EURANDOM PREPRINT SERIES

2021-007

July 9, 2021

Breaking of ensemble equivalence for dense random graphs under a single constraint

F. den Hollander, M Markering
ISSN 1389-2355

Breaking of ensemble equivalence for dense random graphs under a single constraint

F. den Hollander [1]
M. Markering [2]

July 8, 2021

Abstract

Two ensembles are often 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. The latter means that it matters for the scaling properties of the graph whether the constraint is met for every single realisation of the graph or only holds as an ensemble average. Various examples were analysed in the literature, and the specific relative entropy was computed as a function of the constraint. It was found that BEE is the rule rather than the exception for two classes: *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*.

In the present paper we establish BEE for a third class: dense random graphs with a *single* constraint, namely, on the density of a given finite simple graph. We show that BEE occurs only 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 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 *spectral signature* of BEE. Proofs are based on an analysis of the variational formula on the space of graphons for the limiting specific relative entropy derived in [13], in combination with an identification of the minimising graphons and replica symmetry arguments. We show that in the replica symmetric region of the BEE-phase, as the size of the graph tends to infinity, the microcanonical ensemble behaves like an Erdős-Rényi random graph, while the canonical ensemble behaves like a *mixture* of two Erdős-Rényi random graphs. In other words, BEE is due to *coexistence* of two densities.

MSC2020: 05C80, 60C05, 60F10, 82B20.
Keywords: Constrained random graphs, Gibbs ensembles, Relative entropy, Breaking of ensemble equivalence, Graphons, Variational representations, Maximal eigenvalues, Replica symmetry.
Acknowledgement: The research in this paper was supported through NWO Gravitation Grant NETWORKS 024.002.003.

[1]Mathematical Institute, Leiden University, P.O. Box 9512, 2300 RA Leiden, The Netherlands, *denholla@math.leidenuniv.nl*

[2]Mathematical Institute, Leiden University, P.O. Box 9512, 2300 RA Leiden, The Netherlands, *maartenmarkering@outlook.com*

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 [13], 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.

1.1 Background and motivation

In this paper we analyse random graphs that are subject to *constraints*. Statistical physics prescribes what probability distribution on the set of graphs we should choose when we want to model a given type of constraint [11]. Two important choices are:

(1) The *microcanonical ensemble*, where the constraints are *hard* (i.e., are satisfied by each individual graph).

(2) The *canonical ensemble*, where the constraints are *soft* (i.e., hold as ensemble averages, while individual graphs may violate the constraints).

For random graphs that are large but finite, the two ensembles are obviously different and, in fact, represent different empirical situations. Each ensemble represents the unique probability distribution with *maximal entropy* respecting the constraints. In the limit as the size of the graph diverges, the two ensembles are traditionally *assumed* to become equivalent as a result of the expected vanishing of the fluctuations in the soft constraints, i.e., the soft constraints are expected to behave asymptotically like hard constraints. This assumption of *ensemble equivalence* is one of the cornerstones of statistical physics, but it does *not* hold in general. We refer to [20] for more background on this phenomenon.

In a series of papers *breaking of ensemble equivalence* (BEE) for various choices of the constraints was investigated, including the degree sequence and the total number of subgraphs of a specific type. Both the *sparse regime* (where the number of edges per vertex remains bounded) and the *dense regime* (where the number of edges per vertex is of the order of the number of vertices) were considered. Let S_n be the *relative entropy* of the microcanonical ensemble with respect to the canonical ensemble when the graph has n vertices. In the sparse regime the relevant quantity is $s_\infty = \lim_{n\to\infty} n^{-1} S_n$, because n is the scale of the number of *vertices*. In the dense regime the relevant quantity is $s_\infty = \lim_{n\to\infty} n^{-2} S_n$, because n^2 is the scale of the number of *edges*.

- In [19, 9, 10] it was shown that, in the sparse regime, constraining the degrees of *all* the vertices leads to BEE, even when the graph consists of multiple communities. The main result was an explicit formula for s_∞ in terms of the limit of the empirical degree distribution of the constraint. In [18] a formula was put forward that expresses the specific relative entropy in terms of a covariance matrix under the canonical ensemble. This formula is a powerful computational tool.

- In [13] it was shown that, in the dense regime, 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 random graph [3, 5]. The main result was a variational formula for s_∞ in the space of graphons. Also this variational formula is a powerful computational tool. In [14], 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. Erdős-Rényi random graphs with constraints on the number of of edges and triangles have been studied extensively. In [14] this knowledge was crucial for arriving at a detailed analysis of the variational formula for s_∞.

Naively, we might expect that a *single* constraint cannot lead to BEE because 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 the BEE-phase. We analyse how s_∞ tends to zero near the curve that borders the BEE-phase.

In [8] 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 $n \to \infty$. 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 $\delta_\infty = \lim_{n\to\infty} n^{-1}\Delta_n$, 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 [16] 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.

1.2 Definitions and preliminaries

In this section, which is largely lifted from [13], we present the definitions of the main concepts to be used in the sequel, together with some key results from prior work. We consider general vector-valued constraints, even though later we will only focus on scalar-valued constraints.

Section 1.2.1 presents the formal definition of the two ensembles we are interested in and gives our definition of ensemble equivalence in the dense regime. Section 1.2.2 recalls some basic facts about graphons, while Section 1.2.3 recalls some basic properties of the canonical ensemble. Section 1.2.4 looks at convergence of Lagrange multipliers. Section 1.2.5 provides a variational characterisation of ensemble equivalence, proven in [13]. Section 1.2.6 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 [8] 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. Any graph $G \in \mathcal{G}_n$ can be represented by a symmetric $n \times n$ matrix A^G with elements

$$A^G(i,j) := \begin{cases} 1 & \text{if there is an edge between vertex } i \text{ and vertex } j, \\ 0 & \text{otherwise.} \end{cases} \tag{1.1}$$

Let $\vec{C}$ denote a vector-valued function on $\mathcal{G}_n$. We choose a specific vector $\vec{C}^*$, which we assume to be *graphical*, i.e., realisable by at least one graph in $\mathcal{G}_n$. Given $\vec{C}^*$, the *microcanonical ensemble* is the probability distribution $\mathrm{P}_{\mathrm{mic}}$ on $\mathcal{G}_n$ with *hard constraint* $\vec{C}^*$ defined as

$$\mathrm{P}_{\mathrm{mic}}(G) := \begin{cases} 1/\Omega_{\vec{C}^*}, & \text{if } \vec{C}(G) = \vec{C}^*, \\ 0, & \text{otherwise,} \end{cases} \qquad G \in \mathcal{G}_n, \tag{1.2}$$

where

$$\Omega_{\vec{C}^*} := |\{G \in \mathcal{G}_n \colon \vec{C}(G) = \vec{C}^*\}| \tag{1.3}$$

is the number of graphs that realise $\vec{C}^*$. The *canonical ensemble* $\mathrm{P}_{\mathrm{can}}$ is the unique probability distribution on $\mathcal{G}_n$ that maximises the *entropy*

$$S_n(\mathrm{P}) := -\sum_{G \in \mathcal{G}_n} \mathrm{P}(G) \log \mathrm{P}(G) \tag{1.4}$$

subject to the *soft constraint* $\langle \vec{C} \rangle = \vec{C}^*$, where

$$\langle \vec{C} \rangle := \sum_{G \in \mathcal{G}_n} \vec{C}(G)\, \mathrm{P}(G). \tag{1.5}$$

