It is said that *breaking of ensemble equivalence* (BEE) occurs when the specific relative entropy of the two ensembles does not vanish as the size of the graph tends to infinity. Various examples have been analysed in the literature. It was found that BEE is the rule rather than the exception for two classes of constraints: *sparse* random graphs when the *number* of constraints is of the order of the number of vertices, and *dense* random graphs when there are *two or more* constraints that are *frustrated*. We establish BEE for a third class: *dense* random graphs with a *single* constraint on the *density* of a given simple graph. We show that BEE occurs in a certain range of choices for the density and the number of edges of the simple graph, which we refer to as the BEE-phase. We also show that, in part of the BEE-phase, there is a gap between the scaling limits of the averages of the maximal eigenvalue of the adjacency matrix of the random graph under the two ensembles, a property that is referred to as the *spectral signature* of BEE. We further show that in the replica symmetric region of the BEE-phase, BEE is due to the *coexistence* of two densities in the canonical ensemble.

Keywords: Constrained random graph; Gibbs ensemble; relative entropy; breaking of ensemble equivalence; graphon; variational representation; maximal eigenvalues; replica symmetry

2020 Mathematics Subject Classification: Primary 05C80
Secondary 60C05; 60F10; 82B80

1. Introduction and main results

Section 1.1 provides the background and the motivation behind our paper. Section 1.2 states the definition of the microcanonical and the canonical ensemble in the context of constrained random graphs, recalls the notion of ensemble equivalence, lists the key definitions of graphons and subgraph counts, and gives the variational characterisation of the specific relative entropy of the two ensembles for dense random graphs derived in [6], which is the main tool in our paper. Section 1.3 states our main theorems. Section 1.4 identifies the typical graphs under the two ensembles. Section 1.5 offers a brief discussion and an outline of the remainder of the paper.

Received 24 January 2022; revision received 27 October 2022.

^{*} Postal address: Mathematical Institute, Leiden University, P.O. Box 9512, 2300 RA Leiden, The Netherlands. Email: denholla@math.leidenuniv.nl

^{**} Postal address: Pembroke College, Cambridge, CB2 1RF, United Kingdom. Email: mjrm2@cam.ac.uk

© The Author(s), 2023. Published by Cambridge University Press on behalf of Applied Probability Trust.

1.1. Background and motivation

In this paper we analyse random graphs that are subject to *constraints*. Statistical physics prescribes which probability distribution on the set of graphs we should choose when we want to model a given type of constraint [12]. Two important choices are: (i) the *microcanonical ensemble*, where the constraints are satisfied by each individual graph; (ii) the *canonical ensemble*, where the constraints are satisfied as ensemble averages. For random graphs that are large but finite, the two ensembles represent different empirical situations. One of the cornerstones of statistical physics is that the two ensembles become equivalent in the thermodynamic limit, which in our setting corresponds to letting the size of the graph tend to infinity. However, this property does *not* hold in general. We refer to [19] for more background on the phenomenon of *breaking of ensemble equivalence* (BEE).

BEE has been investigated for various choices of constraints, including the degree sequence and the total number of subgraphs of a specific type. The key distinctive object is the *relative entropy* S_n of the microcanonical ensemble with respect to the canonical ensemble when the graph has n vertices. In the *sparse regime*, where the number of edges per vertex remains bounded, the relevant quantity is $s_\infty = \lim_{n \rightarrow \infty} n^{-1} S_n$, because n is the scale of the number of *vertices*. In the *dense regime*, where the number of edges per vertex is of the order of the number of vertices, the relevant quantity is $s_\infty = \lim_{n \rightarrow \infty} n^{-2} S_n$, because n^2 is the scale of the number of *edges*.

Sparse regime: In [10, 11, 18] it was shown that constraining the degrees of *all* the vertices leads to BEE, even when the graph consists of multiple communities. An explicit formula was derived for s_∞ in terms of the limit of the empirical degree distribution of the constraint. In [17] a formula was put forward that expresses the specific relative entropy in terms of a covariance matrix under the canonical ensemble.

Dense regime: In [6] it was shown that constraining the densities of a finite number of subgraphs may lead to BEE. The analysis relied on the *large deviation principle for graphons* associated with the Erdős–Rényi (ER) random graph [3, 5]. The main result was a variational formula for s_∞ in the space of graphons. In [7], for the special case where the constraint is on the densities of the edges and triangles, it was shown that $s_\infty > 0$ when the constraints are *frustrated*, i.e. do not lie on the ER-line where the density of triangles is the third power of the density of edges. Moreover, the asymptotics of s_∞ near the ER-line was identified, and turns out to depend on whether the ER-line is approached from above or below.

It is an open problem whether a *single* constraint may lead to BEE [6]. It was believed that this cannot be the case, because for a single constraint there is no frustration. The goal of the present paper is to show that this intuition is wrong: we condition on the density of a given finite simple graph and prove that BEE occurs in a certain *range* of choices for the density and the number of edges of the simple graph, which we refer to as the BEE-phase. We analyse how s_∞ tends to zero near the curve that borders the BEE-phase. This phase transition is *unlike* any of the phenomena surrounding BEE observed before. In our case, BEE is due to the *coexistence* of two densities in the BEE-phase, similar to the phase transition between water and ice. Thus, our paper provides new insight into the mechanisms causing BEE.

In [9] the gap Δ_n between the *averages* of the maximal eigenvalue of the adjacency matrix of a constrained random graph under the two ensembles was considered. A *working hypothesis* was put forward, stating that BEE is equivalent to this gap not vanishing in the limit as $n \rightarrow \infty$. For a random regular graph with a fixed degree, this equivalence was proved for a range of

degrees that interpolates between the sparse and the dense regime. In the present paper we prove the same for the single constraint. In particular, we compute $\delta_\infty = \lim_{n \rightarrow \infty} n^{-1} \Delta_n$, show that $\delta_\infty \neq 0$ if and only if the density and the number of edges of the simple graph fall in the BEE-phase, and analyse how δ_∞ tends to zero near the curve that borders the BEE-phase.

We will see that the notions of *replica symmetry* and *replica symmetry breaking* highlighted in [15] play an important role. In the regime of replica symmetry we have a complete identification of s_∞ and δ_∞ ; in the regime of replica symmetry breaking some pieces of the characterisation are missing. Furthermore, we establish a direct connection between the region of replica symmetry for regular graphs and the region of ensemble equivalence.

1.2. Definitions and preliminaries

In this section, which is partly lifted from [6], we present the definitions of the main concepts to be used here, together with some key results from prior work. We consider scalar-valued constraints, even though [6] deals with more general vector-valued constraints.

Section 1.2.1 presents the formal definition of the two ensembles and the definition of ensemble equivalence in the dense regime. Section 1.2.2 recalls some basic facts about graphons. Section 1.2.3 recalls the variational characterisation of ensemble equivalence proven in [6]. Section 1.2.4 looks at the average of the maximal eigenvalue value of the adjacency matrix in the two ensembles and recalls a working hypothesis put forward in [9] that links ensemble equivalence to a vanishing gap between the two averages.

1.2.1. *Microcanonical ensemble, canonical ensemble, relative entropy* For $n \in \mathbb{N}$, let \mathcal{G}_n denote the set of all $2^{\binom{n}{2}}$ simple undirected graphs with n vertices. Let T denote a scalar-valued function on \mathcal{G}_n , and T^* a specific scalar that is *graphical*, i.e. realisable by at least one graph in \mathcal{G}_n . Given T^* , the *microcanonical ensemble* is the probability distribution P_{mic} on \mathcal{G}_n with *hard constraint* T^* defined, for $G \in \mathcal{G}_n$, as

$$P_{\text{mic}}(G) := \begin{cases} |\{G \in \mathcal{G}_n : T(G) = T^*\}|^{-1} & \text{if } T(G) = T^*, \\ 0 & \text{otherwise.} \end{cases}$$

The *canonical ensemble* P_{can} is the unique probability distribution on \mathcal{G}_n that maximises the *entropy* $S_n(P) := -\sum_{G \in \mathcal{G}_n} P(G) \log P(G)$ subject to the *soft constraint* $\langle T \rangle := \sum_{G \in \mathcal{G}_n} T(G) P(G) = T^*$. This gives the formula [13]

$$P_{\text{can}}(G) := \frac{1}{Z(\theta^*)} e^{\theta^* T(G)}, \quad G \in \mathcal{G}_n, \tag{1.1}$$

with $Z(\theta^*)$ the *partition function*. In (1.1), the Lagrange multiplier θ^* must be set to the unique value that realises $\langle T \rangle = T^*$; see [6, (2.6) and (2.7)].

The *relative entropy* of P_{mic} with respect to P_{can} is defined as

$$S_n(P_{\text{mic}} | P_{\text{can}}) := \sum_{G \in \mathcal{G}_n} P_{\text{mic}}(G) \log \frac{P_{\text{mic}}(G)}{P_{\text{can}}(G)}.$$

For any $G_1, G_2 \in \mathcal{G}_n$, $P_{\text{can}}(G_1) = P_{\text{can}}(G_2)$ whenever $T(G_1) = T(G_2)$, i.e. the canonical probability is the same for all graphs with the same value of the constraint. We may therefore rewrite the relative entropy as

$$S_n(P_{\text{mic}} | P_{\text{can}}) = \log \frac{P_{\text{mic}}(G^*)}{P_{\text{can}}(G^*)},$$

where G^* is any graph in \mathcal{G}_n such that $T(G^*) = T^*$.

Remark 1.1. Both the constraint T^* and the Lagrange multiplier θ^* in general depend on n , i.e. $T^* = T_n^*$ and $\theta^* = \theta_n^*$. We consider constraints that converge when we pass to the limit $n \rightarrow \infty$, i.e.

$$\lim_{n \rightarrow \infty} T_n^* =: T_\infty^*. \tag{1.2}$$

Consequently, we expect that

$$\lim_{n \rightarrow \infty} \theta_n^* =: \theta_\infty^*. \tag{1.3}$$

Throughout the paper we *assume* that (1.3) holds. If convergence fails, then we may still consider subsequential convergence. The subtleties concerning (1.3) were discussed in detail in [6, Appendix A].

