

Bootstrapping exchangeable random graphs*

Alden Green

*Department of Statistics and Data Science,
Carnegie Mellon University,
Pittsburgh PA 15213, USA*

e-mail: ajgreen@andrew.cmu.edu

and

Cosma Rohilla Shalizi

*Department of Statistics and Data Science, Machine Learning Department,
Carnegie Mellon University,
Pittsburgh PA 15213, USA*

e-mail: cshalizi@stat.cmu.edu

Abstract: We introduce two new bootstraps for exchangeable random graphs. One, the “empirical graphon bootstrap”, is based purely on resampling, while the other, the “histogram bootstrap”, is a model-based “sieve” bootstrap. We show that both of them accurately approximate the sampling distributions of motif densities, i.e., of the normalized counts of the number of times fixed subgraphs appear in the network. These densities characterize the distribution of (infinite) exchangeable networks. Our bootstraps therefore give a valid quantification of uncertainty in inferences about fundamental network statistics, and so of parameters identifiable from them.

MSC2020 subject classifications: Primary 62F40; secondary 62G09, 62G05.

Keywords and phrases: Bootstrap, exchangeable random graph, network, motif density.

Received August 2020.

Contents

1	Introduction and Goal	1059
2	Notation and methodology	1061
	2.1 Resampling procedures	1062
3	Main results	1064
4	Simulations	1068
5	Supporting propositions and lemmas	1074
6	Proofs	1078
	6.1 Proof of Theorem 1	1078

*Our work was supported by NSF grant DMS1418124.

6.2 Proof of Theorem 2 1080
 6.3 Proof of Lemma 1 1085
 6.4 Proof of Lemma 2 1086
 6.5 Proof of Proposition 2 1087
 6.6 Proof of Proposition 3 1087
 6.7 Proof of Lemma 3 1089
 6.8 Proof of Lemma 4 1091
 Acknowledgments 1093
 References 1093

1. Introduction and Goal

By this point, it is a cliché to say that networks matter, and that network data analysis is an increasingly important part of statistics. Statistical work has largely focused on elaborating models and obtaining point-estimates of their parameters [31, 21]; there has been comparatively little progress in quantifying uncertainty in these estimates, though that is essential to their scientific utility. If we had widely-accepted parametric models, we might hope to use standard asymptotics, at least heuristically, but we do not have such models, and we have reason to doubt that standard asymptotics apply to networks [34]¹. In other areas of statistics, bootstrapping has been highly successful at quantifying uncertainty, even in the face of model mis-specification and complicated dependence structures [22]. Accordingly, in this paper, we introduce two bootstraps, one, the “empirical graphon bootstrap”, based purely on resampling, the other, the “histogram bootstrap”, being a model-based “sieve” bootstrap. We prove that they both accurately approximate the sampling distributions of “motif densities”, the normalized count of the number of times any fixed subgraph (or “motif”) appears in the network. Under exchangeability of the nodes, such densities are known to characterize the distribution of (infinite) networks, as well as defining the convergence of sequences of individual (non-random) graphs. Our bootstraps therefore provide theoretically sound ways to quantify the uncertainty in inferences regarding a fundamental class of network statistics, and so of parameters identifiable from these statistics.

As a contribution to network data analysis, our work extends previous proposals for quantifying uncertainty by means of subsampling the network and using plug-in Gaussian approximations [2], and heuristic parametric bootstraps [33] and resampling schemes [10]. However, from work on bootstrapping in other areas, we know that estimation of the distribution via resampling can be more accurate, particularly with small sample sizes. While our proofs rely on asymptotic arguments via normal approximations, we think its reasonable to expect that our estimators will perform well in the small sample setting as well. From

¹Since the standard asymptotics essentially rest on the log-likelihood having a quadratic maximum [15], this says something interesting and strange about network dependence, but that will have to be pursued elsewhere.

a bootstrap-theory perspective, our contribution is to extend the validity of resampling and sieve bootstraps to a new type of dependence structure, joining previous work on time series, spatial data, and U-statistics.

Probabilistic background and general approach Exchangeability of the nodes² is a common assumption on networks; it corresponds to the assumption that any two isomorphic graphs should be equally probable, and that no information we have on individual nodes (other than their location in the network) is useful for predicting their links. As with other probabilistic symmetries, exchangeability is useful, in part, because of representation theorems which state that all (infinite) exchangeable distributions are mixtures of certain extremal distributions with the same symmetry, but stronger independence properties [8, 19, 23]. In the case of exchangeable networks, the relevant extremal distributions, now often called “graphons”, were characterized by Aldous and Hoover, and work as follows [19, ch. 7]. Every node gets an independent, uniformly distributed random variable on $[0, 1]$, say ϵ_i for node i , and there is a link function $h : [0, 1] \times [0, 1] \mapsto [0, 1]$, symmetric in its arguments, such that the probability of an edge between i and j is $h(\epsilon_i, \epsilon_j)$, and all edges are independent (given the ϵ s). Any exchangeable distribution is a mixture of such graphon distributions³, and any one realization of an exchangeable distribution is drawn from a *single* h .

This provides a natural approach to bootstrapping: estimate the link function h , then randomly redraw node variables and reconnect the edges with the corresponding probability. Our task is thus just (!) to estimate the link function sufficiently well. We propose two approaches. One, the “empirical graphon”, takes the adjacency matrix, views it as a binary-valued function on the unit square, and uses that as our estimate of h . Proving the validity of this bootstrap then relies on results about the convergence, in a suitable topology, of exchangeable random graphs to their generating graphon. Our other bootstrap is a histogram-like estimator of the graphon, a special case of stochastic block models, essentially approximating h by a series of simple functions. Its validity rests on some smoothness assumptions regarding h , but, when they hold, it gives a faster rate of convergence. In both approaches, a Berry-Esseen inequality for U-statistics due to Callaert and Janssen [7] provides a crucial technical tool.

Organization Section 2 fixes notation, lays out assumptions, and, in §2.1, formally proposes the two bootstraps. Section 3 gives the main theorems, stating conditions under which our bootstraps consistently approximate the distribution of motif densities. Section 4 shows empirically that our bootstraps perform well even on moderate sized networks. Section 5 collects supporting propositions and lemmas, and Section 6 proves the main results.

²Often called “joint” exchangeability, to indicate that one applies the same permutation to both the rows and the columns of the adjacency matrix of a graph, as opposed to the stricter “separate” exchangeability, where rows and columns can be differently permuted [19, sec. 7.1].

³Naturally, the same distribution can be realized by many different h functions, which leads to some subtleties in a formal statement of the representation theorem. We do not elaborate on this, since it is not relevant to our concerns, but see Kallenberg [19, sec. 7.6].

2. Notation and methodology

We (mostly) follow the notation of Bhattacharyya and Bickel [2]. Unless otherwise noted, by “graph” we will always mean an undirected, simple graph. For any graph G , $V(G)$ will be the set of its vertices, and $E(G)$ the set of its edges; when $i \in V(G)$, $j \in V(G)$, we write (i, j) for the (unordered) pair of nodes, and $(i, j) \in E(G)$ or $(i, j) \notin E(G)$ depending on whether or not there is an edge. We will sometimes abbreviate this as $(i, j) \in G$ when there is no chance of ambiguity. We will use colons to abbreviate sequences, so that $i : j$ stands for $i, i + 1, \dots, j - 1, j$, and (say) $x_{i:j}$ the sequence of variables $x_i, x_{i+1}, \dots, x_{j-1}, x_j$. Given an ordered p -tuple of indices in $1 : n$, $\mathbf{i} = (i_1, i_2, \dots, i_p)$, we let $G(\mathbf{i}) = (\mathbf{i}, E(G) \cap \{\mathbf{i} \times \mathbf{i}\})$ be the induced subgraph of G with those vertices; we will write $S_n(p)$ for the collection of all ordered p -tuples of $1 : n$. Two graphs G_1 and G_2 are isomorphic when their nodes can be put in one-to-one correspondence while preserving both edges and non-edges, i.e., there is an invertible mapping $\sigma : V(G_1) \mapsto V(G_2)$ such that $(i, j) \in E(G_1)$ if and only if $(\sigma(i), \sigma(j)) \in E(G_2)$. When this holds, we write $G_1 \simeq G_2$, and we write $N(G)$ for the number of graphs on $1 : |V(G)|$ which are isomorphic to G . K_p will indicate the complete graph on $1 : p$, i.e., the p -node graph with all possible edges.

Our data G_n is a graph on the vertices $1 : n$, with corresponding $n \times n$ adjacency matrix A . We assume that the graph is exchangeable, and hence was generated as follows:

$$\epsilon_i \stackrel{\text{i.i.d}}{\sim} \text{Uniform}(0, 1) \tag{1}$$

$$A_{ij} | \epsilon_{1:n} \stackrel{\text{ind}}{\sim} \text{Bernoulli}(h_n(\epsilon_i, \epsilon_j)) \tag{2}$$

for a symmetric, measurable link function $h_n : [0, 1] \times [0, 1] \mapsto [0, 1]$. Without loss of generality, we decompose the function h_n as

$$h_n(u, v) = \rho_n w(u, v) \tag{3}$$

where $\int_0^1 \int_0^1 w(u, v) du dv = 1$, so that ρ_n is the marginal probability of an edge between any two nodes, i.e., the (expected) edge density⁴. We will frequently make assume that w is bounded in various $L^p([0, 1]^2)$ norms, and we use $\|w\|_p = (\int_{[0,1]^2} w(u, v)^p du dv)^{1/p}$ to refer to such norms.

Fixing any p -node connected, simple, undirected graph R that we like, we can ask about the probability that the first p nodes of G_n instantiate this motif⁵,

$$P_R(h) = \mathbb{P}(G_n(1 : p) = R) \tag{4}$$

⁴The job of the ρ_n factor is to allow the graph to become sparse as n grows, as in Borgs et al. [6]; otherwise, graph sequences generated by graphons are “dense”, i.e., the number of edges grows quadratically with the number of nodes. If this is not a concern, and this is a point of some debate in the field, one can fix ρ_n to a constant value for all n . All of our results are valid under such dense-graph limits, and indeed most of them would simplify.

⁵The literature typically calls both R and $G_n(1 : p)$ “subgraphs”; to avoid confusion, we borrow the term “motif” from Milo et al. [30] to designate the pattern being matched, though those authors suggested using it for the patterns which were, in some sense, more common than expected by chance.

Of course, by exchangeability, $P_R(h) = \mathbb{P}(G_n(\mathbf{i}) = R)$ for any $\mathbf{i} \in S_n(p)$. These probabilities are thus very much like moments of the distribution of G_n , and indeed it is known from previous work [27] that the collection of these probabilities, over all motifs R , suffice to characterize an exchangeable graph distribution⁶. One can show [27] that

$$P_R(h) = \mathbb{E} \left[\prod_{(i,j) \in E(R)} h_n(\epsilon_i, \epsilon_j) \prod_{(i,j) \notin E(R)} (1 - h_n(\epsilon_i, \epsilon_j)) \right] \quad (5)$$

It is natural to want to relate these moments to their sample counter-parts. It turns out that a good estimate for $P_R(h)$ is simply to count the number of induced subgraphs in G_n which are isomorphic to R :

$$P_R(G_n) = \frac{1}{\binom{n}{p} p! N(R)} \sum_{\mathbf{i} \in S_n(p)} \mathbb{I}(G_n(\mathbf{i}) \simeq R) \quad (6)$$

Unsurprisingly, $\mathbb{E}[P_R(G_n)] = P_R(h)$. Moreover, previous work⁷ on graph limits tells us that, for fixed R , $P_R(G_n) \xrightarrow{p} P_R(h)$.

Finally, we will need a few scaled versions of the above quantities, since we allow the sparsity factor ρ_n to approach 0 as n grows to ∞ . First, let $\hat{\rho}_n = P_{K_2}(G_n)$ be the edge density observed in the graph. Second, let $\tilde{P}_R(h) = \frac{P_R(h)}{\rho_n^{|E(R)|}}$, and its corresponding empirical quantity $\tilde{P}_R(G_n) = \frac{P_R(G_n)}{\hat{\rho}_n^{|E(R)|}}$.

Miscellaneous notation and conventions Unless otherwise noted, all limits are taken as the number of nodes in the graph grows, i.e., as $n \rightarrow \infty$.

2.1. Resampling procedures

Our resampling procedures begin with an estimate of the graphon, $\hat{h} : [0, 1]^2 \rightarrow [0, 1]$, a mapping to be estimated using the graph G_n . Given such an estimate \hat{h} , we then generate m random variables $\epsilon_i^* \stackrel{\text{i.i.d.}}{\sim} \text{Uniform}(0, 1)$; here $m \in \mathbb{N}$ will be the number of nodes in our resampled network. We then simulate from \hat{h} in the way we generate from graphons, forming the bootstrapped network G_m^* . More precisely, we let the bootstrapped adjacency matrix $A^* = (A_{ij}^*)$ where the entries A_{ij}^* are conditionally independent given G_n and $\epsilon_{1:m}^*$, and follow the distribution $A_{ij}^* | (G_n, \epsilon_{1:m}^*) \sim \text{Bern}(\hat{h}(\epsilon_i^*, \epsilon_j^*))$. (Implicitly, $m = m(n)$). The properties of our bootstrapping procedure clearly depend on the graphon estimate \hat{h} , and we now formally define the two estimators we will subsequently analyze.

⁶See Bickel, Chen and Levina [3] for a discussion, and a method-of-moments procedure for estimating h , based on this fact.

⁷See, for instance, Borgs et al. [5, Lemma 4.4], which gives an explicit (though potentially loose) rate of convergence. This rate is fast enough that a Borel-Cantelli argument could strengthen convergence in probability to almost-sure convergence, but this goes beyond what we need here.

Our first approach estimates the graphon by using its empirical counterpart: the adjacency matrix.

Definition 1 (Empirical Graphon). *The **empirical graphon**, denoted \hat{h}_{adj} , is $\hat{h}_{adj}(u, v) := A_{\lceil nu \rceil \lceil nv \rceil}$. We refer to the process of resampling using the empirical graphon as the empirical graphon bootstrap.*

The empirical graphon bootstrap is equivalent to sampling m vertices from G_n (with replacement), and adding in adjacencies exactly as they appear in G_n . Despite the intuitive analogy between this scheme and the classical i.i.d bootstrap — here, we treat the vertices as the units of data to be resampled — and the widespread success of the bootstrap in the i.i.d case, vertex resampling procedures have not heretofore been rigorously analyzed⁸. One potential explanation for this is that when the number of resampled vertices m is sufficiently large, with high probability the resampled graph G_m^* will contain multiple copies of the same vertex in G_n . Since G_n is a simple graph containing no loops, these copies will never be connected in G_m^* . This (non)-adjacency structure between copies of the same vertex does not reflect any underlying feature of the true generative process by which G_n was formed, and therefore induces a bias between the conditional distribution of G_m^* and the distribution of G_n , a fact remarked upon by [10, 25]. However, as we will see, under appropriate conditions this bias is asymptotically negligible, at least with regards to estimating the distribution of motif densities.

