The phases of large networks with edge and triangle constraints

Richard Kenyon^{*} Charles Radin[†] Kui Ren [‡] Lorenzo Sadun[§]

July 26, 2017

Abstract

Based on numerical simulation and local stability analysis we describe the structure of the phase space of the edge/triangle model of random graphs. We support simulation evidence with mathematical proof of continuity and discontinuity for many of the phase transitions. All but one of the many phase transitions in this model break some form of symmetry, and we use this model to explore how changes in symmetry are related to discontinuities at these transitions.

1 Introduction

We use the variational formalism of equilibrium statistical mechanics to analyze large random graphs. More specifically we analyze "emergent phases", which represent the (nonrandom) large-scale structure of typical large graphs under global constraints on subgraph densities. We concentrate on the model with edge and triangle density constraints, ε and τ , which in this context play somewhat similar roles that mass and energy density constraints play in microcanonical models of simple materials. Our goal is to understand the statistical states $g_{\varepsilon,\tau}$ which maximize entropy for given ε and τ . These are the analogues of Gibbs states in statistical mechanics.

Other parametric models of random graphs are widely used, in particular exponential random graph models (analogues of grand canonical models of materials), and the edge/triangle constraints have been studied in that formalism since they were popularized by Strauss in 1986 [24]. There is a short discussion of exponential models in Section 2.

This paper, following on [16, 17, 18, 6, 7, 19], is the first attempt to determine the qualitative features of the whole phase space of the edge/triangle model. We discuss what

^{*}Department of Mathematics, Brown University, Providence, RI 02912; rkenyon@math.brown.edu

[†]Department of Mathematics, University of Texas, Austin, TX 78712; radin@math.utexas.edu

[‡]Department of Mathematics and ICES, University of Texas, Austin, TX 78712; ren@math.utexas.edu

[§]Department of Mathematics, University of Texas, Austin, TX 78712; sadun@math.utexas.edu

phases exist, and give the basic features of the transitions between these phases; see Figure 1 for a schematic diagram. While the topography of Figure 1 is correct, the relative sizes of the phases have been distorted to make smaller phases more visible. (In the actual portrait, the phases F(1,1), A(2,0), and B(1,1) comprise more than 95% of the area, making the other phases difficult to see.) The boundary of the phase space, the "scalloped triangle" of Razborov [20], represented by the blue curve in Figure 1, is accurately represented in Figure 6. More detail can be drawn from the numerics in Section 3.



Figure 1: Schematic sketch of 14 of the phases in the edge/triangle model. Continuous phase transitions are shown in red, and discontinuous phase transitions in black. The boundary of the phase space is in blue.

The phase space Γ is the space of achievable values of the constraints, and a phase is a connected open subset of Γ in which the entropy optimizing states are unique for given constraints, and vary analytically with them. Determining the entropy-optimal states for edge/triangle constraints is a difficult variational problem, one that has not been solved rigorously except in special regions of the phase space [16, 17, 18, 6, 7, 19]. (This basic problem will be made more explicit once we carefully introduce the mathematical formalism in Section 2.) However there is solid evidence that each individual optimizer is *multipodal*, also known as a stochastic block model [13], or rather, the graphon version of it; see definition in Section 2. This evidence is borne out by our numerical studies, and based on them we conjecture that the optimal states occur in three families (plus one extra phase), corresponding to certain equivalences or symmetries, as follows. We will give the evidence for our conjectures as we proceed.

In the region $\tau < \varepsilon^3$ of Γ there are three infinite families of phases, denoted A(n, 0), B(n, 1) and C(n, 2); See Figure 1. These phases are distinguished by the symmetry of their underlying graphon. An A(n, 0) phase has a multipodal graphon g in which the vertex set [0, 1] is divided into n equal-length pieces P_1, \ldots, P_n and for $x_i \in P_i$ and $x_j \in P_j$, the quantity $p_{ij} := g(x_i, x_j)$ only depends on whether j = k or $j \neq k$, that is, these pieces P_i are statistically equivalent. It follows that for an A(n, 0) phase there are only two statistical parameters p_{11} and p_{12} , which are thus easily computable functions of the edge and triangle constraints, ε and τ . See for example Figure 3 for the case A(2, 0), and the left most matrix in Figure 4.



Figure 2: Distorted sketch of Razborov's scalloped triangle. The dotted curve is not part of the region; it just indicates the quadratic curve on which cusps lie.

The A(n,0) phase touches the lower boundary of Γ at the cusp $(\varepsilon,\tau) = \left(\frac{n}{n+1}, \frac{n(n-1)}{(n+1)^2}\right)$.

A graphon in a B(n,1) phase has a partition of [0,1] into n statistically equivalent intervals, $n \ge 1$, plus one other interval, see Figure 4 middle. In the phase diagram, these phases are arranged in stripes coming out of the point $(\varepsilon, \tau) = (1,1)$, and various parts of the boundary of B(n,1) are shared with A(n+1,0), A(n+2,0), B(n-1,1), B(n+1,1), and C(n,2). These are depicted schematically in Figure 1.

A graphon in a C(n, 2) phase has a vertex set partition into n statistically equivalent subsets, $n \ge 1$, plus another statistically equivalent pair of subsets, see Figure 4 right. The phase shares boundary with B(n, 1), A(n + 2, 0), and the scallop connecting the cusps with $\varepsilon = n/(n+1)$ and $\varepsilon = (n+1)/(n+2)$. See Figure 1.

F(1,1) denotes the single phase for the region $\tau > \varepsilon^3$; see Figure 1. A graphon in F(1,1) is bipodal, that is, represented by a 2 × 2 matrix. Although this is the same basic structure as B(1,1), we consider the two phases to be different because they have different structure in a neighborhood of their common boundary $\tau = \varepsilon^3$, as an analyzed in [7]. In particular, across the curve $\tau = \varepsilon^3$ the entropy and pattern densities are not differentiable.

$\varepsilon - (\varepsilon^3 - \tau)^{1/3}$ $\varepsilon + (\varepsilon^3 - \tau)^{1/3}$	y	$\varepsilon + (\varepsilon^3 - \tau)^{1/3}$	$\varepsilon - (\varepsilon^3 - \tau)^{1/3}$
	1/2	$\varepsilon - (\varepsilon^3 - \tau)^{1/3}$	$\varepsilon + (\varepsilon^3 - \tau)^{1/3}$

Figure 3: The optimal graphons in phase A(2,0).

The multipodal structure of the $g_{\varepsilon,\tau}$ has only been proven on certain subsets of the phase space: on the boundary, on the so-called "Erdős-Rényi curve" $\tau = \varepsilon^3$ (consisting of constant graphons, see [13, 8]), on the line segment $(1/2, \tau)$ for $0 \le \tau \le 1/8$, and in a region just above the Erdős-Rényi curve; see [7]. It has however also been proven just above the Erdős-Rényi curve in a wide family of other models [7], for all parameters in edge/k-star models [6], and is supported in our edge/triangle model by extensive simulations for a range of constraints [18].

The support for the specific pattern of phases in Figure 1 is purely from simulation and will be described in Section 3.