All the quantities above depend on n . In order not to burden the notation, we exhibit this n -dependence only in the symbols \mathcal{G}_n and $S_n(\mathbb{P}_{\text{mic}} \mid \mathbb{P}_{\text{can}})$. When we pass to the limit $n \rightarrow \infty$, we need to specify how $T(G)$, T^* , and θ^* are chosen to depend on n . We refer the reader to [6], where this issue was discussed in detail.

Definition 1.1. (*Ensemble equivalence.*) In the dense regime, if

$$s_\infty := \lim_{n \rightarrow \infty} n^{-2} S_n(\mathbb{P}_{\text{mic}} \mid \mathbb{P}_{\text{can}}) = 0,$$

then \mathbb{P}_{mic} and \mathbb{P}_{can} are said to be *equivalent*.

This particular notion of ensemble equivalence is known as *measure equivalence* of ensembles and is standard in the study of ensemble equivalence of networks. Other notions of ensemble equivalence are *thermodynamic equivalence* and *macrostate equivalence*. Under certain hypotheses, the three notions have been shown to be equivalent for physical systems [19]. We refer the reader to [19, 20] for further discussion of different notions of ensemble equivalence.

1.2.2. *Graphons* There is a natural way to embed a simple graph on n vertices in a space of functions called *graphons*. Let \mathcal{W} be the space of functions $h: [0, 1]^2 \rightarrow [0, 1]$ such that $h(x, y) = h(y, x)$ for all $(x, y) \in [0, 1]^2$. A finite simple graph G on n vertices can be represented as a graphon $h^G \in \mathcal{W}$ in a natural way as

$$h^G(x, y) := \begin{cases} 1 & \text{if there is an edge between vertex } \lceil nx \rceil \text{ and vertex } \lceil ny \rceil, \\ 0 & \text{otherwise.} \end{cases} \tag{1.4}$$

The space of graphons \mathcal{W} is endowed with the *cut distance*

$$d_\square(h_1, h_2) := \sup_{S, T \subset [0, 1]} \left| \int_{S \times T} dx dy [h_1(x, y) - h_2(x, y)] \right|, \quad h_1, h_2 \in \mathcal{W}.$$

On \mathcal{W} there is a natural equivalence relation \sim . Let Σ be the space of measure-preserving bijections $\sigma: [0, 1] \rightarrow [0, 1]$. Then $h_1(x, y) \sim h_2(x, y)$ if $\delta_\square(h_1, h_2) = 0$, where δ_\square is the *cut metric* defined by

$$\delta_\square(\tilde{h}_1, \tilde{h}_2) := \inf_{\sigma_1, \sigma_2 \in \Sigma} d_\square(h_1^{\sigma_1}, h_2^{\sigma_2}), \quad \tilde{h}_1, \tilde{h}_2 \in \tilde{\mathcal{W}},$$

with $h^\sigma(x, y) = h(\sigma x, \sigma y)$. This equivalence relation yields the quotient space $(\tilde{\mathcal{W}}, \delta_\square)$. As noted above, we suppress the n -dependence. Thus, by G we denote any simple graph on n

vertices, by h^G its image in the graphon space \mathcal{W} , and by \tilde{h}^G its image in the quotient space $\tilde{\mathcal{W}}$. For a more detailed description of the structure of the space $(\tilde{\mathcal{W}}, \delta_\square)$ we refer to [1, 2, 8].

For $h \in \tilde{\mathcal{W}}$ and F a finite simple graph with m vertices and edge set $E(F)$, define

$$t(F, h) := \int_{[0,1]^m} \prod_{\{i,j\} \in E(F)} h(x_i, x_j) dx_1 \dots dx_m.$$

Then the *homomorphism density* of F in G equals $t(F, h^G)$, where h^G is the empirical graphon defined in (1.4).

In this paper we focus on the special case where the constraint $T(G) = T(h^G) := t(F, h^G)$ is on the *homomorphism density* T_n^* of a *specific subgraph* F . The map T is well-defined on both the space of graphs \mathcal{G}_n for each n and the space of graphons. Rewriting (1.1), we obtain

$$P_{\text{can}}(G) = e^{n^2[\theta_n^* T(G) - \psi_n(\theta_n^*)]}, \quad G \in \mathcal{G}_n,$$

where

$$\psi_n(\theta_n^*) := \frac{1}{n^2} \log \sum_{G \in \mathcal{G}_n} e^{n^2[\theta_n^* T(G)]} = \frac{1}{n^2} \log Z(\theta_n^*).$$

It turns out that, under the scaling n^2 , the function ψ_n converges. Hence, rewriting (1.1) in this form aids us in the analysis of the canonical ensemble.

1.2.3. *Variational characterisation of ensemble equivalence* In order to characterise the asymptotic behaviour of the two ensembles, the entropy function of a Bernoulli random variable is essential. For $u \in [0, 1]$, let

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

Extend the domain of this function to the graphon space \mathcal{W} by defining

$$I(h) := \int_{[0,1]^2} dx dy I(h(x, y))$$

(with the convention that $0 \log 0 := 0$). On the quotient space $(\tilde{\mathcal{W}}, \delta_\square)$, define $I(\tilde{h}) = I(h)$, where h is any element of the equivalence class \tilde{h} . Note that $I(h)$ takes values in $[-\frac{1}{2} \log 2, 0]$. Apart from a shift by $\frac{1}{2} \log 2$, $h \mapsto I(h)$ plays the role of the rate function in the large deviation principle for the empirical graphon associated with the Erdős–Rényi random graph, derived in [5].

The key result in [6] is the following variational formula for s_∞ , where $\tilde{\mathcal{W}}^* := \{\tilde{h} \in \tilde{\mathcal{W}} : T(\tilde{h}) = T_\infty^*\}$ is the subspace of all graphons that meet the constraint T_∞^* . This is a compact set, since T is continuous in the cut metric and $(\tilde{\mathcal{W}}, \delta_\square)$ is compact [14].

Theorem 1.1. (Variational characterisation of ensemble equivalence.) *Subject to (1.2) and (1.3), $\lim_{n \rightarrow \infty} n^{-2} S_n(P_{\text{mic}} | P_{\text{can}}) =: s_\infty$, with*

$$s_\infty = \sup_{\tilde{h} \in \tilde{\mathcal{W}}} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})] - \sup_{\tilde{h} \in \tilde{\mathcal{W}}^*} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})]. \tag{1.5}$$

Theorem 1.1 and the compactness of $\tilde{\mathcal{W}}^*$ give us a variational characterisation of ensemble equivalence: $s_\infty = 0$ if and only if at least one of the maximisers of $\theta_\infty^* T(\tilde{h}) - I(\tilde{h})$ in $\tilde{\mathcal{W}}$ also lies in $\tilde{\mathcal{W}}^* \subset \tilde{\mathcal{W}}$, i.e. satisfies the hard constraint.

We need the following lemma, which relates T_∞^* and θ_∞^* without requiring knowledge of T_n^* and θ_n^* .

Lemma 1.1. *Under the assumptions in (1.2) and (1.3), $\theta_\infty^* = \arg \max_{\theta \in \mathbb{R}} [\theta T_\infty^* - \psi_\infty(\theta)]$, where $\psi_\infty(\theta) := \lim_{n \rightarrow \infty} \psi_n(\theta) = \sup_{\tilde{h} \in \tilde{\mathcal{W}}} [\theta T(\tilde{h}) - I(\tilde{h})]$.*

Proof. For every $n \in \mathbb{N}$,

$$\theta_n^* = \arg \max_{\theta \in \mathbb{R}} [n^2[\theta T_n^* - \psi_n(\theta)]] = \arg \max_{\theta \in \mathbb{R}} [\theta T_n^* - \psi_n(\theta)].$$

Let $f_n(\theta, T_n^*) := \theta T_n^* - \psi_n(\theta)$ and $f_\infty(\theta, T_\infty^*) = \theta T_\infty^* - \psi_\infty(\theta)$. By [6, Theorem 3.2 and Lemma A.1],

$$\begin{aligned} f_n(\theta_n^*, T_n^*) &= \sup_{\theta \in \mathbb{R}} f_n(\theta, T_n^*) = [\theta_n^* T_n^* - \psi_n(\theta_n^*)] \rightarrow [\theta_\infty^* T_\infty^* - \psi_\infty(\theta_\infty^*)] \\ &= \theta_\infty^* T_\infty^* - \sup_{\tilde{h} \in \tilde{\mathcal{W}}} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})] = f_\infty(\theta_\infty^*, T_\infty^*), \quad n \rightarrow \infty. \end{aligned}$$

Furthermore, for every $\theta \in \mathbb{R}$, $f_n(\theta, T_n^*) \leq f_n(\theta_n^*, T_n^*)$, and hence

$$f_\infty(\theta, T_\infty^*) = \lim_{n \rightarrow \infty} f_n(\theta, T_n^*) \leq \lim_{n \rightarrow \infty} f_n(\theta_n^*, T_n^*) = f_\infty(\theta_\infty^*, T_\infty^*),$$

so that θ_∞^* is a maximiser of $f_\infty(\cdot, T_\infty^*)$. □

1.2.4. Maximal eigenvalue of the adjacency matrix In [9] a working hypothesis was put forward, stating that breaking of ensemble equivalence is manifested by a gap between the scaling limits of the averages of the maximal eigenvalue of the adjacency matrix of the random graph under the two ensembles. More precisely, let λ_n denote the maximal eigenvalue of the adjacency matrix of $G \in \mathcal{G}_n$. Then the working hypothesis is that

$$\begin{aligned} \lim_{n \rightarrow \infty} \Delta_n \neq 0 &\implies \text{BEE}, \\ \text{BEE} &\implies \lim_{n \rightarrow \infty} \Delta_n \neq 0 \text{ apart from exceptional constraints,} \end{aligned}$$

with $\Delta_n := E_{\text{can}}[\lambda_n] - E_{\text{mic}}[\lambda_n]$. In [9] this equivalence was proven for the specific example where the constraint is on all the degrees being equal to $d(n)$, with $(\log n)^\beta \leq d(n) \leq n - (\log n)^\beta$ for some $\beta \in (6, \infty)$. It turns out that BEE occurs and that $\lim_{n \rightarrow \infty} \Delta_n = 1 - p$ when $\lim_{n \rightarrow \infty} n^{-1}d(n) = p \in [0, 1]$, i.e. the exceptional constraints correspond to the ultra-dense regime where $p = 1$.

