



Universiteit
Leiden
The Netherlands

Breaking of ensemble equivalence for complex networks

Roccaverde, A.

Citation

Roccaverde, A. (2018, December 5). *Breaking of ensemble equivalence for complex networks*. Retrieved from <https://hdl.handle.net/1887/67095>

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/67095>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The following handle holds various files of this Leiden University dissertation:

<http://hdl.handle.net/1887/67095>

Author: Roccaverde, A.

Title: Breaking of ensemble equivalence for complex networks

Issue Date: 2018-12-05



CHAPTER 5

Ensemble Equivalence for dense graphs

This chapter is based on:

F. den Hollander, M. Mandjes, A. Roccaverde, and N. J. Starreveld. Ensemble equivalence for dense graphs. *Electron. J. Probab.*, 23:Paper No. 12, 26, 2018

Abstract

In this paper we consider a random graph on which topological restrictions are imposed, such as constraints on the total number of edges, wedges, and triangles. We work in the dense regime, in which the number of edges per vertex scales proportionally to the number of vertices n . Our goal is to compare the micro-canonical ensemble (in which the constraints are satisfied for every realization of the graph) with the canonical ensemble (in which the constraints are satisfied on average), both subject to maximal entropy. We compute the relative entropy of the two ensembles in the limit as n grows large, where two ensembles are said to be *equivalent* in the dense regime if this relative entropy divided by n^2 tends to zero. Our main result, whose proof relies on large deviation theory for graphons, is that breaking of ensemble equivalence occurs when the constraints are *frustrated*. Examples are provided for three different choices of constraints.

§5.1 Introduction

Section 5.1.1 gives background and motivation, Section 5.1.2 describes relevant literature, while Section 5.1.3 outlines the remainder of the paper.

§5.1.1 Background and motivation

For large networks a detailed description of the architecture of the network is infeasible and must be replaced by a *probabilistic* description, where the network is assumed to be a random sample drawn from a set of allowed graphs that are consistent with a set of empirically observed features of the network, referred to as *constraints*. Statistical physics deals with the definition of the appropriate probability distribution over the set of graphs and with the calculation of its relevant properties (Gibbs [53]). The two main choices¹ of probability distribution 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 networks that are large but finite, the two ensembles are obviously different and, in fact, represent different empirical situations: they serve as *null-models* for the network after incorporating what is known about the network *a priori* via the constraints. 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 of the soft constraints, i.e., the soft constraints are expected to become asymptotically hard. This assumption of *ensemble equivalence*, which is one of the corner stones of statistical physics, does however *not* hold in general (we refer to Touchette [97] for more background).

In Squartini *et al.* [92] the question of the possible breaking of ensemble equivalence was investigated for two types of constraint:

- (I) The total number of edges.
- (II) The degree sequence.

In the *sparse regime*, where the empirical degree distribution converges to a limit as the number of vertices n tends to infinity such that the maximal degree is $o(\sqrt{n})$, it was shown that the relative entropy of the micro-canonical ensemble w.r.t. the canonical ensemble divided by n (which can be interpreted as the relative entropy per vertex) tends to s_∞ , with $s_\infty = 0$ in case the constraint concerns the total number of edges, and $s_\infty > 0$ in case the constraint concerns the degree sequence. For the latter

¹The microcanonical ensemble and the canonical ensemble work with a fixed number of vertices. There is a third ensemble, the *grandcanonical ensemble*, where also the size of the graph is considered as a soft constraint.

case, an explicit formula was derived for s_∞ , which allows for a *quantitative analysis* of the breaking of ensemble equivalence.

In the present paper we analyse what happens in the *dense regime*, where the number of edges per vertex is of order n . We consider case (I), yet allow for constraints not only on the total number of edges but also on the total number of wedges, triangles, etc. We show that the relative entropy divided by n^2 (which, up to a constant, can be interpreted as the relative entropy per edge) tends to s_∞ , with $s_\infty > 0$ when the constraints are *frustrated*. Our analysis is based on a large deviation principle for graphons.

§5.1.2 Relevant literature

In the past few years, several papers have studied the microcanonical ensemble and the canonical ensemble. Most papers focus on dense graphs, but there are some interesting advances for sparse graphs as well. Closely related to the canonical ensemble are the *exponential random graph model* (Bhamidi *et al.* [12], Chatterjee and Diaconis [29]) and the *constrained exponential random model* (Aristoff and Zhu [3], Kenyon and Yin [67], Yin [102], Zhu [105]).

In Aristoff and Zhu [3], Kenyon *et al.* [63], Radin and Sadun [86], the authors study the microcanonical ensemble, focusing on the constrained entropy density. In [3] directed graphs are considered with a *hard* constraint on the number of directed edges and j -stars, while in [63, 86] the focus is on undirected graphs with a *hard* constraint on the edge density, j -star density and triangle density, respectively. Following the work in Bhamidi *et al.* [12] and in Chatterjee and Diaconis [29], a deeper understanding has developed of how these models behave as the size of the graph tends to infinity. Most results concern the asymptotic behaviour of the partition function (Chatterjee and Diaconis [29], Kenyon, Radin, Ren and Sadun [63]) and the identification of regions where phase transitions occur (Aristoff and Zhu [4], Lubetsky and Zhao [70], Yin [101]). For more details we refer the reader to the recent monograph by Chatterjee [27], and references therein. Significant contributions for sparse graphs were made in Chatterjee and Dembo [28] and in subsequent work of Yin and Zhu [103].

For an overview on random graphs and their role as models of complex networks, we refer the reader to the recent monograph by van der Hofstad [99]. The most important distinction between our paper and the existing literature on exponential random graphs is that in the canonical ensemble we impose a *soft* constraint.

§5.1.3 Outline

The remainder of this paper is organised as follows. Section 5.2 defines the two ensembles, gives the definition of equivalence of ensembles in the dense regime, recalls some basic facts about graphons, and states the large deviation principle for the Erdős-Rényi random graph. Section 5.3 states a key theorem in which we give a *variational representation* of s_∞ when the constraint is on *subgraph counts*, properly normalised. Section 5.4 presents our main theorem for ensemble equivalence, which provides three

examples for which *breaking of ensemble equivalence* occurs when the constraints are *frustrated*. In particular, the constraints considered are on the number of edges, triangles and/or stars. Frustration corresponds to the situation where the canonical ensemble scales like an Erdős-Rényi random graph model with an appropriate edge density but the microcanonical ensemble does not. The proof of the main theorem is given in Sections 5.5–5.6, and relies on various papers in the literature dealing with exponential random graph models. Appendix A discusses convergence of Lagrange multipliers associated with the canonical ensemble.

§5.2 Key notions

In Section 5.2.1 we introduce the model and give our definition of equivalence of ensembles in the dense regime (Definition 5.2.1 below). In Section 5.2.2 we recall some basic facts about graphons (Propositions 5.2.4–5.2.6 below). In Section 5.2.3 we recall the large deviation principle for the Erdős-Rényi random graph (Proposition 5.2.7 and Theorem 5.2.8 below), which is the key tool in our paper.

§5.2.1 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 with elements

$$h^G(i, j) := \begin{cases} 1 & \text{if there is an edge between vertex } i \text{ and vertex } j, \\ 0 & \text{otherwise.} \end{cases} \quad (5.1)$$

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

$$P_{\text{mic}}(G) := \begin{cases} 1/\Omega_{\vec{C}^*}, & \text{if } \vec{C}(G) = \vec{C}^*, \\ 0, & \text{otherwise,} \end{cases} \quad G \in \mathcal{G}_n, \quad (5.2)$$

where

$$\Omega_{\vec{C}^*} := |\{G \in \mathcal{G}_n : \vec{C}(G) = \vec{C}^*\}| \quad (5.3)$$

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

$$S_n(P) := - \sum_{G \in \mathcal{G}_n} P(G) \log P(G) \quad (5.4)$$

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

$$\langle \vec{C} \rangle := \sum_{G \in \mathcal{G}_n} \vec{C}(G) P(G). \quad (5.5)$$

This gives the formula (see Jaynes [61])

$$P_{\text{can}}(G) := \frac{1}{Z(\vec{\theta}^*)} e^{H(\vec{\theta}^*, \vec{C}(G))}, \quad G \in \mathcal{G}_n, \quad (5.6)$$

with

$$H(\vec{\theta}^*, \vec{C}(G)) := \vec{\theta}^* \cdot \vec{C}(G), \quad Z(\vec{\theta}^*) := \sum_{G \in \mathcal{G}_n} e^{\vec{\theta}^* \cdot \vec{C}(G)}, \quad (5.7)$$

denoting the *Hamiltonian* and the *partition function*, respectively. In (5.6)–(5.7) the parameter $\vec{\theta}^*$ (which is a real-valued vector the size of the constraint playing the role of a Lagrange multiplier) must be set to the unique value that realises $\langle \vec{C} \rangle = \vec{C}^*$. The Lagrange multiplier $\vec{\theta}^*$ exists and is unique. Indeed, the gradients of the constraints in (5.5) are linearly independent vectors. Consequently, the Hessian matrix of the entropy of the canonical ensemble in (5.6) is a positive definite matrix, which implies uniqueness of the Lagrange multiplier.

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

$$S_n(P_{\text{mic}} | P_{\text{can}}) := \sum_{G \in \mathcal{G}_n} P_{\text{mic}}(G) \log \frac{P_{\text{mic}}(G)}{P_{\text{can}}(G)}. \quad (5.8)$$

5.2.1 Definition. In the dense regime, if ²

$$s_\infty := \lim_{n \rightarrow \infty} \frac{1}{n^2} S_n(P_{\text{mic}} | P_{\text{can}}) = 0, \quad (5.9)$$

then P_{mic} and P_{can} are said to be *equivalent*.

Before proceeding, we recall an important observation made in Squartini *et al.* [92]. For any $G_1, G_2 \in \mathcal{G}_n$, $P_{\text{can}}(G_1) = P_{\text{can}}(G_2)$ 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 (5.8) as

$$S_n(P_{\text{mic}} | P_{\text{can}}) = \log \frac{P_{\text{mic}}(G^*)}{P_{\text{can}}(G^*)}, \quad (5.10)$$

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). This fact greatly simplifies computations.

5.2.2 Remark. 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 $S_n(P_{\text{mic}} | P_{\text{can}})$. When we pass to the limit $n \rightarrow \infty$, we need to specify how $\vec{C}(G)$, \vec{C}^* and $\vec{\theta}^*$ are chosen to depend on n . This will be done in Section 5.3.1.

²In Squartini *et al.* [92], 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).

§5.2.2 Graphons

There is a natural way to embed a simple graph on n vertices in a space of functions called *graphons*. Let 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 $h^G \in W$ in a natural way as (see Fig. 5.1)

$$h^G(x, y) := \begin{cases} 1 & \text{if there is an edge between vertex } \lceil nx \rceil \text{ and vertex } \lceil ny \rceil, \\ 0 & \text{otherwise.} \end{cases} \quad (5.11)$$




Figure 5.1: An example of a graph G and its graphon representation h^G .