That being said, these conditions, such as on the maximum amount of sparsity tolerated, may be unrealistic depending on the particular problem of interest. Under stronger assumptions on the link function w , it is possible to more accurately estimate the graphon h_n , with resulting improvements to the downstream bootstrap procedure. This motivates our second estimator, which is exactly the restricted least squares histogram estimator set forth in Klopp, Tsybakov and Verzelen [20].

Definition 2 (Histogram). *Fix an integer $r > 1$ which corresponds to the number of bins in the histogram, and a number $s \in (0, 1]$ which corresponds to the maximum value the histogram estimate can take. Define the set $\mathcal{Z}_{n,r}$ to consist of all functions which assign each of the n nodes to one of the r classes. Then we set the histogram estimate \hat{h}_{hist} of h to be the least-squares estimate over functions which are piecewise-constant on partitions over the unit square, and which are bounded above by s . That is, for $Q = (Q_{ab}) \in \mathbb{R}^{r \times r}$, $\|Q\|_\infty \leq s$ and $z \in \mathcal{Z}_{n,r}$, we set*

$$L(Q, z) = \sum_{a,b \in [r]} \sum_{(i,j) \in z^{-1}(a) \times z^{-1}(b)} (A_{ij} - Q_{ab})^2 \tag{7}$$

$$(\hat{Q}, \hat{z}) = \underset{Q,z}{\operatorname{argmin}} L(Q, z) \tag{8}$$

$$\hat{\theta}_{ij} = \hat{Q}_{\hat{z}(i)\hat{z}(j)} \tag{9}$$

⁸Owen and Eckles [32], drawing on McCullagh [29], consider a similar bootstrap for estimating the variance in the mean of a multi-dimensional array of real-valued random variables.

$$\hat{h}_{hist}(u, v) = \theta_{\lceil nu \rceil \lceil nv \rceil} \quad (10)$$

We refer to the process of resampling using the histogram estimate \hat{h}_{hist} as the histogram bootstrap.

The histogram estimator of the graphon is a specific case of the stochastic block model, which itself dates back at least to [12, 17, 11]. In such models, every node is independently and randomly assigned to one of r latent classes or “blocks”, and edges form independently between nodes, with probabilities depending only on the nodes’ block assignments. The histogram estimator used here was introduced by Klopp, Tsybakov and Verzelen [20], though see also [36, 13, 14]. Klopp, Tsybakov and Verzelen [20] derive upper bounds on the mean-squared error of this particular histogram estimator which hold for all $Cn^{-1} \leq \rho_n \leq 1$. In our analysis of the histogram bootstrap, we will make use of these estimates on mean-squared error; interestingly, we will also make use of the fact that the histogram estimate \hat{h} is itself bounded, by construction.

Related work We have already mentioned a few suggested schemes for quantifying uncertainty of network statistics. Of these, the closest to our own approach is that of Bhattacharyya and Bickel [2], who consider a pair of subsampling schemes which they show lead to consistent distributional estimates for the same types of statistics (motif densities) that we analyze. The empirical graphon bootstrap can be viewed as analogous to their uniform subsampling procedure, but with nodes sampled with replacement rather than without replacement; as mentioned previously, it is known in the i.i.d setting that resampling can lead to much more accurate estimates at small sample sizes than subsampling (see e.g. [4]). Our theoretical results also cover general motifs, as opposed to Bhattacharyya and Bickel [2] who study only acyclic motifs and rings.

There also exists some other related work, which appeared after an initial version of this manuscript was made available as a preprint. Levin and Levina [26] study a model-based bootstrap under the assumption that the observed network is a random dot product graph. In a different direction, Lunde and Sarkar [28] establish the consistency of a subsampling approach for more general classes of network statistics. Finally, Zhang and Xia [38] consider a studentized version of our empirical graphon bootstrap, and use the Edgeworth expansion to derive rates of convergence and prove higher-order correctness.

3. Main results

Our main pair of results establish that if one samples G_m^* using either the empirical graphon or histogram bootstraps, then the conditional distribution (after the right scaling and centering) of $P_R(G_m^*)$ converges in probability to the distribution of $P_R(G_n)$, under some assumptions about the sparsity of the graphon, the structure of the motif R , and — in the case of estimation using a histogram — the smoothness of the graphon. For notational convenience, set $\bar{\rho}_n := P_{K_2}(\hat{h})$ to be the expected edge density of the resampled graph G_m^* .

Theorem 1. Let G_m^* be sampled from the empirical graphon bootstrap. Suppose that $w \neq 1$ on a set of strictly positive (Lebesgue) measure in $[0, 1]^2$. For any p -node motif R , if (i) $\int_{[0,1]^2} w^{4|E(R)|}(u_1, u_2) du_{1:2} < \infty$, (ii) either R is acyclic and $\rho_n = \omega(n^{-1})$ or R is general and $\rho_n = \omega(n^{-\frac{1}{2p}})$, and (iii) $m \rightarrow \infty$ and $m = \omega(\rho_n^{-4|E(R)|})$, then

$$\begin{aligned} \sup_x \left| \mathbb{P} \left(\frac{\sqrt{m}}{\hat{\rho}_n^{|E(R)|}} \left(P_R(G_m^*) - P_R(\hat{h}) \right) \leq x \middle| G_n \right) \right. \\ \left. - \mathbb{P} \left(\frac{\sqrt{n}}{\hat{\rho}_n^{|E(R)|}} \left(P_R(G_n) - P_R(h) \right) \leq x \right) \right| \xrightarrow{P} 0. \end{aligned} \tag{11}$$

In practice one would approximate the conditional distribution of $P_R(G_m^*)$ given G_n through Monte Carlo: that is, by repeatedly drawing samples of G_m^* according to the resampling procedure outlined in Section 2.1. Theorem 1 establishes that, asymptotically in n and the number of resamples B , such a procedure accurately approximates the distribution of $P_R(G_n)$.

We discuss the various conditions on w , ρ_n and m in detail in Sections 5 and 6, which is also where the proof of Theorem 1 and all our other theorems can be found. For now, we make only a few basic observations. Theorem 1 shows that the empirical graphon bootstrap is consistent under the relatively modest condition $\|w\|_{4|E(R)|} < \infty$. On the other hand, a more severe limitation of the empirical graphon bootstrap lies in the assumption for general motifs R that $\rho_n = \Omega(n^{-1/2p})$. In contrast, the scaled and centered motif density $\rho_n^{-|E(R)|} \sqrt{n} (P_R(G_n) - P_R(h))$ is known to converge to a non-degenerate Gaussian limit as long as $\rho_n = \Omega(n^{-2/p})$, which allows for much sparser sequences of graphs.

To obtain a consistent distributional estimate for these sparser sequences, we turn to a different graphon estimate. Under appropriate smoothness assumptions on w , estimators such as \hat{h}_{hist} converge in a strong sense (e.g. in an $L^p([0, 1]^2)$ norm) to the graphon function h_n ; by contrast, \hat{h}_{adj} converges to h_n only in a weaker topology. By leveraging this stronger notion of convergence, we can prove that the subsequent bootstrap procedure is consistent under weaker minimal conditions on the sparsity ρ_n .

Before focusing on the histogram estimator \hat{h}_{hist} , we start by considering an arbitrary graphon estimate \hat{h} : this could be formed by binning, smoothing, low-rank reconstruction, etc. In Theorem 2, we show that so long as \hat{h} satisfies a pair of general conditions, the distribution of the resampled graph G_n^* will converge to that of the original graph G_n .

Theorem 2. Let \hat{h} be a graphon estimate, and let G_n^* be a graph resampled from \hat{h} as described in Section 2.1. Suppose $w \neq 1$ on a set of positive Lebesgue measure on $[0, 1]^2$. For a p -node motif R , if (i) $\int_{[0,1]^2} w^{2|E(R)|}(u, v) du dv < \infty$, (ii) either R is acyclic and $\rho_n = \omega(n^{-1})$ or R is general and $\rho_n = \omega(n^{-2/p})$,

and (iii) the graphon estimate \hat{h} satisfies

$$\|\hat{h}\|_{3|E(R)|} = O_p(\rho_n), \quad \text{and} \quad \|\hat{h} - h_n\|_{2|E(R)|} = o_p(\rho_n), \quad (12)$$

then

$$\begin{aligned} \sup_x \left| \mathbb{P} \left(\frac{\sqrt{n}}{\hat{\rho}_n^{|E(R)|}} \left(P_R(G_n^*) - P_R(\hat{h}) \right) \leq x \middle| G_n \right) \right. \\ \left. - \mathbb{P} \left(\frac{\sqrt{n}}{\hat{\rho}_n^{|E(R)|}} \left(P_R(G_n) - P_R(h) \right) \leq x \right) \right| \xrightarrow{p} 0. \end{aligned} \quad (13)$$

Of course, for Theorem 2 to be practically useful, we need an estimator \hat{h} which actually achieves the notions of boundedness and convergence that are assumed in (12). Various difficulties arise in finding such an estimator:

- **Smoothness.** Although in Theorem 2, the only explicit assumption made on the function w is that $w \in L^2|E(R)|([0, 1]^2)$, this is insufficient to guarantee the convergence of $\hat{h} - h_n$ in the sense of (12). In order to obtain such a result, some additional structure must be placed on w . Typical assumptions are that w is piecewise constant (which corresponds to the stochastic block model) or that w is α -Hölder for some $\alpha \in (0, \infty]$. For concreteness, we will stick with the assumption that w is L -Lipschitz, meaning

$$|w(u, v) - w(x, y)| \leq L(|u - x| + |v - y|), \quad \text{for all } (u, v) \text{ and } (x, y) \in [0, 1]^2.$$
- **Sparsity.** The sparsity parameter ρ_n plays various roles in the hardness of graphon estimation. Intuitively, as the graph becomes sparser the whole function h_n gets closer to 0, and is thus easier to estimate; on the other hand, the available data suffers from a worse signal-to-noise ratio. Moreover, we note that as ρ_n decreases the conditions in (12) become stronger; thus it is important to have an estimator which has a fast rate of convergence in the very-sparse regime, where say $\rho_n = n^{-1} \log \log n$.
- **Correct norm.** Assuming appropriate conditions on smoothness and sparsity, various works have considered the problem of graphon estimation. Typically these works study either the squared-error loss $\|\hat{h} - h\|_2^2$ or an in-sample analogue. However, for any motif R except for the edge $R = K_2$, the norm $\|\cdot\|_{2|E(R)|}$ is a stricter norm than the $\|\cdot\|_2$, in the sense that $\|\cdot\|_2 \leq \|\cdot\|_{2|E(R)|}$. So it is not the case that every graphon estimate \hat{h} which accurately approximates h in $L^2([0, 1]^2)$ norm necessarily yields a consistent bootstrapping procedure.

We consider the particular restricted histogram estimator of [20] precisely because, assuming the link function w is Lipschitz, the estimator converges in a sufficiently strong norm at a sufficiently fast rate.

Proposition 1 (Corollary 3.6 of [20]). *Suppose w is L -Lipschitz, $\rho_n = \omega(n^{-1})$, and the histogram estimate \hat{h}_{hist} is computed with $r = \sqrt{n\rho_n}$ and $s = \rho_n$. Then*

$$\|\hat{h}_{\text{hist}}\|_{\infty} \leq \rho_n, \quad \text{and} \quad \|\hat{h}_{\text{hist}} - h_n\|_2 = O_p\left(\sqrt{\rho_n n^{-1} \log(n\rho_n)}\right) = o_p(\rho_n). \quad (14)$$

Using Hölder’s inequality, we see that the guarantees in (14) imply the conditions in (12). Therefore we may apply Theorem 2, and conclude that the histogram bootstrap yields a consistent distributional estimate.

Corollary 1. *Let G_n^* be sampled from the histogram bootstrap, where the histogram estimate \hat{h}_{hist} is computed with $r = \sqrt{n\rho_n}$ and $s = \rho_n$. Suppose w is L -Lipschitz, and $w \neq 1$ on a set of positive Lebesgue measure on $[0, 1]^2$. For a p -node motif R , if either R is acyclic and $\rho_n = \omega(n^{-1})$ or R is general and $\rho_n = \omega(n^{-2/p})$, then*

$$\begin{aligned} \sup_x \left| \mathbb{P} \left(\frac{\sqrt{n}}{\hat{\rho}_n^{|E(R)|}} (P_R(G_n^*) - P_R(\hat{h})) \leq x \middle| G_n \right) \right. \\ \left. - \mathbb{P} \left(\frac{\sqrt{n}}{\hat{\rho}_n^{|E(R)|}} (P_R(G_n) - P_R(h)) \leq x \right) \right| \xrightarrow{p} 0. \end{aligned} \tag{15}$$

Corollary 1 shows that when the link function w is Lipschitz, the histogram bootstrap is consistent for all ranges of sparsity in which the limiting distribution of $\rho_n^{-|E(R)|} \sqrt{n}(P_R(G_n) - P_R(h))$ is known (see [3]). As promised, when R is a general motif, the sparsity assumptions needed for Theorem 2 and Corollary 1 to hold can be much weaker than those required for Theorem 1. For example, suppose the motif of interest R is a triangle, i.e. $R = \{\{1, 2, 3\}, \{(1, 2), (2, 3), (1, 3)\}\}$ is the complete graph on 3 nodes. In this case, for the guarantees of Theorem 1 to hold the sparsity must satisfy $\rho_n = \omega(n^{-1/6})$. On the other hand, assuming w is Lipschitz the conclusions of Corollary 1 hold whenever the sparsity is at least $\omega(n^{-2/3})$, which is much smaller.

There are certainly drawbacks to the histogram bootstrap. First of all, as already mentioned, to obtain stronger results in terms of the sparsity parameter ρ_n , we are forced to make stronger assumptions on the link function w . Additionally, we note that the estimator \hat{h}_{hist} explicitly enforces sparsity in the estimate through the tuning parameter s , which in Theorem 2 is set based on the (typically unknown) sparsity ρ_n . [20] suggest a data-dependent way for choosing s , which will not affect the results of (14) nor the consistency of the resulting bootstrap procedure. Finally, although the estimator of [20] has strong theoretical guarantees, it involves solving a combinatorial optimization problem and is not computationally feasible. We now turn to a discussion of this and other computational issues.

Computational considerations Each of our proposed bootstraps are computationally intensive procedures. Naively computing $P_R(G_m^*)$ even once requires checking $O(m^p)$ separate subgraphs, which may be infeasible except when the sample size n is modest, and the number of vertices in the motif R is small. Repeating this computation for each bootstrap sample only exacerbates these issues. The histogram bootstrap poses the added challenge that one need compute the estimate \hat{h}_{hist} , which as already mentioned poses a serious computational challenge.

Still, the situation is not as bleak as these considerations suggest, for several reasons. First, for certain graphs, such as stars and triangles, there exist fast (indeed, sublinear time) algorithms for approximating the subgraph count [9, 16]. Second, the motifs whose densities are of *scientific* interest in applied problems tend to be small. Third, any graphon estimate \hat{h} which satisfies (14) will meet the condition (12), and will yield a consistent bootstrap. For instance, the spectral method studied by [37] is a computationally reasonable alternative to the histogram estimate which—after appropriate truncation of the estimate \hat{h} to satisfy $\|\hat{h}\|_\infty \leq \rho_n$, and extension of the estimate to $[0, 1]^2$ —should satisfy (14).