For our single constraint in the dense regime, we will be interested in the quantity $\delta_\infty := \lim_{n \rightarrow \infty} n^{-1} \Delta_n$.

1.3. Main results

In what follows, F is any finite simple graph with k edges, and the constraint is on the homomorphism density of F being equal to T_n^* , as defined in Section 1.2.2. Recall from Remark 1.1 that we assume that $(T_n^*)_{n \in \mathbb{N}}$ and $(\theta_n^*)_{n \in \mathbb{N}}$ converge to some constants T_∞^* and θ_∞^* respectively. For the sake of convenience, we write $T^* = T_\infty^*$ and $\theta^* = \theta_\infty^*$. In the four theorems below we allow for $k \in [1, \infty)$, although $k \in \mathbb{N}$ is needed to interpret the constraint in terms of a subgraph density.




FIGURE 1. A numerical picture of the phase diagram. The left and right lines together form the critical curve $(T^*, k_c(T^*))$. In the figure, T^* is denoted by T^* and $k_c(T^*)$ is denoted by $k(T^*)$. The minimum is achieved at $k_0 = 4.591 \dots$ and $T_0 = 0.3237 \dots$

1.3.1. *Parameter regime* Our first two theorems identify the *parameter regime* for BEE.

Theorem 1.2. (Computational criterion for ensemble equivalence.) For $\theta \in [0, \infty)$ and $k \in [1, \infty)$, let $u^*(\theta, k)$ be a maximiser of

$$\sup_{u \in [0, 1]} [\theta u^k - I(u)]. \tag{1.6}$$

- (a) For every $T^* \in [(\frac{1}{2})^k, 1)$ there is ensemble equivalence if and only if there exists a $\theta_0 = \theta_0(T^*, k) \in [0, \infty)$ such that $(u^*(\theta_0, k))^k = T^*$. In that case the Lagrange multiplier $\theta^* = \theta^*(T^*, k)$ equals θ_0 .
- (b) There exists a unique $\hat{\theta} = \hat{\theta}(k) \in [0, \infty)$ such that $\theta^*(T^*, k) = \hat{\theta}$ for all T^* for which there is breaking of ensemble equivalence.

Theorem 1.3. (Phase diagram.)

- (a) There exists a function $k_c : (0, 1) \rightarrow [1, \infty)$ such that for every $T^* \in (0, 1)$ there is ensemble equivalence when $\log_2(1/T^*) \leq k \leq k_c(T^*)$ and breaking of ensemble equivalence when $k > k_c(T^*)$.
- (b) $T^* \mapsto k_c(T^*)$ achieves a unique minimum at the point (T_0, k_0) , with k_0 the unique solution of the equation $((k_0 - 1)/k_0) \log(k_0 - 1) = 1$ and $T_0 = ((k_0 - 1)/k_0)^{k_0}$.
- (c) $T^* \mapsto k_c(T^*)$ is analytic on $(0, 1) \setminus \{T_0\}$.
- (d) $(\frac{1}{2})^{k_c(T^*)} \sim T^*$ as $T^* \downarrow 0$ and $k_c(T^*)(\frac{1}{2})^{k_c(T^*)} \sim 1 - T^*$ as $T^* \uparrow 1$.

A numerical picture of the phase diagram described in Theorem 1.3 is shown in Fig. 1.

Note that the results above only hold in the regime $T^* \in [(\frac{1}{2})^k, 1)$, which corresponds to the regime $\theta^* \geq 0$. This assumption is necessary, since the results from [4] that we use only

hold for non-negative θ^* . For k and T^* not in this regime, we do not know if there is ensemble equivalence.

1.3.2. *Replica symmetry* Our last two theorems quantify the specific relative entropy and the spectral gap in the *replica symmetry regime*. This regime was first defined in [5] and further studied in [15]. Using the theory developed in [15], it is possible to quantify the specific relative entropy s_∞ and the difference of the largest eigenvalue δ_∞ for certain T^* in the BEE-phase.

Definition 1.2. (*Replica symmetry.*) Consider the Erdős–Rényi random graph G on n vertices with retention probability $p \in [0, 1]$ conditioned on $t(F, G) \geq T^*$ for some finite simple graph F . If G converges in the cut metric to a constant graphon, then we say that T^* is in the *replica symmetric region*.

From the theory of large deviations for random graphs developed in [5], we know that T^* is in the replica symmetric region if and only if

$$\inf_{t(F,f) \geq T^*} I_p(f) \tag{1.7}$$

is minimised by a constant graphon, with I_p the rate function given by

$$I_p(f) = \int_{[0,1]^2} dx dy \left(f(x, y) \log \frac{f(x, y)}{p} + [1 - f(x, y)] \log \frac{1 - f(x, y)}{1 - p} \right). \tag{1.8}$$

Note that $I(f) = I_{1/2}(f) - \frac{1}{2} \log 2$. Hence, if T^* is in the replica symmetric region, then there is an explicit solution for the second supremum in (1.5). In [15], it was shown that T^* is in the replica symmetric region when $(T^*, I_p(T^{*1/d}))$ lies on the convex minorant of the function $x \mapsto I_p(x^{1/d})$, with d the maximum degree of the subgraph F . If F is regular, then the converse statement holds as well.

Fix a subgraph F with k edges and maximum degree d . Let $T_1^*(k) \in ((\frac{1}{2})^k, T_0)$ and $T_2^*(k) \in (T_0, 1)$ be the two solutions of the equation $k_c(T^*(k)) = k$, so that $(T_1^*(k), T_2^*(k)) = \text{BEE-phase}$. In Lemma 3.1, we prove that the replica symmetric region contains $[(\frac{1}{2})^k, T_1^*(d)] \cup [T_2^*(d), 1]$. Thus, if $d < k$, then in part of the BEE-phase there is replica symmetry. This allows us to formulate the following two theorems (which are vacuous for $d = k$).

Theorem 1.4. (Specific relative entropy.) *For every T^* in the replica symmetric part of the phase of breaking of ensemble equivalence,*

$$s_\infty = \begin{cases} \hat{\theta}(k)[T_1^*(k) - T^*] + [I(T^{*1/k}) - I(T_1^*(k)^{1/k})] > 0, & T^* \in (T_1^*(k), T_1^*(d)], \\ \hat{\theta}(k)[T_2^*(k) - T^*] + [I(T^{*1/k}) - I(T_2^*(k)^{1/k})] > 0, & T^* \in [T_2^*(d), T_2^*(k)]. \end{cases}$$

Consequently,

$$s_\infty = \begin{cases} C(T_1^*(k), k) [T^* - T_1^*(k)]^2 + O([T - T_1^*(k)]^3), & T^* \downarrow T_1^*(k), \\ C(T_2^*(k), k) [T^* - T_2^*(k)]^2 + O([T - T_2^*(k)]^3), & T^* \uparrow T_2^*(k), \end{cases}$$

with

$$C(T^*, k) = \frac{T^{*(1-2k)/k}}{2k} \left\{ \frac{1}{k} \left(1 + \frac{T^{*1/k}}{1 - T^{*1/k}} \right) + \left(\frac{1}{k} - 1 \right) \log \left(\frac{T^{*1/k}}{1 - T^{*1/k}} \right) \right\}.$$

Theorem 1.5. (Spectral signature.) *For every T^* in the replica symmetric part of the phase of breaking of ensemble equivalence,*

$$\delta_\infty = \frac{T_1^{*1/k}[T_2^*(k) - T^*] + T_2^{*1/k}[T^* - T_1^*(k)]}{T_2^*(k) - T_1^*(k)} - T^{*1/k} < 0,$$

$$T^* \in (T_1^*(k), T_1^*(d)] \cup [T_2^*(d), T_2^*(k)).$$

Consequently,

$$\delta_\infty = \begin{cases} \hat{C}(T_1^*(k), k)[T^* - T_1^*(k)] + O([T^* - T_1^*(k)]^2), & T^* \downarrow T_1^*(k), \\ \hat{C}(T_2^*(k), k)[T^* - T_2^*(k)] + O([T^* - T_2^*(k)]^2), & T^* \uparrow T_2^*(k), \end{cases}$$

with

$$\hat{C}(T^*, k) = \frac{T_2^{*1/k}(k) - T_1^{*1/k}(k)}{T_2^*(k) - T_1^*(k)} - \frac{1}{k} T^{*(1-k)/k}.$$

1.4. Typical graph under the microcanonical and canonical ensemble

The BEE-phase can also be characterised through convergence of the random graph drawn from the two ensembles. In Lemmas 5.1 and 5.2 we show that the random graph drawn from the canonical ensemble converges to the maximiser(s) of the first supremum of (1.5), while the random graph drawn from the microcanonical ensemble converges to the maximiser(s) of the second supremum of (1.5).

