Private Graphon Estimation for Sparse Graphs

Christian Borgs* Jennifer T. Chayes* Adam Smith†

June 23, 2015

Abstract

We design algorithms for fitting a high-dimensional statistical model to a large, sparse network without revealing sensitive information of individual members. Given a sparse input graph G, our algorithms output a node-differentially-private nonparametric block model approximation. By node-differentially-private, we mean that our output hides the insertion or removal of a vertex and all its adjacent edges. If G is an instance of the network obtained from a generative nonparametric model defined in terms of a graphon W, our model guarantees consistency, in the sense that as the number of vertices tends to infinity, the output of our algorithm converges to W in an appropriate version of the L_2 norm. In particular, this means we can estimate the sizes of all multi-way cuts in G.

Our results hold as long as W is bounded, the average degree of G grows at least like the log of the number of vertices, and the number of blocks goes to infinity at an appropriate rate. We give explicit error bounds in terms of the parameters of the model; in several settings, our bounds improve on or match known nonprivate results.

*Microsoft Research New England, Cambridge, MA, USA. Email: {cborgs,jchayes}@microsoft.com

†Pennsylvania State University, University Park, PA, USA. Email: asmith@psu.edu. Supported by NSF award IIS-1447700 and a Google Faculty Award. Part of this work was done while visiting Boston University's Hariri Institute for Computation and Harvard University's Center for Research on Computation and Society.

Contents

1 Introduction

Differential Privacy. Social and communication networks have been the subject of intense study over the last few years. However, while these networks comprise a rich source of information for science, they also contain highly sensitive private information. What kinds of information can we release about these networks while preserving the privacy of their users? Simple measures, such as removing obvious identifiers, do not work; for example, several studies (e.g., [6, 52]) reidentified individuals in the graph of a social network even after all vertex and edge attributes were removed. Such attacks highlight the need for statistical and learning algorithms that provide rigorous privacy guarantees.

Differential privacy [28], which emerged from a line of work started by [25], provides meaningful guarantees in the presence of arbitrary side information. In a traditional statistical data set, where each person corresponds to a single record (or row of a table), differential privacy guarantees that adding or removing any particular person's data will not noticeably change the distribution on the analysis outcome. There is now a rich and deep literature on differentially private methodology for learning and other algorithmic tasks; see [27] for a recent tutorial. By contrast, differential privacy in the context of graph data is much less developed. There are two main variants of graph differential privacy: *edge* and *node* differential privacy. Intuitively, edge differential privacy ensures that an algorithm's output does not reveal the inclusion or removal of a particular edge in the graph, while node differential privacy hides the inclusion or removal of a node together with all its adjacent edges. Edge privacy is a weaker notion (hence easier to achieve) and has been studied more extensively, with particular emphasis on the release of individual graph statistics [53, 55, 39, 51, 47, 40], the degree distribution [32, 33, 38, 44, 37], and data structures for estimating the edge density of all cuts in a graph [31, 9]. Several authors designed edge-differentially private algorithms for fitting generative graph models [51, 40, 47, 37, 60], but these do not appear to generalize to node privacy with meaningful accuracy guarantees.

The stronger notion, node privacy, corresponds more closely to what was achieved in the case of traditional data sets, and to what one would want to protect an individual's data: it ensures that *no matter what an analyst observing the released information knows ahead of time*, she learns the same things about an individual Alice regardless of whether Alice's data are used or not. In particular, no assumptions are needed on the way the individuals' data are generated (they need not even be independent). Node privacy was studied more recently [41, 22, 10, 54], with a focus on on the release of descriptive statistics (such as the number of triangles in a graph). Unfortunately, differential privacy's stringency makes the design of accurate, node-private algorithms challenging.

In this work, we provide the first algorithms for node-private inference of a high-dimensional statistical model that does not admit simple sufficient statistics.

Modeling Large Graphs via Graphons. Traditionally, large graphs have been modeled using various parametric models, one of the most popular being the stochastic block model [35]. Here one postulates that an observed graph was generated by first assigning vertices at random to one of k groups, and then connecting two vertices with a probability that depends on the groups the two vertices are members of.

As the number of vertices of the graph in question grows, we do not expect the graph to be well described by a stochastic block model with a fixed number of blocks. In this paper we consider nonparametric models (where the number of parameters need not be fixed or even finite) given in

terms of a *graphon*. A graphon is a measurable, bounded function $W : [0,1]^2 \to [0,\infty)$ such that $W(x,y) = W(y,x)$, which for convenience we take to be normalized: $\int W = 1$. Given a graphon, we generate a graph on n vertices by first assigning i.i.d. uniform labels in $[0,1]$ to the vertices, and then connecting vertices with labels x, y with probability $\rho_n W(x,y)$, where ρ_n is a parameter determining the density of the generated graph G_n with $\rho_n \|W\|_\infty \leq 1$. We call G_n a W-random graph with target density ρ_n (or simply a $\rho_n W$-random graph).

To our knowledge, random graph models of the above form were first introduced under the name latent position graphs [34], and are special cases of a more general model of "inhomogeneous random graphs" defined in [12], which is the first place were n-dependent target densities ρ_n were considered. For both dense graphs (whose target density does not depend on the number of vertices) and sparse graphs (those for which $\rho_n \to 0$ as $n \to \infty$), this model is related to the theory of convergent graph sequences [13, 14, 16, 17, 18]. For dense graphs it was first explicitly proposed in [46], though it can be implicitly traced back to [36, 5], where models of this form appear as extremal points of two-dimensional exchangeable arrays; see [24] (roughly, their results relate graphons to exchangeable arrays the way de Finetti's theorem relates i.i.d. distributions to exchangeable sequences). For sparse graphs, [19] offers a different nonparametric approach.

Estimation and Identifiability. Assuming that G_n is generated in this way, we are then faced with the task of estimating W from a *single observation* of a graph G_n. To our knowledge, this task was first explicitly considered in [7], which considered graphons describing stochastic block models with a fixed number of blocks. This was generalized to models with a growing number of blocks [56, 23], while the first estimation of the nonparametric model was proposed in [8]. Various other estimation methods were proposed recently, for example [45, 58, 43, 59, 20, 4, 61, 30, 3, 21, 1, 2]. These works make various assumptions on the function W, the most common one being that after a measure-preserving transformation, the integral of W over one variable is a strictly monotone function of the other, corresponding to an asymptotically strictly monotone degree distribution of G_n. (This assumption is quite restrictive: in particular, such results do not apply to graphons that represent block models.) For our purposes, the most relevant works are Wolfe and Olhede [59], Gao et al. [30], Chatterjee [21] and Abbe and Sandon [2], which provide consistent estimators without monotonicity assumptions (see "Comparison to nonprivate bounds", below).

One issue that makes estimation of graphons challenging is *identifiability*: multiple graphons can lead to the same distribution on G_n. Specifically, two graphons W and $\tilde{W}$ lead to the same distribution on W-random graphs if and only if there are measure preserving maps $\phi, \tilde{\phi} : [0,1] \to [0,1]$ such that $W^\phi = \widetilde{W}^{\tilde{\phi}}$, where W^ϕ is defined by $W(x,y) = W(\phi(x), \phi(y))$ [24, 15]. Hence, there is no "canonical graphon" that an estimation procedure can output, but rather an equivalence class of graphons. Some of the literature circumvents identifiability by making strong additional assumptions, such as strict monotonicity, that imply the existence of canonical equivalent class representatives. We make no such assumptions, but instead define consistency in terms of a metric on these equivalence classes, rather than on graphons as functions. We use a variant of the L_2 metric,

$$\delta_2(W, W') = \inf_{\phi:[0,1]\to[0,1]} \|W^\phi - W'\|_2 \,. \tag{1}$$

where ϕ ranges over measure-preserving bijections.