Practical implications To summarize, the theory developed in this section shows that our proposed empirical graphon and histogram bootstraps result in consistent estimates for the distributions of motif densities. Motif densities are statistics of direct scientific interest, for instance in biological [30] and social [1] networks. Additionally, our conclusions apply to sparse graphons, which are one way to model the sparsity often seen in real-world networks.

Finally, the fact that our bootstraps are valid for motif densities suggests (via the delta method) that they should also be valid for functionals which can be expressed as well-behaved functions of motif densities. Precisely because motif densities characterize infinite-exchangeable network distributions, this class of functionals should be quite rich, and include more complicated and scientifically interesting functionals than just motif densities. We thus view our results as evidence that the bootstrap is a statistically reasonable off-the-shelf method of generally quantifying uncertainty in network data analysis.

4. Simulations

Our theory kicks in only as the number of vertices $n \rightarrow \infty$. Of course we would like to be confident that our bootstraps work when n is finite, and even when n is reasonably small. In this section, we provide empirical evidence supporting this conclusion. We consider various graphons h_n and motifs R , and show that confidence intervals formed using either of our proposed bootstraps typically have close to nominal coverage when n is even moderately large. Moreover, the interval widths shrink quickly and at comparable rates for each bootstrap procedure.

Setup We will consider three separate simulation setups. In each, the graph G_n will be sampled as in (1), and the setups will be distinguished only by different choices of the link function w . Within each setup, we will vary $n = 25, 50, \dots, 400$ and $\rho_n = .02, .1$ and $.25$, to investigate the effects of both sample size and sparsity. We will always draw 100 resampled graphs $(G_n^*)^{(1)}, \dots, (G_n^*)^{(100)}$ with resulting motif densities $P_R((G_n^*)^{(1)}), \dots, P_R((G_n^*)^{(100)})$, and take the empirical distribution

$$\frac{1}{100} \sum_{b=1}^{100} \frac{\sqrt{n}}{\bar{\rho}_n^{|E(R)|}} \left(P_R((G_n^*)^{(b)}) - P_R(\hat{h}) \right)$$

as a Monte Carlo estimate of the (centered and scaled) conditional distribution of $P_R(G_n^*)$. Letting $1 - \alpha$ be the nominal coverage, we take $\hat{q}_{\alpha/2}$ and $\hat{q}_{1-\alpha/2}$ to be the $\alpha/2$ and $1 - \alpha/2$ quantiles of this Monte Carlo estimate. Our final confidence interval is then given by

$$\hat{I}_R(G_n) = \left(P_R(G_n) + \hat{q}_{\alpha/2} \frac{\hat{\rho}_n^{|E(R)|}}{\sqrt{n}}, P_R(G_n) + \hat{q}_{1-\alpha/2} \frac{\hat{\rho}_n^{|E(R)|}}{\sqrt{n}} \right)$$

Our theory establishes that, under appropriate conditions on the link function w , the true coverage of $\hat{I}_R(G_n)$ should approach $1 - \alpha$ as $n \rightarrow \infty$ (up to error induced by the Monte Carlo approximation). To examine finite sample behavior, we repeat the above setup over $N = 1000$ draws of the graph G_n — and resulting intervals $\hat{I}_R(G_n)$ — to obtain an estimated coverage and average interval width for each value of n and ρ_n , and for each of our two bootstraps. It is these quantities which we examine below.

For the empirical graphon bootstrap, we implemented our own code. To compute a histogram estimator, we use the `blockmodels` package Leger [24], which approximates the least squares estimator using a variational EM procedure, and chooses the optimal number of blocks by cross-validation. Note that this deviates slightly from our theoretical definition of the histogram estimator, but the deviation is in the direction of being more computationally tractable, i.e., not having to solve a difficult combinatorial optimization problem.

Simulation 1: Gaussian link function In our first simulation, we let

$$w(u, v) \propto \exp(-25\|u - v\|^2/2)$$

be a Gaussian link function with bandwidth $1/5$ (the constant of proportionality is chosen so that $\int w(u, v) du dv = 1$). In Figure 1, we see coverage plotted against sample size, for the motifs $R = \{(1, 2), (2, 3), (1, 3)\}$ (triangle), $R = \{(1, 2), (2, 3)\}$ (two-star), $R = \{(1, 2), (2, 3), (3, 4), (4, 1)\}$ (four-cycle), and $R = \{(1, 2), (1, 3), (1, 4)\}$ (claw). The lines are colored according to the resampling procedure used, and marked according to whether $\rho_n = .02, .1$ or $.25$.

When $\rho_n = .25$, for all of the aforementioned motifs except the four-cycle and for all $n > 300$, the estimated coverage is within $.05$ of nominal coverage for both bootstraps; it is therefore fair to say that both of our bootstraps “work” reasonably well in this setting. When $\rho_n = .1$, the story is more mixed. Resampling using the histogram estimate still gives approximately nominal coverage, but estimated coverage of the empirical bootstrap deviates substantially from nominal coverage. That being, even for the empirical bootstrap estimated coverage appears to be tending towards $1 - \alpha$ as n increases, as we would expect. When $\rho_n = .02$, even the histogram estimator is not particularly accurate for the four-cycle. All of this is in line with our theory, which requires suitable lower bounds on ρ_n as a function of n for both bootstraps, with the lower bound being more stringent for the empirical graphon than for the histogram estimator, and for motifs with more nodes and edges such as the four-cycle.

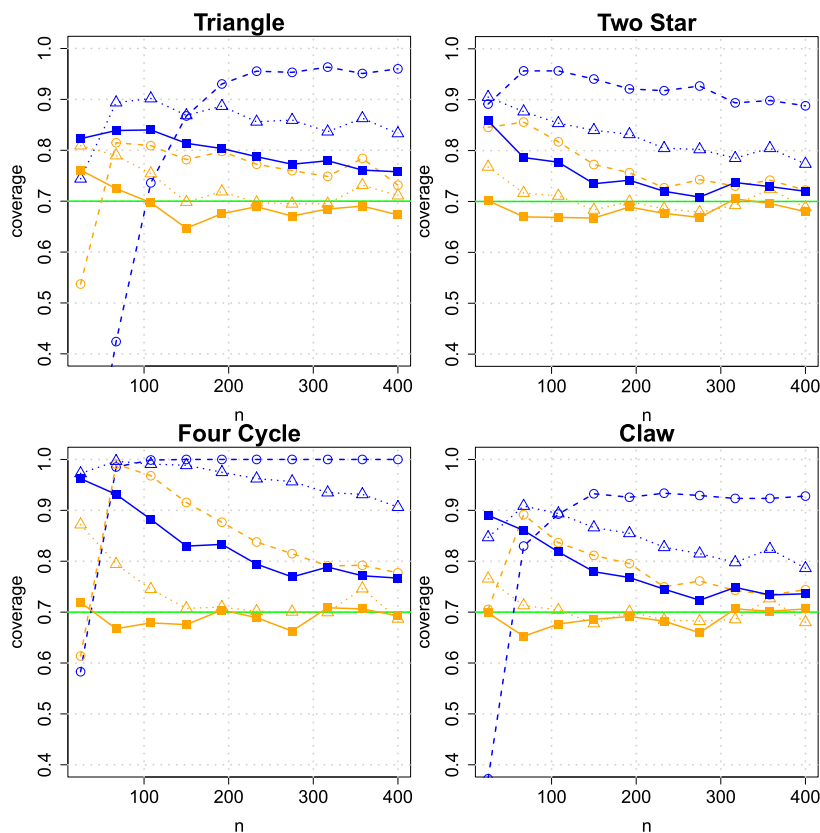


FIG 1. Coverage of bootstrapping methods as a function of sample size for Gaussian link function. Nominal level $1 - \alpha = .7$ denoted by horizontal green line. The empirical graphon procedure is shown in blue, and the histogram graphon procedure in orange. Dashed lines and circular symbols correspond to $\rho_n = .02$; dotted lines and triangles to $\rho_n = .1$; filled lines and squares to $\rho_n = .25$.

Figure 2 shows the average interval widths for each of the four aforementioned motifs when $\rho_n = .2$, plotted on a log-log scale. We see that, in addition to having better coverage, the histogram bootstrap has narrower confidence intervals than the empirical graphon bootstrap. However, as n grows the difference between the two decreases, and the slopes become more and more similar, converging towards the $1/\sqrt{n}$ rate we would expect from Theorems 1 and 2 (since $\rho_n = .25$ does not change with n .)

Simulation 2: stochastic block model In our second simulation, we take

$$w(u, v) \propto \begin{cases} 1, & \text{if } u \leq .5, v \leq .5 \\ .8, & \text{if } u > .5, v > .5 \\ .25, & \text{otherwise;} \end{cases}$$

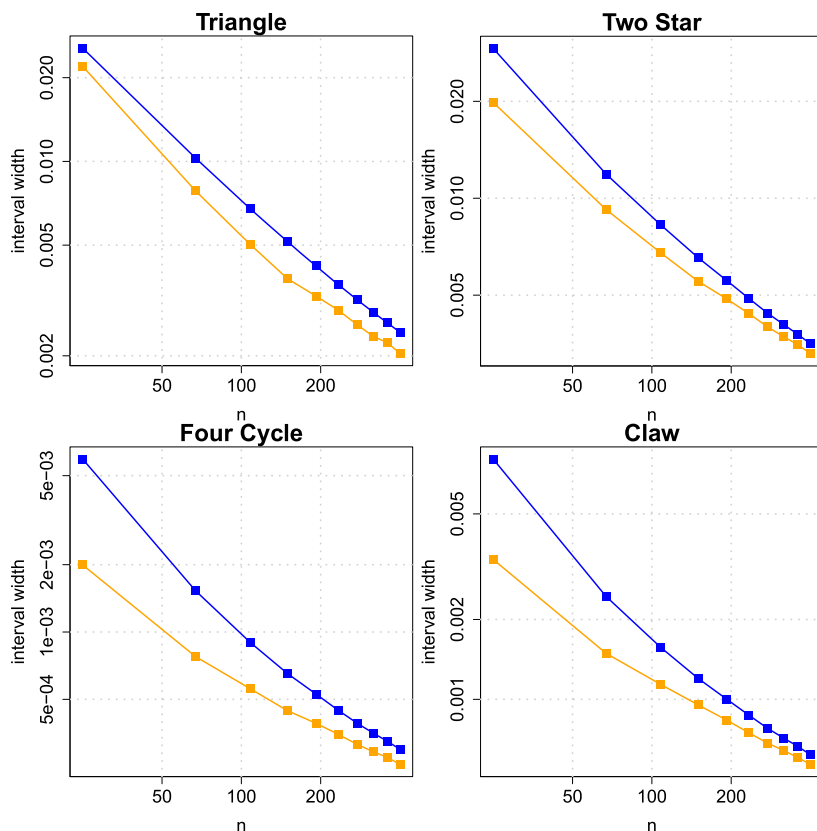


FIG 2. Interval width of bootstrapping methods as a function of sample size for Gaussian link function. Colors and symbols are the same as in Figure 1.

in other words we sample G_n from a stochastic block model with two blocks.

Figure 3 shows estimated coverage as a function of sample size. The conclusions we draw are similar in spirit to those drawn from our first experiment. Both bootstrap methods appear to be approaching $1 - \alpha$ coverage as n increases, when the density parameter ρ_n is sufficiently large (although we note that for the empirical graphon bootstrap, the difference between nominal and actual coverage is larger than it was for the Gaussian link function). The histogram bootstrap has close to nominal coverage across all values of ρ_n , but after all this is a case where the true model is itself a histogram, and we should not be surprised that the histogram bootstrap performs quite well.

Figure 4 shows the average confidence interval width as a function of n when $\rho_n = .25$. Again the conclusions are similar to those of our first experiment. In particular, while resampling using a histogram graphon estimate results in narrower average confidence intervals than resampling using the empirical graphon, the difference between the two decreases as n grows, and the rate at which in-

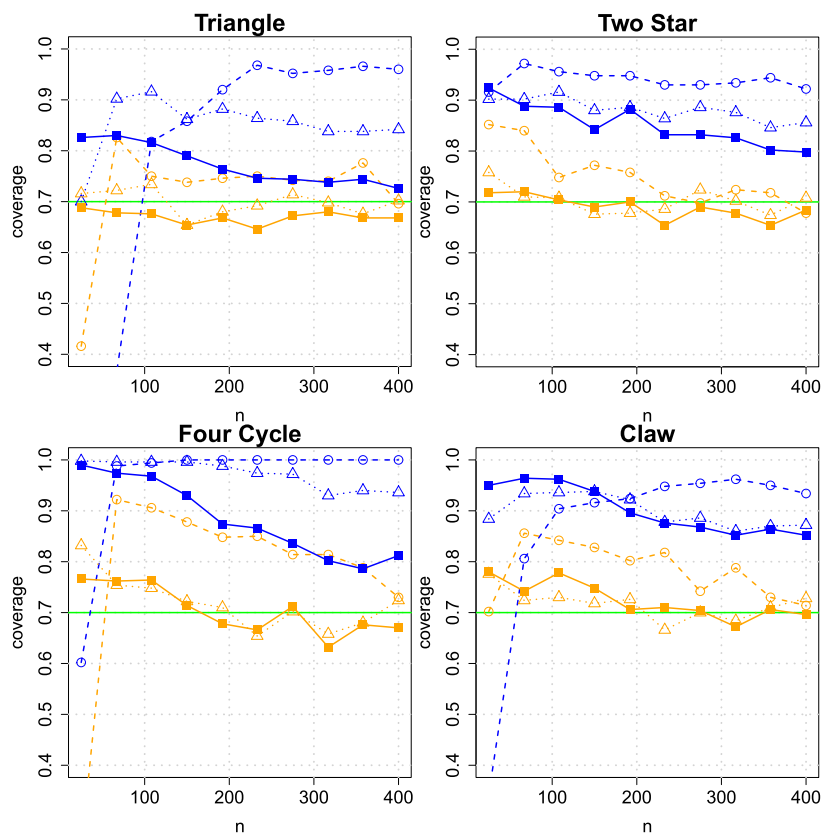


FIG 3. Coverage of bootstrapping methods as a function of sample size for stochastic block model. Colors and symbols are the same as in Figure 1.

terval width shrinks tends towards the expected $1/\sqrt{n}$ rate.

Simulation 3: horseshoe link function In our third and final simulation, we let

$$w(u, v) \propto \exp(-200(u - v^2)^2/2) + \exp(-200(v - u^2)^2/2).$$

We call this the horseshoe link function, since when w is visualized using a heat map it looks like a horseshoe.⁹ Due to its odd shape, w is poorly approximated by a histogram of any reasonable binwidth. Relatedly, the Lipschitz condition required for Corollary 1 is satisfied only for a very large Lipschitz constant L . On the other hand, our theory for the empirical graphon does not depend on the smoothness of the function w , and so it applies equally as well to the horseshoe link function as to the previous — smoother — graphons considered in our first two simulations.

⁹We have borrowed this graphon from Wang [35], where it was used as a challenging example for graphon estimation.