The space of graphons W is endowed with the *cut distance*

$$d_{\square}(h_1, h_2) := \sup_{S, T \subset [0, 1]} \left| \int_{S \times T} dx dy [h_1(x, y) - h_2(x, y)] \right|, \quad h_1, h_2 \in W. \quad (5.12)$$

On W there is a natural equivalence relation \equiv . Let Σ be the space of measure-preserving bijections $\sigma: [0, 1] \rightarrow [0, 1]$. Then $h_1(x, y) \equiv h_2(x, y)$ if $h_1(x, y) = h_2(\sigma x, \sigma y)$ for some $\sigma \in \Sigma$. This equivalence relation yields the quotient space $(\tilde{W}, \delta_{\square})$, where δ_{\square} is the metric defined by

$$\delta_{\square}(\tilde{h}_1, \tilde{h}_2) := \inf_{\sigma_1, \sigma_2} d_{\square}(h_1^{\sigma_1}, h_2^{\sigma_2}), \quad \tilde{h}_1, \tilde{h}_2 \in \tilde{W}. \quad (5.13)$$

To avoid cumbersome notation, throughout the sequel 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 W , and by \tilde{h}^G its image in the quotient space \tilde{W} . Let F and G denote two simple graphs with vertex sets $V(F)$ and $V(G)$, respectively, and let $\text{hom}(F, G)$ be the number of homomorphisms from F to G . The *homomorphism density* is defined as

$$t(F, G) := \frac{1}{|V(G)|^{|V(F)|}} \text{hom}(F, G). \quad (5.14)$$

Two graphs are said to be *similar* when they have similar homomorphism densities.

5.2.3 Definition. A sequence of labelled simple graphs $(G_n)_{n \in \mathbb{N}}$ is left-convergent when $(t(F, G_n))_{n \in \mathbb{N}}$ converges for any simple graph F .

Consider a simple graph F on k vertices with edge set $E(F)$, and let $h \in W$. Similarly as above, define the density

$$t(F, h) := \int_{[0,1]^k} dx_1 \cdots dx_k \prod_{(i,j) \in E(F)} h(x_i, x_j). \quad (5.15)$$

If h^G is the image of a graph G in the space W , then

$$t(F, h^G) = \int_{[0,1]^k} dx_1 \cdots dx_k \prod_{(i,j) \in E(F)} h^G(x_i, x_j) = \frac{1}{|V(G)||V(F)|} \text{hom}(F, G) = t(F, G). \quad (5.16)$$

Hence a sequence of graphs $(G_n)_{n \in \mathbb{N}}$ is left-convergent to $h \in W$ when

$$\lim_{n \rightarrow \infty} t(F, G_n) = t(F, h). \quad (5.17)$$

We conclude this section with three basic facts that will be needed later on. The first gives the relation between left-convergence of sequences of graphs and convergence in the quotient space $(\tilde{W}, \delta_\square)$, the second is a compactness property, while the third shows that the homomorphism density is Lipschitz continuous with respect to the δ_\square -metric.

5.2.4 Proposition (Borgs et al. [20]). *For a sequence of labelled simple graphs $(G_n)_{n \in \mathbb{N}}$ the following properties are equivalent:*

- (i) $(G_n)_{n \in \mathbb{N}}$ is left-convergent.
- (ii) $(h^{G_n})_{n \in \mathbb{N}}$ is a Cauchy sequence in the metric δ_\square .
- (iii) $(t(F, h^{G_n}))_{n \in \mathbb{N}}$ converges for all finite simple graphs F .
- (iv) There exists an $h \in W$ such that $\lim_{n \rightarrow \infty} t(F, h^{G_n}) = t(F, h)$ for all finite simple graphs F .

5.2.5 Proposition (Lovász and Szegedy [69]). $(\tilde{W}, \delta_\square)$ is compact.

5.2.6 Proposition (Borgs et al. [20]). *Let G_1, G_2 be two labelled simple graphs, and let F be a simple graph. Then*

$$|t(F, G_1) - t(F, G_2)| \leq 4|E(F)|\delta_\square(G_1, G_2). \quad (5.18)$$

For a more detailed description of the structure of the space $(\tilde{W}, \delta_\square)$ we refer the reader to Borgs et al. [20, 21] and Diao et al. [39].

§5.2.3 Large deviation principle for the Erdős-Rényi random graph

In this section we recall a few key facts from the literature about rare events in Erdős-Rényi random graphs, formulated in terms of a large deviation principle. Importantly, the scale that is used is n^2 , the order of the number of edges in the graph.

We start by introducing the large deviation rate function. For $p \in (0, 1)$ and $u \in [0, 1]$, let

$$\begin{aligned} I_p(u) &:= \frac{1}{2}u \log \left(\frac{u}{p} \right) + \frac{1}{2}(1-u) \log \left(\frac{1-u}{1-p} \right), \\ I(u) &:= \frac{1}{2}u \log u + \frac{1}{2}(1-u) \log(1-u) = I_{\frac{1}{2}}(u) - \frac{1}{2} \log 2, \end{aligned} \quad (5.19)$$

with the convention that $0 \log 0 = 0$. For $h \in W$ we write, with a mild abuse of notation,

$$I_p(h) := \int_{[0,1]^2} dx dy I_p(h(x, y)), \quad I(h) := \int_{[0,1]^2} dx dy I(h(x, y)). \quad (5.20)$$

On the quotient space $(\tilde{W}, \delta_\square)$ we define $I_p(\tilde{h}) = I_p(h)$, where h is any element of the equivalence class \tilde{h} .

5.2.7 Proposition (Chatterjee and Varadhan [31]). *The function I_p is well-defined on \tilde{W} and is lower semi-continuous under the δ_\square -metric.*

Consider the set \mathcal{G}_n of all graphs on n vertices and the Erdős-Rényi probability distribution $\mathbb{P}_{n,p}$ on \mathcal{G}_n . Through the mappings $G \rightarrow h^G \rightarrow \tilde{h}^G$ we obtain a probability distribution on W (with a slight abuse of notation again denoted by $\mathbb{P}_{n,p}$), and a probability distribution $\tilde{\mathbb{P}}_{n,p}$ on \tilde{W} .

5.2.8 Theorem (Chatterjee and Varadhan [31]). *For every $p \in (0, 1)$, the sequence of probability distributions $(\tilde{\mathbb{P}}_{n,p})_{n \in \mathbb{N}}$ satisfies the large deviation principle on $(\tilde{W}, \delta_\square)$ with rate function I_p defined by (5.20), i.e.,*

$$\begin{aligned} \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \tilde{\mathbb{P}}_{n,p}(\tilde{C}) &\leq - \inf_{\tilde{h} \in \tilde{W}} I_p(\tilde{h}) \quad \forall \tilde{C} \subset \tilde{W} \text{ closed}, \\ \liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \tilde{\mathbb{P}}_{n,p}(\tilde{O}) &\geq - \inf_{\tilde{h} \in \tilde{O}} I_p(\tilde{h}) \quad \forall \tilde{O} \subset \tilde{W} \text{ open}. \end{aligned} \quad (5.21)$$

Using the large deviation principle we can find asymptotic expressions for the number of simple graphs on n vertices with a given property. In what follows a property of a graph is defined through an operator $T: W \rightarrow \mathbb{R}^m$ for some $m \in \mathbb{N}$. We assume that the operator T is continuous with respect to the δ_\square -metric, and for some $\vec{T}^* \in \mathbb{R}^m$ we consider the sets

$$\tilde{W}^* := \{\tilde{h} \in \tilde{W} : T(\tilde{h}) = \vec{T}^*\}, \quad \tilde{W}_n^* := \{\tilde{h} \in \tilde{W}^* : \tilde{h} = \tilde{h}^G \text{ for some } G \text{ on } n \text{ vertices}\}. \quad (5.22)$$

By the continuity of the operator T , the set \tilde{W}^* is closed. Therefore, using Theorem 5.2.8, we obtain the following asymptotics for the cardinality of \tilde{W}_n^* .

5.2.9 Corollary (Chatterjee [26]). *For any measurable set $\tilde{W}^* \subset \tilde{W}$, with \tilde{W}_n^* as defined in (5.22),*

$$- \inf_{\tilde{h} \in \text{int}(\tilde{W}^*)} I(\tilde{h}) \leq \liminf_{n \rightarrow \infty} \frac{\log |\tilde{W}_n^*|}{n^2} \leq \limsup_{n \rightarrow \infty} \frac{\log |\tilde{W}_n^*|}{n^2} \leq - \inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}), \quad (5.23)$$

where $\text{int}(\tilde{W}^*)$ is the interior of \tilde{W}^* .

§5.3 Variational characterisation of ensemble equivalence

In this section we present a number of preparatory results we will need in Section 5.4 to state our theorem on the equivalence between P_{mic} and P_{can} . Our main result is Theorem 5.3.4 below, which gives us a *variational characterisation of ensemble equivalence*. In Section 5.3.1 we introduce our constraints on the subgraph counts. In Section 5.3.2 we rephrase the canonical ensemble in terms of graphons. In Section 5.3.3 we state and prove Theorem 5.3.4.

§5.3.1 Subgraph counts

First we introduce the concept of subgraph counts, and point out how the corresponding canonical distribution is defined. Label the simple graphs in any order, e.g., F_1 is an edge, F_2 is a wedge, F_3 is triangle, etc. Let $C_k(G)$ denote the number of subgraphs F_k in G . In the dense regime, $C_k(G)$ grows like n^{V_k} , where $V_k = |V(F_k)|$ is the number of vertices in F_k . For $m \in \mathbb{N}$, consider the following *scaled vector-valued function* on \mathcal{G}_n :

$$\vec{C}(G) := \left(\frac{p(F_k)C_k(G)}{n^{V_k-2}} \right)_{k=1}^m = n^2 \left(\frac{p(F_k)C_k(G)}{n^{V_k}} \right)_{k=1}^m. \quad (5.24)$$

The term $p(F_k)$ counts the edge-preserving permutations of the vertices of F_k , i.e., $p(F_1) = 2$ for an edge, $p(F_2) = 2$ for a wedge, $p(F_3) = 6$ for a triangle, etc. The term $C_k(G)/n^{V_k}$ represents a subgraph density in the graph G . The additional n^2 guarantees that the full vector scales like n^2 , the scaling of the large deviation principle in Theorem 5.2.8. For a simple graph F_k we define the *homomorphism density* as

$$t(F_k, G) := \frac{\text{hom}(F_k, G)}{n^{V_k}} = \frac{p(F_k)C_k(G)}{n^{V_k}}, \quad (5.25)$$

which does not distinguish between permutations of the vertices. Hence the Hamiltonian becomes

$$H(\vec{\theta}, \vec{T}(G)) = n^2 \sum_{k=1}^m \theta_k t(F_k, G) = n^2 (\vec{\theta} \cdot \vec{T}(G)), \quad G \in \mathcal{G}_n, \quad (5.26)$$

where

$$\vec{T}(G) := (t(F_k, G))_{k=1}^m. \quad (5.27)$$

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

$$P_{\text{can}}(G \mid \vec{\theta}) := e^{n^2 [\vec{\theta} \cdot \vec{T}(G) - \psi_n(\vec{\theta})]}, \quad G \in \mathcal{G}_n, \quad (5.28)$$

where ψ_n replaces the *partition function*:

$$\psi_n(\vec{\theta}) := \frac{1}{n^2} \log \sum_{G \in \mathcal{G}_n} e^{n^2 (\vec{\theta} \cdot \vec{T}(G))}. \quad (5.29)$$

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

$$\langle \vec{T} \rangle = \sum_{G \in \mathcal{G}_n} \vec{T}(G) P_{\text{can}}(G) = \vec{T}^*. \quad (5.30)$$

The canonical probability then becomes

$$P_{\text{can}}(G) = P_{\text{can}}(G \mid \vec{\theta}^*) \quad (5.31)$$

In Section 5.5.1 we will discuss how to find $\vec{\theta}^*$.