This gives the formula [15]

$$\mathrm{P}_{\mathrm{can}}(G) := \frac{1}{Z(\vec{\theta}^*)}\, \mathrm{e}^{H(\vec{\theta}^*, \vec{C}(G))}, \qquad G \in \mathcal{G}_n, \tag{1.6}$$

with

$$H(\vec{\theta}^*, \vec{C}(G)) := \vec{\theta}^* \cdot \vec{C}(G), \qquad Z(\vec{\theta}^*) := \sum_{G \in \mathcal{G}_n} \mathrm{e}^{\vec{\theta}^* \cdot \vec{C}(G)}, \tag{1.7}$$

denoting the *Hamiltonian* and the *partition function*, respectively. In (1.6)–(1.7) the parameter $\vec{\theta}^*$, which is a real-valued vector whose dimension is equal to the number of constraints, must be set to the unique value that realises $\langle \vec{C} \rangle = \vec{C}^*$. As a Lagrange multiplier, $\vec{\theta}^*$ always exists, but uniqueness is non-trivial. In the sequel we will only consider examples where the gradients of the constraints in (1.5) are *linearly independent* vectors. Consequently, the Hessian matrix of the entropy of the canonical ensemble in (1.6) is a positive-definite matrix, which implies uniqueness.

The *relative entropy* of $\mathrm{P}_{\mathrm{mic}}$ with respect to $\mathrm{P}_{\mathrm{can}}$ is defined as

$$S_n(\mathrm{P}_{\mathrm{mic}} \mid \mathrm{P}_{\mathrm{can}}) := \sum_{G \in \mathcal{G}_n} \mathrm{P}_{\mathrm{mic}}(G) \log \frac{\mathrm{P}_{\mathrm{mic}}(G)}{\mathrm{P}_{\mathrm{can}}(G)}. \tag{1.8}$$

For any $G_1, G_2 \in \mathcal{G}_n$, $\mathrm{P}_{\mathrm{can}}(G_1) = \mathrm{P}_{\mathrm{can}}(G_2)$ whenever $\vec{C}(G_1) = \vec{C}(G_2)$, i.e., the canonical probability is the same for all graphs with the same value of the constraint. We may therefore rewrite (1.8) as

$$S_n(\mathrm{P}_{\mathrm{mic}} \mid \mathrm{P}_{\mathrm{can}}) = \log \frac{\mathrm{P}_{\mathrm{mic}}(G^*)}{\mathrm{P}_{\mathrm{can}}(G^*)}, \tag{1.9}$$

where G^* is *any* graph in $\mathcal{G}_n$ such that $\vec{C}(G^*) = \vec{C}^*$ (recall that we assumed that $\vec{C}^*$ is realisable by at least one graph in $\mathcal{G}_n$). The removal of the sum over $\mathcal{G}_n$ constitutes a major simplification.

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 $S_n(\mathrm{P}_{\mathrm{mic}} \mid \mathrm{P}_{\mathrm{can}})$. When we pass to the limit $n \to \infty$, we need to specify how $\vec{C}(G)$, $\vec{C}^*$ and $\vec{\theta}^*$ are chosen to depend on n. We refer the reader to [13], where this issue was discussed in detail.

Definition 1.1. [Ensemble equivalence] In the dense regime, if

$$s_\infty := \lim_{n\to\infty} n^{-2} S_n(\mathrm{P}_{\mathrm{mic}} \mid \mathrm{P}_{\mathrm{can}}) = 0, \tag{1.10}$$

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

Remark 1.2. In [19], which was concerned with the *sparse regime*, the relative entropy was divided by n (the number of vertices). In the *dense regime*, however, it is appropriate to divide by n^2 (the order of the number of edges).

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 $h\colon [0,1]^2 \to [0,1]$ 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 (see Figure 1)

$$h^G(x,y) := \begin{cases} 1 & \text{if there is an edge between vertex } \lceil nx \rceil \text{ and vertex } \lceil ny \rceil, \\ 0 & \text{otherwise.} \end{cases} \tag{1.11}$$

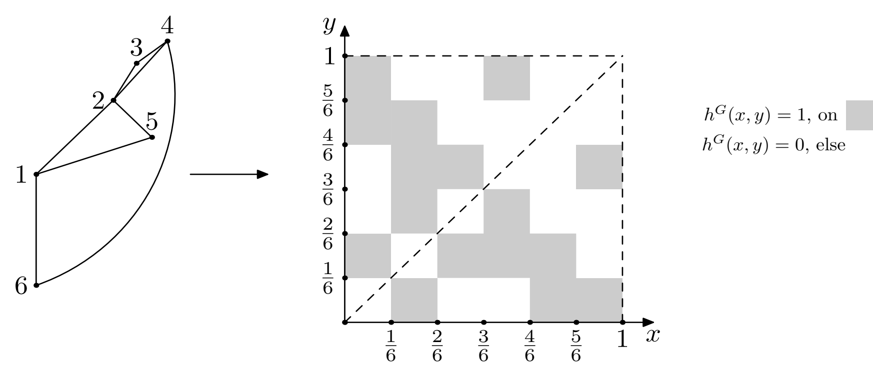

Figure 1: An example of a graph G and its graphon representation h^G.

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

$$d_\square(h_1,h_2) := \sup_{S,T\subset[0,1]} \left| \int_{S\times T} \mathrm{d}x\,\mathrm{d}y\,[h_1(x,y) - h_2(x,y)] \right|, \qquad h_1, h_2 \in \mathcal{W}. \tag{1.12}$$

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

$$\delta_\square(\tilde{h}_1,\tilde{h}_2) := \inf_{\sigma_1,\sigma_2\in\Sigma} d_\square(h_1^{\sigma_1}, h_2^{\sigma_2}), \qquad \tilde{h}_1,\tilde{h}_2 \in \tilde{\mathcal{W}}, \tag{1.13}$$

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, 7].

1.2.3 Subgraph counts

Fix $m \in \mathbb{N}$. Let $(F_i)_{i=1}^m$ be any collection of finite simple graphs. Let $C_i(G)$ denote the number of subgraphs F_i in G. In the dense regime, $C_i(G)$ grows like n^{V_i} as $n \to \infty$, where V_i is the number of vertices in F_i. Consider the following *scaled vector-valued function* on $\mathcal{G}_n$:

$$\vec{C}(G) := \left(\frac{p(F_i)\, C_i(G)}{n^{V_i-2}}\right)_{i=1}^m = n^2 \left(\frac{p(F_i)\, C_i(G)}{n^{V_i}}\right)_{i=1}^m. \tag{1.14}$$