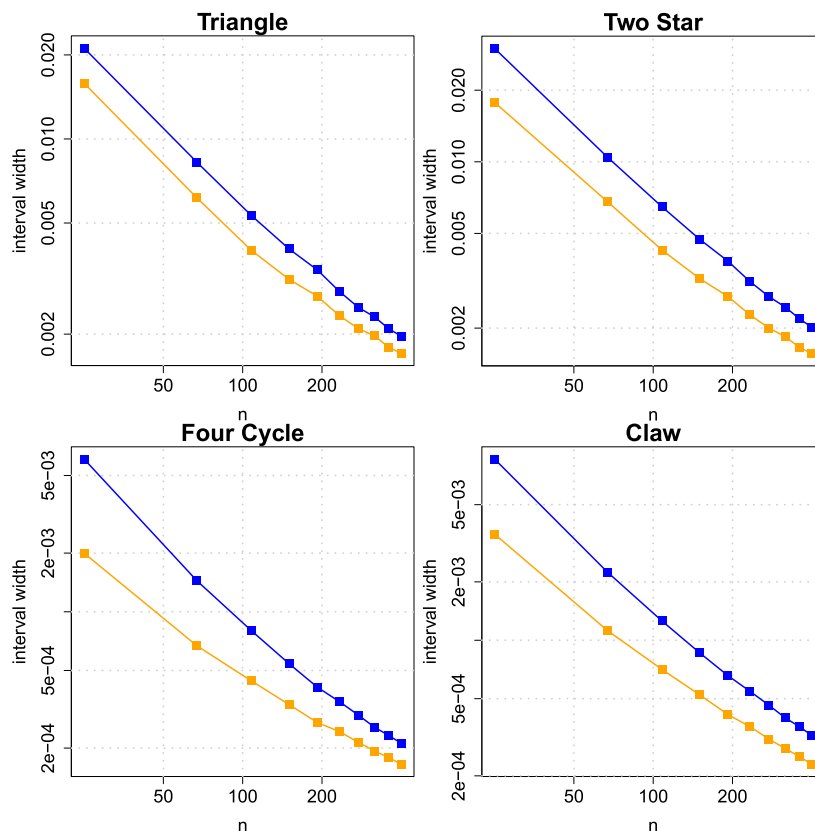


FIG 4. Interval width of bootstrapping methods as a function of sample size for stochastic block model. Colors and symbols are the same as in Figure 2.

Such is the theoretical state of affairs; Figure 5 shows the empirical reality. In contrast to the previous two simulations, when $\rho_n = .25$ it is now the empirical graphon bootstrap which has closer to nominal coverage, consistently across the different possibilities for R and n . (When $\rho_n = .1$ or $\rho_n = .02$, resampling from a histogram estimate still results in coverage closer to $1 - \alpha$). We conclude that for sufficiently non-smooth link functions w and large values of ρ_n , the empirical graphon may outperform the histogram estimate as a bootstrap procedure. Figure 6 shows that average interval widths are similar for both resampling procedures, reiterating the takeaway message of Simulations 1 and 2.

Taken as a whole, our simulations demonstrate that even for networks with a moderate number of vertices, both of our bootstraps result in reasonable estimates of the motif density distribution, across different choices of graphon function and motif. In particular, confidence intervals formed using these bootstraps have approximately nominal coverage when ρ_n is sufficiently large, as suggested by our asymptotic theory, and interval widths shrink rapidly (approaching the

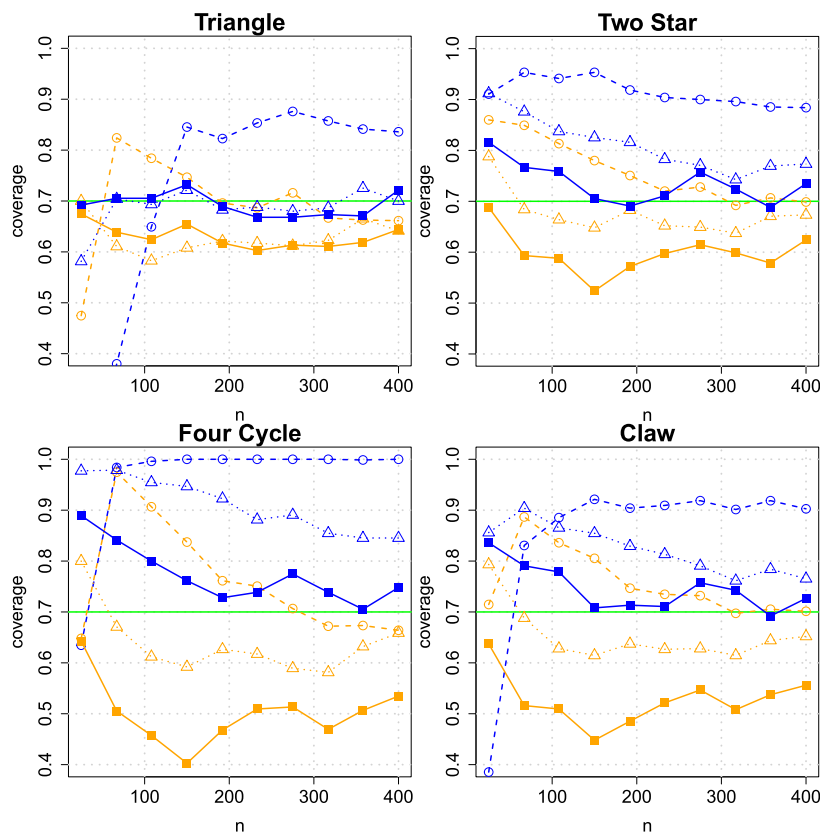


FIG 5. Coverage of bootstrapping methods as a function of sample size for horseshoe graphon. Nominal level $1 - \alpha = .7$ denoted by green line. Colors and symbols are the same as in Figure 1.

anticipated $1/\sqrt{n}$ rate for fixed ρ_n) as n grows.

5. Supporting propositions and lemmas

As mentioned earlier, the proofs for Theorems 1 and 2 each rely on a series of 3 approximations: first, that the distribution of a scaled version of $P_R(G_n)$ is close to Gaussian; second, that the conditional distribution of a scaled version of $P_R(G_m^*)$ is close to Gaussian; and third, that the variances of both distributions are close. Propositions 2 and 3 formalize the first and third of these assertions (the second assertion is covered in the proof of the theorems.)

In order to state Proposition 2, we will need an expression for the (normalized) variance of $P_R(G_n)$. Intuitively, this will involve taking expectation over the product of indicator functions, of the form $\mathbb{E}[\mathbb{I}\{G_n(\mathbf{i}) \simeq R\}\mathbb{I}\{G_n(\mathbf{j}) \simeq R\}]$. As we might expect, this quantity can be related to $P_W(h)$ for a set of motifs

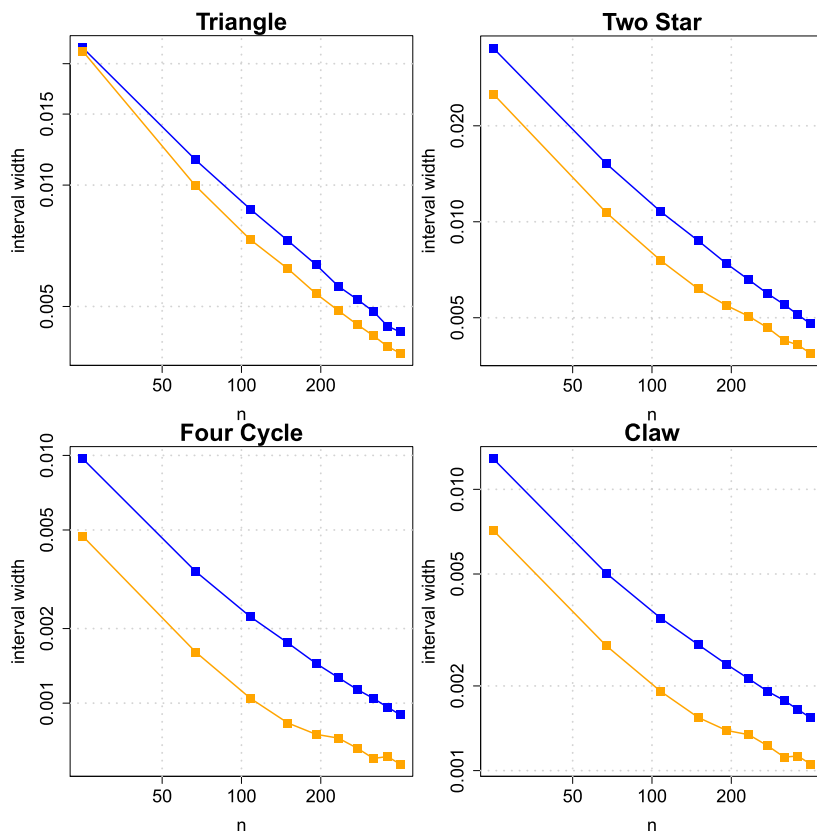


FIG 6. Interval width of bootstrapping methods as a function of sample size for horseshoe graphon. Colors and symbols are the same as in Figure 2.

W . The intuition is that the event of seeing the same motif on two different subgraphs corresponds to seeing one of several particular motifs on the union subgraph. We will call the set of these motifs W the “merged copy set”, because it consists of motifs formed by taking two copies of R and merging some of their vertices.

Definition 3. The *merged copy set* of a motif R on k vertices, denoted $MC(R, k)$, is the set of all W such that

$$MC(R, k) = \{W : |V(W)| = k, \exists \mathbf{i}, \mathbf{j} \text{ s.t. } \mathbf{i} \cup \mathbf{j} = 1 : k, W(\mathbf{i}) \simeq R, W(\mathbf{j}) \simeq R\} \tag{16}$$

Lemma 1, next, relates the double sum present in $\mathbb{E}[P_R(G_n)^2]$ to summing over motifs in the various merged copy sets. The merged copy sets tells us which motifs we need to sum over, but not how many times we’ll need to consider each motif. Luckily, this depends only on $\binom{n}{k}$ and a combinatorial factor, itself

a function of R and W , which counts how many ways W can be formed by merging two copies of R .¹⁰

Lemma 1. *For any p -node motif R ,*

$$\sum_{\mathbf{i}, \mathbf{j} \in S_p(n)} \mathbb{E} [\mathbb{I}\{G_n(\mathbf{i}) \simeq R\} \mathbb{I}\{G_n(\mathbf{j}) \simeq R\}] = \sum_{k=p}^{2p} \binom{n}{k} \sum_{W \in MC(R,k)} C_R(W) P_W(h), \tag{17}$$

where $C_R(W)$ counts the number of ways of forming W by merging two copies of R .

Lemma 2. *Define*

$$\sigma_R^2(h_n) \equiv \text{Var} \left(\frac{\sqrt{n}}{\rho_n^{|E(R)|}} P_R(G_n) \right). \tag{18}$$

Then,

$$\begin{aligned} \sigma_R^2(h_n) &= \frac{n}{(\rho_n^{|E(R)|} \binom{n}{p} |p!N(R)|)^2} \sum_{k=p}^{2p-1} \binom{n}{k} \sum_{S \in MC(R,k)} C_R(S) P_S(h_n) \\ &\quad - \left(1 - \frac{\binom{n-p}{p}}{\binom{n}{p}} \right) \tilde{P}_R(h_n)^2. \end{aligned}$$

Moreover, if $\int_{[0,1]^2} w^{2|E(R)|}(u_1, u_2) du_{1:2} < \infty$, and either R is acyclic and $\rho_n = \omega(n^{-1})$ or R is general and $\rho_n = \omega(n^{-2/p})$, then

$$\sigma_R^2 := \lim_{n \rightarrow \infty} \sigma_R^2(h_n) < \infty. \tag{19}$$

Now that we know that $\rho_n^{-|E(R)|}$ is the appropriate normalization, we can state a central limit theorem from [3] for the empirical subgraph densities.

Proposition 2 (Theorem 1 of [3]). *Suppose $w \neq 1$ on a subset of $[0, 1]^2$ with positive Lebesgue measure. For a p -node motif R , if (i) $\int_{[0,1]^2} w^{2|E(R)|}(u_1, u_2) du_{1:2} < \infty$ and (ii) either R is acyclic and $\rho_n = \omega(n^{-1})$, or R is general and $\rho_n = \Omega(n^{-\frac{2}{p}})$ then*

$$\sqrt{n} \rho_n^{-|E(R)|} \left(P_R(G_n) - P_R(h) \right) \xrightarrow{w} \mathcal{N}(0, \sigma_R^2) \tag{20}$$

Also

$$\tilde{P}_R(G_n) \xrightarrow{p} \tilde{P}_R(h) \tag{21}$$

¹⁰For example, if $R = K_2$, the single edge between two nodes, there are four different ways we can merge two copies of R to get a 2-star.

We reiterate that the requirement $\rho_n = \omega(n^{-2/p})$ in Theorem 2 essentially matches that of Proposition 2, and that when $\rho_n = o(n^{-2/p})$ the limiting distribution of $\sqrt{n}\rho_n^{-|E(R)|}(P_R(G_n) - P_R(h))$ is not known.

As stated above, one step required for the proofs of our main theorems is to show that the bootstrap estimate of variance, for both of our procedures, is close to σ_R^2 . Proposition 3 formalizes this statement.

Proposition 3. *For a motif R , under the conditions of Theorem 1, the empirical graphon $\hat{h} = \hat{h}_{adj}$ satisfies*

$$\sigma_R^2(\hat{h}) \xrightarrow{p} \sigma_R^2, \tag{22}$$

and under the conditions of Theorem 2, for any graphon estimate \hat{h} which satisfies (12),

$$\sigma_R^2(\hat{h}) \xrightarrow{p} \sigma_R^2. \tag{23}$$

Finally, in order to show that Proposition 3 holds, we must have that the $P_S(\hat{h})$ terms in $\sigma_R^2(\hat{h})$ converge at the appropriate rate to the $P_S(h)$ terms in σ_R^2 . Lemmas 3 and 4 establish this, respectively, for the empirical graphon and for an arbitrary estimator that satisfies (12).

Lemma 3. *For any k -node motif S , if $\int_{[0,1]^2} w^{2|E(S)|}(u_1, u_2) du_{1:2} < \infty$, and $\rho_n = \omega(n^{-1})$, then for $\hat{h} = \hat{h}_{adj}$*

$$\left| \frac{P_S(\hat{h}) - P_S(h_n)}{\rho_n^{k-1}} \right| \xrightarrow{p} 0. \tag{24}$$

If additionally $\rho_n = \omega(n^{-\frac{1}{k}})$, then

$$\left| \frac{P_S(\hat{h}) - P_S(h_n)}{\rho_n^{|E(S)|}} \right| \xrightarrow{p} 0. \tag{25}$$

In order to prove (22) in Proposition 3, we will need to invoke Lemma 3 with respect to motifs S in the merged copy set $MC(R, k)$, for $k = p, \dots, 2p - 1$. For such motifs S , the number of edges $|E(S)|$ may be as large as $2|E(R)|$, and the number of nodes as large as $2p - 1$. This explains why in Theorem 1, we require that $w \in L^{4|E(R)|}([0, 1]^2)$ and $\rho_n = \omega(n^{-1/2p})$. On the other hand, the requirements of Lemma 4 are weaker, and translate to weaker requirements in Theorem 2.