Outside the BEE-phase, both suprema are attained by the constant graphon $h \equiv T^{*1/k}$, meaning that for large n both ensembles behave approximately like the Erdős–Rényi random graph with retention probability $p = T^{*1/k}$. Inside the BEE-phase, the first supremum is maximised by the two constant graphons $T_1^*(k)^{1/k}$ and $T_2^*(k)^{1/k}$, neither of which lies in $\tilde{\mathcal{V}}^*$. Consequently, the random graph drawn from the canonical ensemble converges to the random graphon

$$\frac{T_2^*(k) - T^*}{T_2^*(k) - T_1^*(k)} \delta_{T_1^*(k)^{1/k}} + \frac{T^* - T_1^*(k)}{T_2^*(k) - T_1^*(k)} \delta_{T_2^*(k)^{1/k}},$$

meaning that for large n the canonical ensemble behaves approximately like a mixture of two Erdős–Rényi random graphs. If T^* is in the replica symmetric part of the BEE-phase, then the second supremum is still minimised by the constant graphon $h \equiv T^{*1/k}$. Hence, the random graph is asymptotically deterministic under the microcanonical ensemble and random under the canonical ensemble. Thus, BEE occurs due to coexistence of two densities. This is similar in spirit to the coexistence of water and ice at the melting point, at which a first-order phase transition between water and ice occurs.

In the region of replica symmetry breaking, the maximisers of the second supremum are unknown, and it is not even known whether or not there is a unique maximiser (see Fig. 2). In the case of non-uniqueness, also under the microcanonical ensemble the random graph is asymptotically random.

1.5. Discussion and outline

Theorem 1.2 reduces the variational formula on $\tilde{\mathcal{V}}$ to a variational formula on $[0, 1]$, and is an application of a reduction principle explained in [4] (see also [3]). The proof relies on the

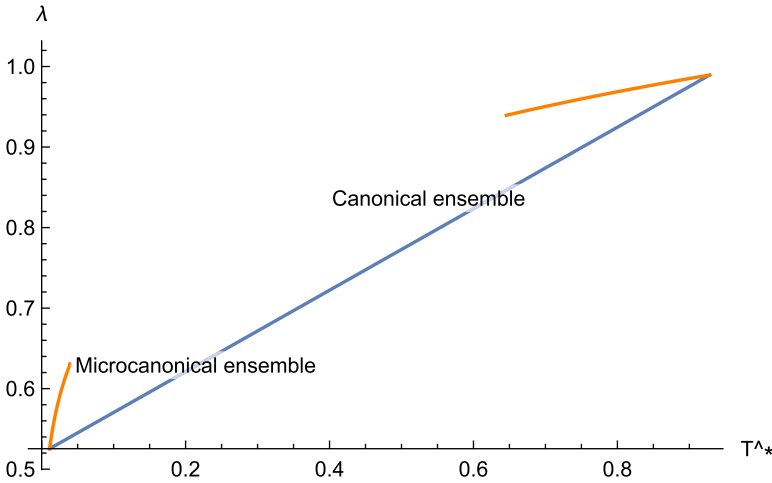


FIGURE 2. A numerical picture of the average largest eigenvalue $\lambda = \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\lambda_n]$ of the adjacency matrix under the microcanonical ensemble (top curve) and the canonical ensemble (bottom curve), as a function of T^* for a subgraph F with $k = 7$ edges and maximum degree $d = 5$. The top curve is shown only for T^* in the replica symmetric region. In the region of replica symmetry breaking we have no explicit expression for λ under the microcanonical ensemble.

variational characterisation in Theorem 1.1. The main difficulty lies in computing the tuning parameter θ^* as a function of the density T^* , which is resolved through Lemma 1.1. The proof follows from an analysis of the two variational expressions, for which we rely in part on the results in [16]. From Theorem 1.2, for each k we can identify the BEE-phase as follows. The expression in (1.6) has at most two local maximisers $u_1^*(\theta) < u_2^*(\theta)$, which are both increasing in θ . For $\theta < \hat{\theta}$, $u_1^*(\theta)$ is the global maximiser; for $\theta > \hat{\theta}$, $u_2^*(\theta)$ is the global maximiser; and for $\theta = \hat{\theta}$, $u_1^*(\theta)$ and $u_2^*(\theta)$ are both global maximisers. Hence, the values $u \in (u_1^*(\theta), u_2^*(\theta))$ can never be a global maximiser, and so the BEE-phase contains $(u_1^*(\theta)^k, u_2^*(\theta)^k)$. Since $u_1^*(0) = \frac{1}{2}$ and $\lim_{\theta \rightarrow \infty} u_2^*(\theta) = 1$, the interval $(u_1^*(\theta)^k, u_2^*(\theta)^k)$ is the entire BEE-phase.

Theorem 1.3 identifies the BEE-phase and captures the main properties of the critical curve bordering this phase. The proof relies on Lemma 3.1, which allows us to use results from [15] and establish a connection between ensemble equivalence and replica symmetry, in the sense that T^* lies in the BEE-phase for a subgraph with k edges if and only if T^* lies in the region of replica symmetry breaking for $p = \frac{1}{2}$ and a k -regular subgraph (recall (1.7) and (1.8)). This connection is purely analytic: it establishes equivalence of variational formulas and implies that the graph in Fig. 1 is a cross-section of the curves in [15, Figure 2] at $p = \frac{1}{2}$. It is not clear, however, how to probabilistically interpret the relationship between replica symmetry for regular subgraphs and breaking of ensemble equivalence for general graphs. Note that we do not require any regularity of the subgraph F , and also the degrees of F do not play any role. It might be easier to use the variational formula in (1.6) (with I_p instead of I) to analyse replica symmetry, rather than the convex minorant of $x \mapsto I_p(x^{1/k})$.

Theorem 1.4 gives an explicit formula for the specific relative entropy s_∞ in part of the BEE-phase. The proof exploits the connection with replica symmetry. If a subgraph has more edges than its maximal degree (i.e. is not a k -star), then the BEE-phase near $T_1^*(k)$ and $T_2^*(k)$ is replica symmetric. This implies that the second supremum in (1.5) also has a constant maximiser,

which allows us to explicitly compute s_∞ . It turns out that the relative entropy undergoes a second-order phase transition as T^* approaches the critical curve.

Theorem 1.5 shows that the working hypothesis put forward in [9] is met in the replica symmetric part of the BEE-phase. A random graph drawn from the canonical ensemble converges to a constant graphon whose height is a random mixture of the two maximisers u_1, u_2 of (1.6). The average largest eigenvalue converges to a value on the line segment connecting $(u_1^{1/k}, u_1)$ and $(u_2^{1/k}, u_2)$. In the region of replica symmetry, a random graph drawn from the microcanonical ensemble converges to the constant graphon whose height is $(T^*)^{1/k}$, as illustrated in Fig. 2. Note that the average largest eigenvalue is larger in the microcanonical ensemble than in the canonical ensemble, contrary to what was found in [9], where the constraint was on the degree sequence. It turns out that δ_∞ undergoes a first-order phase transition as T^* approaches the critical curve.

The numerical picture of the phase diagram in Fig. 1 was made using Mathematica. The computations involve finding an approximate value of $\hat{\theta}(k)$ for each k (up to an accuracy of five digits), and computing $u_1^*(\hat{\theta}(k), k)$ and $u_2^*(\hat{\theta}(k), k)$. The dotted lines are formed by the points $(u_1^*(k)^k, k)$ and $(u_2^*(k)^k, k)$. This is done for k starting at 4.592 and increasing with increments of 0.002.

In [19], BEE for interacting particle systems was studied at three different levels: thermodynamic, macrostate, and measure. It was shown that these levels are in fact equivalent. A general formalism was put forward, based on an abstract large deviation principle, linking the occurrence of BEE to non-convexity of the rate function associated with the microcanonical ensemble as a function of the parameters controlling the constraint. In our context, the large deviation principle for graphons in [5] provides the conceptual basis for identifying the BEE-phase via the variational formula derived in [6], and the link with the convex minorant mentioned above fits in with the picture provided in [19].

Theorems 1.2–1.5 are proved in Sections 2–5, respectively.

2. Proof of Theorem 1.2

Throughout the proof, we fix $k \in \mathbb{N}$, and suppress k from the notation. We analyse the expression

$$\sup_{\tilde{h} \in \mathcal{W}} [\theta T(\tilde{h}) - I(\tilde{h})] \tag{2.1}$$

with $\theta \in [0, \infty)$, and determine for which values of T^* a maximiser of this supremum is in the set $\tilde{\mathcal{W}}^*$. Note that it suffices to consider $\theta \in [0, \infty)$, since $T^* \geq (\frac{1}{2})^k$. This was shown in [6, Lemma 5.1] in the case that F is a triangle, but the proof generalizes to general finite simple graphs.

By [4, Theorem 4.1], the supremum equals the supremum in (1.6), and each maximiser of (2.1) is a constant function, where the constant is a maximiser of (1.6). Furthermore, by Lemma 1.1, θ^* is a maximiser of the supremum

$$\sup_{\theta \geq 0} [\theta T^* - \theta T(u^*(\theta)) + I(u^*(\theta))] = \sup_{\theta \geq 0} [\theta T^* - \theta (u^*(\theta))^k + I(u^*(\theta))],$$

where $u^*(\theta)$ is a maximiser of (1.6). By [16, Proposition 3.2], $l_\theta(u) := \theta u^k - I(u)$ has at most two maxima and there exists a $\hat{\theta}$ such that, for $\theta < \hat{\theta}$, the first local maximum is the unique global maximum and, for $\theta > \hat{\theta}$, the second local maximum is the unique global maximum.




FIGURE 3. Three plots of $l_\theta(u)$ for $k = 7$ and $\theta = 0.3$, $\theta = \hat{\theta}(7)$, and $\theta = 0.4$. For $\theta = 0.3$, $u_1^*(\theta)$ is the global maximiser; for $\theta = \hat{\theta}(7)$, $u_1^*(\theta)$ and $u_2^*(\theta)$ are both global maximisers; and for $\theta = 0.4$, $u_2^*(\theta)$ is the global maximiser. In the figures, the function $l_\theta(u)$ is denoted by $l(u)$ and the local maximisers $u_1^*(\theta)$ and $u_2^*(\theta)$ are denoted by $u1$ and $u2$ respectively. The BEE-phase is $(u_1^*(\hat{\theta})^k, u_2^*(\hat{\theta})^k)$.