The term $p(F_i)$ counts the edge-preserving permutations of the vertices of F_i. The term $C_i(G)/n^{V_i}$ represents the density of F_i in G. The additional n^2 guarantees that the full vector scales like n^2, in line with the scaling of the large deviation principle for graphons in the Erdős-Rényi random graph derived in [5]. Let $\hom(F_i, G)$ be the number of homomorphisms from F_i to G, and define the *homomorphism density* as

$$t(F_i, G) := \frac{\hom(F_i,G)}{n^{V_i}} = \frac{p(F_i)\, C_i(G)}{n^{V_i}}, \tag{1.15}$$

which does not distinguish between permutations of the vertices. In terms of this quantity, the Hamiltonian associated with the constraint of the homomorphism densities of $(F_i)_{i=1}^m$ becomes

$$H(\vec{\theta}, \vec{T}(G)) = n^2 \sum_{i=1}^m \theta_i\, t(F_i, G) = n^2(\vec{\theta}\cdot\vec{T}(G)), \qquad G \in \mathcal{G}_n, \tag{1.16}$$

where $\cdot$ denotes the inner product for vectors, and

$$\vec{T}(G) := (t(F_i,G))_{i=1}^m. \tag{1.17}$$

The canonical ensemble with parameter $\vec{\theta}$ takes the form

$$\mathrm{P}_{\mathrm{can}}(G \mid \vec{\theta}\,) := \mathrm{e}^{n^2\left[\vec{\theta}\cdot\vec{T}(G) - \psi_n(\vec{\theta})\right]}, \qquad G \in \mathcal{G}_n, \tag{1.18}$$

where ψ_n replaces the *partition function* $Z(\vec{\theta})$:

$$\psi_n(\vec{\theta}) := \frac{1}{n^2} \log \sum_{G\in\mathcal{G}_n} \mathrm{e}^{n^2(\vec{\theta}\,\cdot\,\vec{T}(G))} = \frac{1}{n^2}\log Z(\vec{\theta}). \tag{1.19}$$

In the sequel we take $\vec{\theta}$ equal to a specific value $\vec{\theta}^*$ so as to meet the soft constraint, i.e.,

$$\langle \vec{T} \rangle = \sum_{G\in\mathcal{G}_n} \vec{T}(G)\, \mathrm{P}_{\mathrm{can}}(G) = \vec{T}^* \tag{1.20}$$

for some choice of $\vec{T}^* \in [0,1]^m$. With this choice, the canonical probability becomes

$$\mathrm{P}_{\mathrm{can}}(G) = \mathrm{P}_{\mathrm{can}}(G \mid \vec{\theta}^*). \tag{1.21}$$

1.2.4 Convergence of Lagrange multipliers

Both the constraint $\vec{T}^*$ and the *Lagrange multiplier* $\vec{\theta}^*$ in general depend on n, i.e., $\vec{T}^* = \vec{T}^*_n$ and $\vec{\theta}^* = \vec{\theta}^*_n$. We consider constraints that converge when we pass to the limit $n \to \infty$, i.e.,

$$\lim_{n\to\infty} \vec{T}^*_n =: \vec{T}^*_\infty. \tag{1.22}$$

Consequently, we expect that

$$\lim_{n\to\infty} \vec{\theta}^*_n =: \vec{\theta}^*_\infty. \tag{1.23}$$

Throughout the paper we *assume* that (1.23) holds. If convergence fails, then we may still consider subsequential convergence. The subtleties concerning (1.23) were discussed in detail in [13, Appendix A].

In what follows we need the following lemma, which relates T^*_∞ and $\vec{\theta}^*_\infty$ without requiring knowledge of T^*_n and $\vec{\theta}^*_n$.

Lemma 1.3. *Let $\vec{T}\colon \tilde{\mathcal{W}} \to \mathbb{R}^m$ be a bounded continuous function. Then, under the assumptions in* (1.22) *and* (1.23)*,*

$$\vec{\theta}^*_\infty = \arg\max_{\vec{\theta}\in\mathbb{R}^m} [\vec{\theta}\cdot\vec{T}^*_\infty - \psi_\infty(\vec{\theta})], \tag{1.24}$$

where

$$\psi_\infty(\vec{\theta}) := \lim_{n\to\infty} \psi_n(\vec{\theta}) = \sup_{\tilde{h}\in\tilde{\mathcal{W}}} \left[\vec{\theta}\cdot\vec{T}(\tilde{h}) - I(\tilde{h})\right]. \tag{1.25}$$

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

$$\vec{\theta}^*_n = \arg\max_{\vec{\theta}\in\mathbb{R}^m} [n^2\vec{\theta}\cdot\vec{T}^*_n - n^2\psi_n(\vec{\theta})] = \arg\max_{\vec{\theta}\in\mathbb{R}^m} [\vec{\theta}\cdot\vec{T}^*_n - \psi_n(\vec{\theta})]. \tag{1.26}$$

Let $f_n(\vec{\theta},\vec{T}^*_n) := \vec{\theta}\cdot\vec{T}^*_n - \psi_n(\vec{\theta})$ and $f_\infty(\vec{\theta},\vec{T}^*_\infty) = \vec{\theta}\cdot\vec{T}^*_\infty - \psi_\infty(\vec{\theta})$. By [13, Theorem 3.2 and Lemma A.1],

$$\begin{aligned} f_n(\vec{\theta}^*_n,\vec{T}^*_n) &= \sup_{\vec{\theta}\in\mathbb{R}^m} f_n(\vec{\theta},\vec{T}^*_n) = [\vec{\theta}^*_n\cdot\vec{T}^*_n - \psi_n(\vec{\theta}^*_n)] \to [\vec{\theta}^*_\infty\cdot\vec{T}^*_\infty - \psi_\infty(\vec{\theta}^*_\infty)] \\ &= \vec{\theta}^*_\infty\cdot\vec{T}^*_\infty - \sup_{\tilde{h}\in\tilde{\mathcal{W}}} \left[\vec{\theta}^*_\infty\cdot\vec{T}(\tilde{h}) - I(\tilde{h})\right] = f_\infty(\vec{\theta}^*_\infty,\vec{T}^*_\infty), \qquad n\to\infty. \end{aligned} \tag{1.27}$$

(Note that [13, Theorem 3.2 and Lemma A.1] were stated for homomorphism densities, but the proofs generalise to bounded continuous functions $\vec{T}\colon \tilde{\mathcal{W}} \to \mathbb{R}^m$.) Furthermore, for every $\vec{\theta} \in \mathbb{R}^m$, $f_n(\vec{\theta},\vec{T}^*_n) \leq f_n(\vec{\theta}^*_n,\vec{T}^*_n)$, and hence $f_\infty(\vec{\theta},\vec{T}^*_\infty) = \lim_{n\to\infty} f_n(\vec{\theta},\vec{T}^*_n) \leq \lim_{n\to\infty} f_n(\vec{\theta}^*_n,\vec{T}^*_n) = f_\infty(\vec{\theta}^*_\infty,\vec{T}^*_\infty)$, so that $\vec{\theta}^*_\infty$ is a maximiser of $f_\infty(\cdot,\vec{T}^*_\infty)$. □

1.2.5 Variational characterisation of ensemble equivalence

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

$$t(F,h) := \int_{[0,1]^k} \prod_{\{i,j\}\in E(F)} h(x_i,x_j)\, \mathrm{d}x_1 \ldots \mathrm{d}x_k. \tag{1.28}$$

Then the homomorphism density of F in G equals

$$t(F,G) = t(F,h^G), \tag{1.29}$$

where h^G is the empirical graphon defined in (1.11). Therefore the Hamiltonian in (1.16) can be rewritten as

$$H(\vec{\theta},\vec{T}(G)) = n^2 \sum_{i=1}^{m} \theta_i\, t(F_i,h^G). \tag{1.30}$$

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

$$I(u) := \tfrac{1}{2} u \log u + \tfrac{1}{2}(1-u)\log(1-u). \tag{1.31}$$

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

$$I(h) := \int_{[0,1]^2} \mathrm{d}x\,\mathrm{d}y\; I(h(x,y)) \tag{1.32}$$

(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 [13] is the following variational formula for s_∞, where

$$\tilde{\mathcal{W}}^* := \{\tilde{h} \in \tilde{\mathcal{W}} \colon\; \vec{T}(h) = \vec{T}^*_\infty\} \tag{1.33}$$

is the subspace of all graphons that meet the constraint $\vec{T}^*_\infty$.