1.1 Our Contributions

In this paper we construct an algorithm that produces an estimate $\hat{W}$ from a single instance G_n of a W-random graph with target density ρ_n (or simply ρ, when n is clear from the context). We aim for several properties:

1. $\hat{W}$ is differentially private;
2. $\hat{W}$ is consistent, in the sense that $\delta_2(W, \hat{W}) \to 0$ in probability as $n \to \infty$;
3. $\hat{W}$ has a compact representation (in our case, as a matrix with $o(n)$ entries);
4. The procedure works for sparse graphs, that is, when the density ρ is small;
5. On input G_n, $\hat{W}$ can be calculated efficiently.

Here we give an estimation procedure that obeys the first four properties, leaving the question of polynomial-time algorithms for future work. Given an input graph G_n, a privacy-parameter ϵ and a target number k of blocks, our algorithm $\mathcal{A}$ produces a k-block graphon $\hat{W} = \mathcal{A}(G_n)$ such that

- $\mathcal{A}$ is ϵ-differentially node private. The privacy guarantee holds for all inputs, independent of modeling assumptions.
- Assume that (1) W is an arbitrary graphon, normalized so $\int W = 1$; (2) the expected average degree $(n-1)\rho$ grows at least as fast as $\log n$; and (3) k goes to infinity sufficiently slowly with n. Then when G_n is ρW-random, the estimate $\hat{W}$ for W is *consistent* (that is, $\delta_2(\hat{W}, W) \to 0$, both in probability and almost surely).

Combined with the general theory of convergent graphs sequences, these result in particular give a node-private procedure for estimating the edge density of all cuts in a ρW-random graph, see (24) in Section 2.2 below.

The main idea of our algorithm is to use the exponential mechanism of [50] to select a block model which approximately minimizes the ℓ_2 distance to the observed adjacency matrix of G, under the best possible assignment of nodes to blocks (this explicit search over assignments makes the algorithm take exponential time). In order to get an algorithm that is accurate on sparse graphs, we need several nontrivial extensions of current techniques. To achieve privacy, we use a new variation of the Lipschitz extension technique of [41, 22] to reduce the sensitivity of the δ_2 distance. While those works used Lipschitz extensions for noise addition, we use of Lipshitz extensions inside the "exponential mechanism" [50] (to control the sensitivity of the score functions). To bound our algorithm's error, we provide a new analysis of the ℓ_2-minimization algorithm; we show that approximate minimizers are not too far from the actual minimizer (a "stability" property). Both aspects of our work are enabled by restricting the ℓ_2^2-minimization to a set of block models whose density (in fact, L_∞ norm) is not much larger than that of the underlying graph. The algorithm is presented in Section 3.

Our most general result proves consistency for arbitrary graphons W but does not provides a concrete rate of convergence. However, we provide explicit rates under various assumptions on W. Specifically, we relate the error of our estimator to two natural error terms involving the graphon W: the error $\epsilon_k^{(O)}(W)$ of the best k-block approximation to W in the L_2 norm (see (5) below) and an error term $\epsilon_n(W)$ measuring the L_2-distance between the graphon W and the matrix of probabilities $H_n(W)$ generating the graph G_n (see (4) below.) In terms of these error terms,

Theorem 1 shows

$$\delta_2\Big(W, \hat{W}\Big) \leq \epsilon_k^{(O)}(W) + 2\epsilon_n(W) + O_P\left(\sqrt[4]{\frac{\log k}{\rho n}} + \sqrt{\frac{k^2 \log n}{n\epsilon}} + \frac{1}{\rho\epsilon n}\right). \tag{2}$$

Along the way, we provide a novel analysis of a straightforward, nonprivate least-squares estimator, whose error bound has a better dependence on k:

$$\delta_2\Big(W, \hat{W}_{\text{nonprivate}}\Big) \leq \epsilon_k^{(O)}(W) + 2\epsilon_n(W) + O_P\left(\sqrt[4]{\frac{\log k}{\rho n} + \frac{k^2}{\rho n^2}}\right). \tag{3}$$

It follows from the theory of graph convergence that for all graphons W, we have $\epsilon_k^{(O)}(W) \to 0$ as $k \to \infty$ and $\epsilon_n(W) \to 0$ almost surely as $n \to \infty$. As proven in Appendix D, we also have $\epsilon_n(W) = O_P(\epsilon_k^{(O)}(W) + \sqrt[4]{k/n})$, though this upper bound is loose in many cases.

As a specific instantiation of these bounds, let us consider the case that W is exactly described by a k-block model, in which case $\epsilon_k^{(O)}(W) = 0$ and $\epsilon_n(W) = O_P(\sqrt[4]{k/n})$ (see Appendix D). For $k \leq (n/\log^2 n)^{1/3}$, $\rho \geq \log(k)/k$ and constant ϵ, our private estimator has an asymptotic error that is dominated by the (unavoidable) error of $\epsilon_n(W) = \sqrt[4]{k/n}$, showing that we do not lose anything due to privacy in this special case. Another special case is when W is α-Hölder continuous, in which case $\epsilon_k^{(O)}(W) = O(k^{-\alpha})$ and $\epsilon_n(W) = O_P(n^{-\alpha/2})$; see Remark 2 below.

Comparison to Previous Nonprivate Bounds. We provide the first consistency bounds for estimation of a nonparametric graph model subject to node differential privacy. Along the way, for sparse graphs, we provide more general consistency results than were previously known, regardless of privacy. In particular, to the best of our knowledge, *no prior results give a consistent estimator for W that works for sparse graphs without any additional assumptions besides boundedness.*

When compared to results for nonprivate algorithms applied to graphons obeying additional assumptions, our bounds are often incomparable, and in other cases match the existing bounds.

We start by considering graphons which are themselves step functions with a known number of steps k. In the dense case, the nonprivate algorithms of [30] and [21], as well as our nonprivate algorithm, give an asymptotic error that is dominated by the term $\epsilon_n(W) = O(\sqrt[4]{k/n})$, which is of the same order as our private estimator as long as $k = \tilde{o}(n^{1/3})$. [59] provided the first convergence results for estimating graphons in the sparse regime. Assuming that W is bounded above and below (so it takes values in a range $[\lambda_1, \lambda_2]$ where $\lambda_1 > 0$), they analyze an inefficient algorithm (the MLE). The bounds of [59] are incomparable to ours, though for the case of k-block graphons, both their bounds and our nonprivate bound are dominated by the term $\sqrt[4]{k/n}$ when $\rho > (\log k)/k$ and $k \leq \rho n$. A different sequence of works shows how to consistently estimate the underlying block model with a *fixed* number of blocks k in polynomial time for very sparse graphs (as for our non-private algorithm, the only thing which is needed is that $n\rho \to \infty$) [3, 1, 2]; we are not aware of concrete bounds on the convergence rate.

For the case of *dense* α-Hölder-continuous graphons, the results of [30] give an error which is dominated by the term $\epsilon_n(W) = O_P(n^{-\alpha/2})$. For $\alpha < 1/2$, our nonprivate bound matches this bound, while for $\alpha > 1/2$ it is worse. [59] consider the sparse case. The rate of their estimator is incomparable to the rate of our estimator; further, their analysis requires a lower bound on the edge probabilities, while ours does not.

See Appendix A for a more detailed discussion of the previous literature.

2 Preliminaries

2.1 Notation

For a graph G on $[n] = \{1, \dots, n\}$, we use $E(G)$ and $A(G)$ to denote the edge set and the adjacency matrix of G, respectively. The edge density $\rho(G)$ is defined as the number of edges divided by $\binom{n}{2}$. Finally the degree d_i of a vertex i in G is the number of edges containing i. We use the same notation for a weighted graph with nonnegative edge weights β_{ij}, where now $\rho(G) = \frac{2}{n(n-1)} \sum_{i<j} \beta_{ij}$, and $d_i = \sum_{j \neq i} \beta_{ij}$. We use $\mathbb{G}_n$ to denote the set of weighted graphs on n vertices with weights in $[0,1]$, and $\mathbb{G}_{n,d}$ to denote the set of all graphs in $\mathbb{G}_n$ that have maximal degree at most d.