5.3.1 Remark. (i) The constraint \vec{T}^* and the Lagrange multiplier $\vec{\theta}^*$ in general depend on n , i.e., $\vec{T}^* = \vec{T}_n^*$ and $\vec{\theta}^* = \vec{\theta}_n^*$ (recall Remark 5.2.2). We consider constraints that converge when we pass to the limit $n \rightarrow \infty$, i.e.,

$$\lim_{n \rightarrow \infty} \vec{T}_n^* = \vec{T}_\infty^*. \quad (5.32)$$

Consequently, we expect that

$$\lim_{n \rightarrow \infty} \vec{\theta}_n^* = \vec{\theta}_\infty^*. \quad (5.33)$$

Throughout the sequel we *assume* that (5.33) holds. If convergence fails, then we may still consider subsequential convergence. The subtleties concerning (5.33) are discussed in Appendix A.

(ii) In what follows, we suppress the dependence on n and write $\vec{T}^*, \vec{\theta}^*$ instead of $\vec{T}_n^*, \vec{\theta}_n^*$, but we keep the notation $\vec{T}_\infty^*, \vec{\theta}_\infty^*$ for the limit. In addition, throughout the sequel we write $\vec{\theta}, \vec{\theta}_\infty$ instead of $\vec{\theta}^*, \vec{\theta}_\infty^*$ when we view these as parameters that do not depend on n . This distinction is crucial when we take the limit $n \rightarrow \infty$.

§5.3.2 From graphs to graphons

In (5.16) we saw that if we map a finite simple graph G to its graphon h^G , then for each finite simple graph F the homomorphism densities $t(F, G)$ and $t(F, h^G)$ are identical. If $(G_n)_{n \in \mathbb{N}}$ is left-convergent, then

$$\lim_{n \rightarrow \infty} \vec{T}(G_n) = (t(F_k, h))_{k=1}^m \quad (5.34)$$

for some $h \in W$, as an immediate consequence of Theorem 5.2.4. We further see that the expression in (5.26) can be written in terms of graphons as

$$H(\vec{\theta}, \vec{T}(G)) = n^2 \sum_{k=1}^m \theta_k t(F_k, h^G). \quad (5.35)$$

With this scaling the *hard constraint* is denoted by \vec{T}^* , has the interpretation of the *density* of an observable quantity in G , and defines a subspace of the quotient space \tilde{W} , which we denote by \tilde{W}^* , and which consists of all graphons that meet the hard constraint, i.e.,

$$\tilde{W}^* := \{\tilde{h} \in \tilde{W} : \vec{T}(\tilde{h}) = \vec{T}^*\}. \quad (5.36)$$

The *soft constraint* in the canonical ensemble becomes $\langle \vec{T} \rangle = \vec{T}^*$ (recall (5.5)).

§5.3.3 Variational formula for specific relative entropy

In what follows, the limit as $n \rightarrow \infty$ of the partition function $\psi_n(\vec{\theta})$ defined in (5.29) plays an important role. This limit has a variational representation that will be key to our analysis.

5.3.2 Theorem (Chatterjee and Diaconis [29]). *Let $\vec{T}: \tilde{W} \rightarrow \mathbb{R}^m$ be the operator defined in (5.27). For any $\vec{\theta} \in \mathbb{R}^m$ (not depending on n),*

$$\lim_{n \rightarrow \infty} \psi_n(\vec{\theta}) = \sup_{\tilde{h} \in \tilde{W}} \left(\vec{\theta} \cdot \vec{T}(\tilde{h}) - I(\tilde{h}) \right) \quad (5.37)$$

with I and ψ_n as defined in (5.20) and (5.29).

5.3.3 Theorem (Chatterjee and Diaconis [29]). *Let F_1, \dots, F_m be subgraphs as defined in Section 5.3.1. Suppose that $\theta_2, \dots, \theta_m \geq 0$. Then*

$$\lim_{n \rightarrow \infty} \psi_n(\vec{\theta}) = \sup_{0 \leq u \leq 1} \left(\sum_{i=1}^m \theta_i u^{E(F_k)} - I(u) \right), \quad (5.38)$$

where $E(F_k)$ denotes the number of edges in the subgraph F_k .

The key result in this section is the following variational formula for s_∞ defined in Definition 5.2.1. Recall that for $n \in \mathbb{N}$ we write $\vec{\theta}^*$ for $\vec{\theta}_n^*$.

5.3.4 Theorem. *Consider the microcanonical ensemble defined in (5.2) with constraint $\vec{T} = \vec{T}^*$ defined in (5.27), and the canonical ensemble defined in (5.28)–(5.29) with parameter $\vec{\theta} = \vec{\theta}^*$ such that, for every $n \in \mathbb{N}$, (5.30), (5.32) and (5.33) hold. Then*

$$s_\infty = \lim_{n \rightarrow \infty} \frac{1}{n^2} S_n(P_{\text{mic}} | P_{\text{can}}) = \sup_{\tilde{h} \in \tilde{W}} [\vec{\theta}_\infty^* \cdot \vec{T}(\tilde{h}) - I(\tilde{h})] - \sup_{\tilde{h} \in \tilde{W}^*} [\vec{\theta}_\infty^* \cdot \vec{T}(\tilde{h}) - I(\tilde{h})], \quad (5.39)$$

where I is defined in (5.19) and $\tilde{W}^* = \{\tilde{h} \in \tilde{W} : \vec{T}(\tilde{h}) = \vec{T}_\infty^*\}$.

Proof. From (5.10) we have

$$s_\infty = \lim_{n \rightarrow \infty} \frac{1}{n^2} [\log P_{\text{mic}}(G^*) - \log P_{\text{can}}(G^*)], \quad (5.40)$$

where G^* is any graph in \mathcal{G}_n such that $\vec{T}(G^*) = \vec{T}^*$. For the microcanonical ensemble we have

$$\log P_{\text{mic}}(G^*) = -\log \Omega_{\vec{T}^*} = -\log \mathbb{P}_{\frac{1}{2}, n} \left(\{G \in \mathcal{G}_n : \vec{T}(G) = \vec{T}^*\} \right) - \binom{n}{2} \log 2, \quad (5.41)$$

where

$$\Omega_{\vec{T}^*} = |\{G \in \mathcal{G}_n : \vec{T}(G) = \vec{T}^*\}| > 0. \quad (5.42)$$

Define the operator $\vec{T}: W \rightarrow \mathbb{R}^m$, $h \mapsto (t(F_k, h))_{k=1}^m$. This operator can be extended to an operator (with a slight abuse of notation again denoted by \vec{T}) on the quotient space $(\tilde{W}, \delta_\square)$ by defining $\vec{T}(\tilde{h}) = \vec{T}(h)$ with $h \in \tilde{h}$. Define the following sets

$$\tilde{W}^* := \{\tilde{h} \in \tilde{W}: \vec{T}(\tilde{h}) = \vec{T}_\infty^*\}, \quad \tilde{W}_n^* := \{\tilde{h} \in \tilde{W}^*: \tilde{h} = \tilde{h}^G \text{ for some } G \in \mathcal{G}_n\}. \quad (5.43)$$

From the continuity of the operator \vec{T} on \tilde{W} , we see that \tilde{W}^* is a compact subspace of \tilde{W} , and hence is also closed. From Theorem 5.2.6 we have that \vec{T} is a Lipschitz continuous operator on the space $(\tilde{W}, \delta_\square)$. Since \tilde{W} is a compact space, we have that

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}_{\frac{1}{2}, n} \left(\left\{ G \in \mathcal{G}_n: \vec{T}(G) = \vec{T}_\infty^* \right\} \right) = - \inf_{\tilde{h} \in \tilde{W}^*} I_{\frac{1}{2}}(\tilde{h}) = - \inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}) - \frac{1}{2} \log 2. \quad (5.44)$$

The large deviation principle applied to (5.41) yields

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log P_{\text{mic}}(G^*) = \inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}). \quad (5.45)$$

Consider the canonical ensemble and a graph G_n^* on n vertices such that $\vec{T}(G_n^*) = \vec{T}_\infty^*$. By Definition 5.2.3, Proposition 5.2.4, and (5.32) we may suppose that $(G_n^*)_{n \in \mathbb{N}}$ is left-convergent and converges to the graphon h^* . Since \vec{T} is continuous, we have that $\vec{T}(G_n^*)$ converges to $\vec{T}(h^*) = \vec{T}_\infty^*$. From (5.28) we have that

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log P_{\text{can}}(G_n^*) = \vec{\theta}_\infty^* \cdot \vec{T}_\infty^* - \psi_\infty(\vec{\theta}_\infty^*). \quad (5.46)$$

By Theorem 5.3.2,

$$\psi_\infty(\vec{\theta}_\infty^*) = \sup_{\tilde{h} \in \tilde{W}} [\vec{\theta}_\infty^* \cdot \vec{T}(\tilde{h}) - I(\tilde{h})]. \quad (5.47)$$

There is an additional subtlety in proving (5.47) in our setup because $\vec{\theta}_\infty^*$ depends on n . This dependence is treated in Appendix A. Combining (5.45) and (5.47), we get

$$s_\infty = \lim_{n \rightarrow \infty} \frac{1}{n^2} S_n(P_{\text{mic}} | P_{\text{can}}) = \inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}) - \vec{\theta}_\infty^* \cdot \vec{T}_\infty^* + \sup_{\tilde{h} \in \tilde{W}} [\vec{\theta}_\infty^* \cdot \vec{T}(\tilde{h}) - I(\tilde{h})]. \quad (5.48)$$

By definition all elements $\tilde{h} \in \tilde{W}^*$ satisfy $\vec{T}(\tilde{h}) = \vec{T}_\infty^*$. Hence the expression in the right-hand side of (5.48) can be written as

$$\sup_{\tilde{h} \in \tilde{W}} [\vec{\theta}_\infty^* \cdot \vec{T}(\tilde{h}) - I(\tilde{h})] - \sup_{\tilde{h} \in \tilde{W}^*} [\vec{\theta}_\infty^* \cdot \vec{T}(\tilde{h}) - I(\tilde{h})], \quad (5.49)$$

which settles the claim.

5.3.5 Remark. Theorem 5.3.4 and the compactness of \tilde{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{W} also lies in $\tilde{W}^* \subset \tilde{W}$. Equivalently, $s_\infty = 0$ when at least one the maximisers of $\vec{\theta}_\infty^* \cdot \vec{T}(\tilde{h}) - I(\tilde{h})$ satisfies the hard constraint.

§5.4 Main theorem

The variational formula for the relative entropy s_∞ in Theorem 5.3.4 allows us to identify examples where ensemble equivalence holds ($s_\infty = 0$) or is broken ($s_\infty > 0$). We already know that if the constraint is on the edge density alone, i.e., $T(G) = t(F_1, G) = T^*$, then $s_\infty = 0$ (see Garlaschelli *et al.* [48]). In what follows we will look at three models:




Figure 5.2: A 5-star graph and its graphon representation.