Theorem 1.4. [Variational characterisation of ensemble equivalence] *Subject to* (1.20), (1.22) *and* (1.23),

$$\lim_{n\to\infty} n^{-2} S_n(\mathrm{P}_{\mathrm{mic}} \mid \mathrm{P}_{\mathrm{can}}) =: s_\infty \tag{1.34}$$

with

$$s_\infty = \sup_{\tilde{h}\in\tilde{\mathcal{W}}} \left[\vec{\theta}^*_\infty \cdot \vec{T}(\tilde{h}) - I(\tilde{h})\right] - \sup_{\tilde{h}\in\tilde{\mathcal{W}}^*} \left[\vec{\theta}^*_\infty \cdot \vec{T}(\tilde{h}) - I(\tilde{h})\right]. \tag{1.35}$$

Theorem 1.4 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 $\vec{\theta}^*_\infty \cdot \vec{T}(\tilde{h}) - I(\tilde{h})$ in $\tilde{\mathcal{W}}$ also lies in $\tilde{\mathcal{W}}^* \subset \tilde{\mathcal{W}}$, i.e., satisfies the hard constraint.

1.2.6 Maximal eigenvalue of the adjacency matrix

In [8] 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

$$\begin{array}{lll} \lim_{n\to\infty} \Delta_n \neq 0 & \Longrightarrow & \text{BEE}, \\ \text{BEE} & \Longrightarrow & \lim_{n\to\infty} \Delta_n \neq 0 \text{ apart from exceptional constraints}, \end{array} \tag{1.36}$$

with

$$\Delta_n := \mathrm{E}_{\mathrm{can}}[\lambda_n] - \mathrm{E}_{\mathrm{mic}}[\lambda_n]. \tag{1.37}$$

In [8] 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 $\lim_{n\to\infty} \Delta_n = 1 - p$ when $\lim_{n\to\infty} n^{-1} d(n) = p \in [0, 1]$, 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

$$\delta_\infty := \lim_{n\to\infty} n^{-1} \Delta_n. \tag{1.38}$$

1.3 Main results

In what follows, F is any finite simple graph with k edges, and the constraint is on the density of F being equal to T^*_∞. Henceforth 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.

1.3.1 Parameter regime

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

Theorem 1.5. [Computational criterion for ensemble equivalence] *For $\theta \in [0, \infty)$ and $k \in [1, \infty)$, let $u^*(\theta, k)$ be a maximiser of*

$$\sup_{u\in[0,1]} \big[\theta u^k - I(u)\big]. \tag{1.39}$$

(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 1.6. [Phase diagram]
(a) There exists a function $k_c\colon (0, 1) \to [1, \infty)$ 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 $\frac{k_0-1}{k_0} \log(k_0 - 1) = 1$ and $T_0 = (\frac{k_0-1}{k_0})^{k_0}$.
(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 $T^* \uparrow 1$.

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 [16]. Using the theory developed in [16], 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.

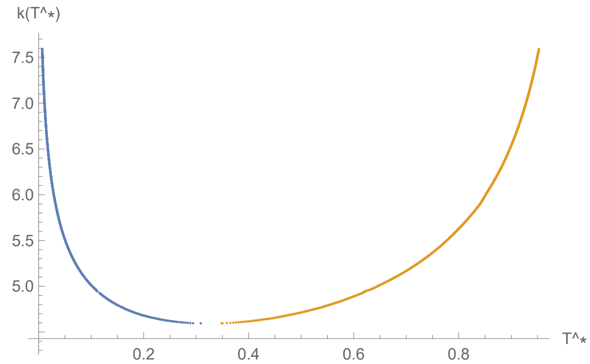

Figure 2: A numerical picture of the phase diagram. The blue and orange lines together form the critical curve $(T^*, k_c(T^*))$. In the figure, T^* is denoted by `T^*` and $k_c(T^*)$ is denoted by `k(T^*)`. The minimum is achieved at $k_0 = 4.591\ldots$ and $T_0 = 0.3237\ldots$.

Definition 1.7. [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

$$\inf_{t(F,f)\geq T^*} I_p(f) \tag{1.40}$$

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

$$I_p(f) = \int_{[0,1]^2} \mathrm{d}x\,\mathrm{d}y \left(f(x,y) \log \frac{f(x,y)}{p} + [1-f(x,y)] \log \frac{1-f(x,y)}{1-p} \right). \tag{1.41}$$

Note that $I(f) = I_{\frac{1}{2}}(f) - \frac{1}{2}\log 2$. Hence, if T^* is in the replica symmetric region, then there is an explicit solution for the second supremum in (1.35). In [16], 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

$$T_1^*(k) \in ((\tfrac{1}{2})^k, T_0), \qquad T_2^*(k) \in (T_0, 1), \tag{1.42}$$

be the two solutions of the equation $k_c(T^*(k)) = k$, so that

$$(T_1^*(k), T_2^*(k)) = \text{BEE-phase}. \tag{1.43}$$

In Lemma 3.1, we prove that the replica symmetric region is contained in $[(\frac{1}{2})^k, T_1^*(d)] \cup [T_2^*(d), 1]$. 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 1.8. [Specific relative entropy] *For every T^* in the replica symmetric part of the phase of breaking of ensemble equivalence,*

$$s_\infty = \begin{cases} \hat{\theta}(k)\,[T_1^*(k) - T^*] + \left[I(T^{*\,1/k}) - I(T_1^*(k)^{1/k})\right] > 0, & T^* \in (T_1^*(k), T_1^*(d)], \\ \hat{\theta}(k)\,[T_2^*(k) - T^*] + \left[I(T^{*\,1/k}) - I(T_2^*(k)^{1/k})\right] > 0, & T^* \in [T_2^*(d), T_2^*(k)). \end{cases} \tag{1.44}$$

Consequently,

$$s_\infty = \begin{cases} C(T_1^*(k),k)\,[T^* - T_1^*(k)]^2 + O([T - T_1^*(k)]^3), & T^* \downarrow T_1^*(k), \\ C(T_2^*(k),k)\,[T^* - T_2^*(k)]^2 + O([T - T_2^*(k)]^3), & T^* \uparrow T_2^*(k), \end{cases} \tag{1.45}$$

with

$$C(T^*,k) = \frac{T^{*\,(1-2k)/k}}{2k} \left\{ \frac{1}{k}\left(1 + \frac{T^{*\,1/k}}{1 - T^{*\,1/k}}\right) + \left(\frac{1}{k} - 1\right) \log\left(\frac{T^{*\,1/k}}{1 - T^{*\,1/k}}\right) \right\}. \tag{1.46}$$

Theorem 1.9. [Spectral signature] *For every T^* in the replica symmetric part of the phase of breaking of ensemble equivalence,*

$$\begin{aligned} &\delta_\infty = \frac{T_1^{*\,1/k}\,[T_2^*(k) - T^*] + T_2^{*\,1/k}\,[T^* - T_1^*(k)]}{T_2^*(k) - T_1^*(k)} - T^{*\,1/k} < 0, \\ &T^* \in (T_1^*(k), T_1^*(d)] \cup [T_2^*(d), T_2^*(k)). \end{aligned} \tag{1.47}$$