From Matrices to Graphons. We define a graphon to be a bounded, measurable function $W : [0,1]^2 \to \mathbb{R}_+$ such that $W(x,y) = W(y,x)$ for all $x, y \in [0,1]$. It will be convenient to embed the set of a symmetric $n \times n$ matrix with nonnegative entries into graphons as follows: let $\mathcal{P}_n = (I_1, \dots I_n)$ be the partition of $[0,1]$ into adjacent intervals of lengths $1/n$. Define $W[A]$ to be the step function which equals A_{ij} on $I_i \times I_j$. If A is the adjacency matrix of an unweighted graph G, we use $W[G]$ for $W[A]$.

Distances. For $p \in [1, \infty)$ we define the L_p norm of an $n \times n$ matrix A by $\|A\|_p^p = \frac{1}{n^2} \sum_{i,j} |A_{ij}|^p$, and the L_p norm of a (Borel)-measurable function $W : [0,1]^2 \to \mathbb{R}$ by $\|f\|_p^p = \int |f(x,y)|^p dxdy$. Associated with the L_2-norm is a scalar product, defined as $\langle A, B \rangle = \frac{1}{n^2} \sum_{i,j} A_{ij} B_{ij}$ for two $n \times n$ matrices A and B, and $\langle U, W \rangle = \int U(x,y) W(x,y) dxdy$ for two square integrable functions $U, W : [0,1]^2 \to \mathbb{R}$. Note that with this notation, the edge density and the L_1 norm are related by $\|G\|_1 = \frac{n-1}{n} \rho(G)$.

Recalling (1), we define the δ_2 distance between two matrices A, B, or between a matrix A and a graphon W by $\delta_2(A,B) = \delta_2(W[A], W[B])$ and $\delta_2(A,W) = \delta_2(W[A], W)$. In addition, we will also use the in general larger distances $\hat{\delta}_2(A,B)$ and $\hat{\delta}_2(A,W)$, defined by taking a minimum over matrices A' which are obtained from A by a relabelling of the indices: $\hat{\delta}_2(A,B) = \min_{A'} \|A' - B\|_2$ and $\hat{\delta}_2(A,W) = \min_{A'} \|A' - W\|_2$.

2.2 W-random graphs and graph convergence

W-random graphs and stochastic block models. Given a graphon W we define a random $n \times n$ matrix $H_n = H_n(W)$ by choosing n "positions" $x_1, \dots, x_n$ i.i.d. uniformly at random from $[0,1]$ and then setting $(H_n)_{ij} = W(x_i, x_j)$. If $\|W\|_\infty \leq 1$, then $H_n(W)$ has entries in $[0,1]$, and we can form a random graph $G_n = G_n(W)$ on n-vertices by choosing an edge between two vertices $i < j$ with probability $(H_n)_{ij}$, independently for all $i < j$. Following [46] we call $G_n(W)$ a W-random graph and $H_n(W)$ a W-weighted random graph. We incorporate a target density ρ_n (or simply ρ, when n is clear from the context) by normalizing W so that $\int W = 1$ and taking G to be a sample from $G_n(\rho W)$. In other words, we set $Q = H_n(\rho W) = \rho H_n(W)$ and then connect i to j with probability Q_{ij}, independently for all $i < j$.

The error from our main estimates measuring the distance between $H_n(W)$ and W is defined as

$$\epsilon_n(W) = \hat{\delta}_2(H_n(W), W) \tag{4}$$

and goes to zero as $n \to \infty$ by the following lemma, which follows easily from the results of [17].

Lemma 1. *Let W be a graphon with $\|W\|_\infty < \infty$. With probability one, $\|H_n(W)\|_1 \to \|W\|_1$ and $\epsilon_n(W) \to 0$.*

Stochastic block models are specific examples of W-random graph in which W is constant on sets of the form $I_i \times I_j$, where $(I_1, \dots, I_k)$ is a partition of $[0,1]$ into intervals of possibly different lengths.

Approximation by block models. In the opposite direction, we can map a function W to a matrix B by the following procedure. Starting from an arbitrary partition $\mathcal{P} = (Y_1, \dots, Y_k)$ of $[0,1]$ into sets of equal Lebesgues measure, define $W/\mathcal{P}$ to be the matrix obtained by averaging over sets of the form $Y_i \times Y_j$,

$$(W/\mathcal{P})_{ij} = \frac{1}{\lambda(Y_i)\lambda(Y_j)} \Big(\int_{Y_i \times Y_j} W(x,y) dx dy \Big)$$

where $\lambda(\cdot)$ denotes the Lebesgue measure. Finally, we will use $W_\mathcal{P}$ to denote the step function

$$W_\mathcal{P} = \sum_{i,j \in [k]} (W/\mathcal{P})_{ij} 1_{Y_i} \times 1_{Y_j}.$$

Using the above averaging procedure, it is easy to see that any graphon W can be well approximated by a block model. Indeed, let

$$\epsilon_k^{(O)}(W) = \min_B \|W - W[B]\|_2 \tag{5}$$

where the minimum goes over all $k \times k$ matrices B. Given that we are minimizing the L_2-distance, the minimimizer can easily be calculated, and is equal to $W_{\mathcal{P}_k}$, where $\mathcal{P}_k$ is a partition of $[0,1]$ into adjacent intervals of lengths $1/k$. It then follows from the Lebesgue density theorem (see, e.g., [17] for details) that $\epsilon_k^{(O)}(W) = \|W - W_{\mathcal{P}_k}\|_2 \to 0$ as $k \to \infty$.

We will take the above approximation as a benchmark for our approach, and consider it the error an "oracle" could obtain (hence the superscript O).

Convergence. The theory of graph convergence was first developed [13, 14, 16], where it was formulated for dense graphs, and then generalized to sparse graphs in [11, 17, 18]. One of the notions of graph convergence considered in these papers is the notion of convergence in metric. The metric in question is similar to the metric δ_2, but instead of the L_2-norm, one starts from the cut-norm $\|\cdot\|_\square$ first defined in [29],

$$\|W\|_\square = \sup_{S,T \subset [0,1]} \Big| \int_{S \times T} W \Big|,$$

where the supremum goes over all measurable sets $S, T \subset [0,1]$. The cut-distance $\delta_\square$ between two integrable functions $U, W; [0,1]^2 \to \mathbb{R}$ is then defined as

$$\delta_\square(U,W) = \inf_\phi \|U^\phi - W\|_\square,$$

where the inf goes over all measure preserving bijections on $[0,1]$. We will also need the following variations: a distance $\hat{\delta}_\square(G,G')$ between two graphs on the same node set, as well as a distance $\hat{\delta}_\square(G,W)$ between a graph G and a graphon W, defined as $\hat{d}_\square(G,G') = \min_{G''} \|W[G''] - W[G]\|_\square$

and $\hat{d}_\square(G,W) = \min_{G''} \|W[G''] - W\|_\square$, respectively, where the minimum goes over graphs G'' isomorphic to G.

Given these notions, we say a (random or deterministic) sequence G_n of graphs converges to a graphon W in the cut metric if, as $n \to \infty$,

$$\delta_\square\left(\frac{1}{\rho(G_n)} W[G_n], W\right) \to 0\,.$$

With this notion of convergence, for any graphon W with $\int W = 1$, a sequence of W-random graphs G_n with target density ρ_n converges to the generating graphon W. This was shown for bounded W and n-independent target densities ρ with $\rho\|W\|_\infty \le 1$ in [14], but the statement is much more general, and in particular holds for arbitrary target densities ρ_n as long as $n\rho_n \to \infty$ and $\limsup \rho_n \|W\|_\infty \le 1$ [17].