P _{1,2}	P _{1,2}	p _{1,1}	p _{1,3}		р _{3,3}	D1 4			p _{4,5}	p _{4,4}
						r= 1,4			p _{4,4}	p 4,5
P _{1,2}	P _{1,1}	P _{1,2}	p _{1,2}	<i>P</i> _{1,1}	– <i>p</i> _{1,3}	<i>p</i> _{1,2}	p _{1,2}	p _{1,1}	P 1,4	
			P _{1,1}	p _{1,2}		p _{1,2}	p _{1,1}	p _{1,2}		
						<i>p</i> _{1,1}	<i>p</i> _{1,2}	<i>p</i> _{1,2}		
1/3	1/3	1/3	c/2	c/2	1–c	c/3	c/3	c/3	1–c	

Figure 4: Example of an A(3,0), a B(2,1) and a C(3,2) graphon.

In Section 4 we use this detailed structure of the optimal graphons to analyze the transitions between phases. In our edge/triangle model all phase transitions involve a change of symmetry. Some of the transitions are continuous (have a unique entropy optimizer) and some discontinuous (have a nonunique entropy optimizer), and we prove some of these results in that section. The role of symmetry in determining these qualitative features is of interest, given the major role that symmetry plays in the understanding of phase transitions in statistical physics [10, 1]. We discuss this further in Section 5.

2 Notation and formalism

The study of large dense graphs uses the mathematical tool of graphons, which we now review [8], making use of the discussion in [15]. We let G_n denote the set of graphs on n nodes, which we label $\{1, \ldots, n\}$. Graphs are assumed simple, i.e. undirected and without multiple edges or loops. A graph G in G_n can be represented by its 0, 1-valued adjacency matrix, or alternatively by the function g_G on the unit square $[0, 1]^2$ with constant value 0 or 1 in each of the subsquares of area $1/n^2$ centered at the points $((j - \frac{1}{2})/n, (k - \frac{1}{2})/n)$; see Figure 5. Note that when we represent the adjacency matrix in this way, we naturally put the origin in the lower left corner of the matrix, so that for example the symmetry of the adjacency matrix becomes the symmetry $g_G(x, y) = g_G(y, x)$.

More generally, a graphon $g \in \mathcal{G}$ is an arbitrary symmetric measurable function on $[0, 1]^2$ with values in [0, 1].

We define the "cut metric" on graphons by

$$d(f,g) \equiv \sup_{S,T \subseteq [0,1]} \Big| \int_{S \times T} [f(x,y) - g(x,y)] \, dx \, dy \Big|. \tag{1}$$

Informally, g(x, y) is the probability of an edge between nodes x and y, and so two graphons are called equivalent if they agree up to a 'node rearrangement', that is, $g(x, y) \sim g(\phi(x), \phi(y))$ where ϕ is a measure-preserving transformation of [0, 1] (see [8] for details). The cut metric



Figure 5: Adjacency matrix of a graph with 14 nodes, with origin in the lower left corner.

on graphons is invariant under the action of ϕ : $d(f \circ \phi, g \circ \phi) = d(f, g)$. We define the cut metric on the quotient space $\tilde{\mathcal{G}} = \mathcal{G}/\sim$ of 'reduced graphons' to be the infimum of (1) over all representatives of the given equivalence classes. $\tilde{\mathcal{G}}$ is compact in the topology induced by this metric [8].

A graphon is said to be *multipodal* if it is piecewise constant with rectangular pieces, that is, for some k there is a partition of the vertex set [0, 1] into k sets P_1, \ldots, P_k , and g is constant on $P_i \times P_j$. A k-podal graphon is a multipodal graphon with k pieces. The analogous notion for finite graphs is a stochastic block model.

The 'large scale' features of a graph G which are relevant for us are the *densities* with which various subgraphs H sit in G. Assume for instance that H is a 4-cycle. We could represent the density of H in G in terms of the adjacency matrix A_G by

$$\frac{1}{n(n-1)(n-2)(n-3)} \sum_{w,x,y,z} A_G(w,x) A_G(x,y) A_G(y,z) A_G(z,w),$$
(2)

where the sum is over distinct nodes $\{w, x, y, z\}$ of G. For large n this can approximated, within O(1/n), as:

$$\int_{[0,1]^4} g_G(w,x) g_G(x,y) g_G(y,z) g_G(z,w) \, dw \, dx \, dy \, dz. \tag{3}$$

It is therefore useful to define the density $t_H(g)$ of this H in a graphon g by

$$t_H(g) = \int_{[0,1]^4} g(w, x)g(x, y)g(y, z)g(z, w) \, dw \, dx \, dy \, dz, \tag{4}$$

The density for other subgraphs is defined analogously. In particular the density of edges is represented by

$$t_{\varepsilon}(g) = \int_{[0,1]} g(x,y) \, dx \, dy \tag{5}$$

and the density for triangles is given by

$$t_{\tau}(g) = \int_{[0,1]^3} g(x,y)g(y,z)g(z,x)\,dx\,dy\,dz.$$
(6)

We note that $t_H(g)$ only depends on the equivalence class of g and is a continuous function of g with respect to the cut metric on reduced graphons. It is an important result of [8] that the densities of subgraphs are separating: any two reduced graphons with the same values for all densities t_H are the same.

Our goal is to analyze typical large graphs with variable edge/triangle constraints in the phase space of Figure 2. Our densities are real numbers, limits of densities which are attainable in large finite systems, so we begin by softening the constraints, considering graphs with $n \gg 1$ nodes and with edge/triangle densities (ε', τ') satisfying $\varepsilon - \delta < \varepsilon' < \varepsilon + \delta$ and $\tau - \delta < \tau' < \tau + \delta$ for some small δ (which will eventually disappear.) It is easy to show that the number of such constrained graphs grows as $\exp(sn^2)$ for some $s = s(\varepsilon, \tau, \delta) > 0$, and by a typical graph we mean one chosen from the uniform distribution on the constrained set.

Getting back to our goal of analyzing constrained uniform distributions, a related step is to determine the cardinality of the set of graphs on n vertices subject to the constraints. Constraints are expressed in terms of a vector α of values of a finite set C of densities, and a small parameter δ . Denoting the cardinality by $Z_n(\alpha, \delta)$, it was proven in [16, 17] that $\lim_{\delta \to 0} \lim_{n \to \infty} (1/n^2) \ln[Z_n(\alpha, \delta)]$ exists; it is called the *constrained entropy* $s(\alpha)$. As in statistical mechanics this can be usefully represented via a variational principle.

Theorem 2.1 (The variational principle for constrained graphs [16, 17]) For any k-tuple $\overline{H} = (H_1, \ldots, H_k)$ of subgraphs and k-tuple $\alpha = (\alpha_1, \ldots, \alpha_k)$ of numbers,

$$s(\alpha) = \max_{t_{H_j}(g) = \alpha_j} S(g),\tag{7}$$

where S is the graphon entropy:

$$S(g) = -\frac{1}{2} \int_{[0,1]^2} g(x,y) \ln[g(x,y)] + [1 - g(x,y)] \ln[1 - g(x,y)] \, dx \, dy. \tag{8}$$

We can now state in detail the basic goal of this paper as the determination of those graphons $g_{\varepsilon,\tau}$, with edge/triangle constraints (ε, τ) , which optimize the graphon entropy S(g).