- (I) The constraint is on the triangle density, i.e., $\vec{T}_2(G) = t(F_3, G) = T_2^*$ with F_3 the triangle. This will be referred to as the *Triangle Model*.
- (II) The constraint is on the edge density and triangle density, i.e., $\vec{T}(G) = (t(F_1, G), t(F_3, G)) = (T_1^*, T_2^*)$ with F_1 the edge and F_3 the triangle. This will be referred to as the *Edge-Triangle Model*.
- (III) The constraint is on the j -star density, i.e., $\vec{T}(G) = t(T[j], G) = T[j]^*$ with $T[j]$ the j -star graph, consisting of 1 root vertex and $j \in \mathbb{N} \setminus \{1\}$ vertices connected to the root but not connected to each other (see Fig. 5.2). This will be referred to as the *Star Model*.

For a graphon $h \in W$ (recall (5.15)), the edge density and the triangle density equal

$$\begin{aligned} T_1(h) &= \int_{[0,1]^2} dx_1 dx_2 h(x_1, x_2), \\ T_2(h) &= \int_{[0,1]^3} dx_1 dx_2 dx_3 h(x_1, x_2) h(x_2, x_3) h(x_3, x_1), \end{aligned} \quad (5.50)$$

while the j -star density equals

$$T[j](h) = \int_{[0,1]} dx \int_{[0,1]^j} dx_1 dx_2 \cdots dx_j \prod_{i=1}^j h(x, x_i). \quad (5.51)$$

5.4.1 Theorem. *For the above three types of constraint:*

- (I) (a) If $T_2^* \geq \frac{1}{8}$, then $s_\infty = 0$.
 (b) If $T_2^* = 0$, then $s_\infty = 0$.
- (II) (a) If $T_2^* = T_1^{*3}$, then $s_\infty = 0$.
 (b) If $T_2^* \neq T_1^{*3}$ and $T_2^* \geq \frac{1}{8}$, then $s_\infty > 0$.
 (c) If $T_2^* \neq T_1^{*3}$, $0 < T_1^* \leq \frac{1}{2}$ and $0 < T_2^* < \frac{1}{8}$, then $s_\infty > 0$.
 (d) If $T_1^* = \frac{1}{2} + \epsilon$ with $\epsilon \in \left(\frac{\ell-2}{2\ell}, \frac{\ell-1}{2\ell+2}\right)$, $\ell \in \mathbb{N} \setminus \{1\}$, and T_2^* is such that (T_1^*, T_2^*) lies on the scallopy curve in Fig. 5.3, then $s_\infty > 0$.
 (e) If $0 < T_1^* \leq \frac{1}{2}$ and $T_2^* = 0$, then $s_\infty = 0$.
- (III) For every $j \in \mathbb{N} \setminus \{1\}$, if $T[j]^* \geq 0$, then $s_\infty = 0$.

Here, $T_1^*, T_2^*, T[j]^*$ are in fact the limits $T_{1,\infty}^*, T_{2,\infty}^*, T[j]_\infty^*$ in (5.32), but in order to keep the notation light we now also suppress the index ∞ .




Figure 5.3: The admissible edge-triangle density region is the region on and between the blue curves (cf. Radin and Sadun [86]).

Theorem 5.4.1, which states our main results on ensemble equivalence and which is proven in Sections 5.5–5.6, is illustrated in Fig. 5.3. The region on and between the blue curves corresponds to the set of all realisable graphs: if the pair (e, t) lies in this region, then there exists a graph with edge density e and triangle density t . The red curves represent ensemble equivalence, the blue curves and the grey region represent breaking of ensemble equivalence, while in the white region between the red curve and the lower blue curve we do not know what happens. Breaking of ensemble equivalence arises from frustration between the edge and the triangle density.

Each of the cases in Theorem 5.4.1 corresponds to typical behavior of graphs drawn from the two ensembles:

- In cases (I)(a) and (II)(a), graphs drawn from both ensembles are asymptotically like Erdős-Rényi random graphs with parameter $p = T_2^{*1/3}$.
- In cases (I)(b) and (II)(e), almost all graphs drawn from both ensembles are asymptotically like bipartite graphs.
- In cases (II)(b), (II)(c) and (II)(d), we do not know what graphs drawn from the canonical ensemble look like. Graphs drawn from the microcanonical ensemble do not look like Erdős-Rényi random graphs. The structure of graphs drawn from the microcanonical ensemble when the constraint is as in (II)(d) has been determined in Pirkhurko and Razborov [83] and Radin and Sadun [86]. The vertex set of a graph drawn from the microcanonical ensemble can be partitioned into ℓ subsets: the first $\ell - 1$ have size $\lfloor cn \rfloor$ and the last has size between $\lfloor cn \rfloor$ and $2\lfloor cn \rfloor$, where c is a known constant depending on ℓ . The graph has the form of a complete ℓ -partite graph on these pieces, plus some additional edges in the last piece that create no additional triangles.
- In case (III), graphs drawn from both ensembles are asymptotically like Erdős-Rényi random graphs with parameter $p = T[j]^{*1/j}$.

5.4.2 Remark. Similar results hold for the Edge-Wedge-Triangle Model and the Edge-Star Model.

Here are three open questions:

- Identify in which cases (5.32) implies (5.33).
- Is $s_\infty = 0$ as soon as the constraint involves a *single* subgraph count only?
- What happens for subgraphs other than edges, wedges, triangles and stars? Is again $s_\infty > 0$ under appropriate frustration?

§5.5 Choice of the tuning parameter

The tuning parameter is to be chosen so as to satisfy the soft constraint (5.30), a procedure that in equilibrium statistical physics is referred to as the *averaging principle*. Depending on the choice of constraint, finding $\tilde{\theta}^*$ may not be easy, neither analytically nor numerically. In Section 5.5.1 we investigate how $\tilde{\theta}^*$ behaves as we vary \tilde{T}^* for fixed n . We focus on the Edge-Triangle Model (a slight adjustment yields the same results for the Triangle Model). In Section 5.5.2 we investigate how averages under the canonical ensemble, like (5.30), behave when $n \rightarrow \infty$. Here we can treat general constraints defined in (5.27).

For the behaviour of our constrained models, the *sign* of the coordinates of the tuning parameter $\tilde{\theta}^*$ is of pivotal importance, both for a fixed $n \in \mathbb{N}$ and asymptotically (see Bhamidi *et al.* [12], Chatterjee and Diaconis [29], Radin and Yin [87], and references therein). We must therefore carefully keep track of this sign. The key results in this direction are Lemmas 5.5.1 and 5.5.2 below.

§5.5.1 Tuning parameter for fixed n

5.5.1 Lemma. *Consider the Triangle Model with the constraint given by the triangle density T_2^* . For every n , $\theta^* \geq 0$ if and only if $T_2^* \geq \frac{1}{8}$.*

Proof. The proof is similar to that of Lemma 5.5.2 below.

5.5.2 Lemma. *Consider the Edge-Triangle Model. For every n , $\theta_2^* \geq 0$ if and only if $T_2^* \geq \frac{1}{8}$, irrespective of T_1^* . Furthermore, $\theta_1^* \geq 0$ if and only if $T_1^* \geq \frac{1}{2}$.*

Proof. Define, for $\theta_1, \theta_2 \in \mathbb{R}$, the function

$$g(\theta_1, \theta_2) := \sum_{G \in \mathcal{G}_n} \exp \left[n^2 \left(\theta_1 (T_1(G) - \frac{1}{2}) + \theta_2 (T_2(G) - \frac{1}{8}) \right) \right]. \quad (5.52)$$

We first prove that g attains a unique global minimum at $(\theta_1, \theta_2) = (0, 0)$. Consider the canonical ensemble P_{can} as defined in (5.28) and (5.31), with \vec{T} as defined above, and the probability distribution P_{hom} on \mathcal{G}_n that assigns probability $2^{-\binom{n}{2}}$ to every graph $G \in \mathcal{G}_n$. Since P_{hom} is absolutely continuous with respect to P_{can} , the relative entropy $S_n(P_{\text{hom}} | P_{\text{can}})$ is well defined:

$$S_n(P_{\text{hom}} | P_{\text{can}}) = \sum_{G \in \mathcal{G}_n} P_{\text{hom}}(G) \log \frac{P_{\text{hom}}(G)}{P_{\text{can}}(G)} \geq 0. \quad (5.53)$$

Using the form of the canonical ensemble we get, after some straightforward calculations, that, for all $\theta_1, \theta_2 \in \mathbb{R}$,

$$\sum_{G \in \mathcal{G}_n} \exp \left[n^2 (\theta_1 T_1(G) + \theta_2 T_2(G)) \right] \geq 2^{\binom{n}{2}} \exp \left[n^2 \left(\theta_1 \frac{1}{2} + \theta_2 \frac{1}{8} \right) \right], \quad (5.54)$$

where the term in the right-hand side comes from the relation

$$\sum_{G \in \mathcal{G}_n} \frac{1}{2^{\binom{n}{2}}} (\theta_1 T_1(G) + \theta_2 T_2(G)) = \theta_1 \frac{1}{2} + \theta_2 \frac{1}{8}. \quad (5.55)$$

Observe that the left-hand side represents the average edge and triangle density, multiplied with θ_1, θ_2 , in an Erdős-Rényi random graph with parameters $(n, \frac{1}{2})$. From (5.54) we find that $g(\theta_1, \theta_2) \geq 2^{\binom{n}{2}} = g(0, 0)$ for all $\theta_1, \theta_2 \in \mathbb{R}$, and so g attains a global minimum at $(0, 0)$. In what follows we show that this global minimum is unique. A straightforward computation shows that $\partial_{\theta_1} g(\theta_1, \theta_2) = \partial_{\theta_2} g(\theta_1, \theta_2) = 0$ if and only if $\langle T_1 \rangle = \frac{1}{2}$ and $\langle T_2 \rangle = \frac{1}{8}$. Furthermore, the Hessian matrix is a covariance matrix and hence is positive semi-definite. For $\vec{\theta} = (\theta_1, \theta_2) = (0, 0)$ we know that $\langle T_1 \rangle = \frac{1}{2}$ and $\langle T_2 \rangle = \frac{1}{8}$. Hence, by uniqueness of the multiplier $\vec{\theta}^*$ for the constraint $T_1^* = \frac{1}{2}$, $T_2^* = \frac{1}{8}$, we obtain that g has a unique global minimum at $(0, 0)$. Moreover, this shows that g has no other stationary points. Consider the parameter $(\theta_1, \theta_2) = (\theta_1^*, \theta_2^*)$. We have

$$\begin{aligned} \partial_{\theta_2} g(\theta_1^*, \theta_2^*) &= (\langle T_2 \rangle - \frac{1}{8}) \exp[-n^2(\theta_1^* \frac{1}{2} + \theta_2^* \frac{1}{8})] \sum_{G \in \mathcal{G}_n} \exp[n^2(\theta_1^* T_1(G) + \theta_2^* T_2(G))] \\ &= (T_2^* - \frac{1}{8}) \exp[-n^2(\theta_1^* \frac{1}{2} + \theta_2^* \frac{1}{8})] \sum_{G \in \mathcal{G}_n} \exp[n^2(\theta_1^* T_1(G) + \theta_2^* T_2(G))]. \end{aligned} \quad (5.56)$$

If $T_2^* \geq \frac{1}{8}$, then $\partial_{\theta_2} g(\theta_1^*, \theta_2^*) \geq 0$. Because g has a unique stationary point at $(0, 0)$, which is a global minimum, we get $\theta_2^* \geq 0$. Similarly, we can show that if $T_2^* < \frac{1}{8}$, then $\theta_2^* < 0$. Suppose that $T_1^* \geq \frac{1}{2}$. For the parameter $(\theta_1, \theta_2) = (\theta_1^*, \theta_2^*)$ we have

$$\begin{aligned} \partial_{\theta_1} g(\theta_1^*, \theta_2^*) &= \left(\langle T_1 \rangle - \frac{1}{2}\right) \exp\left[-n^2\left(\theta_1^* \frac{1}{2} + \theta_2^* \frac{1}{8}\right)\right] \sum_{G \in \mathcal{G}_n} \exp\left[n^2\left(\theta_1^* T_1(G) + \theta_2^* T_2(G)\right)\right] \\ &= \left(T_1^* - \frac{1}{2}\right) \exp\left[-n^2\left(\theta_1^* \frac{1}{2} + \theta_2^* \frac{1}{8}\right)\right] \sum_{G \in \mathcal{G}_n} \exp\left[n^2\left(\theta_1^* T_1(G) + \theta_2^* T_2(G)\right)\right]. \end{aligned} \quad (5.57)$$

Arguing in a similar way as before, we conclude that $\theta_1^* \geq 0$ if and only if $T_1^* \geq \frac{1}{2}$.