Estimation of Multi-Way Cuts. Using the results of [18], the convergence of G_n in the cut-metric $\delta_\square$ implies many interesting results for estimating various quantities defined on the graph G_n. Indeed, a consistent approximation $\hat{W}$ to W in the metric δ_2 is clearly consistent in the weaker metric $\delta_\square$. But this distance controls various quantities of interest to computer scientists, e.g., the size of all multi-way cuts, implying that a consistent estimator for W also gives consistent estimators for all multi-way cuts.

To formalize this, we need some notation. Given a weighted graph G on $[n]$ with node-weights one and edge-weights $\beta_{xy}(G)$, and given a partition $\mathcal{P} = (V_1, \ldots, V_q)$ of $[n]$ into q groups, let $G/\mathcal{P}$ be the weighted graph with weights

$$\alpha_i(G/\mathcal{P}) = |V_i|/|V(G)| \quad \text{and} \quad \beta_{ij}(G/\mathcal{P}) = \frac{1}{n^2\|G\|_1} \sum_{x\in V_i, y\in V_j} \beta_{xy}(G).$$

We call $G/\mathcal{P}$ a q-quotient or q-way cut of G, and denote the set of all q-way cuts by $S_q(G)$:

$$S_q(G) = \{G/\mathcal{P} : \mathcal{P} \text{ is a partition of } [n] \text{ into } q \text{ sets } \}\,.$$

We also consider the set of *fractional q-way cuts*, $\hat{S}_q(G) = \{G/\rho\}$, defined in terms of *fractional q-partitions* ρ. A fractional q-partition of $V(G)$ is a map $\rho : V(G) \to \Delta_q : x \mapsto \rho(x)$, where Δ_q is the simplex $\Delta_q = \{\rho = (\rho_i) \in [0,1]^q \colon \sum_i \rho_i = 1\}$, and the corresponding fractional quotient G/ρ is the weighted graph with weights $\alpha_i(G/\rho) = \frac{1}{|V(G)|}\sum_x \rho_i(x)$ and $\beta_{ij}(G/\rho) = \frac{1}{n^2\|G\|_1}\sum_{x\in V_i, y\in V_j} \beta_{xy}(G)\rho_i(x)\rho_j(x)$.

The set of fractional q-partitions of a graphon W, $\hat{S}_q(W)$, is defined similarly: $\hat{S}_q(W) = \{W/\rho \mid \rho : [0,1] \to \Delta_q\}$, with a fractional partition now a measurable function $\rho : [0,1] \to \Delta_q$, and W/ρ given in terms of the weights

$$\alpha_i(W/\rho) = \int \rho_i(x)dx \quad \text{and} \quad \beta_{ij}(W/\rho) = \frac{1}{\|W\|_1}\int \rho_i(x)\rho_j(y)W(x,y).$$

To measure the distance between the various sets of q-way cuts, we use the Hausdorff distance between sets $S, S' \subset \mathbb{R}^{q+q^2}$,

$$d_\infty^{\text{Haus}}(S,S') = \max\{\sup_{H\in S}\inf_{H'\in S'} \|H - H'\|_\infty,\ \sup_{H'\in S'}\inf_{H\in S}\|H - H'\|_\infty\},$$

where $\|\cdot\|_\infty$ is the L_∞-norm $\|H - H'\|_\infty = \max\{\max_i |\alpha_i(H) - \alpha_i(H')|, \max_{ij} |\beta_{ij}(H) - \beta_{ij}(H')|\}$.

It was shown in [17] that if G_n converges to W in the cut-metric, then $d_\infty^{\text{Haus}}(S_q(G_n), \hat{S}_q(W)) \to 0$. In particular, a consistent estimator $\hat{W}$ for the generating graphon W of a W-random graph G_n leads to a consistent estimator $\hat{S}_q(\hat{W})$ for the cuts $S_q(G_n)$, in the sense that $d_\infty^{\text{Haus}}(S_q(G_n), \hat{S}_q(\hat{W})) \to 0$. With a little more work, we can give quantitative bounds; see Theorem 3 below.

2.3 Differential Privacy for Graphs

The goal of this paper is the development of a differentially private algorithm for graphon estimation. The privacy guarantees are formulated for worst-case inputs — we do not assume that G is generated from a graphon when analyzing privacy. This ensures that the guarantee remains meaningful no matter what an analyst knows ahead of time about G.

In this paper, we consider the notion of node privacy. We call two graphs G and G' *node neighbors* if one can be obtained from the other by removing one node and its adjacent edges.

Definition 1 (ϵ-node-privacy)**.** *A randomized algorithm $\mathcal{A}$ is* ϵ-node-private *if for all events S in the output space of $\mathcal{A}$, and node neighbors G, G',*

$$\Pr[\mathcal{A}(G) \in S] \leq \exp(\epsilon) \times \Pr[\mathcal{A}(G') \in S]\,.$$

We also need the notion of the *node-sensitivity* of a function $f : \mathbb{G}_n \to \mathbb{R}$, defined as maximum $\max_{G,G'} |f(G) - f(g')|$, where the maximum goes over node-neighbors. This constant is often called the Lipshitz constant of f

Finally, we need a lemma concerning the extension of functions $f : \mathbb{G}_{n,d} \to \mathbb{R}$ to functions $\hat{f} : \mathbb{G}_n \to \mathbb{R}$. We say a function on adjacency matrices is *nondecreasing* if adding an edge to the adjacency matrix does not increase the value of the function.

Lemma 2 ([48, 41])**.** *For every function $f : \mathbb{G}_{n,d} \to \mathbb{R}$, there is an extension $\hat{f} : \mathbb{G}_n \to \mathbb{R}$ of f with the same node-sensitivity as f. If f is a nondecreasing linear function of the adjacency matrix, then we can select $\hat{f}$ to be nondecreasing and computable in polynomial-time and so that $\hat{f}(G) \leq f(G)$ for all graphs $G \in \mathbb{G}_n$.*

The lemma is proved in Appendix B.

3 Differentially Private Graphon Estimation

3.1 Least-squares Estimation

Given a graph as input generated by an unknown graphon W, our goal is to recover a block-model approximation to W. The basic nonprivate algorithm we emulate is least squares estimation, which outputs the $k \times k$ matrix B which is closest to the input adjacency matrix A in the distance

$$\hat{\delta}_2(B, A) = \min_\pi \|B_\pi - A\|_2,$$

where the minimum runs over all equipartitions π of $[n]$ into k classes, i.e., over all maps $\pi : [n] \to [k]$ such that all classes have size as close to n/k as possible, i.e., such that $||\pi^{-1}(i)| - n/k| < 1$ for all i, and B_π is the $n \times n$ block-matrix with entries $(B_\pi)_{xy} = B_{\pi(x)\pi(y)}$. If A is the adjacency matrix of a graph G, we write $\hat{\delta}_2(B, G)$ instead of $\hat{\delta}_2(B, A)$. In the above notation, the basic algorithm we would want to emulate is then the algorithm which outputs the least square fit $\hat{B} = \text{argmin}_B\, \hat{\delta}_2(B, G)$, where the argmin runs over all symmetric $k \times k$ matrices B.

3.2 Towards a Private Algorithm

Our main idea to turn the least square algorithm into a private algorithm is to use the so-called exponential mechanism of McSherry and Talwar [50]. Applied naively, we would therefore want to output a random $k \times k$ matrix B according to the probability distribution

$$\Pr(\hat{B} = B) \propto \exp\Big(-C\hat{\delta}_2^2(B, A)\Big),$$

with C chosen small enough to guarantee differential privacy. It is a standard fact from the theory of differential privacy, that C should be at most ϵ over twice the node-sensitivity of the "score function", here $\delta_2^2(B, \cdot)$. But this value of C turns out to be too small to produce an output that is a good approximation to the least square estimator. Indeed, for a given matrix B and equipartition π, the distance $\|G - B_\pi\|_2^2$ can change by as much as $\frac{1}{n}$ when G is replaced by a node-neighbor, regardless of the magnitude of the entries of B. To obtain differential privacy, we then would need to choose $C \geq n\epsilon/2$, which turns out to not produce useful results when the input graph G is sparse, since small values of C will lead to large errors relative to the least square estimator.