Lemma 4. *For any motif S , if $\|w\|_{|E(S)|} < \infty$ and \hat{h} is a graphon estimate which satisfies*

$$\rho_n^{-1} \|\hat{h} - h_n\|_{|E(S)|} \xrightarrow{p} 0, \tag{26}$$

then

$$\left| \frac{P_S(\hat{h}) - P_S(h_n)}{\rho_n^{|E(S)|}} \right| \xrightarrow{p} 0. \tag{27}$$

6. Proofs

6.1. Proof of Theorem 1

We can upper bound (11) by

$$\begin{aligned} & \sup_x \left| \mathbb{P} \left(\frac{\sqrt{m}}{\hat{\rho}_n^{|E(R)|}} \left(P_R(G_m^*) - P_R(\hat{h}) \right) \leq x \middle| G_n \right) - \Phi \left(\frac{x}{\sigma_R(\hat{h})} \right) \right| \\ & + \sup_x \left| \Phi \left(\frac{x}{\sigma_R(\hat{h})} \right) - \Phi \left(\frac{x}{\sigma_R(h)} \right) \right| \\ & + \sup_x \left| \Phi \left(\frac{x}{\sigma_R(h)} \right) - \mathbb{P} \left(\frac{\sqrt{n}}{\hat{\rho}_n^{|E(R)|}} \left(P_R(G_n) - P_R(h) \right) \leq x \right) \right|. \end{aligned}$$

The third term goes to 0 by Proposition 2. The second term goes to 0 by Proposition 3. All that remains is to bound the first term. Note that to do so, we cannot simply invoke Proposition 2, because the empirical graphon \hat{h} is random and changing with n . Instead, we bound the first term using a Berry-Esseen inequality for U-statistics, keeping in mind that $P_R(G_m^*)$ is strictly a U-statistic, conditional on G_n , because once the ϵ_i^* are specified there is no more randomness in $P_R(G_m^*)$. To ease notation, for $\mathbf{i} \in S_n(p)$ define

$$\epsilon_{\mathbf{i}}^* = (\epsilon_{i_1}^*, \dots, \epsilon_{i_p}^*) \tag{28}$$

$$\hat{H}_R(\epsilon_{\mathbf{i}}^*) := \mathbb{E} [\mathbb{I}(G_n^*(\mathbf{i}) = R) | \epsilon^*, G_n] \tag{29}$$

and note that

$$P_R(G_m^*) = \frac{1}{N(R)p! \binom{n}{p}} \sum_{\mathbf{i} \in S_n(p)} \sum_{R_1 \sim R} \hat{H}_R(\epsilon_{\mathbf{i}}^*) \tag{30}$$

Therefore, by Janssen [18], we have that

$$\sup_x \left| \mathbb{P} \left(\frac{\sqrt{m}}{\hat{\rho}_n^{|E(R)|}} \left(P_R(G_m^*) - P_R(\hat{h}) \right) \leq x \middle| G_n \right) - \Phi \left(\frac{x}{\sigma_R(\hat{h})} \right) \right| \leq C \frac{\bar{\nu}_3}{\bar{\sigma}_g^3 m^{\frac{1}{2}}} \tag{31}$$

where

$$\bar{\sigma}_g^2 := \text{Var} \left(\mathbb{E} \left[\sum_{R_1 \sim R} \frac{\hat{H}_{R_1}(\epsilon_{1:p}^*)}{p! N(R) \hat{\rho}_n^{|E(R)|}} \middle| \epsilon_{\mathbf{i}}^*, G_n \right] \middle| G_n \right) \tag{32}$$

$$\bar{\nu}_3 := \mathbb{E} \left[\left| \sum_{R_1 \sim R} \frac{\hat{H}_{R_1}(\epsilon_{1:p}^*) - P_R(\hat{h})}{p! N(R) \hat{\rho}_n^{|E(R)|}} \right|^3 \middle| G_n \right] \tag{33}$$

We first bound $\bar{\nu}_3$, using the fact that $\hat{H}_R(\epsilon_{1:p}^*)^3 = \hat{H}_R(\epsilon_{1:p}^*)$. By Holder's Inequality, we have

$$\mathbb{E} \left[\hat{H}_{R_1}(\epsilon_{1:p}^*) \hat{H}_{R_2}(\epsilon_{1:p}^*) \hat{H}_{R_3}(\epsilon_{1:p}^*) \right] \leq \mathbb{E} \left[\hat{H}_R(\epsilon_{1:p}^*)^3 \right] \tag{34}$$

By Lemma 3, we have that

$$\left| \frac{P_R(\hat{h}) - P_R(h_n)}{\rho_n^{|E(R)|}} \right| \xrightarrow{p} 0 \tag{35}$$

since either $|E(R)| = p - 1$ or $\rho_n = \omega(n^{-\frac{1}{p}})$. Therefore,

$$P_R(\hat{h}) = O_p(P_R(h_n)) = O_p\left(\rho_n^{|E(R)|}\right) \tag{36}$$

where the last statement is implied by the condition $\int_{[0,1]^2} w^{4|E(R)|}(u_1, u_2) du_{1:2} < \infty$. Putting these together, we can upper bound $\bar{\nu}_3$,

$$\bar{\nu}_3 \leq \frac{8}{\bar{\rho}_n^{3|E(R)|}} (\mathbb{E} [\hat{H}_R(\epsilon_{1:p}^*)^3 | G_n] + P_R(\hat{h})^3) \tag{37}$$

$$= \frac{8}{\bar{\rho}_n^{3|E(R)|}} (\mathbb{E} [\hat{H}_R(\epsilon_{1:p}^*) | G_n] + P_R(\hat{h})^3) \tag{38}$$

$$= \frac{8}{\bar{\rho}_n^{3|E(R)|}} (P_R(\hat{h}) + P_R(\hat{h})^3) \tag{39}$$

$$= O_p\left(\rho_n^{-2|E(R)|}\right). \tag{40}$$

Turning to $\bar{\sigma}_g^2$, we have

$$\bar{\sigma}_g^2 = \text{Var} \left(\mathbb{E} \left[\frac{\sum_{R_1 \sim R} \hat{H}_{R_1}(\epsilon_{1:p}^*)}{p! N(R) \bar{\rho}_n^{|E(R)|}} \middle| \epsilon_1^*, G_n \right] \middle| G_n \right) \tag{41}$$

$$= \frac{\sum_{S \in MC(R,1)} C_R(S) (P_S(\hat{h}) - P_R(\hat{h})^2)}{\bar{\rho}_n^{2|\bar{E}(R)|}}, \tag{42}$$

with the second equality following because

$$\mathbb{E} \left[\mathbb{E} [\hat{H}_{R_1}(\epsilon_{1:p}^*) | \epsilon_1^*] \mathbb{E} [\hat{H}_{R_2}(\epsilon_{1:p}^*) | \epsilon_1^*] \middle| G_n \right] = P_S(\hat{h})$$

for some S in $MC(R, 1)$. If R is acyclic, $2|E(R)| = 2p - 2 = |V(S)| - 1$. Otherwise, by assumption, $\rho_n = \omega(n^{-\frac{1}{2p}})$ and $|V(S)| = 2p - 1$. Either way, by Lemma 3, we have that for all $S \in MC(R, 1)$,

$$\left| \frac{P_S(\hat{h}) - P_R(\hat{h})^2 - P_S(h) + P_R(h)^2}{\bar{\rho}_n^{2|E(R)|}} \right| \xrightarrow{p} 0 \tag{43}$$

Note that $\frac{P_S(h) - P_R(h)^2}{\bar{\rho}_n^{2|E(R)|}} = \tilde{P}_S(h) - (\tilde{P}_R(h))^2 = \theta(1)$ by Holder's inequality, unless $w = 1$ almost everywhere. So, $\bar{\sigma}_g^2 = \Omega_p(1)$. Finally, if $m = \omega(\rho_n^{-4|E(R)|})$,

$$\bar{\nu}_3 \bar{\sigma}_g^{-3} m^{-\frac{1}{2}} = o_p(1) \tag{44}$$

which completes the proof.

6.2. Proof of Theorem 2

We begin, similarly to the Proof of Theorem 1, by upper bounding (13) via the triangle inequality by

$$\begin{aligned} & \sup_x \left| \mathbb{P}\left(\frac{\sqrt{n}}{\bar{\rho}_n^{|E(R)|}} \left(P_R(G_n^*) - P_R(\hat{h})\right) \leq x \middle| G_n\right) - \Phi\left(\frac{x}{\sigma_R^2(\hat{h})}\right) \right| \\ & + \sup_x \left| \Phi\left(\frac{x}{\sigma_R^2(\hat{h})}\right) - \Phi\left(\frac{x}{\sigma_R^2(h)}\right) \right| \end{aligned} \tag{45}$$

$$+ \sup_x \left| \Phi\left(\frac{x}{\sigma_R^2(h)}\right) - \mathbb{P}\left(\frac{\sqrt{n}}{\hat{\rho}_n^{|E(R)|}} (P_R(G_n) - P_R(h)) \leq x \right) \right|. \tag{46}$$

The third term goes to 0 by Proposition 2. The second term goes to 0 by Proposition 3. To bound the first term, we split $P_R(G_n^*) - P_R(\hat{h})$ up into two components, based on the randomness from resampling latent variables and the randomness from resampling edges respectively. In other words,

$$P_R(G_n^*) - P_R(\hat{h}) = P_R(G_n^*) - \mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n] + \mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n] - P_R(\hat{h}). \tag{47}$$

Lemma 5 establishes that $\mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n] - P_R(\hat{h})$ obeys, conditional on G_n , a central limit theorem for U-statistics. Lemma 6 establishes that $P_R(G_n^*) - \mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n]$, once appropriately scaled, has asymptotically negligible contribution to the overall randomness. Let

$$\tau_R^2(\hat{h}) := \text{Var}\left(\frac{\sqrt{n}}{\bar{\rho}_n^{|E(R)|}} \mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n] \middle| G_n\right) \tag{48}$$

Lemma 5. *Let R be a fixed motif. Then, if the conditions of Theorem 2 hold,*

$$\sup_x \left| \mathbb{P}\left(\frac{\sqrt{n}}{\bar{\rho}_n^{|E(R)|}} (\mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n] - P_R(\hat{h})) \leq x \middle| G_n\right) - \Phi\left(\frac{x}{\tau_R(\hat{h})}\right) \right| \xrightarrow{p} 0. \tag{49}$$

Proof. We use the same notation as in the Proof of Theorem 1. We can write

$$\mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n] = \frac{1}{\binom{n}{p}} \sum_{i \in S_n(p)} \frac{\mathbb{E}[\mathbb{I}(G_n^*(i) \simeq R)|\epsilon_i^*, G_n]}{p!N(R)\bar{\rho}_n^{|E(R)|}} \tag{50}$$

$$= \frac{1}{\binom{n}{p}} \sum_{\mathbf{i} \in S_n(p)} \sum_{R_1 \sim R} \frac{\hat{H}_{R_1}(\epsilon_i^*)}{p!N(R)\bar{\rho}_n^{|E(R)|}}, \tag{51}$$

which shows that $\mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n]$ is a U-statistic conditional on the graph G_n . The Berry-Esseen theorem for U-statistics [18] therefore tells us

$$\sup_x \left| \mathbb{P} \left(\frac{\sqrt{n}}{\bar{\rho}_n^{|E(R)|}} (\mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n] - P_R(\hat{h})) \leq x \mid G_n \right) - \Phi \left(\frac{x}{\tau_R(\hat{h})} \right) \right| \tag{52}$$

$$\leq \bar{\nu}_3 \bar{\sigma}_g^{-3} n^{-\frac{1}{2}}$$

where $\bar{\nu}_3$ and $\bar{\sigma}_g^{-3}$ are defined as in the proof of Theorem 1. First, we'll upper bound $\bar{\nu}_3$.

$$\bar{\nu}_3 \leq 8 \left(\mathbb{E} \left[\left| \frac{\sum_{R_1 \sim R} \hat{H}_{R_1}(\epsilon_{1:p}^*)}{p!N(R)\bar{\rho}_n^{|E(R)|}} \right|^3 \mid G_n \right] + \mathbb{E} \left[\frac{\sum_{R_1 \sim R} \hat{H}_{R_1}(\epsilon_{1:p}^*)}{p!N(R)\bar{\rho}_n^{|E(R)|}} \mid G_n \right]^3 \right) \tag{53}$$

$$\leq \frac{16}{\bar{\rho}_n^{3|E(R)|}} \mathbb{E} \left[\hat{H}_R(\epsilon_{1:p}^*)^3 \mid G_n \right] \tag{54}$$

where the second line follows from Holder's inequality. Then,

$$\mathbb{E} \left[\hat{H}_R(\epsilon_{1:p}^*)^3 \mid G_n \right] \tag{55}$$

$$= \int_{[0,1]^p} \prod_{(i,j) \in R} \hat{h}(u_i, u_j)^3 \prod_{(i,j) \in K_n \setminus R} (1 - \hat{h}(u_i, u_j))^3 du_{1:p}$$

$$\leq \int_{[0,1]^p} \prod_{(i,j) \in R} \hat{h}(u_i, u_j)^3 du_{1:p} \tag{56}$$

$$\leq \int_{[0,1]^2} (\hat{h}(u_1, u_2)^{3|E(R)|}) du_{1:2} \tag{57}$$

$$= O \left(\rho_n^{3|E(R)|} \right), \tag{58}$$

where the final equality follows by the boundedness in norm assumed in (12). This implies an upper bound on the skewness,

$$\bar{\nu}_3 = \frac{O_p \left(\rho_n^{3|E(R)|} \right)}{\bar{\rho}_n^{3|E(R)|}} = O_p(1). \tag{59}$$

Turning to $\bar{\sigma}_g^2$, just as in Theorem 1, we have

$$\bar{\sigma}_g^2 = \text{Var} \left(\mathbb{E} \left[\frac{\sum_{R_1 \sim R} \hat{H}_{R_1}(\epsilon_{1:p}^*)}{p!N(R)\bar{\rho}_n^{|E(R)|}} \mid \epsilon_1^*, G_n \right] \mid G_n \right) \tag{60}$$

$$= \frac{\sum_{S \in MC(R, 2p-1)} C_R(S) \left(P_S(\hat{h}) - P_R(\hat{h})^2 \right)}{\bar{\rho}_n^{2|E(R)|}}. \tag{61}$$

Observe that $S \in MC(R, 2p - 1)$ satisfies $|E(S)| = 2|E(R)|$. Consequently, the assumptions $\|w\|_{|E(S)|} < \infty$ and $\rho_n \|h_n - \hat{h}\|_{|E(S)|} = o_p(\rho_n)$ of Lemma 4 hold for all $S \in MC(R, 2p - 1)$ as well as for $S = R$, and so do the conclusions of Lemma 4. From here, reasoning exactly the same as in the proof of Theorem 1, except using Lemma 4 rather than Lemma 3, implies that $\bar{\sigma}_g^2 = \Omega_p(1)$. Combining this with our upper bound on $\bar{\nu}_3$, we conclude that

$$\bar{\nu}_3 \bar{\sigma}_g^{-3} n^{-\frac{1}{2}} = O_p(1) O_p(1) n^{-\frac{1}{2}} \xrightarrow{p} 0, \tag{62}$$

which concludes the proof of Lemma 5. □

Lemma 6. *Let R be a fixed motif. Then, if the conditions of Theorem 2 hold,*

$$\text{Var} \left(\frac{\sqrt{n}}{p!N(R)\bar{\rho}_n^{|E(R)|}} (P_R(G_n^*) - \mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n]) \Big| G_n \right) \xrightarrow{p} 0, \tag{63}$$

and

$$\frac{\sqrt{n}}{p!N(R)\bar{\rho}_n^{|E(R)|}} (P_R(G_n^*) - \mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n]) \xrightarrow{p} 0. \tag{64}$$