Hence, for all $\theta \neq \hat{\theta}$, $u^*(\theta)$ is well-defined. For $\theta = \hat{\theta}$, both maxima are a global maximum. In that case, we let $u^*(\theta)$ denote either of the two maximisers.

Let $m(\theta) = \theta T^* - l_\theta(u^*(\theta)) = \theta T^* - \theta(u^*(\theta))^k + I(u^*(\theta))$. In Fig. 3, plots of l_θ are shown for several values of θ . Write

$$u := u^*(\theta), \quad u' := \frac{\partial u}{\partial \theta}(\theta).$$

Then

$$\begin{aligned} l'_\theta(u) &= \theta k u^{k-1} - \frac{1}{2} \log u + \frac{1}{2} \log(1-u) = 0, \\ m'(\theta) &= T^* - u^k - \theta k u^{k-1} u' + \frac{1}{2} u' \log(u) - \frac{1}{2} u' \log(1-u) \\ &= T^* - u^k - u' \left(\frac{1}{2} \log u - \frac{1}{2} \log(1-u) - \theta k u^{k-1} \right) \\ &= T^* - u^k. \end{aligned}$$

Hence, if there exists a $\theta_0 \geq 0$ such that $(u^*(\theta_0))^k = T^*$, then $m'(\theta_0) = 0$ and so $\theta^* = \theta_0$. In that case $(u^*(\theta^*))^k = T^*$, so there is ensemble equivalence. If such a θ_0 does not exist, then there is breaking of ensemble equivalence.

Let $u_1^*(\theta)$ and $u_2^*(\theta)$ be the first and second local maximum of l_θ , respectively. Then $\theta \mapsto u_1^*(\theta)$ and $\theta \mapsto u_2^*(\theta)$ are increasing. Furthermore, for all $\theta < \hat{\theta}$, $u_1^*(\theta)$ is the unique global maximum, while for all $\theta > \hat{\theta}$, $u_2^*(\theta)$ is the unique global maximum. Hence, if there is breaking of ensemble equivalence, then $m'(\theta) > 0$ for $\theta < \hat{\theta}$ and $m'(\theta) < 0$ for $\theta > \hat{\theta}$. We conclude that $\theta^* = \hat{\theta}$.

3. Proof of Theorem 1.3

We first fix some notation. For given k and θ , let $u_1^*(\theta, k)$ and $u_2^*(\theta, k)$ be the first and second local maximum respectively of $l_{\theta,k}(u) = \theta u^k - I(u)$. Let $\hat{\theta}(k)$ be the unique value of θ such that $u_1^*(\hat{\theta}(k), k) = u_2^*(\hat{\theta}(k), k)$. Define $J_k(x) = I(x^{1/k})$ and $T_1(k) = u_1^*(\hat{\theta}(k), k)^k$, $T_2(k) = u_2^*(\hat{\theta}(k), k)^k$.

3.1. Existence of k_c

Lemmas 3.1 and 3.2 establish the existence of the critical curve. Lemma 3.1 shows the connection between replica symmetry and ensemble equivalence as discussed in Section 1.5, since T is in the region of replica symmetry for $p = \frac{1}{2}$ if and only if $(T, I(T^{1/k}))$ lies on the convex minorant of J_k .

Lemma 3.1. (Connection with replica symmetry.) *Let $k \geq 1$ and $T \in \left[\left(\frac{1}{2}\right)^k, 1\right)$. There is ensemble equivalence for $T^* = T$ if and only if $(T, I(T^{1/k}))$ lies on the convex minorant of the function J_k .*

Proof. Note that $I(x) = I_{1/2}(x) - \frac{1}{2} \log 2$ (recall (1.8)), so $(T, I(T^{1/k}))$ lies on the convex minorant of J_k if and only if $(T, I_{1/2}(T^{1/k}))$ lies on the convex minorant of the function $x \mapsto I_{1/2}(x^{1/k})$.

In [15, Appendix A], it is shown that there exist $q_1, q_2 \in (0, 1)$ such that $(q^k, I(q))$ is not on the convex minorant of J if and only if $q^k \in (q_1^k, q_2^k)$. The values q_1, q_2 are defined as the unique values in $[0, 1]$ such that the tangent lines of J at q_1^k and q_2^k are the same, i.e. $J'(q_1^k) = J'(q_2^k) =: D$ and $J(q_1^k) + D(q_2^k - q_1^k) = J(q_2^k)$, or, equivalently, $Dq_1^k - J(q_1^k) = Dq_2^k - J(q_2^k)$.

Recall from Section 1.5 that there is breaking of ensemble equivalence for $T^* = T \in \left[\left(\frac{1}{2}\right)^k, 1\right)$ if and only if $T \in (u_1^k, u_2^k)$, where $u_1 = u_1^*(\hat{\theta}(k), k)$ and $u_2 = u_2^*(\hat{\theta}(k), k)$. Since u_1, u_2 are the maximisers of $x \mapsto \hat{\theta}x^k - I(x)$ and $x \mapsto x^k$ is monotone, we have that $T_1 := u_1^k$ and $T_2 := u_2^k$ are the maximisers of $x \mapsto \hat{\theta}x - I(x^{1/k}) = \hat{\theta}x - J(x)$. Hence, $J'(T_1) = J'(T_2) = \hat{\theta}$. Furthermore, $\hat{\theta}$ was defined such that $\hat{\theta}u_1^k - I(u_1) = \hat{\theta}u_2^k - I(u_2)$, so $\hat{\theta}T_1 - J(T_1) = \hat{\theta}T_2 - J(T_2)$.

From the above, we conclude that $u_1 = q_1$ and $u_2 = q_2$, which completes the proof. □

There is ensemble equivalence for $T^* \leq T_1(k)$ and $T^* \geq T_2(k)$, and ensemble inequivalence for $T^* \in (T_1(k), T_2(k))$. By [15, Lemma A.5], $k \mapsto u_1^*(\hat{\theta}, k)$ is decreasing and $k \mapsto u_2^*(\hat{\theta}, k)$ is increasing. Although $k \mapsto (u_1^*(\hat{\theta}, k))^k$ is clearly decreasing, it is not a priori obvious whether $k \mapsto (u_2^*(\hat{\theta}, k))^k$ is increasing, since $u_2^*(\hat{\theta}, k) < 1$. If the latter is the case, then for all $k > k_c(T^*)$ there is breaking of ensemble equivalence, and for all $k \leq k_c(T^*)$ there is ensemble equivalence, where $k_c(T^*)$ is chosen such that $T^* = T_1(k_c)$ or $T^* = T_2(k_c)$. This proves the first part of Theorem 1.3. Also, since $T_1(k) \geq \left(\frac{1}{2}\right)^k$, this also shows that $k_c \geq \log_2(1/T^*)$. The following lemma fills in the gap.

Lemma 3.2. (Monotonicity.) *The function $k \mapsto T_1(k)$ is decreasing and $k \mapsto T_2(k)$ is increasing.*

Proof. The function $\partial J_k / \partial k$ is a concave function for every k . Because the line segment connecting $(T_1(k), J_k(T_1(k)))$ with $(T_2(k), J_k(T_2(k)))$ lies below the curve $(x, J_k(x))$, we have that, for all $\alpha \in [0, 1]$ and $k' \downarrow k$,

$$\begin{aligned} J_{k'}(\alpha T_1(k) + (1 - \alpha)T_2(k)) &= J_k(\alpha T_1(k) + (1 - \alpha)T_2(k)) \\ &\quad + (k' - k) \frac{\partial J_k}{\partial k}(\alpha T_1(k) + (1 - \alpha)T_2(k)) + o(k' - k) \\ &\geq \alpha J_k(T_1(k)) + (1 - \alpha)J_k(T_2(k)) \\ &\quad + (k' - k) \left(\alpha \frac{\partial J_k}{\partial k}(T_1(k)) + (1 - \alpha) \frac{\partial J_k}{\partial k}(T_2(k)) \right) + o(k' - k) \\ &= \alpha J_{k'}(T_1(k)) + (1 - \alpha)J_{k'}(T_2(k)) + o(k' - k). \end{aligned}$$

Hence, for $k' > k$ small enough, the line segment connecting the points $(T_1(k), J_{k'}(T_1(k)))$ and $(T_2(k), J_{k'}(T_2(k)))$ lies below the curve $(x, J_{k'}(x))$, and is not tangent to the curve at any of the end points. Thus, by [15, Lemma A.3], $T_1(k') < T_1(k) < T_2(k) < T_2(k')$. □

3.2. Minimum of k_c

By [16, Proposition 3.2], for all $k \leq k_0$, $l_{\theta,k}$ has a unique maximiser for all $\theta \geq 0$. For all $k > k_0$, there exists a $\theta \geq 0$ such that $l_{\theta,k}$ has two maximisers. Hence, the minimum value of $k_c(T^*)$ is k_0 . In the proof of [16, Proposition 3.2] it is shown that $\hat{\theta}(k_0) = k_0^{k_0-1}/(2(k_0 - 1)^{k_0})$, and so

$$\begin{aligned} l'_{\hat{\theta}(k_0),k_0} \left(\frac{k_0 - 1}{k_0} \right) &= \frac{(k_0)^{k_0-1}}{2(k_0 - 1)^{k_0}} k_0 \left(\frac{k_0 - 1}{k_0} \right)^{k_0-1} - \frac{1}{2} \log(k_0 - 1) \\ &= \frac{k_0}{2(k_0 - 1)} - \frac{1}{2} \log(k_0 - 1) = 0. \end{aligned}$$

Hence, $u^*(\hat{\theta}(k_0), k_0) = (k_0 - 1)/k_0$, and so for $T^* = ((k_0 - 1)/k_0)^{k_0}$ we have $k_c(T^*) = k_0$. We conclude that k_c has a unique minimum at the point $((k_0 - 1)/k_0)^{k_0}, k_0$.