Consider the Edge-Triangle Model and suppose that the constraint (T_1^*, T_2^*) is such that $T_2^* = T_1^{*3}$. Then $\theta_2^* = 0$ and θ_1^* matches the constraint on the edge density only. The following lemma shows that in this case the canonical ensemble behaves like the Erdős-Rényi model with parameter T_1^* , a fact that will be needed later to prove equivalence.

5.5.3 Lemma. *Consider the Edge-Triangle Model with the constraint given by the edge-triangle densities $\vec{T}^* = (T_1^*, T_2^*)$ with $T_2^* = T_1^{*3}$. Consider the canonical ensemble as defined in (5.31). Then, for every $n \in \mathbb{N}$,*

$$\theta_1^* = \frac{1}{2} \log \frac{T_1^*}{1 - T_1^*}, \quad \theta_2^* = 0. \quad (5.58)$$

Proof. From the definition of the canonical ensemble we have that, for $G \in \mathcal{G}_n$,

$$P_{\text{can}}(G) = P_{\text{can}}(G \mid \vec{\theta}^*) = e^{n^2[\theta_1^* T_1(G) + \theta_2^* T_2(G) - \psi_n(\vec{\theta}^*)]}, \quad (5.59)$$

where $\psi_n(\vec{\theta}^*)$ is the partition function defined in (5.29). For the specific value $\vec{\theta} = \vec{\theta}^*$ we have that (recall (5.30))

$$\langle T_1 \rangle = T_1^*, \quad \langle T_2 \rangle = T_2^* = T_1^{*3}. \quad (5.60)$$

We claim that the correct parameter is $\vec{\theta}^* = (\frac{1}{2} \log \frac{T_1^*}{1 - T_1^*}, 0)$. The average fraction of edges is T_1^* (see Park and Newman [81]). The average number of triangles is

$$\begin{aligned} \langle T_2 \rangle &= \frac{\sum_{G \in \mathcal{G}_n} T_2(G) \exp\left[n^2\left(\frac{1}{2} \log \frac{T_1^*}{1 - T_1^*} T_1(G)\right)\right]}{\sum_{G \in \mathcal{G}_n} \exp\left[n^2\left(\frac{1}{2} \log \frac{T_1^*}{1 - T_1^*} T_1(G)\right)\right]} \\ &= \frac{\sum_{G \in \mathcal{G}_n} T_2(G) (T_1^*)^{E(G)} (1 - T_1^*)^{\binom{n}{2} - E(G)}}{\sum_{G \in \mathcal{G}_n} (T_1^*)^{E(G)} (1 - T_1^*)^{\binom{n}{2} - E(G)}} = T_1^{*3}, \end{aligned}$$

where the last equation comes from the fact we are calculating the average number of triangles in an Erdős-Rényi model with probability T_1^* . Since the multiplier θ^* is unique, the proof is complete.

§5.5.2 Tuning parameter for $n \rightarrow \infty$

In Lemma 5.5.4 below we show how averages under the canonical ensemble behave asymptotically when $\vec{\theta}$ does not depend on n . In Lemma A.2 we will look at what happens when $\vec{\theta}$ is a one-dimensional multiplier and depends on n .

5.5.4 Lemma. *Suppose that the operator $\vec{T}: W \rightarrow \mathbb{R}^m$ is bounded and continuous with respect to the δ_\square -norm as defined in (5.13). For $\vec{\theta} \in \mathbb{R}^m$ independent of n , consider the variational problem*

$$\sup_{\tilde{h} \in \tilde{W}} [\vec{\theta} \cdot \vec{T}(\tilde{h}) - I(\tilde{h})], \quad (5.61)$$

where I is defined in (5.19). Suppose that the supremum is attained at a unique point, denoted by $\tilde{h}^*(\vec{\theta})$. Then

$$\lim_{n \rightarrow \infty} \sum_{G \in \mathcal{G}_n} T_k(G) P_{\text{can}}(G \mid \vec{\theta}) = T_k(\tilde{h}^*(\vec{\theta})), \quad k = 1, \dots, m. \quad (5.62)$$

Proof. The average of T_k under the canonical probability distribution is equal to

$$\sum_{G \in \mathcal{G}_n} T_k(G) P_{\text{can}}(G \mid \vec{\theta}) = \sum_{G \in \mathcal{G}_n} T_k(G) e^{n^2 [\vec{\theta} \cdot \vec{T}(G) - \psi_n(\vec{\theta})]} =: \mathcal{J}_n. \quad (5.63)$$

Pick $\delta > 0$ and consider the δ -ball $B_\delta(\tilde{h}^*)$ around the maximiser \tilde{h}^* in the quotient space $(\tilde{W}, \delta_\square)$, i.e.,

$$B_\delta(\tilde{h}^*) := \{\tilde{h} \in \tilde{W} : \delta_\square(\tilde{h}, \tilde{h}^*) < \delta\}. \quad (5.64)$$

We denote by G^δ a graph on n vertices whose graphon is a representative element of the class \tilde{h}^G . With a slight abuse of notation, we denote by G^δ both the graph and the corresponding graphon, and by \tilde{h}^G the corresponding equivalence class in the quotient space $(\tilde{W}, \delta_\square)$. Since $(\tilde{W}, \delta_\square)$ is compact space (recall Proposition 5.2.5), and the graphons associated with finite graphs form a countable family that is dense in $(\tilde{W}, \delta_\square)$ (see Diaconis *et al.* [39], Lovász and Szegedy [69]), there exists a sequence $(\tilde{h}^{G_n})_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} \delta_\square(\tilde{h}^{G_n}, \tilde{h}^*) = 0$. For n large enough the neighbourhood $B_\delta(\tilde{h}^*)$ contains elements of the sequence $(\tilde{h}^{G_n})_{n \in \mathbb{N}}$ and, due to the Lipschitz property (recall Proposition 5.2.6), $\delta_\square(\tilde{h}^{G_n}, \tilde{h}^*) < \delta$ implies $|T_k(\tilde{h}^{G_n}) - T_k(\tilde{h}^*)| < C_k \delta$ for some constant $C_k > 0$ and $k = 1, \dots, m$.

Upper bound for \mathcal{J}_n . We decompose the sum over $G \in \mathcal{G}_n$ into two parts: the first over G whose graphon lies in $B_\delta(\tilde{h}^*)$, the second over G whose graphon lies in $B_\delta(\tilde{h}^*)^c =: \tilde{W}^{\delta, \#}$. We further denote by

$$\mathcal{G}_n^\delta := \{G \in \mathcal{G}_n : |T_k(\tilde{h}^G) - T_k(\tilde{h}^*)| < \delta, \quad k = 1, \dots, m\}, \quad (5.65)$$

the set of all graphs whose subgraph densities $T_k(G)$ are δ -close to $T_k(\tilde{h}^*)$. A graph from this set is denoted by G^δ . We define the set

$$\mathcal{G}_n^{\delta, \#} := \{G \in \mathcal{G}_n : \tilde{h}^G \in \tilde{W}^{\delta, \#}\} \quad (5.66)$$

and, for $k = 1, \dots, m$, obtain the following upper bound:

$$\begin{aligned}
 \mathcal{J}_n &= \sum_{G \in \mathcal{G}_n^\delta} T_k(G) e^{n^2 [\vec{\theta} \cdot \vec{T}(G) - \psi_n(\vec{\theta})]} + \sum_{G \in \mathcal{G}_n^{\delta, \#}} T_k(G) e^{n^2 [\vec{\theta} \cdot \vec{T}(G) - \psi_n(\vec{\theta})]} \\
 &\leq \frac{(T_k(G^\delta) + \delta) \sum_{G \in \mathcal{G}_n^\delta} e^{n^2 \vec{\theta} \cdot \vec{T}(G)}}{\sum_{G \in \mathcal{G}_n^\delta} e^{n^2 \vec{\theta} \cdot \vec{T}(G)}} + \sum_{G \in \mathcal{G}_n^{\delta, \#}} T_k(G) e^{n^2 [\vec{\theta} \cdot \vec{T}(G) - \psi_n(\vec{\theta})]} \\
 &= (T_k(G^\delta) + \delta) + \frac{\sum_{G \in \mathcal{G}_n^{\delta, \#}} T_k(G) e^{n^2 \vec{\theta} \cdot \vec{T}(G)}}{\sum_{G \in \mathcal{G}_n} e^{n^2 \vec{\theta} \cdot \vec{T}(G)}}. \tag{5.67}
 \end{aligned}$$

Next, we further bound the second term in (5.67). By definition, for every $n \in \mathbb{N}$ the range of the operator \vec{T} is a finite set

$$R_n := \{\vec{g} \in [0, \infty)^m : \vec{T}(G) = \vec{g}, G \in \mathcal{G}_n\}. \tag{5.68}$$

For the set R_n we observe that $|R_n| = o(n^{m^2})$. In addition, introduce the sets

$$\begin{aligned}
 \mathcal{G}_n^{\vec{g}} &:= \{G \in \mathcal{G}_n : \vec{T}(G) = \vec{g}\}, \\
 R_n^{\delta, \#} &:= \{\vec{g} \in [0, \infty)^m : \vec{T}(G) = \vec{g}, G \in \mathcal{G}_n^{\delta, \#}\} \subset R_n.
 \end{aligned} \tag{5.69}$$

The operator \vec{T} is bounded, and so there exists an $M > 0$ such that $\|\vec{T}(G)\| \leq M$ for all $G \in \mathcal{G}_n$. Hence, the second term in (5.67) can be bounded from above by

$$\frac{\sum_{G \in \mathcal{G}_n^{\delta, \#}} T_k(G) e^{n^2 \vec{\theta} \cdot \vec{T}(G)}}{\sum_{G \in \mathcal{G}_n} e^{n^2 \vec{\theta} \cdot \vec{T}(G)}} \leq \frac{|R_n^{\delta, \#}| M \exp \left[n^2 \sup_{\vec{g} \in R_n^{\delta, \#}} (\vec{\theta} \cdot \vec{g} + \frac{1}{n^2} \log |\mathcal{G}_n^{\vec{g}}|) \right]}{\exp \left[n^2 \sup_{\vec{g} \in R_n} (\vec{\theta} \cdot \vec{g} + \frac{1}{n^2} \log |\mathcal{G}_n^{\vec{g}}|) \right]}. \tag{5.70}$$