Proof. We start by rewriting

$$\begin{aligned} & \frac{\sqrt{n}}{p!N(R)\bar{\rho}_n^{|E(R)|}} (P_R(G_n^*) - \mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n]) \\ &= \frac{\sqrt{n}}{p!N(R)\binom{n}{p}\bar{\rho}_n^{|E(R)|}} \sum_{\mathbf{i} \in S_n(p)} \sum_{R_1 \sim R} \{\mathbb{I}(G_n^*(\mathbf{i}) = R_1) - \hat{H}_{R_1}(\epsilon_{\mathbf{i}}^*)\}. \end{aligned} \tag{65}$$

By the definition of \hat{H}_R ,

$$\mathbb{E} \left[\mathbb{I}(G_n^*(\mathbf{i}) = R) - \hat{H}_R(\epsilon_{\mathbf{i}}^*) \Big| \epsilon^*, G_n \right] = 0 \tag{66}$$

and so by the law of total variance,

$$\text{Cov} \left[\mathbb{I}(G_n^*(\mathbf{i}) = R_1) - \hat{H}_{R_1}(\epsilon_{\mathbf{i}}^*), \mathbb{I}(G_n^*(\mathbf{j}) = R_2) - \hat{H}_{R_2}(\epsilon_{\mathbf{j}}^*) \Big| G_n \right] \tag{67}$$

$$\begin{aligned} &= \mathbb{E} \left[\text{Cov} \left[\mathbb{I}(G_n^*(\mathbf{i}) = R_1) - \hat{H}_{R_1}(\epsilon_{\mathbf{i}}^*), \mathbb{I}(G_n^*(\mathbf{j}) = R_2) - \hat{H}_{R_2}(\epsilon_{\mathbf{j}}^*) \Big| \epsilon^*, G_n \right] \right] \\ &= \mathbb{E} [\text{Cov} [\mathbb{I}(G_n^*(\mathbf{i}) = R_1), \mathbb{I}(G_n^*(\mathbf{j}) = R_2) | \epsilon^*, G_n]] \end{aligned} \tag{68}$$

Therefore,

$$\text{Var} \left(\frac{\sqrt{n}}{p!N(R)\bar{\rho}_n^{|E(R)|}} (P_R(G_n^*) - \mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n]) \Big| G_n \right) \tag{69}$$

$$\begin{aligned} &= \frac{n}{(N(R)p!\binom{n}{p})^2 \bar{\rho}_n^{2|E(R)|}} \\ &\cdot \sum_{\mathbf{i}, \mathbf{j} \in S_n(p)} \sum_{R_1, R_2 \sim R} \mathbb{E} [\text{Cov} [\mathbb{I}(G_n^*(\mathbf{i}) = R_1), \mathbb{I}(G_n^*(\mathbf{j}) = R_2) | \epsilon^*, G_n]] \end{aligned} \tag{70}$$

Let us fix \mathbf{i}, \mathbf{j} and denote by $k = |\mathbf{i} \cap \mathbf{j}|$ the number of nodes that \mathbf{i} and \mathbf{j} have in common, and $k' = |\mathbf{i} \cup \mathbf{j}| = 2p - k$. Note that if $k < 2$, then $G_n^*(\mathbf{i})$ and $G_n^*(\mathbf{j})$ share no dyads, and are thus independent once we condition on ϵ^* . Otherwise if $k = 2, \dots, p$, we can bound the expected conditional covariance in terms of the moment $P_W(\hat{h})$ for some motif $W \in MC(R, k')$ ¹¹ as follows:

$$\left| \mathbb{E} \left[\text{Cov} [\mathbb{I}(G_n^*(\mathbf{i}) = R_1), \mathbb{I}(G_n^*(\mathbf{j}) = R_2) | \epsilon^*, G_n] \middle| G_n \right] \right| \tag{71}$$

$$\leq \mathbb{E} \left[\mathbb{I}(G_n^*(\mathbf{i}) = R_1), \mathbb{I}(G_n^*(\mathbf{j}) = R_2) \middle| G_n \right] \tag{72}$$

$$= P_W(\hat{h})$$

$$= O_p \left(\rho_n^{|E(W)|} \right), \tag{73}$$

where the last equality follows from Lemma 4, which we may invoke because (12) implies $\|\hat{h} - h\|_{|E(W)|} \leq \|\hat{h} - h\|_{2|E(R)|} = o_p(\rho_n)$ for all $k' = p, \dots, 2p - 1$ and $W \in MC(R, k')$.

From here, we divide our analysis into cases, based on whether (a) R is acyclic and $\rho_n = \omega(n^{-1})$, or (b) R is general and $\rho_n = \omega(n^{-2/p})$. Assuming (a), we have only the lower bound $|E(W)| \geq k' - 1$ for each $W \in MC(R, k')$, because W must be connected. Fortunately, this lower bound is enough, since in this case $2|E(R)| = 2(p - 1)$. Noting that there will be on the order of $n^{k'}$ valid choices for \mathbf{i} and \mathbf{j} which yield two subgraphs with k vertices in common, it follows from (69), (73), and the facts we have just observed that

$$\text{Var} \left(\frac{\sqrt{n}}{p!N(R)\bar{\rho}_n^{|E(R)|}} (P_R(G_n^*) - \mathbb{E} [P_R(G_n^*) | \epsilon^*, G_n]) \middle| G_n \right) \tag{74}$$

$$= \sum_{k=2}^p O(n^{1-k}) O_p \left(\rho_n^{(k'-1)-2|E(R)|} \right) \tag{75}$$

$$= \sum_{k=2}^p O(n^{1-k}) O_p \left(\rho_n^{-(k-1)} \right) \tag{76}$$

$$= o_p(1), \tag{77}$$

which implies (63).

Otherwise we assume (b), that R is general and $\rho_n = \omega(n^{-2/p})$. By the properties of the merged copy set, if $W \in MC(R, k')$ then $|E(W)| \geq 2|E(R)| - \binom{k}{2}$. It follows from (69) and (73) that

$$\text{Var} \left(\frac{\sqrt{n}}{p!N(R)\bar{\rho}_n^{|E(R)|}} (P_R(G_n^*) - \mathbb{E} [P_R(G_n^*) | \epsilon^*, G_n]) \middle| G_n \right) \tag{78}$$

$$= \sum_{k=2}^p O(n^{1-k}) O_p \left(\rho_n^{-\binom{k}{2}} \right) \tag{79}$$

¹¹In particular, the motif W on nodes $1, \dots, k'$ such that $W(\mathbf{i}) = R_1$ and $W(\mathbf{j}) = R_2$

$$= \sum_{k=2}^p o(\rho_n^{(1-k)p/2}) O_p\left(\rho_n^{-\binom{k}{2}}\right) \quad (80)$$

$$= o_p(1), \quad (81)$$

again implying (63).

Equation (64) then follows via an application of (conditional) Chebyshev's Inequality, which we make precise in Lemma 7, along with the fact that

$$\mathbb{E}\left[P_R(G_n^*) - \mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n] \mid G_n\right] = 0.$$

Thus the proof of Lemma 6 follows upon proving Lemma 7. \square

Lemma 7. *Let $(X_n), (Y_n)$ be two sequences of random variables.*

Suppose $\mathbb{E}[X_n|Y_n] = 0$ and $\text{Var}(X_n|Y_n) = o_p(1)$. Then $X_n = o_p(1)$.

Proof. It suffices to show that for any $a, \delta > 0$, $\mathbb{P}(X_n > a) < \delta$ for all n sufficiently large. To begin with, we have that for any $b > 0$:

$$\mathbb{P}(X_n > a) \leq \mathbb{P}(X_n > a | \text{Var}(X_n|Y_n) \leq b) + \mathbb{P}(\text{Var}(X_n|Y_n) > b).$$

By the law of iterated expectation, the conditional zero-mean property $\mathbb{E}[X_n|Y_n] = 0$ and Chebyshev's inequality,

$$\begin{aligned} \mathbb{P}(X_n > a | \text{Var}(X_n|Y_n) \leq b) &= \mathbb{E}\left[\mathbb{P}(X_n > a | Y_n) \mid \text{Var}(X_n|Y_n) \leq b\right] \\ &\leq \mathbb{E}\left[\text{Var}(X_n|Y_n) a^{-2} \mid \text{Var}(X_n|Y_n) \leq b\right] \\ &\leq \frac{b}{a^2}; \end{aligned}$$

note that the second line follows because by assumption $\mathbb{E}[X_n|Y_n] = 0$. Taking $b = \delta a^2/2$, we have by assumption that $\mathbb{P}(\text{Var}(X_n|Y_n) > b) \leq \delta/2$ for all n sufficiently large, and so for all such n we obtain that $\mathbb{P}(X_n > a) \leq \delta$, as desired. \square

The proof of Theorem 2 now follows straightforwardly. Putting Lemmas 5 and 6 together via Slutsky's Theorem yields

$$\sup_x \left| \mathbb{P}\left(\frac{\sqrt{n}}{\hat{\rho}_n^{|\mathbb{E}(R)|}} (P_R(G_n^*) - P_R(\hat{h})) \leq x \mid G_n\right) - \Phi\left(\frac{x}{\tau_R(\hat{h})}\right) \right| \xrightarrow{P} 0 \quad (82)$$

Finally, we have by the definition of conditional expectation that

$$\sigma_R^2(\hat{h}) = \tau_R^2(\hat{h}) + \text{Var}\left(\frac{\sqrt{n}}{\hat{\rho}_n^{|\mathbb{E}(R)|}} (P_R(G_n^*) - \mathbb{E}[P_R(G_n^*)|\epsilon^*, G_n]) \mid G_n\right) \quad (83)$$

By Lemma 6, we therefore have that

$$\left| \sigma_R^2(\hat{h}) - \tau_R^2(\hat{h}) \right| \xrightarrow{P} 0 \quad (84)$$

Since $\sigma_R^2(\hat{h}) = \theta_p(1)$ and thus $\tau_R^2(\hat{h}) = \theta_p(1)$, this in turn implies

$$\left| \Phi\left(\frac{x}{\tau_R(\hat{h})}\right) - \Phi\left(\frac{x}{\sigma_R(\hat{h})}\right) \right| \xrightarrow{p} 0 \tag{85}$$

which completes the proof of Theorem 2.

6.3. Proof of Lemma 1

This proof will be made slightly easier by introducing a second motif S , also on p nodes. (We can think of S as being an isomorphic copy of R .) Let $\mathbf{i} \cup \mathbf{j} = \mathbf{l}$, where (with slight notational mangling) \mathbf{l} is an ordered k -tuple. Then, define $R_{\mathbf{i}} = \{(c, d) : (a, b) \in R, (i_a, i_b) = (l_c, l_d)\}$ and similarly $S_{\mathbf{j}} = \{(c, d) : (a, b) \in S, (j_a, j_b) = (l_c, l_d)\}$. (Here, we've done nothing more than taken the two motifs and sent them to the right vertices as defined by the joint vertex set \mathbf{l} .) Now, let $W = R_{\mathbf{i}} \cup S_{\mathbf{j}}$. We would like to relate $\mathbb{I}\{G_n(\mathbf{i}) \simeq R\}\mathbb{I}\{G_n(\mathbf{j}) \simeq S\}$ and $\mathbb{I}\{G_n(\mathbf{l}) \simeq W\}$. Unfortunately, they are not quite equal. After all, if there are some edges *between* the vertices only in $G_n(\mathbf{i})$ and those only in $G_n(\mathbf{j})$ the LHS can still be 1, but the RHS will clearly be 0. To fix this, we sum over all these possible fuller motifs. Let $C_{V(R_{\mathbf{i}}), V(S_{\mathbf{j}})}$ be the set of dyads between vertices only in $R_{\mathbf{i}}$ and those only in $S_{\mathbf{j}}$. Then,

$$\mathbb{I}\{G_n(\mathbf{i}) = R\}\mathbb{I}\{G_n(\mathbf{j}) = S\} = \sum_{\substack{W: W=R_{\mathbf{i}} \cup S_{\mathbf{j}} \cup Q \\ Q \subseteq C_{V(R_{\mathbf{i}}), V(S_{\mathbf{j}})}}} \mathbb{I}\{G_n(\mathbf{l}) = W\} \tag{86}$$

and the relationship between seeing two motifs on different subsets of nodes and seeing one merged motif on the union of the subsets is established.¹²

These manipulations allow us to write the double sum over \mathbf{i} and \mathbf{j} , with the product of indicators of seeing the motifs R and S on the induced subgraphs $G(\mathbf{i})$ and $G(\mathbf{j})$, as a sum over \mathbf{l} , with the sum of indicators of seeing the motif W on $G(\mathbf{l})$.

$$\begin{aligned} & \sum_{\mathbf{i}, \mathbf{j} \in S_p(n)} \mathbb{I}(G_n(\mathbf{i}) = R)\mathbb{I}(G_n(\mathbf{j}) = S) \\ &= \sum_{k=p}^{2p} \sum_{\mathbf{l} \in S_k(n)} \sum_{\substack{\mathbf{i}, \mathbf{j} \in S_p(n) \\ \mathbf{i} \cup \mathbf{j} = \mathbf{l}}} \sum_{\substack{W: W=R_{\mathbf{i}} \cup S_{\mathbf{j}} \cup Q \\ Q \subseteq C_{V(S_{\mathbf{i}}), V(S_{\mathbf{j}})}}} \mathbb{I}(G_n(\mathbf{l}) = W). \end{aligned} \tag{87}$$

(87) makes clear how we can leverage the assumption of exchangeability, since $\mathbb{E}[\mathbb{I}(G_n(\mathbf{l}) = W)] = P_W(h)$ remains unchanged for all choices of \mathbf{l} , and so we

¹²Notice that we have replaced equality up to isomorphism with strict equality. This will simplify the following algebra, and returning to the isomorphism relationship can be established with one line at the end.

can simplify (17) to

$$\sum_{k=p}^{2p} \binom{n}{k} k! \sum_{\substack{\mathbf{i}, \mathbf{j} \in S_p(n) \\ \mathbf{i} \cup \mathbf{j} = 1:k}} \sum_{\substack{W: W=R_i \cup S_j \cup Q \\ Q \subseteq C_{V(R_i), V(S_j)}}} P_W(h). \tag{88}$$

Now, let us specify that $S \simeq R$. Then, for every choice of \mathbf{i}, \mathbf{j} and Q , by definition $R_i, S_j \simeq R$ and so $W \in MC(R, k)$. Moreover, for a given k the number of choices of \mathbf{i}, \mathbf{j} and Q are clearly fixed in n , and so the number of times each W in $MC(R, k)$ appears in the sum must also be fixed in n . Finally, to return to isomorphism note that

$$\sum_{R_1, R_2 \simeq R} \mathbb{I}(G_n(\mathbf{i}) = R_1) \mathbb{I}(G_n(\mathbf{i}) = R_2) = \mathbb{I}(G_n(\mathbf{i}) \simeq R) \mathbb{I}(G_n(\mathbf{i}) \simeq R). \tag{89}$$

and of course the number of $S \simeq R$, $N(R)$ is fixed in n as well. Denote the number of times each W appears as $C_W(R)$, where

$$C_{W_0}(R_0) = \sum_{\substack{\mathbf{i}, \mathbf{j} \in 1:k \\ \mathbf{i} \cup \mathbf{j} = 1:k}} \sum_{R, S \simeq R_0} \sum_{\substack{W: W=R_i \cup S_j \cup Q \\ Q \subseteq C_{V(R_i), V(S_j)}}} \mathbb{I}(W = W_0), \tag{90}$$

and the expression reduces to exactly the desired form.