3.3. Analyticity of k_c

Analyticity of k_c follows from a straightforward application of the implicit function theorem. Let $f : (0, \infty) \times (0, 1)^2 \rightarrow \mathbb{R}^2$ be given by

$$f(k, x, y) = (J'_k(x) - J'_k(y), J'_k(x)x - J'_k(y)y + J(y) - J(x)).$$

Recall from the proof of Lemma 3.1 that, for each k , $T_1(k)$ and $T_2(k)$ are defined such that $f(k, T_1(k), T_2(k)) = 0$. Note that f is analytic, and its Jacobian,

$$\begin{pmatrix} \frac{\partial f_1}{\partial x} & \frac{\partial f_1}{\partial y} \\ \frac{\partial f_2}{\partial x} & \frac{\partial f_2}{\partial y} \end{pmatrix} (T_1(k), T_2(k)) = \begin{pmatrix} J''_k(T_1(k)) & -J''_k(T_2(k)) \\ T_1(k)J''_k(T_1(k)) & -T_2(k)J''_k(T_2(k)) \end{pmatrix},$$

is invertible if $T_1(k) \neq T_2(k)$. Hence, for all $k > k_0$, T_1 and T_2 are analytic functions of k , so k_c is an analytic function of T^* outside its minimum.

Next, consider the behaviour of k_c near T_0 , so as $T_2 - T_1 \downarrow 0$. By implicit differentiation, as $k \downarrow k_0$, the derivative of $T_1(k)$ is given by

$$\begin{aligned} T'_1(k) &= \frac{1}{(T_1 - T_2)J''_k(T_1)J''_k(T_2)} \\ &\quad \times \left[(T_2 - T_1)J''_k(T_2) \frac{\partial J'_k(T_1)}{\partial k} + J''_k(T_2) \left(\frac{\partial J_k(T_1)}{\partial k} - \frac{\partial J_k(T_2)}{\partial k} \right) \right] \\ &= \frac{1}{J''_k(T_1)} \left(\frac{\partial J'_k(T_1)}{\partial k} + \frac{\frac{\partial J_k(T_1)}{\partial k} - \frac{\partial J_k(T_2)}{\partial k}}{T_2 - T_1} \right) \\ &= \frac{1}{J''_k(T_1)} O(T_2 - T_1). \end{aligned}$$

It is not difficult to show that, for $k = k_0$, the function J''_{k_0} has a zero that is also a minimum at $T = T_0$. Hence, as $k \downarrow k_0$, $J''_k(T_1(k)) = O((T_2 - T_1)^2)$, which implies that the derivative of $T'_1(k)$ diverges as $k \downarrow k_0$. In a similar fashion, we can show that the derivative of $T'_2(k)$ diverges as $k \downarrow k_0$. Hence, at T_0 , k_c is at least differentiable and has derivative zero.

3.4. Scaling of k_c near the boundary

In order to identify the asymptotics of k_c for T^* near the edges of the interval $(0,1)$, we first compute the limit of $\hat{\theta}$ as $k \rightarrow \infty$. In the following, we suppress the dependence of $\hat{\theta}$ on k . By Taylor expansion,

$$l_\theta(u_1^*) \leq l_\theta\left(\frac{1}{2}\right) + \left(u_1^* - \frac{1}{2}\right)l'_\theta\left(\frac{1}{2}\right) \leq \theta\left(\frac{1}{2}\right)^k + \frac{1}{2} \log 2 + \theta k\left(\frac{1}{2}\right)^k = \theta\left(\frac{1}{2}\right)^k(1+k) + \frac{1}{2} \log 2,$$

and $l_\theta(1) = \theta < l_\theta(u_2^*)$. This implies that

$$\hat{\theta} < \frac{\log 2}{2\left[1 - \left(\frac{1}{2}\right)^k(1+k)\right]}.$$

Also, $u_2^*(\theta, k) \in ((k-1)/k, 1)$ by [16, Proposition 3.2]. Hence,

$$\begin{aligned} l_{\theta,k}(u_2^*(\theta, k)) &\leq \theta - \frac{k-1}{2k} \log\left(\frac{k-1}{k}\right) - \frac{1}{2}\left(1 - \frac{k-1}{k}\right) \log\left(1 - \frac{k-1}{k}\right) \\ &= \theta - \frac{1}{2} \log\left(1 - \frac{1}{k}\right) - \frac{1}{2k} \log\left(\frac{1}{k-1}\right), \end{aligned}$$

and $l_{\theta,k}\left(\frac{1}{2}\right) = \theta\left(\frac{1}{2}\right)^k + \frac{1}{2} \log 2 < l_{\theta,k}(u_1^*(\theta, k))$. This implies that

$$\hat{\theta} > \frac{\log 2 + \log\left(1 - \frac{1}{k}\right) + \frac{1}{k} \log\left(\frac{1}{k-1}\right)}{2\left[1 - \left(\frac{1}{2}\right)^k\right]}.$$

Combining the bounds above, we obtain that $\hat{\theta} \rightarrow \frac{1}{2} \log 2$ as $k \rightarrow \infty$.

For $T^* \downarrow 0$, Let $y \in \left(\frac{1}{2}, 1\right)$. Then

$$\begin{aligned} l'_{\hat{\theta},k}\left(\frac{1}{2} + y^k\right) &= \hat{\theta}k\left(\frac{1}{2} + y^k\right)^{k-1} - \frac{1}{2} \log\left(\frac{1+2y^k}{1-2y^k}\right) \\ &= \hat{\theta}k\left(\frac{1}{2} + y^k\right)^{k-1} - \frac{1}{2} \log\left(1 + \frac{4y^k}{1-2y^k}\right) \\ &\leq \frac{\log 2}{2\left[1 - \left(\frac{1}{2}\right)^k(1+k)\right]}k\left(\frac{1}{2}\right)^{k-1} - 2y^k + o\left(k\left(\frac{1}{2}\right)^k\right) + o(y^k) < 0 \end{aligned}$$

as $k \rightarrow \infty$. Thus, $u_1^*(\hat{\theta}, k) < \frac{1}{2} + y^k$ for all $y \in \left(\frac{1}{2}, 1\right)$ and k large enough. Hence, $\left(\frac{1}{2} + y^k\right)^{k_c} \geq T^*$ for T^* small enough. We also have $T^* \geq \left(\frac{1}{2}\right)^k$ for all k . Since this holds for all $y \in \left(\frac{1}{2}, 1\right)$ and $\left(\frac{1}{2} + y^k\right)^k \sim \left(\frac{1}{2}\right)^k$, we have $T^* \sim \left(\frac{1}{2}\right)^{k_c}$.

For $T^* \uparrow 1$, let $x \in (0, 1)$. Then

$$l'_{\hat{\theta},k}(1 - x^k) = k(\hat{\theta}(1 - x^k)^{k-1} + \frac{1}{2} \log x) - \frac{1}{2} \log(1 - x^k).$$

As $k \rightarrow \infty$, $(1 - x^k)^{k-1} \rightarrow 1$ and $\log(1 - x^k) \rightarrow 0$. Hence, if $-\frac{1}{2} \log x \geq \hat{\theta}$, then $l'_{\hat{\theta},k}(1 - x^k) < 0$ for k large enough, which implies that $u_2^*(\hat{\theta}, k) < 1 - x^k$. If $-\frac{1}{2} \log x < \hat{\theta}$, then $l'_{\hat{\theta},k}(1 - x^k) > 0$, which implies that $u_2^*(\hat{\theta}, k) > 1 - x^k$. Recall that $\hat{\theta} \rightarrow \frac{1}{2} \log 2$. Thus, choosing $x = \frac{1}{2}$, we get $\left(1 - \left(\frac{1}{2}\right)^{k_c}\right)^{k_c} \sim T^*$, and so $k_c\left(\frac{1}{2}\right)^{k_c} \sim 1 - T^*$.

4. Proof of Theorem 1.4

If $d = k$, then the statement of the theorem is vacuous, so we may assume that $d < k$. Let T^* denote either $T_1^*(k)$ or $T_2^*(k)$. In this proof, we will often use that fact that $I(f) = I_{1/2}(f) - \frac{1}{2} \log 2$. Any reference to the theory of replica symmetry is made with the implicit assumption that $p = \frac{1}{2}$.

Since there is ensemble equivalence for T^* , $(T^*, I((T^*)^{1/k}))$ lies on the convex minorant of $x \mapsto I(x^{1/k})$, and so $T^* \notin (q_1(k)^k, q_2(k)^k)$, where $q_1(k), q_2(k)$ are defined as in the proof of Lemma 3.1. By [15, Lemma A.5], $q_1(k) < q_1(d) < q_2(d) < q_2(k)$, because $d < k$, with d the largest degree of H . Hence, for all $T \in (T_1^*(k), q_1(d))$ and $T \in [q_2(d), T_2^*(k))$, $(T, I(T^{1/d}))$ lies on the convex minorant of $x \mapsto I(x^{1/d})$, but T is not in the region of ensemble equivalence. Thus, by [15, Lemma 3.3], T is in the region of replica symmetry for $t(H, \cdot)$. This implies that $h \equiv T^{1/k}$ is the unique minimiser of

$$\inf\{I(\tilde{h}) : \tilde{h} \in \tilde{\mathcal{W}}, t(H, \tilde{h}) \geq T\} = \inf\{I(\tilde{h}) : \tilde{h} \in \tilde{\mathcal{W}}, t(H, \tilde{h}) = T\} = \inf_{\tilde{h} \in \tilde{\mathcal{W}}^*} I(\tilde{h}).$$

Furthermore, since T is in the BEE-phase, we have $\theta^* = \hat{\theta}$. We conclude that

$$\begin{aligned} s_\infty &= \sup_{\tilde{h} \in \tilde{\mathcal{W}}} [\theta^* T(\tilde{h}) - I(\tilde{h})] - \sup_{\tilde{h} \in \tilde{\mathcal{W}}^*} [\theta^* T(\tilde{h}) - I(\tilde{h})] \\ &= [\hat{\theta} T^* - I(T^{*1/k})] - [\hat{\theta} T - I(T^{1/k})] \\ &= \hat{\theta}(T^* - T) + [I(T^{1/k}) - I(T^{*1/k})] \\ &= [J'_k(T^*) - \hat{\theta}](T - T^*) + J''_k(T^*)(T - T^*)^2 + O((T - T^*)^3) \\ &= \frac{T^{*1/k-2}}{2k} \left\{ \frac{1}{k} \left(1 + \frac{T^{*1/k}}{1 - T^{*1/k}} \right) + \left(\frac{1}{k} - 1 \right) \log \left(\frac{T^{*1/k}}{1 - T^{*1/k}} \right) \right\} (T - T^*)^2 + O((T - T^*)^3) \end{aligned}$$

as $T \rightarrow T^*$. The last equality follows from the fact that $J'_k(T^*) = \hat{\theta}$ (see the proof of Lemma 3.1).