By the large deviation principle in Theorem 5.2.8, we have

$$\frac{1}{n^2} \log |\mathcal{G}_n^{\vec{g}}| = \inf_{\tilde{h} \in \tilde{W}^{\vec{g}}} I(\tilde{h}) + o(1), \tag{5.71}$$

where $\tilde{W}^g = \{\tilde{h} \in \tilde{W} : \vec{T}(\tilde{h}) = \vec{g}\}$. As a consequence, (5.70) is majorised by

$$\begin{aligned}
 &M |R_n^*| e^{o(n^2)} \exp \left[n^2 \left(\sup_{\vec{g} \in R_n^{\delta, \#}} [\vec{\theta} \cdot \vec{g} - \inf_{\tilde{h} \in \tilde{W}^{\vec{g}}} I(\tilde{h})] - \sup_{\vec{g} \in R_n} [\vec{\theta} \cdot \vec{g} - \inf_{\tilde{h} \in \tilde{W}^{\vec{g}}} I(\tilde{h})] \right) \right] \\
 &= M |R_n^*| e^{o(n^2)} \\
 &\exp \left[n^2 \left(\sup_{\vec{g} \in R_n^{\delta, \#}} \sup_{\tilde{h} \in \tilde{W}^{\vec{g}}} [\vec{\theta} \cdot \vec{T}(\tilde{h}) - I(\tilde{h})] - \sup_{\vec{g} \in R_n} \sup_{\tilde{h} \in \tilde{W}^{\vec{g}}} [\vec{\theta} \cdot \vec{T}(\tilde{h}) - I(\tilde{h})] \right) \right] \\
 &= M |R_n^*| e^{o(n^2)} \exp \left[n^2 \left(\sup_{\tilde{h} \in \tilde{W}^{\delta, \#}} [\vec{\theta} \cdot \vec{T}(\tilde{h}) - I(\tilde{h})] - \sup_{\tilde{h} \in \tilde{W}} [\vec{\theta} \cdot \vec{T}(\tilde{h}) - I(\tilde{h})] \right) \right]. \tag{5.72}
 \end{aligned}$$

The last equation can be justified as follows. Define the sets

$$\tilde{W}_n = \{\tilde{h} \in \tilde{W} : \tilde{h} = \tilde{h}^G \text{ for some graph } G \in \mathcal{G}_n\}, \quad \tilde{W}_n^{\delta, \#} = \tilde{W}^{\delta, \#} \cap \tilde{W}_n. \quad (5.73)$$

Since the graphons associated with finite graphs form a countable set that is dense in $(\tilde{W}, \delta_\square)$, we have that

$$\tilde{W} = \text{cl} \left(\bigcup_{n \in \mathbb{N}} \tilde{W}_n \right), \quad \tilde{W}^{\delta, \#} = \text{cl} \left(\bigcup_{n \in \mathbb{N}} \tilde{W}_n^{\delta, \#} \right), \quad (5.74)$$

where cl denotes closure. Using (5.74), and recalling that \vec{T} is continuous and I is lower-semicontinuous, we get

$$\lim_{n \rightarrow \infty} \sup_{\vec{g} \in R_n^{\delta, \#}} \sup_{\tilde{h} \in \tilde{W}^{\vec{g}}} [\vec{\theta} \cdot \vec{T}(\tilde{h}) - I(h)] = \sup_{\tilde{h} \in \tilde{W}^{\delta, \#}} [\vec{\theta} \cdot \vec{T}(\tilde{h}) - I(\tilde{h})], \quad (5.75)$$

and a similar result can be established for the second supremum in the exponent in (5.72). The exponent in (5.72) is negative for all $\delta > 0$ and is independent of n . Moreover, by the left-continuity of the graph sequence $(G_n^\delta)_{n \in \mathbb{N}}$, we have that $\lim_{n \rightarrow \infty} T_k(G_n^\delta) = T_k(\tilde{h}^*)$ for every $k = 1, \dots, m$ and every $\delta > 0$. Combined with the inequality in (5.67), we obtain, for $k = 1, \dots, m$,

$$\lim_{n \rightarrow \infty} \sum_{G \in \mathcal{G}_n} T_k(G) e^{n^2 [\vec{\theta} \cdot \vec{T}(G) - \psi_n(\vec{\theta})]} \leq T_k(\tilde{h}^*). \quad (5.76)$$

Lower bound for \mathcal{J}_n . We distinguish two cases: $T_k(\tilde{h}^*) = 0$ and $T_k(\tilde{h}^*) > 0$. For the first case we trivially get the lower bound

$$\lim_{n \rightarrow \infty} \sum_{G \in \mathcal{G}_n} T_k(G) e^{n^2 \vec{\theta} \cdot \vec{T}(G)} \geq 0 = T_k(\tilde{h}^*). \quad (5.77)$$

For the second case we show the equivalent upper bound for the inverse, i.e.,

$$\lim_{n \rightarrow \infty} \frac{\sum_{G \in \mathcal{G}_n} e^{n^2 \vec{\theta} \cdot \vec{T}(G)}}{\sum_{G \in \mathcal{G}_n} T_k(G) e^{n^2 \vec{\theta} \cdot \vec{T}(G)}} \leq \frac{1}{T_k(\tilde{h}^*)}. \quad (5.78)$$

Using the fact that $T_k(\tilde{h}^*) \neq 0$ is bounded, and using a similar reasoning as for the upper bound on \mathcal{J}_n , the latter is easily verified.

5.5.5 Remark. The convergence in (5.62) is not necessarily uniform in $\vec{\theta}$. Our results in Theorem (5.4.1) (II)(b)-(II)(d) indicate that breaking of ensemble equivalence manifests itself through non uniform convergence in (5.62). In Lemma (A.2) we show that uniform convergence holds when the constraint is on the triangle density only, which explains our result in Theorem (5.4.1) (I).

5.5.6 Remark. The analogue of Lemma 5.5.4 when the supremum in (5.61) has multiple maximisers in \tilde{W} is considerably more involved.

As observed in Remark 5.2.2, in general the tuning parameter $\vec{\theta}^*$ depends on n . We discuss this dependence in Appendix A.

§5.6 Proof of the Main Theorem 5.4.1

We proceed by computing the relative entropy s_∞ . In Sections 5.6.1, 5.6.3, 5.6.4, 5.6.5, 5.6.6 and 5.6.8 we treat the limiting regime where all constraints and parameters are the limiting parameters as in (5.32) and (5.33). In Sections 5.6.2 and 5.6.7 we write $T_{\infty,1}^*, T_{\infty,2}^*, \theta_{\infty,1}^*$ for the limiting regime.

§5.6.1 Proof of (I)(a) (Triangle model $T_2^* \geq \frac{1}{8}$)

Proof. Theorem 5.3.4 says that

$$s_\infty = \sup_{\tilde{h} \in \tilde{W}} [\theta^* T_2(\tilde{h}) - I(\tilde{h})] - \sup_{\tilde{h} \in \tilde{W}^*} [\theta^* T_2(\tilde{h}) - I(\tilde{h})]. \quad (5.79)$$

Consider the first term in the right-hand side (5.79). From Lemma 5.5.1 we know that $\theta^* \geq 0$ if and only if $T_2^* \geq \frac{1}{8}$. From Theorem 5.3.3 it follows that if $\theta^* \geq 0$, then

$$\sup_{\tilde{h} \in \tilde{W}} [\theta^* T_2(\tilde{h}) - I(\tilde{h})] = \sup_{u \in [0,1]} [\theta^* u^3 - I(u)] = \sup_{u \in [0,1]} \ell_3(u; \theta^*). \quad (5.80)$$

From Radin and Yin [87, Proposition 3.2] we know that $\ell_3(u; \theta^*)$ attains a unique global maximum. Let $u^*(\theta^*) = \arg \sup_{u \in [0,1]} \ell_3(u; \theta^*)$ be the unique global maximiser. Using Lemma A.2, we obtain that $u^*(\theta^*) = T_2^{*1/3}$, which leads to

$$\sup_{u \in [0,1]} \ell_3(u; \theta^*) = \theta^* u^*(\theta^*)^3 - I(u^*(\theta^*)) = \theta^* T_2^* - I(T_2^{*1/3}). \quad (5.81)$$

As to the second term in the right-hand side of (5.79), we use Chatterjee and Varadhan [31, Proposition 4.2], which states that, for $T_2^* \in (\frac{1}{8}, 1]$,

$$\inf_{\tilde{h} \in \tilde{W}} I(\tilde{h}) := \inf \{I(\tilde{h}) : \tilde{h} \in \tilde{W}, T_2(\tilde{h}) = T_2^*\} = \inf \{I(\tilde{h}) : \tilde{h} \in \tilde{W}, T_2(\tilde{h}) \geq T_2^*\}. \quad (5.82)$$

Moreover, I is convex at the point $x = T_2^{*1/3}$, and hence from Chatterjee and Varadhan [31, Theorem 4.3] we have that $\inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}) = I(T_2^{*1/3})$. Combining this with (5.81), we conclude that $s_\infty = 0$.

§5.6.2 Proof of (I)(b) ($T_2^* = 0$)

Consider the Triangle Model with the constraint given by the triangle density $T^* = 0$. It was proven by Erdős *et al.* [42] that almost all triangle-free graph have a bipartite structure. For the case of dense graphs, the condition $T^* = 0$ means that the number of triangles in the graph is of order $o(n^2)$. In the proof we will see that the two ensembles are equivalent and that graphs drawn from the two ensembles have a bipartite structure.

Proof. From the construction of the canonical ensemble P_{can} in Section 5.1.3, we observe that $P_{\text{can}}(G) = 0$ when $T(G) > 0$. This is a direct consequence of (5.5). We write

$$\mathcal{G}_n^0 := \{G \in \mathcal{G}_n : T(G) = 0\} \quad (5.83)$$

for the collection of all graphs with triangle density equal to zero. From (5.6) we obtain that $P_{\text{can}}(G) = 0$ if $G \notin \mathcal{G}_n^0$ and $P_{\text{can}}(G) = |\mathcal{G}_n^0|^{-1}$ if $G \in \mathcal{G}_n^0$. Hence $P_{\text{can}}(G) = P_{\text{mic}}(G)$ when the constraint is given by $T^* = 0$, which yields

$$S_n(P_{\text{mic}} \mid P_{\text{can}}) = 0 \quad \forall n \in \mathbb{N} \quad (5.84)$$

and hence $s_\infty = 0$.

§5.6.3 Proof of (II)(a) (Edge-Triangle model $T_2^* = T_1^{*3}$)

For the case $T_1^* = T_2^{*\frac{1}{3}}$ we have shown in Lemma 5.5.3 that the canonical ensemble essentially behaves like an Erdős-Rényi model with parameter $p = T_1^*$. Furthermore, the microcanonical ensemble also has an explicit expression, which is found by using the following lemma.

5.6.1 Lemma. *If $T_1^* = T_2^{*\frac{1}{3}}$, then*

$$\inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}) = I(T_2^{*\frac{1}{3}}) = I(T_1^*). \quad (5.85)$$

Proof. Consider an element $\tilde{h} \in \tilde{W}^*$ with $\tilde{W}^* := \{\tilde{h} \in \tilde{W} : T_1(\tilde{h}) = T_1^* = T_2^{*\frac{1}{3}}, T_2(\tilde{h}) = T_2^*\}$. Using the convexity of I on \tilde{W} and Jensen's inequality, we get

$$I(\tilde{h}) = \int_{[0,1]^2} dx dy I(h(x, y)) \geq I\left(\int_{[0,1]^2} dx dy h(x, y)\right) = I(T_1(\tilde{h})) = I(T_1^*). \quad (5.86)$$

Hence $I(\tilde{h}) \geq I(T_2^{*\frac{1}{3}})$ for every $\tilde{h} \in \tilde{W}^*$, which proves the claim.