Variational principles such as Theorem 2.1 are well known in statistical mechanics [21, 22, 23]. One aspect of such a variational principle in the current context is not well understood, and that is the uniqueness of the optimizer. In all known graph examples the optimizers in the variational principle are unique in $\tilde{\mathcal{G}}$ except on a lower dimensional set of constraints where there is a phase transition. Assuming there is a unique optimizer for (7), this optimizer is the limiting constrained uniform distribution. We will discuss the question of uniqueness of entropy optimizers in Section 5.

Finally we contrast the above formalism with the formalism of exponential random graph models (ERGM's) mentioned in the Introduction. The latter are widely used, especially in the social sciences, to model graphs on a fixed, small number of nodes [13]. The graphon formalism was partly developed to deal with asymptotically large graphs [8], and in particular the proof of Theorem 2.1 is a modification of one in [4]. However, as was pointed out in [3], for large numbers of nodes the parameters in ERGM's can become redundant. In this circumstance the notions of phase and phase transition are affected by this defect. For instance in the ERGM model based on edge/triangle constraints, as discussed for instance in [3, 2, 11], there is an open region in the two dimensional parameter space in which the optimizing graphon is an Erdős-Rényi graphon, which has only one free parameter, representing edge density. In this case the ERGM model does not capture the behavior of large graphs under variation of edge/triangle constraints; see for instance the discussions in [6, 7].

3 Numerical simulation

We now briefly describe the computational algorithms that we used to obtain the phase portrait sketched in Figure 1. The details of the algorithms, as well as their benchmark with known analytical results, can be found in [18]. In a nutshell, the algorithm works as follows. For each given (ε, τ) value in the phase space, we assume that the entropy maximizing graphons are at most 16-podal, the largest number of podes that our computational power can handle. We draw random samples of symmetric 16-podal graphons. We then compute the entropy, as well as the edge and triangle densities, of these graphons. Among the graphons that satisfy the constraints on edge and triangle densities, we take the ones with maximal entropy as the optimizing graphons.

The computational complexity of such a procedure is extremely high, as we have to generate a large number of random samples in order to ensure (i) there are sufficient number of graphons that satisfy the density constraints and (ii) the optimizing graphons we obtain are "close" to the true optimizers. For a given accuracy parameter ϵ , the size of the random samples needed to ensure that the optimizing graphons we obtain have entropy values that are within ϵ accuracy to the true values is determined by the asymptotic analysis in [18]. We follow this analysis in our numerical computation. For each optimizing graphon which we determined with the sampling algorithm, we use it as the initial guess for a Netwon-type local optimization algorithm, more precisely a sequential quadratic programming (SQP) algorithm, to search for local improvements. Detailed analysis of the computational complexity of the sampling algorithm and the implementation of the SQP algorithm are documented in [18]. The results of the SQP algorithms are the candidate optimizing graphons that we show in the rest of this section.

The A(n,0) phase boundaries. In the first group of numerical simulations we try to determine the boundaries of two of the completely symmetric phases, the A(3,0) phase and the A(4,0) phase. As we pointed out in the Introduction, for a given (ε, τ) value in



Figure 6: Boundaries of the A(2,0), A(3,0) and the A(4,0) phases, in solid blue lines, as determined by numerical simulations. The black dashed lines are the upper and the lower (i.e. the Razborov scallops) boundaries of the phase space. The dash-dotted line in the middle is the Erdős-Rényi curve.

A(n,0) we know analytically the unique expression for the optimal A(n,0) graphon (see, for instance, equation (24)) and the corresponding entropy. Therefore, a given point (ε, τ) is outside of the A(n,0) phase if we can find a graphon that gives a larger entropy value, and is presumed to be within the A(n,0) phase if we cannot find a graphon that has a larger entropy value. We use this strategy to determine the boundaries of the A(3,0) and the A(4,0) phases. In Figure 6 we show the boundaries we determined numerically. Note that since these boundaries are determined using a mesh on the (ε, τ) plane, they are only accurate up to the mesh size in each direction. For better visualization, we reproduced Figure 6 in a rescaled coordinate system, $(\varepsilon, \tau' = \tau - \varepsilon(2\varepsilon - 1))$, in Figure 7.

Phases along line segment $(0.735, \tau)$. In the second group of numerical simulations, we provide some numerical evidence to support our conjecture on phases depicted in Figure 1. Consider the optimizing graphons in the phases which the vertical line segment $(0.735, \tau)$: $\tau \in (0.3504, 0.3971)$, cuts through. (Figure 1 is only schematic: for an accurate representation of the boundaries of A(3,0) and A(4,0) see Figure 6.) From top to bottom, the line cuts through the B(1,1), B(2,1), A(3,0), B(2,1), C(2,2), A(4,0), and C(2,2) phases. In Figure 8, we show the optimizing graphons in phases before and after each transition along the line. Let us mention here two obvious discontinuous transitions, the first being the transition from B(1,1) to B(2,1) at about $\tau = 0.3737$, and the second being the transition from B(2,1) to C(2,2) at about $\tau = 0.3574$. The derivative $\partial s/\partial \tau$ of the entropy with respect



Figure 7: Same as in Figure 6 except that the coordinates $(\varepsilon, \tau' = \tau - \varepsilon(2\varepsilon - 1))$ are used, and ε is restricted to [0.5, 0.8]. Note that the upper boundary of the phase space is outside of the region shown here.

to τ exhibits jumps at these transitions, as seen in Figure 9. A theoretical analysis of the transitions is presented in the next section.

Phases along line segment $(0.845, \tau)$. We performed similar simulations to reveal phases along the line segment $(0.845, \tau)$, $\tau \in (0.5842, 0.6034)$. From all evidence we expect the phases, from top to bottom, to be, respectively B(1,1), B(2,1), B(3,1), B(4,1), A(6,0), B(4,1), B(5,1), and C(5,2). Our simultation did not capture the last two phases, which lie extremely close to the lower boundary of the phase space; representive optimizing graphons for the others are shown in Figure 10. Also, due to limitations in computational power we were not able to resolve the transitions between the phases as accurately as in the previous case, that is, those in Figure 8.

Remark on our conjecture on the phase space structure. We now briefly summarize the evidence behind our conjecture of the phase portrait in Figure 1. First of all, by simulations and proofs in several models, in particular this one, we had found that in the interior of phase spaces optimizing graphons are always multipodal. (It might seem that our simulations assume the multipodality they claim to demonstrate. However the simulations in fact optimize among *n*-podal graphons successively for n = 2, 3, ..., 16 [18], originally expecting to find better and better approximations with increasing *n*, and discovered that for most of the phase space no improvement was found beyond n = 3.) In this model in particular, the only place we find more than 3 podes is for $\varepsilon > 0.7$ near the scalloped bound-



Figure 8: The optimizing graphons at the phase boundaries, cut as τ decreases at fixed $\varepsilon = 0.735$. The first two columns are before the transition and the last two columns are after the transition. Rows, top to bottom, show respectively the following transitions: (i) $B(1,1) \rightarrow B(2,1)$, (ii) $B(2,1) \rightarrow A(3,0)$, (iii) $A(3,0) \rightarrow B(2,1)$, (iv) $B(2,1) \rightarrow C(2,2)$, (v) $C(2,2) \rightarrow A(4,0)$, and (vi) $A(4,0) \rightarrow C(2,2)$. We discuss in the text the τ values surrounding the transitions, corresponding to the middle two columns.