5. Proof of Theorem 1.5

We first show that a graph sampled from the canonical ensemble converges to a probability distribution on a finite set of constant graphons. In [4, Theorems 3.2 and 4.2] this is shown for the exponential random graph model with a fixed parameter θ^* . We adapt the proof to the case where we have a sequence of parameters $(\theta_n^*)_{n \in \mathbb{N}}$ converging to some θ^* .

Lemma 5.1. *Let G_n be a random graph drawn from the canonical ensemble \mathbb{P}_{can} with parameter θ_n^* . Let $U(\theta)$ be the set of maximisers of (1.6) for some parameter θ . Then, recalling (1.4), $\min_{u \in U(\theta_\infty^*)} \delta_{\square}(\tilde{h}^{G_n}, \tilde{u}) \rightarrow 0$ as $n \rightarrow \infty$ in probability.*

Proof. Let $\eta > 0$ and define $\tilde{A}(\theta, \eta) := \{\tilde{h} \in \tilde{\mathcal{W}} \mid \delta_{\square}(\tilde{h}, \tilde{U}(\theta)) \geq \eta\}$. Recall from the proof of Theorem 1.2 that $U(\theta)$ consists of a single point for $\theta \neq \hat{\theta}$ and two points for $\theta = \hat{\theta}$. Also recall the definition of the function l_θ from the proof of Theorem 1.2. We first prove the case that $\theta_\infty^* \neq \hat{\theta}$. Then $l_{\theta_n^*}$ converges to $l_{\theta_\infty^*}$ uniformly as $n \rightarrow \infty$, so $U(\theta_n^*)$ converges to $U(\theta_\infty^*)$. Here we assume without loss of generality that $\theta_n^* \neq \hat{\theta}$ and let $U(\theta)$ denote the single maximiser of l_θ by a slight abuse of notation. Hence,

$$\tilde{A}(\theta_n^*, \eta) \subset \tilde{A}(\theta_\infty^*, \eta/2) \tag{5.1}$$

for all n large enough by the triangle inequality. We now adapt the arguments from the proof of [4, Theorem 3.2].

By the compactness of $\tilde{\mathcal{W}}$ and $\tilde{U}(\theta)$, and upper semi-continuity of $\theta_\infty^* T - I$, it follows that

$$2\varepsilon := \sup_{\tilde{h} \in \tilde{\mathcal{W}}} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})] - \sup_{\tilde{h} \in \tilde{A}(\theta_\infty^*, \eta/2)} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})] > 0.$$

Since the $\theta_n^* T$ are all bounded functions and the sequence $(\theta_n^*)_{n \in \mathbb{N}}$ is bounded, there exists a finite set R such that the intervals $\{(a, a + \varepsilon) \mid a \in R\}$ cover the range of $\theta_n^* T$ and $\theta_\infty^* T$ for all n large enough. For each $a \in R$, let $\tilde{F}^a(\theta_n^*) := (\theta_n^* T)^{-1}([a, a + \varepsilon])$. Now define $\tilde{A}^a(\theta_n^*, \eta) := \tilde{A}(\theta_n^*, \eta) \cap \tilde{F}^a(\theta_n^*)$ and $\tilde{A}_n^a(\theta_n^*, \eta) = \tilde{A}^a(\theta_n^*, \eta) \cap \tilde{\mathcal{G}}_n$. Choose $\delta = \frac{1}{2}\varepsilon$. Since $\theta_n^* \rightarrow \theta_\infty^*$, we have

$$(\theta_n^* T)^{-1}([a, a + \varepsilon]) \subset (\theta_\infty^* T)^{-1}([a - \delta, a + \varepsilon + \delta]) =: \tilde{G}^a \tag{5.2}$$

for all n large enough. Now define $\tilde{B}^a := \tilde{A}(\theta_\infty^*, \eta/2) \cap \tilde{G}^a$ and $\tilde{B}_n^a := \tilde{B}^a \cap \tilde{\mathcal{G}}_n$.

Using (5.1) and (5.2), we obtain $\tilde{A}_n^a(\theta_n^*, \eta) \subset \tilde{B}_n^a$. Hence,

$$\begin{aligned} \mathbb{P}_{\text{can}}(G_n \in \tilde{A}(\theta_n^*, \eta)) &\leq e^{-n^2 \psi_n(\theta_n^*)} \sum_{a \in R} e^{n^2(a+\varepsilon)} |\tilde{A}_n^a(\theta_n^*, \eta)| \\ &\leq e^{-n^2 \psi_n(\theta_n^*)} \sum_{a \in R} e^{n^2(a+\varepsilon)} |\tilde{B}_n^a| \\ &\leq e^{-n^2 \psi_n(\theta_n^*)} |R| \sup_{a \in R} e^{n^2(a+\varepsilon)} |\tilde{B}_n^a|. \end{aligned}$$

Using the large deviation principle for the Erdős–Rényi random graph in [4, (8.1)], we obtain

$$\limsup_{n \rightarrow \infty} \frac{\log |\tilde{B}_n^a|}{n^2} \leq - \inf_{\tilde{h} \in \tilde{B}^a} I(\tilde{h}).$$

Also, by [6, Lemma A.1], we have

$$\lim_{n \rightarrow \infty} \psi_n(\theta_n^*) = \psi_\infty(\theta_\infty^*) = \sup_{\tilde{h} \in \tilde{\mathcal{W}}} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})].$$

Combining these two results, we conclude that

$$\limsup_{n \rightarrow \infty} \frac{\log \mathbb{P}_{\text{can}}(G_n \in \tilde{A}(\theta_n^*, \eta))}{n^2} \leq \sup_{a \in R} \left[a + \varepsilon - \inf_{\tilde{h} \in \tilde{B}^a} I(\tilde{h}) \right] - \sup_{\tilde{h} \in \tilde{\mathcal{W}}} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})]. \tag{5.3}$$

The remainder of the proof now follows exactly as in [4]. Indeed, for each $\tilde{h} \in \tilde{B}^a$, we have $\theta_\infty^* T(\tilde{h}) \geq a - \delta$. Hence,

$$\sup_{\tilde{h} \in \tilde{B}^a} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})] \geq a - \delta - \inf_{\tilde{h} \in \tilde{B}^a} I(\tilde{h}).$$

Substituting this into (5.3), we get

$$\begin{aligned} \limsup_{n \rightarrow \infty} \frac{\log \mathbb{P}_{\text{can}}(G_n \in \tilde{A}(\theta_n^*, \eta))}{n^2} &\leq \varepsilon + \delta + \sup_{a \in R} \sup_{\tilde{h} \in \tilde{B}^a} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})] - \sup_{\tilde{h} \in \tilde{\mathcal{W}}} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})] \\ &\leq \varepsilon + \delta + \sup_{\tilde{h} \in \tilde{A}(\theta_\infty^*, \eta/2)} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})] - \sup_{\tilde{h} \in \tilde{\mathcal{W}}} [\theta_\infty^* T(\tilde{h}) - I(\tilde{h})] \\ &\leq \varepsilon + \delta - 2\varepsilon = -\frac{\varepsilon}{2}. \end{aligned}$$

We have thus shown that $\delta_{\square}(\tilde{h}^{G_n}, \tilde{U}(\theta_n^*)) \rightarrow 0$ as $n \rightarrow \infty$ in probability. Since $U(\theta_n^*) \rightarrow U(\theta_{\infty}^*)$, this concludes the proof in the case $\theta_{\infty} \neq \hat{\theta}$.

Now assume that $\theta_{\infty} = \hat{\theta}$. Then (5.1) may no longer hold, since $U(\theta_{\infty}^*)$ now consists of two points, whereas $U(\theta_n^*)$ may consist of only one point. However, if we define $U'(\theta)$ as the set consisting of the two local maxima of l_{θ} , and define $\tilde{A}'(\theta, \eta)$ analogously, then the analogue of (5.1) does hold for all n large enough. Here we use that the two local maxima of $l_{\theta_n^*}$ converge to the two local maxima of $l_{\theta_{\infty}^*}$. The rest of the proof then goes through as before to show that $\delta_{\square}(\tilde{h}^{G_n}, \tilde{U}'(\theta_n^*)) \rightarrow 0$ in probability. Again using convergence of the local maxima, we obtain $\delta_{\square}(\tilde{h}^{G_n}, \tilde{U}'(\theta_{\infty}^*)) \rightarrow 0$ in probability. However, for $\theta_{\infty}^* = \hat{\theta}$, we have that $\tilde{U}'(\theta_{\infty}^*) = \tilde{U}(\theta_{\infty}^*)$, which concludes the proof. \square

Corollary 5.1. *Assume that T^* is in the BEE-phase. Let G_n be a random graph drawn from the canonical ensemble P_{can} . Then h^{G_n} converges weakly to*

$$\frac{u_2^k - T^*}{u_2^k - u_1^k} \delta_{u_1} + \frac{T^* - u_1^k}{u_2^k - u_1^k} \delta_{u_2},$$

with $u_1 < u_2$ the two maximisers of (1.6) for $\theta = \hat{\theta}$.