Proof of (II)(a). Consider the relative entropy s_∞ as defined in (5.9) and (5.10). Using Lemma 5.5.3, we obtain the expression

$$s_\infty = -\frac{1}{2}T_1^* \log(T_1^*) - \frac{1}{2}(1 - T_1^*) \log(1 - T_1^*) + \inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}). \quad (5.87)$$

From Lemma 5.6.1 we have that $\inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}) = I(T_1^*)$, which yields $s_\infty = 0$.

§5.6.4 Proof of (II)(b) ($T_2^* \neq T_1^{*3}$ and $T_2^* \geq \frac{1}{8}$)

Proof. From Lemma 5.5.2 we know that if $T_1^* \geq \frac{1}{2}$ and $T_2^* \geq \frac{1}{8}$, then $\theta_1^* \geq 0$ and $\theta_2^* \geq 0$ while if $T_1^* < \frac{1}{2}$ and $T_2^* \geq \frac{1}{8}$, then $\theta_1^* < 0$ and $\theta_2^* \geq 0$. An argument similar as above yields

$$\sup_{\tilde{h} \in \tilde{W}} \left[\theta_1^* T_1(\tilde{h}) + \theta_2^* T_2(\tilde{h}) - I(\tilde{h}) \right] = \sup_{u \in [0,1]} \ell_3(u; \vec{\theta}^*), \quad (5.88)$$

where for $\theta_1^* \geq 0$ and $\theta_2^* \geq 0$ the last supremum has a unique solution (see Radin and Yin [87, Proposition 3.2]), while for $\theta_1^* < 0$ and $\theta_2^* \geq 0$ it either has a unique solution or two solutions. We treat these two cases separately.

Unique solution. Because of the uniqueness of the solution, not all realisable hard constraints can be met in the limit (see Lemma 5.5.4). We observe that, if $T_2^* \geq \frac{1}{8}$ and $T_2^* \neq T_1^{*3}$, in the limit as $n \rightarrow \infty$ the canonical ensemble becomes Erdős-Rényi with parameter p . This regime is known as the *high-temperature* regime (see Bhamidi *et al.* [12] and Chatterjee and Diaconis [29]). In what follows we determine the parameter p of the canonical ensemble in the limit. From Bhamidi *et al.* [12, Theorem 7] we have that $p = u^*(\bar{\theta}^*)^{\frac{1}{3}}$ with $u^*(\bar{\theta}^*)^{\frac{1}{3}}$ the unique maximiser of (5.88). The expression in (5.88) thus takes the form

$$\begin{aligned} & \sup_{\tilde{h} \in \tilde{W}} \left[\theta_1^* T_1(\tilde{h}) + \theta_2^* T_2(\tilde{h}) - I(\tilde{h}) \right] \\ &= \sup_{u \in [0,1]} \ell_3(u; \bar{\theta}^*) = \theta_1^* u^*(\bar{\theta}^*)^{\frac{1}{3}} + \theta_2^* u^*(\bar{\theta}^*) - I(u^*(\bar{\theta}^*)^{\frac{1}{3}}). \end{aligned} \quad (5.89)$$

Consider the second term in the right-hand side of (5.39). From the definition of \tilde{W}^* it is straightforward to see that

$$\sup_{\tilde{h} \in \tilde{W}^*} \left[\theta_1^* T_1(\tilde{h}) + \theta_2^* T_2(\tilde{h}) - I(\tilde{h}) \right] = \theta_1^* T_1^* + \theta_2^* T_2^* - \inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}), \quad (5.90)$$

where $\tilde{W}^* = \{\tilde{h} \in \tilde{W} : T_1(\tilde{h}) = T_1^*, T_2(\tilde{h}) = T_2^*\}$. We observe that, due to $T_2^* \neq T_1^{*3}$, the constant function $h \equiv u^*(\bar{\theta}^*)^{\frac{1}{3}}$ does not lie in \tilde{W}^* . This shows that $s_\infty > 0$.

Two solutions. The regime in which the right-hand side of (5.88) has two solutions is known as the *low-temperature* regime. In this case the hard constraints (T_1^*, T_2^*) , with $T_1^* \in [\frac{1}{4}, \frac{1}{2})$, $T_2^* \geq \frac{1}{8}$, lie on a curve on the (T_1, T_2) -plane in such a way such that the tuning parameters (θ_1^*, θ_2^*) lie on the *phase transition* curve found in Chatterjee and Diaconis [29] and Radin and Yin [87]. Denote the two solutions of (5.88) by u_1^*, u_2^* . Because of the constraint we are considering, we have that neither of them lies in \tilde{W}^* . From the compactness of the latter space we see that $s_\infty > 0$.

§5.6.5 Proof of (II)(c) ($T_2^* \neq T_1^{*3}$, $0 < T_1^* \leq \frac{1}{2}$ and $0 < T_2^{*3} < \frac{1}{8}$)

For the case $0 < T_1^* \leq \frac{1}{2}$, $T_2^* < \frac{1}{8}$ we know from Lemma 5.5.2 that $\theta_1^* \leq 0$ and $\theta_2^* < 0$ for every n . Hence, because of (5.33), we have that $\theta_1^* \leq 0$ and $\theta_2^* < 0$. This regime is significantly harder to analyse than the previous regimes. Consider the relative entropy s_∞ and the variational representation given in (5.39). We consider two cases: $T_2^* > T_1^{*3}$ and $T_2^* \leq T_1^{*3}$.

Case $T_2^* > T_1^{*3}$. In this case we have the straightforward inequality

$$s_\infty \geq \theta_2^* (T_1^{*3} - T_2^*) - I(T_1^*) + \inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}). \quad (5.91)$$

Since $T_1^{*3} < T_2^*$, we have $\theta_2^* (T_1^{*3} - T_2^*) > 0$. We show that

$$\inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}) = \inf\{I(\tilde{h}) : \tilde{h} \in \tilde{W}, T_1(\tilde{h}) = T_1^*, T_2(\tilde{h}) = T_2^*\} > I(T_1^*). \quad (5.92)$$

Using the convexity of I on \tilde{W} and Jensen's inequality, we obtain that $I(\tilde{h}) \geq I(T_1^*)$ for all $\tilde{h} \in \tilde{W}^*$. Hence

$$\inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}) > \inf\{I(\tilde{h}) : \tilde{h} \in \tilde{W}, T_1(\tilde{h}) = T_1^*\} = I(T_1^*), \quad (5.93)$$

which settles (5.92). Hence $s_\infty > 0$.

Case $T_2^* \leq T_1^{*3}$. We argue similarly as above. We have the straightforward inequality

$$s_\infty \geq \theta_1^* \left(T_2^{*\frac{1}{3}} - T_1^* \right) - I(T_2^{*\frac{1}{3}}) + \inf_{\tilde{h} \in \tilde{W}} I(\tilde{h}). \quad (5.94)$$

We have seen above that $\inf_{\tilde{h} \in \tilde{W}} I(\tilde{h}) > I(T_1^*)$. We further now that I is decreasing on $[0, \frac{1}{2}]$, and so $I(T_1^*) \geq I(T_2^{*1/3})$. Hence $s_\infty > 0$.

§5.6.6 Proof of (II)(d) ((T_1^*, T_2^*) on the scallopy curve)

We show that if (T_1^*, T_2^*) lies on the lower blue curve in Fig. 5.3 (referred to as the scallopy curve), then $s_\infty > 0$. The case where $T_2^* \geq \frac{1}{8}$ can be dealt with directly via Theorem (II)(b). The proof below deals with the case $T_2^* < \frac{1}{8}$.

Proof. We give the proof for $\ell = 2$, the extension to $\ell > 2$ being similar.

Suppose that $T_1^* = \frac{1}{2} + \epsilon$ with $\epsilon \in (0, \frac{1}{6})$, and that T_2^* is chosen as small as possible. It is known that graphs with a relatively high edge density and with a triangle density that is as small as possible have a d -partite structure with edges added in a suitable way so that the desired triangle density is obtained (see Radin and Sadun [86] and Pikhurko and Raborov [83]). Consider a graph on n vertices, denoted by G , with edge density $T_1 \in (\frac{1}{2}, \frac{2}{3})$ and triangle density as small as possible. The structure of such graphs has been described above before Section 5.5. The graphon counterpart of such graphs is the optimiser of the second supremum in the right-hand side of the variational formula for s_∞ . Using Radin and Sadun [86, Theorem 4.2], we obtain

$$\sup_{\tilde{h} \in \tilde{W}^*} [\theta_1^* T_1(\tilde{h}) + \theta_2^* T_2(\tilde{h}) - I(\tilde{h})] = \theta_1^* T_1^* + \theta_2^* T_2^* - \frac{(1 - c(\epsilon))^2}{2} I(p(\epsilon)), \quad (5.95)$$

where

$$c(\epsilon) = \frac{2 + \sqrt{1 - 6\epsilon}}{6}, \quad p(\epsilon) = \frac{4c(\epsilon)(1 - 2c(\epsilon))}{(1 - c(\epsilon))^2}. \quad (5.96)$$

In order to lighten the notation, we drop the dependence of c and p on ϵ . Furthermore, the optimising graphon has the form

$$h_\epsilon^*(x, y) = \begin{cases} 1 & \text{if } x < c < y \text{ or } y < c < x, \\ p & \text{if } c < x < \frac{1+c}{2} < y \text{ or } c < y < \frac{1+c}{2} < x, \\ 0 & \text{otherwise,} \end{cases} \quad (x, y) \in [0, 1]^2, \quad (5.97)$$

which has triangle density

$$T_2(h_\epsilon^*) = \frac{(2 + \sqrt{1 - 6\epsilon})^2}{36} \frac{1 - \sqrt{1 - 6\epsilon}}{3} = T(\epsilon). \quad (5.98)$$

Let $\tilde{\mathcal{F}}_\epsilon$ be the set of all maximisers of $\theta_1^* T_1(\tilde{h}) + \theta_2^* T_2(\tilde{h}) - I(\tilde{h})$ on \tilde{W} . We show that $h_\epsilon^* \notin \tilde{\mathcal{F}}_\epsilon$, which yields $s_\infty > 0$. From Chatterjee and Diaconis [29, Theorem 6.1] we know that if $\tilde{h} \in \tilde{W}$ maximises $\theta_1^* T_1(\tilde{h}) + \theta_2^* T_2(\tilde{h}) - I(\tilde{h})$ on \tilde{W} , then it must satisfy the Euler-Lagrange equations and it must be bounded away from 0 and 1. Hence we see that \tilde{h}_ϵ^* cannot be a stationary point of $\theta_1^* T_1(\tilde{h}) + \theta_2^* T_2(\tilde{h}) - I(\tilde{h})$ on \tilde{W} , and hence cannot be a maximiser.