ary. Second, for each point (ε, τ) in the phase space except for $\varepsilon > 0.75$ and close to the

scallops), our algorithm, even though computationally expensive, could determine the optimal graphons up to relatively high accuracy. The computational cost is in general tolerable; in the exceptional region it is too expensive to find enough graphons which satisfy all the constraints to get the accuracy we wanted. (In more detail, we use the techniques explained in [18] with edge and triangle constraint intervals of size 10^{-9} , which determines entropy to order 10^{-6} , from which we determine our optimal graphons.) Third, within each phase away from the scallops our computational power allows us to perform simulations on relatively fine meshes of (ε, τ) . This is how we determined the boundary of the A(3,0) and A(4,0)phases as well as the transitions shown in Figure 9, for instance. In more detail, consider the middle two graphons in each row in Figure 8, which straddle the transitions. They all have edge density $\varepsilon = 0.735$. In the first row they have triangle densities 0.37373880 and 0.37368265; in the second row the triangle densities are 0.37171725 and 0.37154879; in the third row they are 0.36138489 and 0.36132874; in the fourth row they are 0.35745410 and 0.35739795; in the fifth they are 0.35706102 and 0.35694871; and in the last row they are 0.35475870 and 0.35318639.



Figure 9: Left: entropy s as a function of τ along the line segment $(0.735, \tau), \tau \in (0.3504, 0.3971)$; Middle: zoom of $s(\tau)$ around the $B(2, 1) \to C(2, 2)$ phase transition at $\tau = 0.3574$; Right: zoom of $s(\tau)$ around the $B(1, 1) \to B(2, 1)$ phase transition at $\tau = 0.3737$.

4 Analysis of transitions

As noted above, the numerical simulations indicate that all phase transitions below the Erdős-Rényi curve, with the exception of $A(n,0) \leftrightarrow B(n-1,1)$, occur discontinuously. That is, at the transition there are two (or more) inequivalent entropy maximizing graphons, representing limits of the optimizing graphons in each phase and yielding different densities of certain subgraphs. By contrast, at a continuous transition there is a unique entropy maximizing graphon and the densities of all subgraphs change continuously as we cross the phase boundary. In this section, we augment the numerical results by proving analytically that certain of these transitions can only occur discontinuously.



Figure 10: Typical optimizing graphons in the phases that the line segment $(0.845, \tau)$, $\tau \in (0.5842, 0.6034)$ cuts through. From top left to bottom right are respectively graphons in phases B(1, 1), B(2, 1), B(3, 1), B(4, 1), A(6, 0) and B(4, 1), respectively.

Theorem 4.1 Except at a finite number of values of the edge density ε , there cannot be a continuous transition from a B(1,1) bipodal phase to a B(2,1) or C(1,2) tripodal phase.

Theorem 4.1 can be generalized to consider all $B \leftrightarrow C$ transitions:

Theorem 4.2 Except at a finite number of values of the edge density ε , there cannot be a continuous transition from a B(n, 1) to a C(n, 2) phase.

Assuming that our conjectured phase portrait is correct, this leaves the transitions $B(n,1) \leftrightarrow B(n-1,1), C(n-2,2) \leftrightarrow A(n,0), A(n,0) \leftrightarrow B(n-2,1)$ and $A(n,0) \leftrightarrow B(n-1,1)$ for us to consider.

Theorem 4.3 For n > 2, there cannot be a continuous transition from a B(n, 1) phase to a B(n-1, 1) phase, and there cannot be a continuous transition from an A(n, 0) phase to a B(n-2, 1) phase.

To summarize (see Figure 1) we prove below that transitions cannot be continuous between: any two *B* phases; any A(n,0) and B(n-2,1) phases; any *B* and *C* phases except perhaps at finitely many values of ε . As noted earlier, transitions between *A* and *C* phases appear to always be discontinuous, and transitions between any A(n,0) and B(n-1,1)phases appear to always be continuous. However, we currently lack a proof for these last two claims, although in [19] there is a possible path to prove continuity for $A(2,0) \leftrightarrow B(1,1)$. For completeness we note that the two transitions across the Erdős-Rényi curve, $A(2,0) \leftrightarrow F(1,1)$ and $B(1,1) \leftrightarrow F(1,1)$, are continuous: for (ε, τ) on the Erdős-Rényi curve the graphon with constant value ε is easily seen to be the unique entropy-optimizer, and the optimizers from each side must approach it in L^2 and therefore in cut metric.

Proof of Theorem 4.1. The C(1,2) phase is in the *C* family and appears to the left of the A(3,0) phase, while B(2,1) is in the *B* family and appears to the right of A(3,0). For purposes of this proof, however, they are indistinguishable. All that matters is that there are three "podes", with a symmetry swapping two of them.

The proof has three steps:

- 1. Showing that the Lagrange multipliers (see definition below) would have to diverge at a continuous transition.
- 2. Showing that for a bipodal optimizing graphon, divergent Lagrange multipliers can only occur at the "natural boundary" of the B(1,1) phase, namely at the minimum possible value of τ achievable by a bipodal graphon for the given value of ε .
- 3. Showing that on the natural boundary, there exists a tripodal graphon with the given values of (ε, τ) and with higher entropy than any bipodal graphon. That is, showing that the natural boundary of the B(1,1) phase actually lies within a tripodal phase. This step requires that a certain analytic function of ε be nonzero. Since analytic functions (that aren't identically zero) can only have finitely many roots in a compact interval, this step of the proof can break down at finitely many values of ε . (Numerical examination of the analytic function reveals that it does not have any roots at all for relevant values of ε , making the "all but finitely many" caveat moot in practice.)

We establish some notation. For any graphon g, we let $\varepsilon(g)$ and $\tau(g)$ denote the densities $t_H(g)$ where H is an edge and a triangle, respectively. For our B(1,1) graphon, we let

$$a = p_{11}, \qquad d = p_{12}, \qquad b = p_{22},$$
(9)

and let c be the size of the first pode. For a B(2, 1) graphon, we assume that the first two podes are interchangable, each of size c/2, and set $a_+ = p_{12}$, $a_- = p_{11} = p_{22}$, $d = p_{13} = p_{23}$, and $b = p_{33}$. That is, the B(2, 1) graphon is obtained from a B(1, 1) graphon by splitting the first pode in half (and renumbering the last pode), and by making p_{11} and p_{12} distinct variables. The only way that a B(1, 1) graphon can be a limit of B(2, 1) graphons is if $a_+ - a_$ goes to zero (note that if $c \to 1$ it becomes a type A(2, 0) graphon, not a B(1, 1)). So to prove our theorem, we must show that a sequence of B(2, 1) entropy maximizers cannot approach a limiting graphon with $a_+ = a_-$.