To address this, we first note that we can work with an equivalent score that is much less sensitive. Given B and π, we subtract off the squared norm of G to obtain the following:

$$\begin{aligned} score(B, \pi; G) &= \|G\|_2^2 - \|G - B_\pi\|_2^2 = 2\langle G, B_\pi\rangle - \|B_\pi\|^2, \text{ and} && (6)\\ score(B; G) &= \max_\pi score(B, \pi; G), && (7)\end{aligned}$$

where the max ranges over equipartitions $\pi : [n] \to [k]$. For a fixed input graph G, maximizing the score is the same as minimizing the distance, i.e. $\operatorname{argmin}_B \hat{\delta}_2(B, G) = \operatorname{argmax}_B score(B; G)$. The sensitivity of the new score is then bounded by $\frac{2}{n^2} \cdot \|B\|_\infty$ times the maximum degree in G (since G only affects the score via the inner product $\langle G, B_\pi\rangle$). But this is still problematic since, a priori, we have no control over either the size of $\|B\|_\infty$ or the maximal degree of G.

To keep the sensitivity low, we make two modifications: first, we only optimize over matrices B whose entries are of order ρ_n (in the end, we expect that a good estimator will have entries which are not much larger than $\|\rho_n W\|_\infty$, which is of order ρ_n), and second we restrict ourselves to graphs G whose maximum degree is not much larger than one would expect for graphs generated from a bounded graphon, namely a constant times the average degree. While the first restriction is something we can just implement in our algorithm, unfortunately the second is something we have no control over: We need to choose C small enough to guarantee privacy for all input graphs, and we have set out to guarantee privacy in the worst case, which includes graphs with maximal degree $n - 1$. Here, we employ an idea from [10, 41]: we first consider the restriction of $score(B, \pi; \cdot)$ to $\mathbb{G}_{n,d_n}$ where d_n will be chosen to be of the order of the average degree of G, and then extend it back to all graphs while keeping the sensitivity low.

3.3 Private Estimation Algorithm

After these motivations, we are now ready to define our algorithm. It takes as input the privacy parameter ϵ, the graph G, a number k of blocks, and a constant $\lambda \geq 1$ that will have to be chosen large enough to guarantee consistency of the algorithm. It outputs a matrix B from the set of matrices

$$\mathcal{B}_\mu = \{B \in [0, \mu]^{k\times k} : \text{all entries } B_{i,j} \text{ are multiples of } \tfrac{1}{n}\}.$$

Inside our algorithm, we use an $\epsilon/2$-private algorithm to get an estimate $\hat{\rho}$ for the edge density of G. We do so by setting $\hat{\rho} = \rho(G) + \mathrm{Lap}(4/n\epsilon)$, where $\mathrm{Lap}(\kappa)$ is a Laplace random variable with density $h(z) = \frac{1}{2\kappa}e^{-|z|/\kappa}$. The existence of the Lipschitz extension used in the algorithm follows from Lemma 2.

Algorithm 1: Private Estimation Algorithm

Input: $\epsilon > 0$, $\lambda \geq 1$, an integer k and graph G on n vertices.

Output: $k \times k$ block graphon (represented as a $k \times k$ matrix) estimating ρW

1 Compute an $(\epsilon/2)$-node-private density approximation $\hat{\rho} = \rho(G) + \mathrm{Lap}(4/n\epsilon)$;

2 $d = \lambda\hat{\rho}n$ (the target maximum degree) ;

3 $\mu = \lambda\hat{\rho}$ (the target L_∞ norm for the matrix B) ;

4 For each B and π, let $\widehat{\mathrm{score}}(B, \pi; \cdot)$ denote a nondecreasing Lipschitz extension of $score(B, \pi; \cdot)$ from $\mathbb{G}_{n,d}$ to $\mathbb{G}_n$ such that for all matrices A, $\widehat{\mathrm{score}}(B, \pi; A) \leq score(B, \pi; A)$, and define

$$\widehat{\mathrm{score}}(B; A) = \max_{\pi} \widehat{\mathrm{score}}(B, \pi; A)$$

5 **return** $\hat{B}$, sampled from the distribution

$$\Pr(\hat{B} = B) \propto \exp\Big(\frac{\epsilon}{4\Delta}\widehat{\mathrm{score}}(B; A)\Big),$$

where B ranges over matrices in $\mathcal{B}_\mu$ and $\Delta = \dfrac{4d\mu}{n^2} = \dfrac{4\lambda^2\hat{\rho}^2}{n}$;

Lemma 3. *Algorithm 1 is ϵ-node private.*

Proof. By Lemma 10 from Appendix B, the estimate $\hat{\rho}$ is $\epsilon/2$-private, so we want to prove that the exponential mechanism itself is $\epsilon/2$-private as well. In view of Lemma 11 from Appendix B, all we need to show is that the the vertex sensitivity of $\widehat{\mathrm{score}}(B; \cdot)$ is at most Δ. To this end, we first bound the vertex sensitivity of the original score when restricted to graphs with degree d. Let $G, G' \in \mathbb{G}_{n,d}$ be node neighbors. From (6), we see that

$$score(B, \pi; G) - score(B, \pi; G') = \frac{2}{n^2}\sum_{x,y\in[n]} (A_{xy} - A'_{xy})B_{\pi(x)\pi(y)}\,,$$

where A, A' are the adjacency matrices of G and G'. Since A and A' differ in at most $2d$ entries, the score differs by at most $4d\|B_\pi\|_\infty/n^2$. This is at most Δ, since $B \in \mathcal{B}_\mu$. Since $\widehat{\mathrm{score}}$ is a Lipschitz extension of *score*, the vertex sensitivity of $\widehat{\mathrm{score}}$ (over *all* neighboring graphs) is at most Δ, as required. □

Theorem 1 (Performance of the Private Algorithm)**.** *Let $W : [0,1]^2 \to [0, \Lambda]$ be a normalized graphon, let $0 < \rho\Lambda \leq 1$, let $G = G_n(\rho W)$, $\lambda \geq 1$, and k be an integer. Assume that $\rho n \geq 6\log n$ and $8\Lambda \leq \lambda \leq \sqrt{n}$, $2 \leq k \leq \min\{n\sqrt{\frac{\rho}{2}}, e^{\frac{\rho n}{2}}\}$. Then the Algorithm 1 outputs an approximation $(\hat{\rho}, \hat{B})$ such that*

$$\delta_2\Big(W, \frac{1}{\hat{\rho}}W[\hat{B}]\Big) \leq \epsilon_k^{(O)}(W) + 2\epsilon_n(W) + O_P\left(\sqrt[4]{\frac{\lambda^2\log k}{\rho n}} + \lambda\sqrt{\frac{k^2\log n}{n\epsilon}} + \frac{\lambda}{n\rho\epsilon}\right).$$

The theorem will be proven in Section 5.

In the course of the proof, we will prove results on the performance of a non-private algorithm, which is a variant of the standard least square algorithm, the main difference being that instead of minimizing $\hat{\delta}_2(B, A)$ over all matrices B, we only optimize it over matrices whose entries are bounded by a constant times the density of G.

Algorithm 2: Nonprivate Algorithm

Input: $\lambda \geq 1$, an integer k and graph G on n vertices.

Output: $k \times k$ block graphon (represented as a $k \times k$ matrix B) estimating ρW

$\mu \leftarrow \lambda\rho(G)$ (the target L_∞ norm for the matrix B) ;

return $\hat{B} \in \underset{B \in \mathcal{B}_\mu}{\operatorname{argmin}} \hat{\delta}_2(B; G)$.