§5.6.7 Proof of (II)(e) ($0 < T_1^* \leq \frac{1}{2}$ and $T_2^* = 0$)

Proof. Consider the Edge-Triangle Model with constraint given by the edge and triangle densities $T_1^* \in (0, \frac{1}{2}]$ and $T_2^* = 0$. Working as in Section 5.6.2, we find that the canonical ensemble assigns positive probability only to graphs satisfying the constraint $T_2^* = 0$. Defining \mathcal{G}_n^0 as in (5.83) we obtain

$$P_{\text{can}}(G \mid \vec{\theta}) = \begin{cases} e^{n^2[\theta_1 T_1(G) - \psi_n(\vec{\theta})]} & \text{if } G \in \mathcal{G}_n^0, \\ 0 & \text{else,} \end{cases} \quad (5.99)$$

where $\psi_n(\vec{\theta}) = \sum_{G \in \mathcal{G}_n^0} e^{n^2 \theta_1 T_1(G)}$ is the partition function. From (5.99) we observe that the canonical probability distribution depends only on the edge parameter θ_1 . The parameter θ_1 is chosen equal to θ_1^* that matches the soft constraint, i.e.,

$$\sum_{G \in \mathcal{G}_n^0} T_1(G) P_{\text{can}}(G \mid \vec{\theta}^*) = T_1^*. \quad (5.100)$$

Arguing as in the proof of Chatterjee and Diaconis [29, Theorem 3.1] we find that the relative entropy equals

$$s_\infty = \sup_{\tilde{h} \in \tilde{W}^0} [\theta_{\infty,1}^* T_1(\tilde{h}) - I(\tilde{h})] - \sup_{\tilde{h} \in \tilde{W}^*} [\theta_{\infty,1}^* T_1(\tilde{h}) - I(\tilde{h})], \quad (5.101)$$

where

$$\tilde{W}^0 := \{\tilde{h} \in \tilde{W} : T_2(\tilde{h}) = 0\}, \quad \tilde{W}^* := \{\tilde{h} \in \tilde{W} : T_1(\tilde{h}) = T_{\infty,1}^*, T_2(\tilde{h}) = 0\}. \quad (5.102)$$

Using Chatterjee and Diaconis [29, Theorem 7.1 and Theorem 8.2], we obtain that $s_\infty = 0$.

§5.6.8 Proof of (III) (Star model $T[j]^* \geq 0$)

Proof. From Chatterjee and Diaconis [29, Theorem 6.4] we have that, for all $\theta_\infty^* \in \mathbb{R}$,

$$\sup_{\tilde{h} \in \tilde{W}} [\theta^* W(\tilde{h}) - I(\tilde{h})] = \sup_{u \in [0,1]} [\theta^* u^2 - I(u)], \quad (5.103)$$

which by Radin and Yin [87, Proposition 3.1] has a unique solution, which we denote by $u^*(\theta^*)$. Using Theorem 5.3.4 we get that

$$s_\infty = \theta^* u^*(\theta^*)^2 - I(u^*(\theta^*)) - \theta^* T^* + \inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}), \quad (5.104)$$

where, by Lemma A.2, we have that $u^*(\theta^*) = T^{*\frac{1}{2}}$. This yields

$$s_\infty = -I(T^{*\frac{1}{2}}) + \inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}). \quad (5.105)$$

We show that $\inf_{\tilde{h} \in \tilde{W}^*} I(\tilde{h}) = I(T^{*\frac{1}{2}})$. This is done by slightly modifying the proof of Chatterjee and Diaconis [29, Theorem 6.4]. Indeed, observe that

$$T[j](h) = \int_{[0,1]} dx M(x)^j, \quad M(x) = \int_{[0,1]} dy h(x, y). \quad (5.106)$$

Since I is convex we have

$$\int_{[0,1]^2} dx dy I(h(x, y)) \geq \int_{[0,1]} dx I(M(x)), \quad h \in W, \quad (5.107)$$

with equality if and only if $h(x, y)$ is the same for almost all y . Since h is a symmetric function, we get that equality holds if and only if h is constant. For the constant function $h \equiv (T_j)^{1/j} \in W^* := \{h \in W : T_j(h) = T_j\}$, (5.107) is an equality. Hence, for any minimiser of I on \tilde{W}^* the inequality must be an equality, and thus any minimiser must be constant. This shows that $s_\infty = 0$.

§A Appendix

In this appendix we elaborate on the assumption made in (5.33), i.e., the multiplier $\vec{\theta}_n^*$ converges to a limit $\vec{\theta}_\infty^*$ as $n \rightarrow \infty$. In order to get a meaningful limit, we consider constraints \vec{T}_n^* such that

$$\lim_{n \rightarrow \infty} \vec{T}_n^* = \vec{T}_\infty^*. \quad (5.108)$$

It is straightforward to deduce from Corollary 5.2.9 and (5.26)–(5.30) that if $\{\vec{T}_n^*\}$ is bounded away from 0 and 1 component-wise, then $(\vec{\theta}_n^*)_{n \in \mathbb{N}}$ is bounded away from $-\infty$ and $+\infty$ component-wise. Such a sequence contains a converging subsequence, say, $(\vec{\theta}_{n_k}^*)_{k \in \mathbb{N}}$, which in general need not be unique. Thus, as long as the constraint is component-wise bounded away from 0 and 1, the asymptotic expressions derived in this paper exist, but their values may depend on the subsequence we choose. The value of s_∞ depends on the chosen subsequence, but whether it is positive or zero (i.e., whether there is equivalence) does *not*. A deeper investigation of the behaviour of $\{\vec{\theta}_n^*\}_{n \in \mathbb{N}}$ is interesting, but is beyond the scope of this paper.

We first extend Theorem 5.3.4 for the case when the tuning parameter $\vec{\theta}^*$ depends on n .

A.1 Lemma. *Consider the microcanonical ensemble defined in (5.2) with constraint $\vec{T} = \vec{T}_n^*$ defined as in (5.27), and the canonical ensemble defined in (5.28)–(5.29) with parameter $\vec{\theta} = \vec{\theta}_n^*$ such that (5.30) holds. If the conditions in Remark 5.3.1 hold, then (5.39) holds too.*

Proof. The proof of Theorem 5.3.4 carries over to the setting in which the parameter $\vec{\theta}^*$ depends on n , i.e., $\vec{\theta}^* = \vec{\theta}_n^*$. The only non-trivial step is to show that

$$\lim_{n \rightarrow \infty} \psi_n(\vec{\theta}_n^*) = \psi_\infty(\vec{\theta}_\infty^*). \quad (5.109)$$

In the proof of Theorem 5.3.4 we have shown the pointwise convergence

$$\lim_{n \rightarrow \infty} \psi_n(\vec{\theta}) = \psi_\infty(\vec{\theta}), \quad (5.110)$$

for every $\vec{\theta} \in \mathbb{R}^m$, independently of n . A straightforward computation shows that $\nabla \psi_n(\vec{\theta}) = (\langle T_1 \rangle, \dots, \langle T_m \rangle)$, recall (5.30). Observe that for the specific choice of the parameter $\vec{\theta} = \vec{\theta}_n^* = \vec{\theta}^*$, we have that $\nabla \psi_n(\vec{\theta}_n^*) = (T_1^*, \dots, T_m^*)$, which yields $\|\nabla \psi_n(\vec{\theta})\| \leq m$ for all $n \in \mathbb{N}$ and $\vec{\theta} \in \mathbb{R}^m$. We prove (5.109) under the assumptions made in Remark 5.3.1,

$$\begin{aligned} |\psi_n(\vec{\theta}_n^*) - \psi_\infty(\vec{\theta}_\infty^*)| &\leq |\psi_n(\vec{\theta}_n^*) - \psi_n(\vec{\theta}_\infty^*)| + |\psi_n(\vec{\theta}_\infty^*) - \psi_\infty(\vec{\theta}_\infty^*)| \\ &\leq \|\nabla \psi_n(\vec{\eta})\| \|\vec{\theta}_n^* - \vec{\theta}_\infty^*\| + |\psi_n(\vec{\theta}_\infty^*) - \psi_\infty(\vec{\theta}_\infty^*)| \\ &\leq m \|\vec{\theta}_n^* - \vec{\theta}_\infty^*\| + |\psi_n(\vec{\theta}_\infty^*) - \psi_\infty(\vec{\theta}_\infty^*)| \rightarrow 0, \quad n \rightarrow \infty, \end{aligned} \quad (5.111)$$

where the second inequality follows from the mean-value theorem for some $\vec{\eta} = c\vec{\theta}_n^* + (1-c)\vec{\theta}_\infty^*$, $c \in (0, 1)$. The rest of the proof of Theorem 5.3.4 carries over intact.

In the following lemma we extend the result of Lemma 5.5.4 for the case the operator \vec{T} is the triangle density T_2 . This extension is needed in the proof of Theorem 5.4.1 (I).

A.2 Lemma. *Consider the operator $T_2: \tilde{W} \rightarrow \mathbb{R}$ which is bounded and continuous with respect to the δ_\square -norm as defined in (5.13). For $n \in \mathbb{N}$, consider the tuning parameter θ_n^* according to (5.30), i.e.,*

$$\sum_{G \in \mathcal{G}_n} T_2(G) P_{\text{can}}(G) = T_2^*. \quad (5.112)$$

Suppose that $T_2^ \geq \frac{1}{8}$ and that the limits $T_\infty^*, \theta_\infty^*$ in (5.33) exists. Then*

$$\lim_{n \rightarrow \infty} \sum_{G \in \mathcal{G}_n} T_2(G) P_{\text{can}}(G) = \lim_{n \rightarrow \infty} \frac{\sum_{G \in \mathcal{G}_n} T_2(G) e^{n^2 \theta_n^* T_2(G)}}{\sum_{G \in \mathcal{G}_n} e^{n^2 \theta_n^* T_2(G)}} = u^*(\theta_\infty^*), \quad (5.113)$$

where

$$u^*(\theta) = \arg \sup_{0 \leq u \leq 1} [\theta u^3 - I(u)]. \quad (5.114)$$

Proof. From Lemma 5.5.2, since $T_2^* \geq \frac{1}{8}$ we have that $\theta_n^* \geq 0$ for all n . Consequently, $\theta^* \geq 0$. Define, for $\theta \geq 0$, the function

$$f_n(\theta) := \sum_{G \in \mathcal{G}_n} T_2(G) P_{\text{can}}(G | \vec{\theta}) = \frac{\sum_{G \in \mathcal{G}_n} T_2(G) e^{n^2 \theta T_2(G)}}{\sum_{G \in \mathcal{G}_n} e^{n^2 \theta T_2(G)}} \quad (5.115)$$

and consider the variational problem in (5.61). From Chatterjee and Diaconis [29] we have that, for $\theta \geq 0$,

$$\psi_\infty(\theta) := \sup_{\tilde{h} \in \tilde{W}} [\theta T(\tilde{h}) - I(\tilde{h})] = \sup_{0 \leq u \leq 1} [\theta u^3 - I(u)]. \quad (5.116)$$

From Radin and Sadun [86, Theorem 2.1] we have that the function $\theta \rightarrow u^*(\theta)$ is differentiable on $[0, \infty)$. We also observe that

$$u^*(0) = \frac{1}{2}, \quad \lim_{\theta \rightarrow \infty} u^*(\theta) = 1. \quad (5.117)$$

Moreover, for very n , $\theta \mapsto f_n(\theta)$ is continuous on $[0, \infty)$. Hence, combining Lemma 5.5.4, the continuity of f_n for every n , the analyticity of the limiting function $\theta \mapsto u^*(\theta)$ and (5.117), we obtain that if the limit θ_∞ in (5.33) exists, then

$$\lim_{n \rightarrow \infty} f_n(\theta_n^*) = u^*(\theta_\infty^*) = T_\infty^*, \quad (5.118)$$

which proves the claim.