The Euler-Lagrange equations for maximizing entropy (see [6]) say that there exist constants α and β such that, for all $(x, y) \in [0, 1]^2$,

$$\frac{\delta S(g)}{\delta g(x,y)} = \alpha \frac{\delta \varepsilon(g)}{\delta g(x,y)} + \beta \frac{\delta \tau(g)}{\delta g(x,y)},\tag{10}$$

or more explicitly

$$S_0'(g(x,y)) = \alpha + 3\beta \int_0^1 g(x,z)g(y,z)dz,$$
(11)

where

$$S(g) = \int S_0(g(x,y)) \, dx \, dy, \quad \text{and} \quad S_0(u) = -u \log(u) - (1-u) \log(1-u). \tag{12}$$

For a B(2,1) graphon, the integral $\int_0^1 g(x,z)g(y,z)dz$ equals

$$\begin{cases} \frac{c}{2}(a_{+}^{2} + a_{-}^{2}) + (1 - c)d^{2}, & \text{for } x, y < \frac{c}{2} \text{ or } \frac{c}{2} < x, y < c; \\ ca_{+}a_{-} + (1 - c)d^{2}, & \text{for } x < \frac{c}{2} < y < c \text{ or } y < \frac{c}{2} < x < c; \\ \frac{c}{2}d(a_{+} + a_{-}) + (1 - c)bd, & \text{for } x < c < y \text{ or } y < c < x; \\ cd^{2} + (1 - c)b^{2}, & \text{for } x, y > c. \end{cases}$$
(13)

The Lagrange multiplier β equals $\partial s/\partial \tau$ and is always positive, since we can increase both the entropy and τ by linearly interpolating between our given graphon and a constant graphon.

Subtracting equation (11) for $(x, y) = (\frac{c}{3}, \frac{2c}{3})$ from equation (11) for $(x, y) = (\frac{c}{3}, \frac{c}{3})$ gives:

$$S_0'(a_-) - S_0'(a_+) = \frac{3\beta c}{2}(a_+ - a_-)^2.$$
(14)

By the mean value theorem, the left hand side equals $-S_0''(u)(a_+ - a_-)$ for some u between a_- and a_+ , and so

$$\beta = \frac{-2S_0''(u)}{3(a_+ - a_-)}.$$
(15)

Since $\beta > 0$ and $S_0''(u) = \frac{-1}{u(1-u)}$ is negative and bounded away from zero, a_+ must be greater than a_- and β must diverge to $+\infty$ as $a_+ - a_-$ approaches zero. To compensate, α must diverge to $-\infty$. This completes step 1.

To understand the effect of divergent Lagrange multipliers, we must consider all of the variational equations for bipodal graphons, including those related to changes in c. The edge, triangle and entropy densities are:

$$\begin{aligned} \varepsilon(g) &= c^2 a + 2c(1-c)d + (1-c)^2 b \\ \tau(g) &= c^3 a^3 + 3c^2(1-c)ad^2 + 3c(1-c)^2 bd^2 + (1-c)^3 b^3 \\ S(g) &= c^2 S_0(a) + 2c(1-c)S_0(d) + (1-c)^2 S_0(d). \end{aligned} \tag{16}$$

Taking gradients with respect to the four parameters (a, b, c, d) and setting $\nabla S = \alpha \nabla \varepsilon + \beta \nabla \tau$ gives the four equations:

$$S_{0}'(a) = \alpha + 3\beta(ca^{2} + (1 - c)d^{2}),$$

$$S_{0}'(b) = \alpha + 3\beta(cd^{2} + (1 - c)b^{2}),$$

$$2cS_{0}(a) + 2(1 - 2c)S_{0}(d) - 2(1 - c)S_{0}(b) = \alpha(2ca + 2(1 - 2c)d - 2(1 - c)b) + 3\beta(c^{2}a^{3} + (2c - 3c^{2})ad^{2} + (3c^{2} - 4c + 1)bd^{2} - (1 - c)^{2}b^{3}),$$

$$S_{0}'(d) = \alpha + 3\beta(cad + (1 - c)bd).$$
(17)

Note that the left hand side of the third equation is always finite, and that the left hand sides of the other equations only diverge if the relevant parameter a, b, or d approaches 0 or 1. Otherwise, in the $\beta \to \infty$ limit the ratios of the coefficients of β and α must be the same for all equations. In other words, there is a constant $\lambda = -\beta/\alpha$ such that

$$\nabla \tau = \lambda \nabla \varepsilon, \tag{18}$$

where we restrict attention to parameters that are not 0 or 1, and such that a parameter $i \in \{a, b, d\}$ equals 0 if $\partial_i \tau > \lambda \partial_i \varepsilon$ and equals 1 if $\partial_i \tau < \lambda \partial_i \varepsilon$. (The signs of the inequalities come from the sign of $S'_0(u)$ as $u \to 0$ or $u \to 1$.)

However, these are precisely the same equations that describe finding a stationary point for τ for fixed ε , without regard to the entropy. For B(1, 1) graphons, such stationary points occur only for Erdős-Rényi graphons, with a = b = d, or for minimizers of τ , with d = 1and b = 0 and c satisfying the algebraic condition

$$0 = a^{3}c^{2} + ac^{2} + 2a^{2}c + 2(1-c) - 4a^{2}c^{2} - 4c(1-c).$$
(19)

This completes step 2.

Finally, we must show that a B(1, 1) graphon with b = 0 and d = 1, and satisfying (19), is not an entropy maximizer. Among bipodal graphons with b = 0 and d = 1 and a fixed value of ε , minimizing τ and minimizing the entropy s give different analytic equations for a and c. For all but finitely many values of ε , these equations have distinct roots, implying that the graphon that minimizes τ does not maximize the entropy. If we start at the bipodal graphon that minimizes τ for fixed ε and change c to $c + \delta_c$, while adjusting a to keep ε fixed, we can increase the entropy to first order in δ_c while only increasing τ to second order in δ_c .

To compensate for this increase in τ , we can split the first pode in half, yielding a B(2,1) tripodal graphon with d = 1 and b = 0, with $a_{+} = a + \delta_{a}$ and $a_{-} = a - \delta_{a}$. This decreases τ by $c^{3}\delta_{a}^{3}$, while decreasing the entropy by $O(\delta_{a}^{2})$. By taking δ_{a} to be of order $\delta_{c}^{2/3}$, we can restore the initial value of τ at an entropy cost of $O(\delta_{c}^{4/3})$.

For δ_c sufficiently small and of the correct sign, the $O(\delta_c^1)$ gain in entropy from changing c and a is greater than the $O(\delta_a^2) = O(\delta_c^{4/3})$ cost in entropy from having $a_+ \neq a_-$, so the resulting B(2,1) graphon has higher entropy, but the same values of ε and τ , than the B(1,1) graphon that minimizes τ (among B(1,1) graphons). This completes step 3.

Proof of Theorem 4.2. The proof follows the same strategy as that of the n = 1 case. In step 1 we show that the Lagrange multipliers must diverge at a continuous transition. In step

2 we show that divergent Lagrange multipliers force a B(n, 1) graphon to be a stationary point of τ for fixed ε . In step 3 we show that we can perturb such a stationary B(n, 1)graphon into a C(n, 2) graphon with the same values of (ε, τ) and more entropy, implying that we are not actually at the phase boundary, which is a contradiction.