Theorem 2 (Performance of the Nonprivate Algorithm)**.** *Let $W : [0,1]^2 \to [0, \Lambda]$ be a normalized graphon, let $0 < \rho\Lambda \leq 1$, let $G = G_n(\rho W)$, $\lambda \geq 1$, and k be an integer. If $\hat{B}$ is the least-squares estimator (Algorithm 2), $2\Lambda \leq \lambda \leq \sqrt{n}$, $2 \leq k \leq \min\{n\sqrt{\frac{\rho}{2}}, e^{\frac{\rho n}{2}}\}$, then*

$$\delta_2\Big(W, \frac{1}{\rho(G)} W[\hat{B}]\Big) \leq \epsilon_k^{(O)}(W) + 2\epsilon_n(W) + O_p\left(\sqrt[4]{\lambda^2\left(\frac{\log k}{\rho n} + \frac{k^2}{\rho n^2}\right)}\right).$$

In particular, $\delta_2\Big(W, \frac{1}{\rho(G)} W[\hat{B}]\Big) \to 0$ in probability if $k \to \infty$ and $\lambda^2\Big(\frac{\log k}{\rho n} + \frac{k^2}{\rho n^2}\Big) \to 0$.

Theorem 2 is proven in Section 4.

Remark 1. *While Theorem 1 and Theorem 2 are stated in term of bounds which hold in probability, our proofs give slightly more, and allow us in particular to prove statements which hold almost surely as $n \to \infty$. Namely, they show that under the assumptions of Theorem 2, the output $\hat{B}$ of the nonprivate algorithm is such that*

$$\delta_2\Big(W, \frac{1}{\rho(G)} W[\hat{B}]\Big) \leq \epsilon_k^{(O)}(W) + O\left(\sqrt[4]{\frac{\lambda^2 \log k}{\rho n} + \frac{\lambda^2 k^2}{\rho n^2}}\right) + o(1);$$

they also show that if we replace the assumption $n\rho \geq 6 \log n$ in Theorem 1 by the stronger assumption $n\rho\epsilon / \log n \to \infty$, then the output $\hat{B}$ of the private algorithm is such that

$$\delta_2\Big(W, \frac{1}{\hat{\rho}} W[\hat{B}]\Big) \leq \epsilon_k^{(O)}(W) + O\left(\sqrt[4]{\frac{\lambda^2 \log k}{\rho n}} + \lambda\sqrt{\frac{k^2 \log n}{n\epsilon}} + \frac{\sqrt{\lambda}}{n\rho\epsilon}\right) + o(1)$$

where in both expressions, $o(1)$ is a term which goes to zero with probability one as $n \to \infty$.

Thus for both algorithm, as long as k grows sufficiently slowly with n, with probability one, the asymptotic error is of the form $\epsilon_k^{(O)}(W) + o(1)$, which is best possible, since we can't do better than the best oracle block model approximation.

Remark 2. *Under additional assumptions on the graphon W, we can say a little more. For example, if we assume that W is Hölder continuous, i.e, if we assume that the exists constants $\alpha \in (0, 1]$ and $C < \infty$ such that $|W(x, y) - W(x', y')| \leq C\delta^\alpha$ whenever $|x - x'| + |y - y'| \leq \delta$, then we have that $\epsilon_k^{(O)}(W) = O(k^{-\alpha})$ and $\epsilon_n(W) = O_P(n^{-\alpha/2})$. See Appendix E for details.*

Theorems 1 and 2 imply that the sets of fractional q-way cuts of the estimator $\hat{B}$ from these theorems provide good approximations to the q-way cuts of the graph G (as defined in Section 2.2). Specifically:

Theorem 3. *Let $q \geq 2$ be an integer.*

(i) Under the assumptions of Theorem 2,

$$d_{\infty}^{Haus}(S_q(G), \hat{S}_q(\hat{B}_{nonprivate})) = O_p\left(\epsilon_k^{(O)}(W) + \epsilon_n(W) + \sqrt[4]{\lambda^2\left(\frac{\log k}{\rho n} + \frac{k^2}{\rho n^2}\right)}\right).$$

(ii) Under the assumptions of Theorem 1,

$$d_{\infty}^{Haus}(S_q(G), \hat{S}_q(\hat{B}_{private})) = O_p\left(\epsilon_k^{(O)}(W) + \epsilon_n(W) + \sqrt[4]{\frac{\lambda^2 \log k}{\rho n}} + \lambda\sqrt{\frac{k^2 \log n}{n\epsilon}} + \frac{\lambda}{n\rho\epsilon}\right).$$

The proof of the theorem relies on the theory of graph convergence, in particular the results of [16, 17, 18], and is given in Appendix F.

Remark 3. *When considering the "best" block model approximation to W, one might want to consider block models with unequal block sizes; in a similar way, one might want to construct a private algorithm that outputs a block model with unequal size blocks, and produces a bound in terms of this best block model approximation instead of $\epsilon_k^{(O)}(W)$. With more cumbersome notation, this can be easily proved with our methods, with the minimal block size taking the role of $1/k$ in all our proofs. We leave the details to a journal version.*

4 Estimation Error of the Least Square Algorithm

At a high level, our proofs of Theorems 1 and of 2 follow from the fact that for all B and π, the expected score $\mathbb{E}[Score(B, \pi; G)]$ is equal to the score $Score(B, \pi; Q)$, combined with a concentration argument. As a consequence, the maximizer $\hat{B}$ of $Score(B; G)$ will approximately minimize the L_2-distance $\hat{\delta}_2(B, Q)$, which in turn will approximately minimize $\|\frac{1}{\rho}W[B] - W\|_2$, thus relating the L_2-error of our estimator $\hat{B}$ to the "oracle error" $\epsilon_k^{(O)}(W)$ defined in (5).

In this section we present the analysis of exact and approximate least squares. This allows us to analyze the nonprivate algorithm. The analysis of the private algorithm (Theorem 1) requires additional arguments relating the private approximate maximizer to the nonprivate one; we present these in Section 5).

Our main concentration statement is contained in the following proposition, which we prove in Section 4.1 below. To state it, we define, for every symmetric $n \times n$ matrix Q with vanishing diagonal, $Bern_0(Q)$ to be the distribution over symmetric matrices A with zero diagonal such that the entries $\{A_{ij} : i < j\}$ are independent Bernouilli random variables with $\mathbb{E}A_{ij} = Q_{ij}$.

Proposition 1. *Let $\mu > 0$, $Q \in [0,1]^{n \times n}$ be a symmetric matrix with vanishing diagonal, and $A \sim Bern_0(Q)$. If $2 \leq k \leq \min\{n\sqrt{\rho(Q)}, e^{\rho(Q)n}\}$ and $\hat{B} \in \mathcal{B}_\mu$ is such that*

$$Score(\hat{B}; A) \geq \max_{B \in \mathcal{B}_\mu} Score(B; A) - \nu^2$$

for some $\nu > 0$, then with probability at least $1 - 2e^{-n}$,

$$\hat{\delta}_2(\hat{B}, Q) \leq \min_{B \in \mathcal{B}_\mu} \hat{\delta}_2(B, Q) + \nu + O\left(\sqrt[4]{\mu^2 \rho(Q)\left(\frac{k^2}{n^2} + \frac{\log k}{n}\right)}\right) \tag{8}$$

and in particular

$$\begin{aligned} \|\hat{B}\|_2 &\leq (2\|Q\|_2 + \nu)\left(1 + \frac{2k}{n}\right) + O\left(\sqrt[4]{\mu^2 \rho(Q)\left(\frac{k^2}{n^2} + \frac{\log k}{n}\right)}\right) \\ &\leq (2\|Q\|_2 + \nu)\left(1 + \frac{2k}{n}\right) + O\left(\sqrt{\mu\rho(Q)}\right) \end{aligned} \tag{9}$$

Morally, the proposition contains almost all that is needed to establish the bound (3) proving consistency of the standard least squares algorithm (which, in fact, only involves the case $\nu = 0$), even though there are several additional steps needed to complete the proof (see Sections 4.2 and 4.3 below).