Consequently,

$$\delta_\infty = \begin{cases} \hat{C}(T_1^*(k),k)\,[T^* - T_1^*(k)] + O([T^* - T_1^*(k)]^2), & T^* \downarrow T_1^*(k), \\ \hat{C}(T_2^*(k),k)\,[T^* - T_2^*(k)] + O([T^* - T_2^*(k)]^2), & T^* \uparrow T_2^*(k), \end{cases} \tag{1.48}$$

with

$$\hat{C}(T^*,k) = \frac{T_2^{*\,1/k}(k) - T_1^{*\,1/k}(k)}{T_2^*(k) - T_1^*(k)} - \frac{1}{k} T^{*\,(1-k)/k}. \tag{1.49}$$

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 5.3 below we show that the random graph drawn from the canonical ensemble converges to the maximiser(s) of the first supremum of (1.35), while the random graph drawn from the microcanonical ensemble converges to the maximiser(s) of the second supremum of (1.35).

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-Rènyi 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 $\tilde{\mathcal{W}}^*$. Consequently, the random graph drawn from the canonical ensemble converges to the random graphon

$$\frac{T_2^*(k) - T^*}{T_2^*(k) - T_1^*(k)}\,\delta_{T_1^*(k)^{1/k}} + \frac{T^* - T_1^*(k)}{T_2^*(k) - T_1^*(k)}\,\delta_{T_2^*(k)^{1/k}}, \tag{1.50}$$

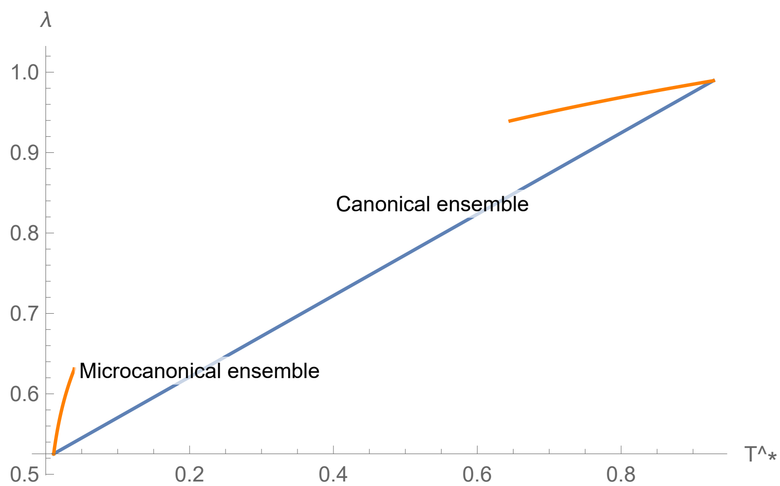

Figure 3: A numerical picture of the average largest eigenvalue $\lambda = \lim_{n\to\infty} \frac{1}{n}\mathbb{E}[\lambda_n]$ of the adjacency matrix under the microcanonical ensemble (top curve) and the canonical ensemble (bottom curve), as a function of T^* for a subgraph F with $k = 7$ edges and maximum degree $d = 5$. The top curve is shown only for T^* in the replica symmetric region. In the region replica symmetry breaking we have no explicit expression for λ under the microcanonical ensemble.

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. In case of non-uniqueness, also under the microcanonical ensemble the random graph is asymptotically random.

1.5 Discussion and outline

1. Theorem 1.5 reduces the variational formula on $\tilde{\mathcal{W}}$ 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 variational characterisation in Theorem 1.4. The main difficulty lies in computing the tuning parameter θ^* as a function of the density T^*, which is resolved through Lemma 1.3. The proof follows from an analysis of the two variational expressions, for which we rely in part on the results in [17]. From Theorem 1.5, for each k we can identify the BEE-phase as follows. The expression in (1.39) 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 $\lim_{\theta\to\infty} u_2^*(\theta) = 1$, the interval $(u_1^*(\theta)^k, u_2^*(\theta)^k)$ is the entire BEE-phase.

2. Theorem 1.6 identifies the BEE-phase and captures the main properties of the critical curve bordering this phase. The proof relies on Lemma 3.1 below, which allows us to use results from [16] 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 (1.40)–(1.41)). This connection is purely analytic: it establishes equivalence of variational formulas and implies that the graph in Figure 2 is a cross-section of the curves in [16, 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 (1.39) (with I_p instead of I) to analyse replica symmetry, rather than the convex minorant of $x \mapsto I_p(x^{1/k})$.

3. Theorem 1.8 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 (1.35) 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.

4. Theorem 1.9 shows that the working hypothesis put forward in [8] is met in the replica symmetric part of the BEE-phase. A random graph drawn from the canonical ensemble converges to a constant graphon who height is a random mixture of the two maximisers u_1, u_2 of (1.39). 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 Figure 3. Note that the average largest eigenvalue is larger in the microcanonical ensemble than in the canonical ensemble, contrary to what was found in [8], where the constraint was on the degree sequence. It turns out that the relative entropy undergoes a first-order phase transition as T^* approaches the critical curve.

5. The numerical picture of the phase diagram in Figure 2 was made using Mathematica. The computations involve finding an approximate value of $\hat{\theta}(k)$ for each k (up to an accuracy of 5 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.

6. In [20], 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 [13], and the link with the convex minorant mentioned in item 2 fits in with the picture provided in [20].

Outline. The remainder of the paper is organised as follows. Theorems 1.5–1.9 are proved in Sections 2–5, respectively.

2 Proof of Theorem 1.5

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

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

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 [13, 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 (1.39), and each maximiser of (2.1) is a constant function, where the constant is a maximiser of (1.39). Furthermore, by Lemma 1.3, θ^* is a maximiser of the supremum

$$\sup_{\theta \geq 0} \left[\theta T^* - \theta T(u^*(\theta)) + I(u^*(\theta))\right] = \sup_{\theta \geq 0} \left[\theta T^* - \theta (u^*(\theta))^k + I(u^*(\theta))\right], \tag{2.2}$$

where $u^*(\theta)$ is a maximiser of (1.39). By [17, Proposition 3.2], $l_\theta(u) := \theta u^k - I(u)$ has at most 2 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. 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 Figure 4, plots of l_θ are shown for several values of θ. Write $u := u^*(\theta^*)$ and $u' := \frac{\partial u}{\partial \theta}(\theta^*)$. Then

$$l'_\theta(u) = \theta k u^{k-1} - \tfrac{1}{2}\log u + \tfrac{1}{2}\log(1-u) = 0 \tag{2.3}$$

and

$$\begin{aligned} m'(\theta) &= T^* - u^k - \theta k u^{k-1} u' + \tfrac{1}{2} u' \log(u) - \tfrac{1}{2} u' \log(1-u) \\ &= T^* - u^k - u'(\tfrac{1}{2}\log u - \tfrac{1}{2}\log(1-u) - \theta k u^{k-1}) \\ &= T^* - u^k. \end{aligned} \tag{2.4}$$

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 1.6

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 $T_1(k) = u_1^*(\hat{\theta}(k), k)^k$, $T_2(k) = u_2^*(\hat{\theta}(k), k)^k$.