We parametrize C(n, 2) graphons as follows: There are two interchangable podes, each of size c/2, and n podes of size (1 - c)/n. Let p_{ij} denote the value of g(x, y) when x is in the *i*-th pode and y is in the *j*-th. We define parameters $\{a_+, a_-, b, d, p\}$ such that

$$p_{ij} = \begin{cases} a_{+} & \text{if } (i,j) = (1,2) \text{ or } (2,1), \\ a_{-} & \text{if } (i,j) = (1,1) \text{ or } (2,2), \\ d & \text{if } i \leq 2 < j \text{ or } j \leq 2 < i, \\ b & \text{if } i = j > 2, \\ p & \text{if } i \neq j \text{ and } i, j > 2. \end{cases}$$

$$(20)$$

The transition that we are trying to rule out is one where a_{\pm} approach a common value a, resulting in a B(n, 1) graphon with one pode of size c and n podes of size (1-c)/n, with

$$p_{ij} = \begin{cases} a & \text{if } (i,j) = (1,1), \\ d & \text{if } i = 1 < j \text{ or } j = 1 < i, \\ b & \text{if } i = j > 1, \\ p & \text{if } i \neq j \text{ and } i, j > 1. \end{cases}$$
(21)

Let ∇S , $\nabla \varepsilon$ and $\nabla \tau$ be the gradients of the functionals S, ε and τ with respect to the parameters (a, b, c, d, p) or (a_+, a_-, b, c, d, p) , depending on the phase we are considering. Maximizing the entropy for fixed (ε, τ) means finding Lagrange multipliers α and β such that

$$\nabla S = \alpha \nabla \varepsilon + \beta \nabla \tau. \tag{22}$$

In the C(n, 2) phase, one checks that $\partial S/\partial a_{-} - \partial S/\partial a_{+}$ is first order in $(a_{+} - a_{-})$, while $\partial \tau/\partial a_{-} - \partial \tau/\partial a_{+}$ is second order. This forces β to diverge to $+\infty$ as $a_{+} - a_{-} \rightarrow 0^{+}$, which in turn forces α to diverge to $-\infty$, exactly as in the proof of Theorem 4.1. This concludes step 1.

Next we consider equation (22) in the limit of divergent α and β . Restricting attention to those parameters for which ∇S does not diverge (i.e. those that are not 0 or 1 in the limiting graphon), we have that $\nabla \varepsilon$ and $\nabla \tau$ must be collinear. That is, there is a constant λ such that, for each parameter $q \in (a, b, d, p)$ taking values in (0, 1), we must have that $\partial \tau / \partial q = \lambda \partial \varepsilon / \partial q$. Furthermore, if in the limit q goes to 0 we must have $\partial \tau / \partial q \geq \lambda \partial \varepsilon / \partial q$ and if q approaches 1 we must have $\partial \tau / \partial q \leq \lambda \partial \varepsilon / \partial q$. That is,

$$\nabla \tau = \lambda \nabla \varepsilon, \tag{23}$$

where we restrict attention to parameters that are not 0 or 1, and we have 1-sided inequalities for those remaining parameters. This is *precisely* the set of equations obtained by ignoring entropy and looking for stationary points of τ for fixed ε , with λ being the Lagrange multiplier of this process. Seeking entropy maximizers with divergent (α, β) is equivalent to seeking stationary points of τ for fixed τ . This concludes step 2.

Now we consider a 1-parameter family of graphons, with a fixed value of ε , that satisfy all of (23) except the equation relating $\partial \tau / \partial c$ and $\partial \varepsilon / \partial c$. Since by assumption we start at a stationary point of τ , moving along this family will only change τ to second order or slower in the change in c, but for all but finitely many values of ε will change S to first order. Move a distance δ_c in the direction of increasing S. A priori we do not know that the resulting change in τ will be positive, but negative changes in τ can be compensated for while *increasing* S, since β is positive. If the change in τ is positive, then we can compensate by splitting the first pode in half, with $a_+ - a_-$ of order $(\delta_c^{2/3})$, restoring τ at an entropy cost of $O(\delta_c^{4/3})$. Since $O(\delta_c^1) > O(\delta_c^{4/3})$, by picking δ_c small enough we can always find a C(n, 2) graphon that does better than the B(n, 1) graphon that was purportedly the entropy maximizer at the phase boundary, which is a contradiction.

Proof of Theorem 4.3. This proof does not require any consideration of entropy. The symmetries are simply incompatible, in that there is no way to approximate a B(n-1,1) graphon with a B(n,1) graphon, and there is no way to approximate a B(n-2,1) graphon with an A(n,0) graphon.

We now turn to determining the locations of the phase transitions. For each of the discontinuous transitions, this is a difficult problem. The optimizing graphons on each side of the transition line are very different, so it is impossible to use perturbation theory to understand the behavior near the line. Instead, one must study each phase separately and approximate the entropy in each phase as an analytic function of ε and τ (e.g., by doing a polynomial fit to numerical data). These functions can then be continued over a larger region and compared. The phase transition line is the locus where the two functions are equal. Using such techniques, we can localize the transition lines and the triple points with considerable accuracy, but in the end the results remain grounded in numerical simulation, and cannot provide independent confirmation of our numerics.

For the continuous $A(n,0) \leftrightarrow B(n-1,1)$ transitions, however, it is possible to obtain an analytic equation satisfied along the transition line. This was already done for the $A(2,0) \leftrightarrow B(1,1)$ transition in [18, 19]. Here we extend the results to n > 2.

This calculation, although elementary, is too long to present here in its entirety; we give the method here. For each fixed n > 2, the space of B(n-1,1) graphons is 5-dimensional. As in Figure 3, we imagine n-1 intervals of size $\frac{c}{n-1}$ and one of size 1-c, and must specify $a = p_{1,1}, b = p_{1,2}, d = p_{1,n}, p = p_{n,n}$ and c. An A(n,0) graphon is a special case of this with p = a, d = b and $c = \frac{n-1}{n}$, and for such graphons the parameters a and b are easily computed from the edge and triangle densities:

$$a = \varepsilon - (n-1) \left(\frac{\varepsilon^3 - \tau}{n-1}\right)^{1/3}; \qquad b = \varepsilon + \left(\frac{\varepsilon^3 - \tau}{n-1}\right)^{1/3}.$$
 (24)

Returning to a general B(n-1,1) graphon, we use the constraints on ε and τ to eliminate two variables, expressing d, p as functions of a, b, c: d = d(a, b, c), p = p(a, b, c). In fact

we don't need to solve explicitly; we only need the first and second partial derivatives of d(a, b, c), p(a, b, c) (evaluated at the parameter values of the A(n, 0) graphon), which can be obtained by implicit differentiation of the equations for ε, τ . Then the entropy S is a function of a, b, c: S = S(a, b, c, d(a, b, c), p(a, b, c)). Computing its Hessian H(S) at the A(n, 0) phase when d = b, p = a yields the matrix

$$H(S) = \begin{pmatrix} \frac{(n-1)\left((a-b)^2 - 4(a-1)a^2X\right)}{(a-1)a(a-b)^2n} & -\frac{2b(n-2)(n-1)X}{(a-b)^2n} & -\frac{2(a+b)X}{a-b} \\ -\frac{2b(n-2)(n-1)X}{(a-b)^2n} & \frac{(n-2)(n-1)\left((a-b)^2 - 2(b-1)b(2a+b(n-4))X\right)}{2(a-b)^2(b-1)bn} & -\frac{2b(n-2)X}{a-b} \\ -\frac{2(a+b)X}{a-b} & -\frac{2b(n-2)X}{a-b} & \frac{2n(\log(1-b)-\log(1-a))}{n-1} \end{pmatrix}$$
(25)

where $X = \tanh^{-1}(a) - \tanh^{-1}(b)$. Within the A(n,0) phase, this second variation is negative-definite, while within the B(n-1,1) phase the matrix has a positive eigenvalue. The boundary is thus defined (locally) by the analytic equation det H(S) = 0.