The proposition also contains an extra ingredient which is a crucial input for the analysis of the private algorithm: it states that if instead of an optimal, least square estimator, we output an estimator whose score is only approximately maximal, then the excess error introduced by the approximation is small. To apply the proposition, we then establish a a lemma which gives us a lower bound on the score of the output $\hat{B}$ in terms of the maximal score and an excess error ν.

There are several steps needed to execute this strategy, the most important ones involving a rigorous control of the error introduced by the Lipschitz extension inside the exponential algorithm (which in turn requires estimating the deviation of the maximal degree from the expected degree, a step where the condition that ρn has to grow like $\log n$ is needed). The excess error ν eventually turns into the second to last error term in (2), while the difference between our private estimator $\hat{\rho}$ for the edge density and the actual edge density of G is responsible for the last one.

The analysis of the private algorithm is presented in Section 5; the remainder of this section presents the detailed analysis of the least squares estimator.

Remark 4. *Note that for both the non-private algorithm and the private algorithm, the above proposition naturally gives a bound for the L_2 estimation error for matrix of probabilities Q. In fact, our proofs provide error bounds on $\hat{\delta}_2(\hat{B}, Q)$ which differ from* (3), (2) *and the bounds in Theorem 3 in that (i) the error term $2\epsilon_n(W)$ is absent, and (ii) the oracle error $\epsilon_k^{(O)}(W)$ is replaced by an oracle error $\hat{\epsilon}_k^{(O)}(H_n)$ for H_n, see Theorems 4, 5 and 6. Converting these bounds into bounds on $\delta_2(W, \frac{1}{\hat{\rho}}W[\hat{B}])$ and expressing the result in terms of $\epsilon_k^{(O)}(W)$ instead of $\hat{\epsilon}_k^{(O)}(H_n)$ then introduces the error term $2\epsilon_n(W)$ in* (3), (2) *and the bounds in Theorem 3.*

4.1 Expectation and Concentration of Scores

The following two lemmas contain the core of the argument outlined at the beginning of this section.

Lemma 4 (Expected scores)**.** *Let $Q \in [0,1]^{n \times n}$ be a symmetric matrix with vanishing diagonal, let $A \sim Bern_0(Q)$, and let B, B' be $k \times k$ matrices. Then*

$$\hat{\delta}_2^2(Q, B) - \hat{\delta}_2^2(Q, B') = \max_{\pi'} \mathbb{E}[Score(B', \pi'; G)] - \max_{\pi} \mathbb{E}[Score(B, \pi; G)],$$

where the two max's go over equipartitions $\pi, \pi' : [n] \to [k]$.

Proof. By linearity of expectation, we have

$$\begin{aligned}\mathbb{E}score(B,\pi;A) &= \mathbb{E}\big(2\langle A, B_\pi\rangle - \|B_\pi\|^2\big) = 2\langle Q, B_\pi\rangle - \|B_\pi\|_2^2 \\ &= \|Q\|_2^2 - \|Q - B_\pi\|_2^2 \,.\end{aligned}$$

Taking into account the definition of $\hat{\delta}_2(B,Q)$, the lemma follows. □

Our second lemma states that the realized scores are close to their expected values. The proof is based on a careful application of the concentration bounds. The argument is delicate because we must take advantage of the low density (when ρ is small).

Lemma 5 (Concentration of scores)**.** *Let $\mu > 0$, let $Q \in [0,1]^{n\times n}$ be a symmetric matrix with vanishing diagonal and let $A \sim Bern_0(Q)$. If $2 \le k \le \min\{n\sqrt{\rho(Q)}, e^{\rho(Q)n}\}$, then, with probability at least $1 - 2e^{-n}$*

$$|score(B,\pi;A) - \mathbb{E}[score(B,\pi;A)]| = O\left(\mu\sqrt{\rho(Q)\left(\frac{k^2}{n^2} + \frac{\log k}{n}\right)}\right)$$

for all equipartitions π and all $B \in [0,\mu]^{k\times k}$.

Proof. First, consider a specific pair B, π. Recall that

$$score(B,\pi;A) - \mathbb{E}[score(B,\pi;A)] = 2\langle A - Q, B_\pi\rangle \,.$$

We wish to bound the deviation of $score(B,\pi;A)$ from its mean. Set $\rho(Q) = \tilde{\rho}$. The quantity $S = \frac{n^2}{2\mu}\cdot\langle A, B_\pi\rangle = \sum_{i<j}\frac{B_{\pi(i)\pi(j)}}{\mu}A_{ij}$ is a sum of $\binom{n}{2}$ independent random variables in $[0,1]$ with expectation $\mathbb{E}S \le \tilde{\rho}\binom{n}{2}$. Using a slight variation on the standard Chernoff bound, which we state in Lemma 17, we will bound the probability that S deviates from its mean by at most $\beta\mu_0$, where $\mu_0 \ge \mathbb{E}[S]$ will be chosen in a moment. Setting $\eta = 2e^{-n}$ and

$$\beta = \sqrt{\frac{k^2 + n\log k + \log(2/\eta)}{3\tilde{\rho}n^2}} = O\left(\sqrt{\frac{k^2}{\tilde{\rho}n^2} + \frac{\log k}{\tilde{\rho}n}}\right)$$

the assumption $k \le \min\{n\sqrt{\rho(Q)}n, e^{\rho(Q)n}\}$ implies $\beta \le 1$, and setting $\mu_0 = 9n^2\tilde{\rho}$, the bound from Lemma 17 becomes

$$2e^{-3\beta^2\tilde{\rho}n^2} = e^{-k^2}k^{-n}\eta \le 2^{-k^2}k^{-n}\eta,$$

implying that

$$\Pr\left(|2\langle A - Q, B_\pi\rangle| \ge \frac{4\mu}{n^2}\beta\mu_0\right) \le \frac{\eta}{k^n 2^{k^2}} \,.$$

Finally, we observe that for any A, the maximum of $|\langle Q - A, B_\pi\rangle|$ over all $B \in [0,\mu]^{k\times k}$ is the same as the maximum over all $B \in \{0,\mu\}^{k\times k}$. Taking a union bound over the (at most $2^{k^2}k^n$) pairs B, π and observing that $\frac{4\mu}{n^2}\beta\mu_0 = O\left(\mu\sqrt{\tilde{\rho}\left(\frac{k^2}{n^2} + \frac{\log k}{n}\right)}\right)$, we get the statement of the lemma. □

Proof of Proposition 1. Let $\hat{B} \in \mathcal{B}_\mu$ be as specified, let $B' \in \mathcal{B}_\mu$ arbitrary, and let $\pi, \pi' : n \to k$ be two equipartitions. By Lemmas 4 and 5,

$$\begin{aligned}\hat{\delta}_2^2(Q, \hat{B}) - \hat{\delta}_2^2(Q, B') &= \max_{\pi'} \mathbb{E}[Score(B', \pi'; G)] - \max_{\pi} \mathbb{E}[Score(\hat{B}, \pi; G)] \\ &\leq Score(B'; G) - Score(\hat{B}; G) + O\left(\mu \sqrt{\rho(Q)\left(\frac{k^2}{n^2} + \frac{\log k}{n}\right)}\right) \\ &\leq \nu^2 + O\left(\sqrt{\mu^2 \rho(Q)\left(\frac{k^2}{n^2} + \frac{\log k}{n}\right)}\right)\end{aligned}$$

which implies the bound (8). (Taking square roots works since $\sqrt{\sum_i C_i^2} \leq \sum_i C_i$ as long as $C_i \geq 0$.) To prove (9), we use that for an arbitrary equipartition π $\|\hat{B}_\pi\|_2^2 \geq \left(1 - \frac{k}{n}\right)\|\hat{B}\|_2^2$ and that $\|\hat{B}_\pi\|_2 \leq \|Q\|_2 + \|B_\pi - Q\|_2$. Inserting the definition of $\hat{\delta}_2(\hat{B}, Q)$ and using the main statement plus the assumptions $2 \leq k \leq \min\{n\sqrt{\rho(Q)}, e^{\rho(Q)n}\}$, we obtain (9). □

4.2 Estimation of the edge-probability matrix Q

Up to technical details, Proposition 1 contains all that is needed to prove consistency of the least square algorithm. As indicated in Remark 4, we will first prove that the algorithm gives a consistent estimator for the matrix Q, and then use this prove that the output also gives a consistent estimator for W. The first statement is formalized in the following theorem.