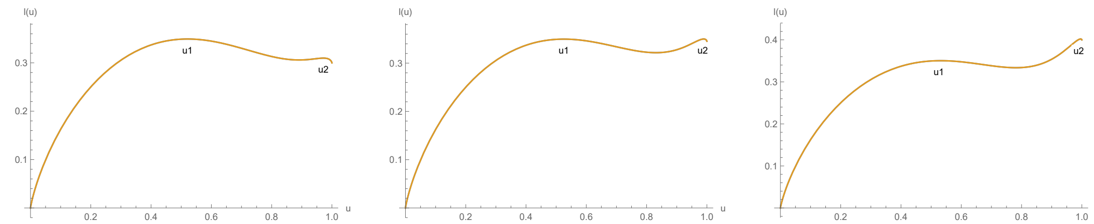

Figure 4: Three plots of $l_\theta(u)$ for $k = 7$ and $\theta = 0.3$, $\theta = \hat{\theta}(7)$ and $\theta = 0.4$, respectively. For $\theta = 0.3$, $u_1^*(\theta)$ is the global maximiser, for $\theta = \hat{\theta}(7)$, $u_1^*(\theta)$ and $u_2^*(\theta)$ are both global maximisers, and for $\theta = 0.4$, $u_2^*(\theta)$ is the global maximiser. In the figures, the function $l_\theta(u)$ is denoted by `l(u)`. The local maximisers $u_1^*(\theta)$ and $u_2^*(\theta)$ are denoted by `u1` and `u2` respectively. The BEE-phase is $(u_1^*(\hat{\theta})^k, u_2^*(\hat{\theta})^k)$.

Existence of k_c. Lemmas 3.1–3.2 below 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 $T \in [(\frac{1}{2})^k, 1)$. 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 (1.41)), 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 [16, 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 $T^* = T \in [(\frac{1}{2})^k, 1)$ 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 $u_2 = q_2$. This 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 [16, 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 $k \mapsto (u_2^*(\hat{\theta}))^k$ is increasing. 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 1.6. 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. Consider the function $\dot{J}_k(x) := \frac{\partial}{\partial k} J_k(x)$. This is a concave function in x 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 $k' \downarrow k$,

$$\begin{aligned}
&J_{k'}\big(\alpha T_1(k) + (1-\alpha)T_2(k)\big) \\
&= J_k\big(\alpha T_1(k) + (1-\alpha)T_2(k)\big) + (k'-k)\dot{J}_k\big(\alpha T_1(k) + (1-\alpha)T_2(k)\big) + o(k'-k) \\
&\geq \alpha J_k(T_1(k)) + (1-\alpha)J_k(T_2(k)) + (k'-k)(\alpha \dot{J}_k(T_1(k)) + (1-\alpha)\dot{J}_k(T_2(k))) + o(k'-k) \\
&= \alpha J_{k'}(T_1(k)) + (1-\alpha)J_{k'}(T_2(k)) + o(k'-k).
\end{aligned} \tag{3.1}$$

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 [16, Lemma A.3], $T_1(k') < T_1(k) < T_2(k) < T_2(k')$. □

Minimum of k_c. By [17, 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 exist 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 [17, Proposition 3.2] it is shown that $\hat{\theta}(k_0) = \frac{k_0^{k_0-1}}{2(k_0-1)^{k_0}}$, and so

$$l'_{\hat{\theta}(k_0),k_0}(\tfrac{k_0-1}{k_0}) = \tfrac{(k_0)^{k_0-1}}{2(k_0-1)^{k_0}}\, k_0 \left(\tfrac{k_0-1}{k_0}\right)^{k_0-1} - \tfrac{1}{2}\log(k_0-1) = \tfrac{k_0}{2(k_0-1)} - \tfrac{1}{2}\log(k_0-1) = 0. \tag{3.2}$$

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

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

$$f(k,x,y) = \big(J'_k(x) - J'_k(y), J'_k(x)x - J'_k(y)y + J(y) - J(x)\big). \tag{3.3}$$

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

$$\begin{pmatrix} \frac{\partial f_1}{\partial x} & \frac{\partial f_1}{\partial y} \\ \frac{\partial f_2}{\partial x} & \frac{\partial f_2}{\partial y} \end{pmatrix} (T_1(k), T_2(k)) = \begin{pmatrix} J''_k(T_1(k)) & -J''_k(T_2(k)) \\ T_1(k)J''_k(T_1(k)) & -T_2(k))J''_k(T_2(k)) \end{pmatrix}, \tag{3.4}$$

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

$$\begin{aligned}
T'_1(k) &= \frac{1}{(T_1 - T_2)J''_k(T_1)J''_k(T_2)} \left[(T_2 - T_1)J''_k(T_2)\frac{\partial J'_k}{\partial k}(T_1) + J''_k(T_2)\left(\frac{\partial J_k}{\partial k}(T_1) - \frac{\partial J_k}{\partial k}(T_2)\right)\right] \\
&= \frac{1}{J''_k(T_1)}\left(\frac{\partial J'_k}{\partial k}(T_1) + \frac{\frac{\partial J_k}{\partial k}(T_1) - \frac{\partial J_k}{\partial k}(T_2}{T_2 - T_1}\right) \\
&= \frac{1}{J''_k(T_1)} O(T_2 - T_1).
\end{aligned} \tag{3.5}$$

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.

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 $k \to \infty$. In the following, we suppress the dependence of $\hat{\theta}$ on k. By Taylor expansion,

$$\begin{aligned} l_\theta(u_1^*) &\le l_\theta\left(\tfrac{1}{2}\right) + \left(u_1^* - \tfrac{1}{2}\right) l'_\theta\left(\tfrac{1}{2}\right) \\ &\le \theta\left(\tfrac{1}{2}\right)^k + \tfrac{1}{2}\log 2 + \theta k\left(\tfrac{1}{2}\right)^k = \theta\left(\tfrac{1}{2}\right)^k(1+k) + \tfrac{1}{2}\log 2, \end{aligned} \tag{3.6}$$

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

$$\hat{\theta} < \frac{\log 2}{2[1-(\frac{1}{2})^k(1+k)]}. \tag{3.7}$$

Also, $u_2^*(\theta,k) \in (\frac{k-1}{k}, 1)$ by [17, Proposition 3.2]. Hence

$$\begin{aligned} l_{\theta,k}(u_2^*(\theta,k)) &\le \theta - \tfrac{k-1}{2k}\log(\tfrac{k-1}{k}) - \tfrac{1}{2}(1-\tfrac{k-1}{k})\log(1-\tfrac{k-1}{k}) \\ &= \theta - \tfrac{1}{2}\log(1-\tfrac{1}{k}) - \tfrac{1}{2k}\log(\tfrac{1}{k-1}), \end{aligned} \tag{3.8}$$

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

$$\hat{\theta} > \frac{\log 2 + \log(1-\frac{1}{k}) + \frac{1}{k}\log(\frac{1}{k-1})}{2[1-(\frac{1}{2})^k]}. \tag{3.9}$$

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

- **Scaling for $T^* \downarrow 0$.** Let $y \in (\frac{1}{2}, 1)$. Then

$$\begin{aligned} l'_{\hat{\theta},k}(\tfrac{1}{2}+y^k) &= \hat{\theta}k(\tfrac{1}{2}+y^k)^{k-1} - \tfrac{1}{2}\log\left(\tfrac{1+2y^k}{1-2y^k}\right) \\ &= \hat{\theta}k(\tfrac{1}{2}+y^k)^{k-1} - \tfrac{1}{2}\log\left(1+\tfrac{4y^k}{1-2y^k}\right) \\ &\le \frac{\log 2}{2[1-(\frac{1}{2})^k(1+k)]}k\left(\tfrac{1}{2}\right)^{k-1} - 2y^k + o(k\left(\tfrac{1}{2}\right)^k) + o(y^k) < 0 \end{aligned} \tag{3.10}$$