One can also consider continuous changes from an A(n,0) graphon to a graphon with a symmetry other than B(n-1,1). The local stability condition for such changes works out to be exactly the same as for $A(n,0) \leftrightarrow B(n-1,1)$. The upshot is that the analytic curve det H(S) = 0 defines the boundary of the region where the A(n,0) graphon is stable against small perturbations.

In Figure 11 we plot the regions where the A(3,0) and A(4,0) graphons are stable against small perturbations. Comparing to Figures 6 and 7, we see that the stability regions are larger than the actual A(3,0) and A(4,0) phases, and even overlap! Without assuming *anything* about the phase portrait (beyond the existence of A(3,0) and A(4,0) phases) this proves that some of the transitions from A(3,0) or A(4,0) to other phases must not be governed by the local stability condition, and so must be discontinuous.



Figure 11: The regions of local stability for A(3,0) and A(4,0) graphons, showing overlap. The left plot uses coordinates (ε, τ) and the right plot uses $(\varepsilon, \tau' = \tau - \varepsilon(2\varepsilon - 1))$.

The transition $A(2,0) \leftrightarrow B(1,1)$ is a supercritical pitchfork bifurcation [19]. The analytic family of entropy maximizers in the A(2,0) phase has a natural analytic continuation

into the B(1,1) region, but no longer maximizes entropy there. The analytic family of entropy maximizers in the B(1,1) phase doubles back on itself and cannot be continued into the A(2,0) region. Power series analysis suggests that something similar happens in the transition from A(n,0) to B(n-1,1). The family of graphons representing A(n,0) entropy maximizers can be continued into the B(n-1,1) region but no longer maximizes entropy. The family of graphons representing B(n-1,1) cannot be analytically continued into the A(n,0) region, but instead doubles back into the B(n-1,1) region. Unlike in the case of B(1,1), the "doubling back" branch is not related by symmetry to the original branch, and so represents a different set of reduced graphons, all stationary points of the entropy, but with presumably lower entropy than graphons of the original branch with the same values of (ε, τ) .

It is a classical result (Mantel's theorem [12], generalized by Turán [25]) that the only way to satisfy the constraints of edge density $\varepsilon = 1/2$ and triangle density $\tau = 0$ is with the complete, balanced bipartite graph, which implies that those values of those density constraints determine the values of the densities of all other subgraphs. Put another way, there is a *unique* reduced graphon g with $\tau(g) = 0$ and $\varepsilon(g) = 1/2$. Likewise, for each ε there is a unique reduced graphon with the maximum possible value of $\tau = \varepsilon^{3/2}$. This phenomenon is called 'finite forcing'; see [9]. However, this phenomenon only occurs on the boundary of the phase space:

Theorem 4.4 For each pair (ε_0, τ_0) in the interior of the space of achievable values (see Figure 2), there exist multiple inequivalent graphons g with $\varepsilon(g) = \varepsilon_0$ and $\tau(g) = \tau_0$.

Proof: For fixed ε_0 , let g_0 be a graphon that minimizes $\tau(g)$ given $\varepsilon(g) = \varepsilon_0$, and let g_1 be a graphon that maximizes $\tau(g)$, and let $\phi : [0,1] \to [0,1]$ be a general measure-preserving homeomorphism. The graphon g_0 is always *n*-podal for some *n* (see [16], Theorem 4.2), while g_1 is bipodal. For $t \in (0,1)$, let $g_{t,\phi} = tg_0 + (1-t)g_1 \circ \phi$. This will be a multipodal graphon, generically with 2n podes whose sizes depend on the details of ϕ but not on the value of *t*. In particular, we can choose homeomorphisms ϕ_1 and ϕ_2 such that the podal structure of g_{t,ϕ_1} is different from that of g_{t,ϕ_2} , and hence is different from that of g_{t',ϕ_2} for arbitrary $t' \in (0, 1)$.

It is easy to check that $\varepsilon(g_{t,\phi}) = t\varepsilon(g_0) + (1-t)\varepsilon(g_1 \circ \phi) = \varepsilon_0$, and that for given ϕ , $\tau(g_{t,\phi})$ is a continuous function of t. By the intermediate value theorem, we can thus find graphons g_{t_1,ϕ_1} and g_{t_2,ϕ_2} such that $\varepsilon(g_{t_1,\phi_1}) = \varepsilon(g_{t_2,\phi_2}) = \varepsilon_0$ and $\tau(g_{t_1,\phi_1}) = \tau(g_{t_2,\phi_2}) = \tau_0$. But g_{t_1,ϕ_1} and g_{t_2,ϕ_2} have different podal structures, and so are inequivalent.

In contrast to the graphons not being determined by ε and τ , simulations indicate that maximizing the entropy for fixed (ε, τ) does give a unique reduced graphon throughout the interior, except on a lower dimensional set, the exceptions being constraint values associated with the discontinuous phase transitions. (The uniqueness of the entropy maximizer was also proven analytically for $(\varepsilon, \tau) = (1/2, \tau)$ for any $0 \le \tau \le 1/8$ in [16], and for an open subset of the F(1, 1) phase in [7].) We do not yet have a theoretical understanding of this fundamental issue, sometimes called the Gibbs phase rule in physics; see however [6] for k-star graph models, and [22, 5] for weak versions in physics.

5 Symmetry

All aspects of this paper relate to the symmetry of phases, referring to the symmetry of the unique entropy-optimizing graphons $g_{\varepsilon,\tau}$ for points (ε,τ) in the phase. These symmetries occur at two levels.

The first and most significant level of symmetry is a consequence of the multipodality of the $g_{\varepsilon,\tau}$, which means that the set of all nodes is composed of a finite number of equivalence classes: the probability of an edge, between a node v_j in class j with a node v_k in class k, is independent of v_j and v_k , only depending on j and k.

The second level of symmetry concerns the equivalence classes of nodes: certain equivalence classes have the same sizes and edge probabilities, and others don't. These size and probability parameters are used to distinguish distinct phases, that is, maximal open regions in the parameter space where the entropy-maximizing graphon $g_{\varepsilon,\tau}$ is unique and varies analytically.

Because of multipodality the function $g_{\varepsilon,\tau}$ restricted to a given phase can be considered a smooth vector valued function of (ε, τ) , the coordinates being the probabilities of edges between node equivalence classes, and the relative sizes of those equivalence classes. By the symmetry of a phase we refer to the symmetries among these coordinates, with the following caveat, illustrated through an example. Within the phase F(1, 1) there is a curve such that the two node equivalence classes have the same size. In a narrow sense this might have signalled a higher symmetry. However there is no singular behavior as (ε, τ) crosses this curve so the curve is simply part of F(1, 1) and does not affect the 'symmetry' of the phase.

Each of the phases has a different symmetry and we conjecture that they all fall into the three families: A(n,0), B(n,1) and C(n,2), or F(1,1), in which the notation completely specifies the symmetry (except for the pairs F(1,1), B(1,1), and C(1,2), B(2,1)). We do not have a theoretical argument for the existence of these phases. From our simulation we are confident there are no others except possibly for constraints near (1,1) where it is not feasible to simulate.

An important result of this paper is the conjectured phase diagram, Figure 1. The other goal is to show how knowledge of the structure of optimizing graphons in phases can be helpful in understanding the role of symmetry in specific features of phase transitions. Consider the following, paraphrased from P.W. Anderson [1], a picture he attributes to Landau [10]:

The First Theorem of solid-state physics states that it is impossible to change symmetry gradually. A symmetry element is either there or it is not; there is no way for it to grow imperceptibly.