Proof. From Lemma 5.1 it is clear that the laws of $(G_n)_{n \in \mathbb{N}}$ form a tight sequence of probability measures. Hence, by Prokhorov’s theorem, for every subsequence $(n_k)_{k \in \mathbb{N}}$ there exists a further subsequence $(n_{k_l})_{l \in \mathbb{N}}$ such that $(G_{n_{k_l}})_{l \in \mathbb{N}}$ converges weakly to the random graphon $p\delta_{u_1} + (1 - p)\delta_{u_2}$ for some $p \in [0, 1]$. Since the homomorphism density is continuous and bounded, this implies that $(E_{\text{can}}[t(H, G_{n_{k_l}})])_{l \in \mathbb{N}}$ converges to $pu_1^k + (1 - p)u_2^k$. However, by the definition of the canonical ensemble, this sequence also converges to T^* . Hence,

$$T^* = \lim_{l \rightarrow \infty} E_{\text{can}}[t(H, G_{n_{k_l}})] = pu_1^k + (1 - p)u_2^k.$$

Solving for p , we obtain that $(G_{n_{k_l}})_{l \in \mathbb{N}}$ converges weakly to

$$\frac{u_2^k - T^*}{u_2^k - u_1^k} \delta_{u_1} + \frac{T^* - u_1^k}{u_2^k - u_1^k} \delta_{u_2}.$$

Since the subsequence $(n_k)_{k \in \mathbb{N}}$ is arbitrary and the expression above does not depend on the chosen subsequence, we conclude that weak convergence holds for the sequence $(G_n)_{n \in \mathbb{N}}$. \square

We can also show convergence of the microcanonical ensemble.

Lemma 5.2. *Let G_n be a random graph drawn from the microcanonical ensemble P_{mic} . Then \tilde{h}^{G_n} converges in probability to \tilde{F}^* , with \tilde{F}^* the set of minimisers in $\tilde{\mathcal{W}}^*$ of I .*

Proof. The proof is similar to the proof of [5, Theorem 3.1]. Fix $\varepsilon > 0$ and let

$$\begin{aligned} \tilde{F}^{\varepsilon} &:= \{\tilde{h} \in \tilde{\mathcal{W}}^* \mid \delta_{\square}(\tilde{h}, \tilde{F}^*) > \varepsilon\}, \\ \tilde{F}_n^{\varepsilon} &:= \{\tilde{h} \in \tilde{F}_n^{\varepsilon} \mid \delta_{\square}(\tilde{h}, \tilde{F}^*) > \varepsilon, \tilde{h} = \tilde{G} \text{ for some } G \in \mathcal{G}_n\}. \end{aligned}$$

Then, by [6, (3.22) and Corollary 2.9],

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n^2} \log P_{\text{mic}}(\tilde{F}^\varepsilon) &= \lim_{n \rightarrow \infty} \frac{1}{n^2} \log (|\tilde{F}_n^\varepsilon| P_{\text{mic}}(G_n = G_n^*)) \\ &= \inf_{\tilde{h} \in \tilde{\mathcal{V}}^*} I(\tilde{h}) + \lim_{n \rightarrow \infty} \frac{1}{n^2} \log |\tilde{F}_n^\varepsilon| \\ &= \inf_{\tilde{h} \in \tilde{\mathcal{V}}^*} I(\tilde{h}) - \inf_{\tilde{h} \in \tilde{F}^\varepsilon} I(\tilde{h}), \end{aligned}$$

where G_n^* is any graph in \mathcal{G}_n such that $\tilde{G}_n^* \in \tilde{\mathcal{V}}^*$. Since $\tilde{\mathcal{V}}^*$ is a compact set and \tilde{F}^ε does not contain any minimisers of $\inf_{\tilde{h} \in \tilde{\mathcal{V}}^*} I(\tilde{h})$, we conclude that the expression above is negative, which implies that $\lim_{n \rightarrow \infty} P_{\text{mic}}(\tilde{F}^\varepsilon) = 0$. \square

We next turn our attention to the largest eigenvalue. For a graph G_n on n vertices, $n^{-1} \lambda_n(G_n)$ equals the operator norm $\|h^{G_n}\|_{\text{op}}$ of the empirical graphon of G_n . The operator norm is continuous and bounded, so we have

$$\lim_{n \rightarrow \infty} n^{-1} E_{\text{can}}[\lambda_n] = pu_1 + (1 - p)u_2 = \frac{T^*(u_2 - u_1) + u_1u_2(u_2^{k-1} - u_1^{k-1})}{u_2^k - u_1^k} =: f(T^*).$$

If T^* is in the region of replica symmetry for the subgraph H , then $h \equiv (T^*)^{1/k}$ is the unique minimiser of I in $\tilde{\mathcal{V}}^*$. So, in this case,

$$\lim_{n \rightarrow \infty} n^{-1} E_{\text{mic}}[\lambda_n] = (T^*)^{1/k} > f(T^*),$$

since the function $x \mapsto x^{1/k}$ is concave, f is affine in T^* , and we have $f(u_1^k) = u_1 = (u_1^k)^{1/k}$ and $f(u_2^k) = u_2 = (u_2^k)^{1/k}$.

The second part of the theorem follows from a simple Taylor expansion.

Funding Information

The research in this paper was supported through NWO Gravitation Grant NETWORKS 024.002.003.

Competing Interests

There were no competing interests to declare which arose during the preparation or publication process of this article.

References

- [1] BORGES, C., CHAYES, J. T., LOVÁSZ, L., SÓS, V. T. AND VESZTERGOMBI, K. (2008). Convergent graph sequences I: Subgraph frequencies, metric properties, and testing. *Adv. Math.* **219**, 1801–1851.
- [2] BORGES, C., CHAYES, J. T., LOVÁSZ, L., SÓS, V. T. AND VESZTERGOMBI, K. (2012). Convergent sequences of dense graphs II: Multiway cuts and statistical physics. *Ann. Math.* **176**, 151–219.
- [3] CHATTERJEE, S. (2015). *Large Deviations for Random Graphs*, École d'Été de Probabilités de Saint-Flour XLV.
- [4] CHATTERJEE, S. AND DIACONIS, P. (2013). Estimating and understanding exponential random graph models. *Ann. Statist.* **5**, 2428–2461.
- [5] CHATTERJEE, S. AND VARADHAN, S. R. S. (2011). The large deviation principle for the Erdős–Rényi random graph. *Europ. J. Combinatorics* **32**, 1000–1017.
- [6] DEN HOLLANDER, F., MANDJES, M., ROCCAVERDE, A. AND STARREVELD, N. J. (2018). Ensemble equivalence for dense graphs. *Electron. J. Probab.* **23**, 12.

- [7] DEN HOLLANDER, F., MANDJES, M., ROCCAVERDE, A. AND STARREVELD, N. J. (2018). Breaking of ensemble equivalence for perturbed Erdős–Rényi random graphs. Preprint, [arXiv:1807.07750](https://arxiv.org/abs/1807.07750).
- [8] DIAO, P., GUILLOT, D., KHARE, A. AND RAJARATNAM, B. (2015). Differential calculus on graphon space. *J. Combinatorial Theory Ser A* **133**, 183–227.
- [9] DIONIGI, P., GARLASCHELLI, D., DEN HOLLANDER, F. AND MANDJES, M. (2021). A spectral signature of breaking of ensemble equivalence for constrained random graphs. *Electron. Commun. Probab.* **26**, 1–15.
- [10] GARLASCHELLI, G., DEN HOLLANDER, F. AND ROCCAVERDE, A. (2018). Covariance structure behind breaking of ensemble equivalence. *J. Statist. Phys.* **173**, 644–662.
- [11] GARLASCHELLI, G., DEN HOLLANDER, F. AND ROCCAVERDE, A. (2021). Ensemble equivalence in random graphs with modular structure. *J. Phys. A* **50**, 015001.
- [12] GIBBS, J. W. (1902). *Elementary Principles of Statistical Mechanics*. Yale University Press, New Haven, CT.
- [13] JAYNES, E. T. (1957). Information theory and statistical mechanics. *Phys. Rev.* **106**, 620–630.
- [14] LOVÁSZ, L. AND SZEGEDY, B. (2007). Szémerédi’s lemma for the analyst. *Geom. Funct. Anal.* **17**, 252–270.
- [15] LUBETZKY, E. AND ZHAO, Y. (2015). On replica symmetry of large deviations in random graphs. *Random Structures Algorithms* **47**, 109–146.
- [16] RADIN, C. AND YIN, M. (2013). Phase transitions in exponential random graphs, *Ann. Appl. Probab.* **6**, 2458–2471.
- [17] SQUARTINI, T. AND GARLASCHELLI, D. (2017). Reconnecting statistical physics and combinatorics beyond ensemble equivalence. Preprint, [arXiv:1710.11422](https://arxiv.org/abs/1710.11422).
- [18] SQUARTINI, T., DE MOL, J., DEN HOLLANDER, F. AND GARLASCHELLI, D. (2015). Breaking of ensemble equivalence in networks. *Phys. Rev. Lett.* **115**, 268701.
- [19] TOUCHETTE, H. (2015). Equivalence and nonequivalence of ensembles: Thermodynamic, macrostate, and measure levels. *J. Statist. Phys.* **159**, 987–1016.
- [20] TOUCHETTE, H., ELLIS, R. S. AND TURKINGTON, B. (2004). An introduction to the thermodynamic and macrostate levels of nonequivalent ensembles. *Physica A* **340**, 138–146.