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

- **Scaling for $T^* \uparrow 1$.** Let $x \in (0,1)$. Then

$$l'_{\hat{\theta},k}(1-x^k) = k(\hat{\theta}(1-x^k)^{k-1} + \tfrac{1}{2}\log x) - \tfrac{1}{2}\log(1-x^k). \tag{3.11}$$

As $k \to \infty$, $(1-x^k)^{k-1} \to 1$ and $\log(1-x^k) \to 0$. Hence, if $-\frac{1}{2}\log x \ge \hat{\theta}$, 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 $\hat{\theta} \to \frac{1}{2}\log 2$. Thus, choosing $x = \frac{1}{2}$, we get $(1-(\frac{1}{2})^{k_c})^{k_c} \sim T^*$, and so $k_c(\frac{1}{2})^{k_c} \sim 1 - T^*$.

4 Proof of Theorem 1.8

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 $T_2^*(k)$. 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 [16, 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 [16, 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

$$\inf\{I(\tilde{h})\colon\ \tilde{h} \in \tilde{\mathcal{W}},\, t(H, \tilde{h}) \geq T\} = \inf\{I(\tilde{h})\colon\ \tilde{h} \in \tilde{\mathcal{W}},\, t(H, \tilde{h}) = T\} = \inf_{\tilde{h} \in \tilde{\mathcal{W}}^*} I(\tilde{h}). \tag{4.1}$$

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

$$\begin{aligned} s_\infty &= \sup_{\tilde{h} \in \tilde{\mathcal{W}}} \left[\theta^* T(\tilde{h}) - I(\tilde{h})\right] - \sup_{\tilde{h} \in \tilde{\mathcal{W}}^*} \left[\theta^* T(\tilde{h}) - I(\tilde{h})\right] \\ &= [\hat{\theta} T^* - I(T^{*\,1/k})] - [\hat{\theta} T - I(T^{1/k})] \\ &= \hat{\theta}(T^* - T) + [I(T^{1/k}) - I(T^{*\,1/k})] \\ &= [J_k'(T^*) - \hat{\theta}](T - T^*) + J_k''(T^*)(T - T^*)^2 + O((T - T^*)^3) \\ &= \frac{T^{*\,1/k-2}}{2k} \left\{ \frac{1}{k} \left(1 + \frac{T^{*\,1/k}}{1 - T^{*\,1/k}}\right) + \left(\frac{1}{k} - 1\right) \log\left(\frac{T^{*\,1/k}}{1 - T^{*\,1/k}}\right) \right\} (T - T^*)^2 + O((T - T^*)^3) \end{aligned} \tag{4.2}$$

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

5 Proof of Theorem 1.9

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, Theorem 3.2] 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 with parameter θ_n^*. Let U be the set of maximisers of* (1.39) *with $\theta = \theta_\infty^*$. Then (recall* (1.11)*)*

$$\min_{u \in U} \delta_\square(\widetilde{h}^{G_n}, \widetilde{u}) \to 0. \tag{5.1}$$

Proof. Let $\eta > 0$ and define

$$\widetilde{A}(\theta_\infty^*) := \{\tilde{h} \in \tilde{\mathcal{W}} \mid \delta_\square(\tilde{h}, \widetilde{U}) \geq \eta\}. \tag{5.2}$$

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

$$2\gamma := \sup_{\tilde{h} \in \tilde{\mathcal{W}}} \left[\theta_\infty^* T(\tilde{h}) - I(\tilde{h})\right] \sup_{\tilde{h} \in \widetilde{A}(\theta_\infty^*)} \left[\theta_\infty^* T(\tilde{h}) - I(\tilde{h})\right] > 0. \tag{5.3}$$

Let $\varepsilon = \gamma$. 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 θ_∞^* for all n large enough. For each $a \in R$, let $\widetilde{F}^a(\theta_n^*) := (\theta_n^* T)^{-1}([a, a+\varepsilon])$. Now define $\widetilde{A}^a(\theta_n^*)$ analogously to $\widetilde{A}(\theta_\infty^*)$ and $\widetilde{A}_n^a(\theta_n^*) = \widetilde{A}(\theta_n^*) \cap \widetilde{\mathcal{G}}_n$. Choose $\delta = \frac{1}{2}\varepsilon$. Since $\theta_n^* \to \theta_\infty^*$ and T is continuous, we have that $(\theta_n^* T)^{-1}([a, a+\varepsilon]) \subseteq (\theta_\infty^* T)^{-1}([a-\delta, a+\varepsilon+\delta]) =: \widetilde{G}^a$ for all n large enough. Now define $\widetilde{B}^a$ and $\widetilde{B}_n^a$ analogously to $\widetilde{A}^a$ and $\widetilde{A}_n^a$.

We have

$$\begin{aligned}
\mathrm{P}_{\mathrm{can}}(G_n \in \widetilde{A}) &\le \mathrm{e}^{-n^2 \psi_n(\theta_n^*)} \sum_{a\in R} e^{n^2(a+\varepsilon)} |\widetilde{A}_n^a(\theta_n^*)| \\
&\le \mathrm{e}^{-n^2 \psi_n(\theta_n^*)} \sum_{a\in R} e^{n^2(a+\varepsilon+\delta)} |\widetilde{B}_n^a| \\
&\le \mathrm{e}^{-n^2 \psi_n(\theta_n^*)} |R| \sup_{a\in R} e^{n^2(a+\varepsilon+\delta)} |\widetilde{B}_n^a|.
\end{aligned} \tag{5.4}$$

As in the proof of [4, Theorem 3.1], we obtain

$$\limsup_{n\to\infty} \frac{\log \mathrm{P}_{\mathrm{can}}(G_n \in \widetilde{A})}{n^2} \le \sup_{a\in R} \left[a + \varepsilon + \delta - \inf_{\tilde{h}\in \widetilde{B}^a} I(\tilde{h}) \right] - \sup_{\tilde{h}\in\tilde{\mathcal{W}}} \left[\theta_\infty^* T(h) - I(h)\right], \tag{5.5}$$

where for the second supremum, we use that

$$\lim_{n\to\infty} \psi_n(\theta_n^*) = \psi_\infty(\theta_\infty^*) = \sup_{\tilde{h}\in\tilde{\mathcal{W}}} \left[\theta_\infty^* T(h) - I(h)\right]. \tag{5.6}$$

This was shown in [13, Lemma A.1]. The remainder of the proof now follows exactly as in [4], with ε replaced by $\varepsilon + \delta = \frac{3}{2}\varepsilon$. □

Corollary 5.2. *Assume that T^* is in the BEE-phase. Let G_n be a random graph drawn from the canonical ensemble. Then h^{G_n} converges in probability to*

$$\frac{u_2^k - T^*}{u_2^k - u_1^k} \delta_{u_1} + \frac{T^* - u_1^k}{u_2^k - u_1^k} \delta_{u_2}, \tag{5.7}$$

with $u_1 < u_2$ the two maximisers of (1.39) *for $\theta = \hat{\theta}$.*

Proof. From the lemma above, it is clear that G_n converges in probability to the random graphon $p\delta_{u_1} + (1-p)\delta_{u_2}$ for some $p \in (0,1)$. It remains to determine p. Since the homomorphism density is continuous, $t(H, G_n)$ converges in probability to $p\delta_{u_1^k} + (1-p)\delta_{u_2^k}$. The homomorphism density is bounded, so the convergence also holds in mean. Hence, by the definition of the canonical ensemble,

$$T^* = \lim_{n\to\infty} E[t(H, G_n)] = p u_1^k + (1-p) u_2^k. \tag{5.8}$$

Solving this equation for p we conclude the proof. □

We can also show convergence of the microcanonical ensemble.