This intuitive picture has been applied, for instance by Landau, to understand why there is no critical point for the fluid/solid transition [1], though it has been difficult to make the argument rigorous:

This is the theoretical argument, which has appeared to some to be a little too straightforward to be absolutely convincing [14].

We suggest that network models such as the edge/triangle model of this paper provide a useful framework for enabling a rigorous study of such symmetry principles. This was done in [19] specifically for the issue of existence of a critical point.

Landau's symmetry principles are commonly applied to the issue of whether phases are continuous at a transition [10], which is also related to uniqueness of entropy optimizers.

We have proven that in the edge/triangle model certain transitions cannot be continuous, and evidence suggests that certain other transitions are continuous. It is worthwhile discussing how these transitions are approached at the micro-level, that is, in terms of the multipodal parameters, the probabilities of edges between various type of nodes. In this regard Figure 1 and Figure 8 are useful.

For all the discontinuous transitions, those proven and those only seen in simulation, we believe from simulation that the transitions can be visualized as the intersection of a pair of two dimensional smooth surfaces both of which exist beyond the intersection but only represent entropy-optimizers on one side. See rows (i), (iv), (v), (vi) in Figure 8.

For the continuous transitions there is more variety. By simulation the continuous $B(2,1) \leftrightarrow A(3,0)$ transition is achieved through directly acquiring the higher symmetry of A(3,0). See rows (*ii*) and (*iii*) in Figure 8. An analogue in statistical mechanics would be a transition between crystal phases in which a rhombohedral unit cell becomes, and remains, cubic. On the other hand the $A(2,0) \leftrightarrow F(1,1)$ transition occurs by the different bipodal symmetries on both sides rising to the full symmetry of the constant graphons at the Erdős-Rényi curve. It is noteworthy that the full symmetry of the constant graphon is incompatible with any possible phase in our sense: since $\varepsilon = \tau^3$, there is not a two-dimensional family of parameter values.

6 Conclusion

The edge/triangle model in this paper was built by analogy with microcanonical mean-field models in statistical mechanics. Mean-field models are useful because frequently the free energy can be determined analytically as a function of the thermodynamic parameters. An important distinction for the edge/triangle and related random graph models is that not only the free energy (entropy in this case) but also the entropy-optimizing graphons (which are the analogues of the Gibbs states) can sometimes be determined for a range of parameters. For instance in [19, Section 3.7] this control of the optimizing states is used to compute how some global quantity changes with the constraint parameters, a level of analysis never possible in short or mean-field models in statistical mechanics.

The main result of this paper is the conjectured phase diagram, Figure 1, for edge/triangle constraints, based largely on simulation, including continuity/discontinuity of all the transitions and the structure of the entropy-optimizing states within each phase.

The secondary goal is to show how knowledge of the structure of the optimizing graphons can be helpful in understanding the role of symmetry in the continuity/discontinuity of phase transitions. Interesting subjects for further investigation include the triple points where phases F(1, 1), A(2, 0) and B(1, 1) meet, all the pairwise transitions being continuous, and where B(1, 1), A(3, 0) and B(2, 1) meet, with $B(1, 1) \leftrightarrow A(3, 0)$ and $B(2, 1) \leftrightarrow B(1, 1)$ discontinuous and $A(3, 0) \leftrightarrow B(2, 1)$ continuous.

Acknowledgments

The main computational results were obtained on the computational facilities in the Texas Super Computing Center (TACC). We gratefully acknowledge this computational support. This work was also partially supported by NSF grants DMS-1208191, DMS-1208941, DMS-1612668, DMS-1509088, DMS-1321018 and DMS-1620473, and Simons Investigator grant 327929.

References

- P.W. Anderson, Basic Notions of Condensed Matter Physics, Benjamin/Cummings, Menlo Park, 1984.
- [2] D. Aristoff and C. Radin, First order phase transition in a model of quasicrystals, J. Phys. A: Math. Theor. 44 (2011) 255001.
- [3] S. Chatterjee and P. Diaconis, Estimating and understanding exponential random graph models, Ann. Stat. **41** (2013) 2428–2461.
- [4] S. Chatterjee and S.R.S. Varadhan, The large deviation principle for the Erdős-Rényi random graph, Eur. J. Combin. 32 (2011) 1000–1017.
- [5] R. Israel, *Convexity in the Theory of Lattice Gases*, Princeton University Press, Princeton, 1979.
- [6] R. Kenyon, C. Radin, K. Ren and L. Sadun, Multipodal structure and phase transitions in large constrained graphs, J. Stat. Phys., 168 (2017), 233–258.
- [7] R. Kenyon, C. Radin, K. Ren and L. Sadun, Bipodal structure in oversaturated random graphs, *Int. Math. Res. Notices*, **2016** (2016), rnw261.
- [8] L. Lovász, *Large Networks and Graph Limits*, American Mathematical Society, Providence, 2012.
- [9] L. Lovász and B. Szegedy, Finitely forcible graphons, J. Combin. Theory Ser. B 101 (2011) 269-301.
- [10] L.D. Landau and E.M. Lifshitz, Statistical Physics, Pergamon Press, London 1958, trans. E. Peierls and R.F. Peierls, chpt. XIV.

- [11] E. Lubetzky and Y. Zhao, On replica symmetry of large deviations in random graphs, Random Structures and Algorithms 47 (2015) 109-146.
- [12] W. Mantel, Problem 28, Wiskundige Opgaven, 10 (1906) 60-61.
- [13] M.E.J. Newman, Networks: an Introduction, Oxford University Press, Oxford, 2010.
- [14] A. Pippard, The Elements of Classical Thermodynamics, Cambridge University Press, Cambridge, 1979, p. 122.
- [15] C. Radin, Phases in large combinatorial systems, arXiv:1601.04787v2.
- [16] C. Radin and L. Sadun, Phase transitions in a complex network, J. Phys. A: Math. Theor. 46 (2013) 305002.
- [17] C. Radin and L. Sadun, Singularities in the entropy of asymptotically large simple graphs, J. Stat. Phys. 158 (2015) 853–865.
- [18] C. Radin, K. Ren and L. Sadun, The asymptotics of large constrained graphs, J. Phys. A: Math. Theor. 47 (2014) 175001.
- [19] C. Radin, K. Ren and L. Sadun, A symmetry breaking transition in the edge/triangle network model, arXiv:1604.07929v1 (2016).
- [20] A. Razborov, On the minimal density of triangles in graphs, Combin. Probab. Comput. 17 (2008) 603-618.
- [21] D. Ruelle, Correlation functionals, J. Math. Phys. 6 (1965) 201-220.
- [22] D. Ruelle, A variational formulation of equilibrium statistical mechanics and the Gibbs phase rule, *Commun. Math Phys.* 5 (1967) 324-329.
- [23] D. Ruelle, Statistical Mechanics; Rigorous Results, Benjamin, New York, 1969.
- [24] D. Strauss, On a general class of models for interaction, SIAM Rev. 28 (1986) 513-527.
- [25] P. Turán, On an Extremal Problem in Graph Theory. Mat. Fiz. Lapok., 48 (1941), 436-452.