6.4. Proof of Lemma 2

To get the expression for the first statement in the lemma, we expand the square and use Lemma 1.

$$\mathbb{E} [P_R(G_n)^2] = \frac{n}{\rho_n^{2|E(R)|} \binom{n}{p} p! N(R)^2} \sum_{\mathbf{i}, \mathbf{j} \in S_p(n)} \mathbb{E} [\mathbb{I}(G_n(\mathbf{i}) \simeq R) \mathbb{I}(G_n(\mathbf{j}) \simeq R)] \tag{91}$$

$$= \frac{n}{\rho_n^{2|E(R)|} \binom{n}{p} p! N(R)^2} \sum_{k=p}^{2p} \binom{n}{k} \sum_{W \in MC(R, k)} C_R(W) P_W(h) \tag{92}$$

Subtracting $P_R(h)^2$ from this, and doing some basic algebraic rearrangement, yields the desired result.

Now, we turn to the second statement in the lemma. Since the set $MC(R, k)$ is finite for any given R and k , and we are summing over a finite number of k , the problem reduces to showing that

$$\lim \frac{n \binom{n}{k} P_W(h)}{(\rho_n^{|E(R)|} \binom{n}{p})^2} < \infty \tag{93}$$

for all k in $p, \dots, 2p - 1$ and all W in $MC(R, k)$. But then,

$$P_W(h) \leq \int_{[0,1]^k} \prod_{(i,j) \in E(W)} h_n(u_i, u_j) du_{1:k} \tag{94}$$

$$\leq \int_{[0,1]^2} h_n(u_1, u_2)^{|E(W)|} du_{1:2} \tag{95}$$

$$= O\left(\rho_n^{|E(W)|}\right) \tag{96}$$

and so

$$\lim_{n \rightarrow \infty} \frac{n \binom{n}{k} P_W(h)}{(\rho_n^{|R|} \binom{n}{p} p! N(R))^2} = O\left(n^{k+1-2p} \rho_n^{|E(W)|-2|E(R)|}\right) \xrightarrow{p} 0 \tag{97}$$

where the last statement follows because either R is acyclic (and thus $|E(R)| = p - 1$) or $\rho_n = O(1/p)$.

6.5. Proof of Proposition 2

Both (20) and (21) come from Bickel, Chen and Levina [3].

6.6. Proof of Proposition 3

By Lemma 2, $\sigma_R^2(h_n) \rightarrow \sigma_R^2$; thus to prove Proposition 3 it suffices to show that for either the specific estimator $\hat{h} = \hat{h}_{adj}$, or for an arbitrary estimator \hat{h} that satisfies (12),

$$|\sigma_R^2(\hat{h}) - \sigma_R^2(h_n)| \xrightarrow{p} 0. \tag{98}$$

Lemma 2 also gives an expression for the normalized variance $\sigma_R^2(h)$, when either $h = h_n$ is the true graphon, or $h = \hat{h}$ is a graphon estimate. Taking the difference between these two gives

$$\begin{aligned} \sigma_R^2(h_n) - \sigma_R^2(\hat{h}) = & \tag{99} \\ & \frac{n}{\left(\binom{n}{p} p! N(R)\right)^2} \sum_{k=p}^{2p-1} \binom{n}{k} \sum_{S \in MC(R,k)} C_R(S) \left[\frac{P_S(h_n)}{\rho_n^{2|E(R)|}} - \frac{P_S(\hat{h})}{\bar{\rho}_n^{2|E(R)|}} \right] \\ & - \frac{1}{p!^2 N(R)^2} \left(1 - \frac{\binom{n-p}{p}}{\binom{n}{p}} \right)^2 \left[\frac{P_R(h_n)}{\rho_n^{|E(R)|}} - \frac{P_R(\hat{h})}{\bar{\rho}_n^{|E(R)|}} \right]. \end{aligned}$$

The above expression makes clear that in order to establish (98), we want to suitably upper bound $|P_S(\hat{h}) - P_S(h_n)|$ for each motif $S \in MC(R, k), k = p, \dots, 2p - 1$, as well as for $S = R$.

We now collect the relevant upper bounds on $|P_S(\hat{h}) - P_S(h_n)|$, under various assumptions and for each of $\hat{h} = \hat{h}_{adj}$ or \hat{h} an arbitrary estimator satisfying (12). We then show that these upper bounds imply (98).

Rates of convergence for motif densities. We summarize the results of Lemmas 3 and 4, insofar as they apply to the proof of Proposition 3

- Suppose R is a general motif, and $S \in MC(R, k)$ for some $k = p, \dots, 2p-1$. Then by assumption $\|w\|_{|E(S)|} \leq \|w\|_{2|E(R)|} < \infty$, and either (a) $\hat{h} = \hat{h}_{adj}$ and $\rho_n = \omega(n^{-1/2p})$, or (b) $\hat{h} = \hat{h}$ and $\|\hat{h} - h_n\|_{|E(S)|} \leq \|\hat{h} - h_n\|_{2|E(R)|} = o_p(\rho_n)$. Either way,

$$|P_S(\hat{h}) - P_S(h_n)| = o_p(\rho_n^{|E(S)|}). \tag{100}$$

- Suppose R is an acyclic motif, and $S \in MC(R, k)$ for some $k = p, \dots, 2p-1$. Then by assumption $\|w\|_{|E(S)|} < \|w\|_{2|E(R)|} < \infty$, and either (a) $\hat{h} = \hat{h}_{adj}$ and $\rho_n = \omega(n^{-1})$, or (b) $\hat{h} = \hat{h}$ and $\|\hat{h} - h_n\|_{|E(S)|} \leq \|\hat{h} - h_n\|_{2|E(R)|} = o_p(\rho_n)$. Either way,¹³ for any connected k -node motif S ,

$$|P_S(\hat{h}) - P_S(h_n)| = o_p(\rho_n^{k-1}). \tag{101}$$

Note that the above pair of statements hold because the assumptions of Proposition 3 subsume those of Lemma 3 and Lemma 4, respectively. Now we turn to establishing (98); we will show the results for $\hat{h} = \hat{h}_{adj}$ and \hat{h} an arbitrary estimator satisfying (12) at the same time, since the proof relies only on the above properties.

Proof of (98). To begin, we note that $\bar{\rho}_n = \hat{P}_S(\hat{h})$ for $S = K_2$, and so we know that

$$|\bar{\rho}_n - \rho_n| = o_p(\rho_n). \tag{102}$$

Next, we observe that

$$\left(1 - \frac{\binom{n-p}{p}}{\binom{n}{p}}\right)^2 |P_R(\hat{h}) - P_R(h_n)| = o\left(|P_R(\hat{h}) - P_R(h_n)|\right) = o_p(\rho_n^{|E(R)|}); \tag{103}$$

the last equality follows immediately from (100) when R is general, and from (101) in the special case where R is acyclic and $|E(R)| = p - 1$.

It remains to show that for all $S \in MC(R, k)$ and $k = p, \dots, 2p - 1$,

$$\frac{n \binom{n}{k}}{\binom{n}{p}^2 \rho_n^{2|E(R)|}} \left[P_S(h_n) - P_S(\hat{h}) \right] \xrightarrow{p} 0. \tag{104}$$

We now separate our analysis into the case where R is a general motif, and the special case of R acyclic.

General motif. When R is general, we can use (100) to reduce the left hand side of (104) to

$$\frac{n \binom{n}{k}}{\binom{n}{p}^2 \rho_n^{2|E(R)| - |E(S)|}} \frac{\left[P_S(h_n) - P_S(\hat{h}) \right]}{\rho_n^{|E(S)|}} = O(n^{k+1-2p}) o_p(1) \rho_n^{|E(S)| - 2|E(R)|}. \tag{105}$$

¹³Of course, Lemma 4 implies something stronger than (101) in case (b), but we will not need this stronger result.

To upper bound $2|E(R)| - |E(S)|$, we leverage the fact that S is a member of the merged copy set of R . The key is to notice that edges which are lost in S when vertices are merged can only be edges between two merged vertices. There are $2p - k$ such merged vertices, so there must be at least $2|E(R)| - \binom{2p-k}{2}$ edges in S , and $2|E(R)| - |E(S)| \leq \frac{(2p-k)(2p-k-1)}{2} =: \frac{k'(k'-1)}{2}$ for $k' = 2p - k$. Then plugging back into (105) gives

$$\frac{n \binom{n}{k}}{\binom{n}{p}^2 \rho_n^{2|E(R)|}} \left[P_S(h_n) - P_S(\hat{h}) \right] = O(n^{1-k'}) \rho_n^{-k'(k'-1)/2} o_p(1) \tag{106}$$

$$= O(n^{1-k'}) O(n^{k'-1}) o_p(1) = o_p(1), \tag{107}$$

with the penultimate equality following because $\rho_n = \Omega(n^{-2/p}) = \Omega(n^{-2/k'})$.

Acyclic motif. Otherwise if R is acyclic, then $|E(R)| = p - 1$, and we use (101) to deduce that

$$\frac{n \binom{n}{k}}{\binom{n}{p}^2 \rho_n^{2|E(R)|}} \left[P_S(h_n) - P_S(\hat{h}) \right] = \frac{n \binom{n}{k}}{\binom{n}{p}^2 \rho_n^{2p-k-1}} \left(\frac{P_S(h_n) - P_S(\hat{h})}{\rho_n^{k-1}} \right) \tag{108}$$

$$= \frac{n \binom{n}{k}}{\binom{n}{p}^2 \rho_n^{2p-k-1}} o_p(1) \tag{109}$$

$$= O\left((n\rho_n)^{-(2p-k-1)} \right) o_p(1) \tag{110}$$

$$= O(1) o_p(1), \tag{111}$$

with the last line following because $\rho_n = \Omega(n^{-1})$.

Thus we have established (104), which concludes the proof of Proposition 3.

6.7. Proof of Lemma 3

Throughout the proof of this lemma, $\hat{h} = \hat{h}_{adj}$. First, we will show that when $\rho_n = \omega(n^{-\frac{1}{k}})$, where $|V(S)| = k$, then

$$\left| \frac{P_S(\hat{h}) - P_S(h)}{\rho_n^{|E(S)|}} \right| \xrightarrow{p} 0 \tag{112}$$

We begin by bounding (112) by

$$\left| \frac{P_S(\hat{h}) - P_S(h)}{\rho_n^{|E(S)|}} \right| \leq \left| \frac{P_S(\hat{h}) - P_S(G_n)}{\rho_n^{|E(S)|}} \right| + \left| \frac{P_S(G_n) - P_S(h)}{\rho_n^{|E(S)|}} \right| \tag{113}$$

The second of these is $o_p(1)$ by Proposition 2, along with the assumption that $\int_{[0,1]^2} w(u, v)^{2|E(S)|}$ is finite. To bound the first term, we will make use of the following combinatorial identity, which relates the V -statistic $P_S(\hat{h})$ to the

U -statistic $P_S(G_n)$ and can be verified by standard counting arguments:

$$\left| \frac{P_S(\hat{h}) - P_S(G_n)}{\rho_n^{|E(S)|}} \right| = \left[1 - \frac{\binom{n}{k}}{n^k} \right] \frac{P_S(G_n)}{\rho_n^{|E(S)|}} + O\left(\sum_{j=1}^{k-1} n^{j-k} \sum_{W \in M(S,j)} \frac{P_W(G_n)}{\rho_n^{|E(S)|}} \right). \tag{114}$$

Here $M(S, j)$ is the set of motifs W on j vertices which can be formed by merging vertices in S , and is not the same as the merged copy set $MC(S, j)$.

W being formed by merging vertices in S restricts how many fewer edges it may have than S . The first merger of two vertices can have merged at most $k - 1$ edges, the second merger can have merged at most $k - 2$ edges, and so forth. As a result, if $W \in M(S, j)$,

$$|E(S)| - |E(W)| \leq \frac{(k + j - 1)(k - j)}{2}. \tag{115}$$

We will also use the fact that $\frac{P_W(G_n)}{\rho_n^{|E(W)|}} = O_p(1)$, again by Proposition 2 along with the fact that W has fewer edges than S . Putting these two together, we have

$$O(n^{j-k}) \frac{P_W(G_n)}{\rho_n^{|E(S)|}} = O_p\left(\frac{n^{j-k}}{\rho_n^{\frac{(k+j-1)(k-j)}{2}}} \right) \tag{116}$$

$$= o_p\left(n^{j-k + \frac{(k+j-1)(k-j)}{2k}} \right) \tag{117}$$

$$= o_p\left(n^{-\frac{k-j}{2k}} \right) = o_p(1) \tag{118}$$

since $j < k$. We must now deal with the leading term, but this is merely

$$\left[1 - \frac{\binom{n}{k} k!}{n^k} \right] \frac{P_S(G_n)}{\rho_n^{|E(S)|}} = O\left(\frac{1}{n} \right) O_p(1) \tag{119}$$

and so we have shown (112) in the case where $\rho_n = \omega(n^{-1/k})$.

Now we turn to the setting where the only restriction on ρ_n is that $n\rho_n \xrightarrow{n} \infty$. What we must show is that

$$\left| \frac{P_S(h) - P_S(\hat{h})}{\rho_n^{k-1}} \right| \xrightarrow{p} 0. \tag{120}$$

Similar to (114), we have

$$\left| \frac{P_S(\hat{h}) - P_S(G_n)}{\rho_n^{|k-1|}} \right| = \left[1 - \frac{\binom{n}{k}}{n^k} \right] \frac{P_S(G_n)}{\rho_n^{|k-1|}} + O\left(\sum_{j=1}^{k-1} n^{j-k} \sum_{W \in M(S,j)} \frac{P_W(G_n)}{\rho_n^{|k-1|}} \right) \tag{121}$$

By Proposition 2, we have that for all j and all $W \in MC(S, j)$, $\frac{P_W(G_n)}{\rho_n^{|E(W)|}} = O_p(1)$. Since $|E(W)| \geq j - 1$, this implies $\frac{P_W(G_n)}{\rho_n^{|k-1|}} = O_p(\rho_n^{j-k})$. Therefore, along

with the fact $n\rho_n \rightarrow \infty$, it follows that

$$O(n^{j-k}) \sum_{W \in M(S,j)} \frac{P_W(G_n)}{\rho_n^{|k-1|}} = o_p(1). \tag{122}$$

Finally, $\left[1 - \frac{\binom{n}{k}}{n^k}\right] \frac{P_S(G_n)}{\rho_n^{|k-1|}} = o_p(1)$, since $\frac{P_S(G_n)}{\rho_n^{|E(S)|}} = O_p(1)$ and $|E(S)| \geq k - 1$. Thus, we have shown (120).