Lemma 5.3. *Let G_n be a random graph drawn from the microcanonical ensemble. Then $\widetilde{h}^{G_n}$ converges in probability to $\widetilde{F}^*$, with $\widetilde{F}^*$ 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

$$\widetilde{F}^{\varepsilon} := \{\widetilde{h} \in \tilde{\mathcal{W}}^* \mid \delta_{\square}(\tilde{h}, \widetilde{F}^*) > \varepsilon\} \tag{5.9}$$

and

$$\widetilde{F}_n^{\varepsilon} := \{\widetilde{h} \in \widetilde{F}_{\varepsilon} \mid \delta_{\square}(\tilde{h}, \widetilde{F}^*) > \varepsilon,\, \tilde{h} = \widetilde{G} \text{ for some } G \in \mathcal{G}_n\}. \tag{5.10}$$

Then, by [13, (3.22) and Corollary 2.9],

$$\begin{aligned} \lim_{n\to\infty} \frac{1}{n^2} \log \mathrm{P}_{\mathrm{mic}}(\widetilde{F}^{\varepsilon}) &= \lim_{n\to\infty} \frac{1}{n^2} \log(|\widetilde{F}_n^{\varepsilon}| \mathrm{P}_{\mathrm{mic}}(G_n = G_n^*)) \\ &= \inf_{\tilde{h}\in\tilde{\mathcal{W}}^*} I(\tilde{h}) + \lim_{n\to\infty} \frac{1}{n^2} \log |\widetilde{F}_n^{\varepsilon}| \\ &= \inf_{\tilde{h}\in\tilde{\mathcal{W}}^*} I(\tilde{h}) - \inf_{\tilde{h}\in\widetilde{F}^{\varepsilon}} I(\tilde{h}), \end{aligned} \tag{5.11}$$

where G_n^* is any graph in $\mathcal{G}_n$ such that $\widetilde{G}_n^* \in \tilde{\mathcal{W}}^*$. Since $\tilde{\mathcal{W}}^*$ is a compact set and $\widetilde{F}^{\varepsilon}$ does not contain any minimisers of $\inf_{\tilde{h}\in\tilde{\mathcal{W}}^*} I(\tilde{h})$, we conclude that the expression above is negative, which implies that

$$\lim_{n\to\infty} \mathrm{P}_{\mathrm{mic}}(\widetilde{F}^{\varepsilon}) = 0. \tag{5.12}$$

□

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 $\|h^{G_n}\|_{\mathrm{op}}$ of the empirical graphon of G_n. The operator norm is continuous, so $\|h^{G_n}\|_{\mathrm{op}}$ converges to $\|p\delta_{u_1} + (1-p)\delta_{u_2}\|_{\mathrm{op}} = p\delta_{u_1} + (1-p)\delta_{u_2}$ in probability, with p as in Corollary 5.2. Since the operator norm is bounded, we also have

$$\lim_{n\to\infty} n^{-1} E_{\mathrm{can}}[\lambda_n] = pu_1 + (1-p)u_2 = \frac{T^*(u_2 - u_1) + u_1 u_2 (u_2^{k-1} - u_1^{k-1})}{u_2^k - u_1^k} =: f(T^*). \tag{5.13}$$

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

$$\lim_{n\to\infty} n^{-1} E_{\mathrm{mic}}[\lambda_n] = (T^*)^{1/k} > f(T^*), \tag{5.14}$$

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.

References

[1] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós and K. Vesztergombi, Convergent graph sequences I: Subgraph frequencies, metric properties, and testing, Adv. Math. 219 (2008) 1801–1851.

[2] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós and K. Vesztergombi, Convergent sequences of dense graphs II: Multiway cuts and statistical physics, Ann. Math. 176 (2012) 151–219.

[3] S. Chatterjee, *Large Deviations for Random Graphs*, École d'Été de Probabilités de Saint-Flour XLV, 2015.

[4] S. Chatterjee, P. Diaconis, Estimating and understanding exponential random graph models, Ann. Statist. 5 (2013) 2428–2461.

[5] S. Chatterjee and S.R.S. Varadhan, The large deviation principle for the Erdős-Rényi random graph, European J. Comb. 32 (2011) 1000–1017.

[6] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications* (2nd edition), Springer, 1998.

[7] P. Diao, D. Guillot, A. Khare and B. Rajaratnam, Differential calculus on graphon space, J. Combin. Theory Ser. A 133 (2015) 183–227.

[8] P. Dionigi, D. Garlaschelli, F. den Hollander and M. Mandjes, A spectral signature of breaking of ensemble equivalence for constrained random graphs, to appear in Electr. J. Probab., [arXiv:2009.05155]

[9] G. Garlaschelli, F. den Hollander and A. Roccaverde, Ensemble equivalence in random graphs with modular structure, J. Phys. A: Math. Theor. 50 (2017) 015001.

[10] G. Garlaschelli, F. den Hollander and A. Roccaverde, Covariance structure behind breaking of ensemble equivalence, J. Stat. Phys. 173 (2018) 644–662.

[11] J.W. Gibbs, *Elementary Principles of Statistical Mechanics*, Yale University Press, New Haven, Connecticut, 1902.

[12] F. den Hollander, *Large Deviations*, Fields Institute Monographs 14, American Mathematical Society, Providence, RI, 2000.

[13] F. den Hollander, M. Mandjes, A. Roccaverde and N.J. Starreveld, Ensemble equivalence for dense graphs, Electronic J. Prob. 23 (2018), Paper no. 12, 1–26.

[14] F. den Hollander, M. Mandjes, A. Roccaverde and N.J. Starreveld, Breaking of ensemble equivalence for perturbed Erdős-Rényi random graphs, [arXiv:1807.07750].

[15] E.T. Jaynes, Information theory and statistical mechanics, Phys. Rev. 106 (1957) 620–630.

[16] E. Lubetzky and Y. Zhao, On replica symmetry of large deviations in random graphs, Random Struct. Algor. 47 (2015) 109–146.

[17] C. Radin, M. Yin, Phase transitions in exponential random graphs, Ann. Appl. Probab. 6 (2013) 2458–2471.

[18] T. Squartini and D. Garlaschelli, Reconnecting statistical physics and combinatorics beyond ensemble equivalence, [arXiv:1710.11422].

[19] T. Squartini, J. de Mol, F. den Hollander and D. Garlaschelli, Breaking of ensemble equivalence in networks, Phys. Rev. Lett. 115 (2015) 268701.

[20] H. Touchette, Equivalence and nonequivalence of ensembles: Thermodynamic, macrostate, and measure levels, J. Stat. Phys. 159 (2015) 987–1016.