6.8. Proof of Lemma 4

Throughout the proof of this Lemma, \hat{h} is a graphon estimate which satisfies (26) but is otherwise arbitrary, and $q = |V(S)|$. We can rewrite $P_S(\hat{h}) - P_S(h_n)$ using their definitions to yield

$$\left| \int_{[0,1]^q} \left\{ \prod_{(i,j) \in E(S)} \hat{h}(u_i, u_j) \prod_{(i,j) \notin E(S)} (1 - \hat{h}(u_i, u_j)) \right. \right. \tag{123}$$

$$\left. \left. - \prod_{(i,j) \in E(S)} h_n(u_i, u_j) \prod_{(i,j) \notin E(S)} (1 - h_n(u_i, u_j)) \right\} du_{1:q} \right|. \tag{124}$$

We now have the difference of products, but want the product of differences. We would also like to ignore the contribution of non-edges. So we use the following Lemma.

Lemma 8. *Let $a_1, \dots, a_\ell, b_1, \dots, b_\ell \in [0, 1]$ for a positive integer ℓ . Then, for any $k \in 1 : \ell$,*

$$\begin{aligned} & \left| \prod_{i=1}^k a_i \prod_{i=k+1}^\ell (1 - a_i) - \prod_{i=1}^k b_i \prod_{i=k+1}^\ell (1 - b_i) \right| \\ & \leq \prod_{i=1}^{k-1} a_i \cdot \sum_{j=k+1}^\ell |a_j - b_j| + \sum_{j=1}^k \left\{ \prod_{i=1}^{j-1} a_i \cdot |a_j - b_j| \cdot \prod_{i=j+1}^k |b_i| \right\}. \end{aligned} \tag{125}$$

In (125) we have adopted the convention that products which run over empty index sets are one—i.e. $\prod_{i=2}^1 a_i = \prod_{i=k+1}^k b_i = 1$ —and sums which run over empty index sets are zero—i.e. $\sum_{j=\ell+1}^\ell |a_j - b_j| = 0$ —for notational conciseness.

Proof. Lemma 8 To begin, suppose $k = \ell$, so that our goal is to show

$$\left| \prod_{i=1}^k a_i - \prod_{i=1}^k b_i \right| \leq \sum_{j=1}^k \left\{ \prod_{i=1}^{j-1} a_i \cdot |a_j - b_j| \cdot \prod_{i=j+1}^k |b_i| \right\}. \tag{126}$$

When $k = \ell = 1$, the claim is obvious. For a general $k = \ell$ it follows by induction; letting $A = \prod_{i=1}^{k-1} a_i$ and $B = \prod_{i=1}^{k-1} b_i$, we have

$$|Aa_k - Bb_k| \leq A|a_k - b_k| + b_k|A - B|$$

$$\begin{aligned} &\leq A|a_k - b_k| + b_k \sum_{j=1}^{k-1} \left\{ \prod_{i=1}^{j-1} a_i \cdot |a_j - b_j| \cdot \prod_{i=j+1}^{k-1} |b_i| \right\} \\ &= \sum_{j=1}^k \left\{ \prod_{i=1}^{j-1} a_i \cdot |a_j - b_j| \cdot \prod_{i=j+1}^k |b_i| \right\}. \end{aligned}$$

Now if $k < \ell$, then we have

$$\begin{aligned} &\left| \prod_{i=1}^k a_i \prod_{i=k+1}^{\ell} (1 - a_i) - \prod_{i=1}^k b_i \prod_{i=k+1}^{\ell} (1 - b_i) \right| \\ &\leq \prod_{i=1}^k a_i \left| \prod_{i=k+1}^{\ell} (1 - a_i) - \prod_{i=k+1}^{\ell} (1 - b_i) \right| + \prod_{i=k+1}^{\ell} (1 - b_i) \left| \prod_{i=1}^k a_i - \prod_{i=1}^k b_i \right| \\ &\leq \prod_{i=1}^k a_i \sum_{j=k+1}^{\ell} |a_j - b_j| + \left| \prod_{i=1}^k a_i - \prod_{i=1}^k b_i \right|, \end{aligned}$$

with the second inequality following because $1 - b_i \in [0, 1]$, along with the sum-product inequality which holds for numbers in $[0, 1]$. Then the claim of the Lemma follows from (126). \square

For notational convenience, let us adopt an arbitrary ordering $e_1, \dots, e_{|E(S)|}$ of the edges in $E(S)$, and also write $\hat{\Delta}(u, v) := |\hat{h}(u, v) - h_n(u, v)|$. From (123) and Lemma 8, we can upper bound $|P_S(\hat{h}) - P_S(h_n)|$ by the sum of two terms,

$$\begin{aligned} &|P_S(\hat{h}) - P_S(h_n)| \\ &\leq \sum_{(i', j') \notin E(S)} \int_{[0,1]^q} \prod_{i=1}^{|E(S)|-1} h_n(u_{e_i(1)}, u_{e_i(2)}) \hat{\Delta}(u_{i'}, u_{j'}) du_{1:q} \\ &\quad + \sum_{j=1}^{|E(S)|} \int_{[0,1]^q} \prod_{i=1}^{j-1} h_n(u_{e_i(1)}, u_{e_i(2)}) \hat{\Delta}(u_{e_i(1)}, u_{e_i(2)}) \prod_{i=j+1}^k \hat{h}(u_{e_i(1)}, u_{e_i(2)}) \\ &=: T_1 + T_2. \end{aligned}$$

Thus it remains to show that $T_1, T_2 = o_p(\rho_n^{|E(S)|})$. This is accomplished in a similar manner for each term, by using Hölder’s inequality, the assumption $\|w\|_{|E(S)|} < \infty$, and the rate of convergence assumed in (26). For T_1 , we have

$$\begin{aligned} T_1 &= \sum_{(i', j') \notin E(S)} \int_{[0,1]^q} \prod_{i=1}^{|E(S)|-1} h_n(u_{e_i(1)}, u_{e_i(2)}) \hat{\Delta}(u_{i'}, u_{j'}) du_{1:q} \\ &\leq \sum_{(i', j')} \|h_n\|_{|E(S)|}^{|E(S)|-1} \|\hat{\Delta}\|_{|E(S)|} \\ &= o_p(\rho_n^{|E(S)|}), \end{aligned}$$

and for T_2 ,

$$\begin{aligned} T_2 &= \sum_{j=1}^{|E(S)|} \int_{[0,1]^q} \prod_{i=1}^{j-1} h_n(u_{e_i(1)}, u_{e_i(2)}) \hat{\Delta}(u_{e_i(1)}, u_{e_i(2)}) \prod_{i=j+1}^k \hat{h}(u_{e_i(1)}, u_{e_i(2)}) \\ &\leq \sum_{j=1}^{|E(S)|} \|h_n\|_{|E(S)|}^{j-1} \|\hat{\Delta}\|_{|E(S)|} \|\hat{h}\|_{|E(S)|-j} \\ &= o_p(\rho_n^{|E(S)|}). \end{aligned}$$

This concludes the proof of Lemma 4.

Acknowledgments

We are grateful to the participants in the CMU Networkshop for valuable suggestions on the content and presentation of our results; to Prof. Paul Janssen for directing us to Janssen [18]; and to Profs. Carl T. Bergstrom, Peter J. Bickel, Dean Eckles, Jennifer Neville, Art B. Owen and Alessandro Rinaldo for valuable discussions, over the years, about network bootstraps.

References

- [1] BEARMAN, P. S., MOODY, J. and STOVEL, K. (2004). Chains of Affection: The Structure of Adolescent Romantic and Sexual Networks. *American Journal of Sociology* **110** 44–91.
- [2] BHATTACHARYYA, S. and BICKEL, P. J. (2015). Subsampling Bootstrap of Count Features of Networks. *Annals of Statistics* **43** 2384–2411. [MR3405598](#)
- [3] BICKEL, P. J., CHEN, A. and LEVINA, E. (2011). The method of moments and degree distributions for network models. *Annals of Statistics* **39** 38–59. [MR2906868](#)
- [4] BICKEL, P. J., GÖTZE, F. and VAN ZWET, W. R. (1997). Resampling fewer than n observations: gains, losses, and remedies for losses. *Statistica Sinica* 1–31. [MR1441142](#)
- [5] BORGS, C., CHAYES, J. T., LOVÁSZ, L., SÓS, V. T. and VESZTERGOMBI, K. (2008). Convergent Sequences of Dense Graphs I: Subgraph Frequencies, Metric Properties and Testing. *Advances in Mathematics* **219** 1801–1851. [MR2455626](#)
- [6] BORGS, C., CHAYES, J. T., COHN, H. and ZHAO, Y. (2019). An L^p Theory of Sparse Graph Convergence I: Limits, Sparse Random Graph Models, and Power Law Distributions. *Transactions of the American Mathematical Society* **372** 3019–3062. [MR3988601](#)
- [7] CALLAERT, H. and JANSSEN, P. (1978). The Berry-Esseen Theorem for U Statistics. *Annals of Statistics* **6** 417–421. [MR0464359](#)
- [8] DYNKIN, E. B. (1978). Sufficient Statistics and Extreme Points. *Annals of Probability* **6** 705–730. [MR0518321](#)

- [9] EDEN, T., LEVI, A., RON, D. and SESHADHRI, C. (2017). Approximately Counting Triangles in Sublinear Time. *SIAM Journal on Computing* **46** 1603–1646. [MR3709896](#)
- [10] ELDARDIRY, H. and NEVILLE, J. (2008). A Resampling Technique for Relational Data Graphs. In *Proceedings of the 2nd SNA Workshop, 14th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. Association for Computing Machinery, New York.
- [11] FIENBERG, S. E., MEYER, M. M. and WASSERMAN, S. S. (1985). Statistical Analysis of Multiple Sociometric Relations. *Journal of the American Statistical Association* **80** 51–67.
- [12] FIENBERG, S. E. and WASSERMAN, S. (1981). Categorical Data Analysis of Single Sociometric Relations. In *Sociological Methodology 1981* (S. Leinhardt, ed.) 156–192. Jossey-Bass, San Francisco.
- [13] GAO, C., LU, Y. and ZHOU, H. H. (2015). Rate-Optimal Graphon Estimation. *Annals of Statistics* **43** 2624–2652. [MR3405606](#)
- [14] GAO, C., LU, Y., MA, Z. and ZHOU, H. H. (2016). Optimal Estimation and Completion of Matrices with Biclustering Structures. *Journal of Machine Learning Research* **17** 1–29. [MR3569248](#)
- [15] GEYER, C. J. (2013). Asymptotics of Maximum Likelihood without the LLN or CLT or Sample Size Going to Infinity. In *Advances in Modern Statistical Theory and Applications: A Festschrift in honor of Morris L. Eaton* (G. JONES and X. SHEN, eds.) 1–24. Institute of Mathematical Statistics, Beachwood, Ohio. [MR3586936](#)
- [16] GONEN, M., RON, D. and SHAVITT, Y. (2011). Counting Stars and Other Small Subgraphs in Sublinear-Time. *SIAM Journal on Discrete Mathematics* **25** 1365–1411. [MR2837605](#)
- [17] HOLLAND, P. W., LASKEY, K. B. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Social Networks* **5** 109–137. [MR0718088](#)
- [18] JANSSEN, P. (1981). Rate of Convergence in the Central Limit Theorem and in the Strong Law of Large Numbers for von Mises Statistics. *Metrika* **28** 35–46. [MR0614539](#)
- [19] KALLENBERG, O. (2005). *Probabilistic Symmetries and Invariance Principles*. Springer-Verlag, New York. [MR2161313](#)
- [20] KLOPP, O., TSYBAKOV, A. B. and VERZELEN, N. (2017). Oracle inequalities for network models and sparse graphon estimation. *The Annals of Statistics* **45** 316–354. [MR3611494](#)
- [21] KOLACZYK, E. D. (2009). *Statistical Analysis of Network Data*. Springer-Verlag, New York. [MR2724362](#)
- [22] LAHIRI, S. N. (2003). *Resampling Methods for Dependent Data*. Springer-Verlag, New York. [MR2001447](#)
- [23] LAURITZEN, S. L. (1984). Extreme Point Models in Statistics. *Scandinavian Journal of Statistics* **11** 65–91. With discussion and response. [MR0754971](#)
- [24] LEGER, J.-B. (2016). Blockmodels: A R-package for estimating in Latent Block Model and Stochastic Block Model, with various probability functions, with or without covariates. E-print, arxiv:1602.07587.

- [25] LEVIN, K. and LEVINA, E. (2019a). Bootstrapping Networks with Latent Space Structure. *E-print*, *arXiv:1907.10821*.
- [26] LEVIN, K. and LEVINA, E. (2019b). Bootstrapping networks with latent space structure. *arXiv preprint arXiv:1907.10821*.
- [27] LOVÁSZ, L. (2009). Very Large Graphs. In *Current Developments in Mathematics, 2008* (D. JERISON, B. MAZUR, T. MROWKA, W. SCHMID, R. P. STANLEY and S.-T. YAU, eds.) 67–128. International Press, Somerville, Massachusetts. [MR1690997](#)
- [28] LUNDE, R. and SARKAR, P. (2019). Subsampling sparse graphons under minimal assumptions. *arXiv preprint arXiv:1907.12528*.
- [29] MCCULLAGH, P. (2000). Resampling and exchangeable arrays. *Bernoulli* **6** 285–301. [MR1748722](#)
- [30] MILO, R., SHEN-ORR, S., ITZKOVITZ, S., KASHTAN, N., CHKLOVSKII, D. and ALON, U. (2002). Network Motifs: Simple Building Blocks of Complex Networks. *Science* **298** 824–827.
- [31] OLDING, B. P. and WOLFE, P. J. (2014). Inference for graphs and networks: Extending classical tools to modern data. In *Data Analysis for Network Cyber-Security* (N. Adams and N. Heard, eds.) 1–31. World Scientific, Singapore. [MR3287662](#)
- [32] OWEN, A. B. and ECKLES, D. (2012). Bootstrapping Data Arrays of Arbitrary Order. *Annals of Applied Statistics* **6** 895–927. [MR3012514](#)
- [33] ROSVALL, M. and BERGSTROM, C. T. (2010). Mapping Change in Large Networks. *PLoS ONE* **5** e8694.
- [34] SHALIZI, C. R. and RINALDO, A. (2013). Consistency Under Sampling of Exponential Random Graph Models. *Annals of Statistics* **41** 508–535. [MR3099112](#)
- [35] WANG, L. (2016). Network Comparisons using Sample Splitting, PhD thesis, Carnegie Mellon University.
- [36] WOLFE, P. J. and OLHEDE, S. C. (2013). Nonparametric graphon estimation. E-print, *arxiv:1309.5936*.
- [37] XU, J. (2018). Rates of convergence of spectral methods for graphon estimation. In *International Conference on Machine Learning* 5433–5442. PMLR.
- [38] ZHANG, Y. and XIA, D. (2020). Edgeworth expansions for network moments. *arXiv preprint arXiv:2004.06615*.