Theorem 4. *Under the assumptions of Theorem 2,*

$$\hat{\delta}_2\Big(\frac{1}{\rho(G)}\hat{B}, H_n(W)\Big) \leq \hat{\epsilon}_k^{(O)}(H_n(W)) + O_p\left(\sqrt{\lambda}\Big(\frac{\log k}{\rho n} + \frac{k^2}{\rho n^2}\Big)^{1/4}\right) \tag{10}$$

where

$$\hat{\epsilon}_k^{(O)}(H) = \inf_B \hat{\delta}_2(B, H),$$

with the inf *going over all symmetric $k \times k$ matrices B. Moreover, a.s. as $n \to \infty$,*

$$\hat{\delta}_2\left(\frac{1}{\rho(G)}\hat{B}, H_n(W)\right) \leq \hat{\epsilon}_k^{(O)}(H_n(W)) + O\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)}\right) + o(1).$$

Proof. As a first step, we will bound the left hand side of (10) by conditioning on the event that

$$\frac{\rho}{2} \leq \rho(Q) \leq 2\rho. \tag{11}$$

By a concentration argument very similar to the proof of Lemma 5 above (in fact, it is easier, see Lemma 12 (part 3) in Appendix C), we have that, with probability at least $1 - 2e^{-n}$,

$$\frac{\rho(G)}{\rho(Q)} = 1 + O\left(\frac{1}{\sqrt{n\rho(Q)}}\right) = 1 + O\left(\frac{1}{\sqrt{n\rho}}\right).$$

We can now apply Proposition 1 with $\nu = 0$ (since the nonprivate algorithm returns an exact minimizer). Recall that $H_n(W) = \frac{Q}{\rho}$ and $\mu = \lambda\rho(G) = \Theta(\lambda\rho)$. We get that, with probability at least $1 - 4e^{-n}$,

$$\hat{\delta}_2\left(\frac{1}{\rho}\hat{B}, H_n(W)\right) \leq \min_{B \in \mathcal{B}_\mu} \hat{\delta}_2\left(\frac{1}{\rho}B, H_n(W)\right) + O\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)}\right). \tag{12}$$

In the remainder of the proof, we bound the first term on the left-hand side above by relating it to the "oracle error" $\hat{\epsilon}_k^{(O)}(H_n(W))$. Let B' and π be such that $\hat{\epsilon}_k^{(O)}(H_n(W)) = \|H_n(W) - B'_\pi\|_2$. It is easy to see that then B'_π is obtained from $H_n(W)$ by averaging over the classes of π, which in turn implies that $\|B'_\pi\|_\infty = \|B'\|_\infty \leq \|H_n(W)\|_\infty \leq \|W\|_\infty \leq \Lambda$ and $\|B'_\pi\|_2 \leq \|H_n(W)\|_2 = \rho^{-1}\|Q\|_2$. Define B by rounding all entries of $\rho(G)B'$ down to the nearest multiple of $1/n$, adding a rounding error of at most $1/n$, so that $\|B - \rho(G)B'\|_\infty \leq 1/n$. Note that B' is on the scale of W and $H_n(W)$ (that is, we expect $\|B'\| = \Theta(1)$), while B is on the scale of ρW and Q; hence, $\|B_\pi\|_2 \leq \frac{\rho(G)}{\rho}\|Q\|_2$. Now $\mu \geq \rho(G)\Lambda$, $\|B\|_\infty \leq \rho(G)\|B'\|_\infty \leq \rho(G)\Lambda$, and $\Lambda \leq \lambda/2$. Thus, B is in the set $\mathcal{B}_\mu$ that the algorithm searches over. We can bound the first term in the left-hand side of (12) by

$$\begin{aligned}
\hat{\delta}_2\left(H_n(W), \frac{1}{\rho}B\right) &\leq \left\|H_n(W) - \frac{1}{\rho}B_\pi\right\|_2 \\
&\leq \left\|H_n(W) - \frac{1}{\rho(G)}B_\pi\right\|_2 + \left\|\frac{1}{\rho(G)}B_\pi - \frac{1}{\rho}B_\pi\right\|_2 \\
&\leq \hat{\epsilon}_k^{(O)}(H_n(W)) + \left\|B'_\pi - \frac{1}{\rho(G)}B_\pi\right\|_2 + \left\|\frac{1}{\rho(G)}B_\pi - \frac{1}{\rho}B_\pi\right\|_2 \\
&\leq \hat{\epsilon}_k^{(O)}(H_n(W)) + \frac{1}{n\rho(G)} + \left|1 - \frac{\rho(G)}{\rho}\right|\frac{\|Q\|_2}{\rho}
\end{aligned}$$

Combined with our previous two bounds and the fact that by (11), we can bound $\|Q\|_2$ by $\|Q\|_2 \leq \sqrt{\|Q\|_1\|Q\|_\infty} \leq \sqrt{\rho(Q)\Lambda\rho} \leq \sqrt{2\Lambda\rho} \leq \sqrt{\lambda}\rho$, this implies that

$$\begin{aligned}
\hat{\delta}_2\left(\frac{1}{\rho}\hat{B}, H_n(W)\right) &\leq \hat{\epsilon}_k^{(O)}(H_n(W)) + \sqrt{\lambda}\left|1 - \frac{\rho(Q)}{\rho}\right| + O\left(\sqrt{\frac{\lambda}{n\rho}}\right) + O\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)}\right) \\
&\leq \hat{\epsilon}_k^{(O)}(H_n(W)) + \sqrt{\lambda}\left|1 - \frac{\rho(Q)}{\rho}\right| + O\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)}\right).
\end{aligned}$$

We can now use (9) to bound $\|\hat{B}\|_2$ by $O(\|Q\|_2 + \sqrt{\mu\rho(Q)}) = O(\sqrt{\lambda}\rho)$ and thus $\delta_2(\hat{B}/\rho, \hat{B}/\rho(G))$ by $O(\sqrt{\lambda})|1 - \rho(Q)/\rho|$ plus an error which can be absorbed into the error term above. We obtain that, conditioned on (11), with probability at least $1 - 4e^{-n}$, we have

$$\hat{\delta}_2\left(\frac{1}{\rho(G)}\hat{B}, H_n(W)\right) \leq \hat{\epsilon}_k^{(O)}(H_n(W)) + O\left(\sqrt[4]{\lambda^2\left(\left|1 - \frac{\rho(Q)}{\rho}\right|^4 + \frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)}\right). \tag{13}$$

By Lemma 12 from Appendix C, $|\rho(Q) - \rho|^2 = O_P(\lambda\rho^2/n)$, implying that

$$\lambda^2\left|1 - \frac{\rho(Q)}{\rho}\right|^4 = O_P\left(\frac{\lambda^4}{n^2}\right) = O_P\left(\frac{\lambda^2}{n}\right) = O_P\left(\frac{\lambda^2 \log k}{\rho n}\right).$$

On the other hand, again by Lemma 12 from Appendix C, the probability that (11) does not hold is $O(\lambda/n)$, showing that with probability $1 - 4e^{-n} - O(\lambda/n) = 1 - O(\lambda/n)$,

$$\hat{\delta}_2\left(\frac{1}{\rho(G)}\hat{B}, H_n(W)\right) \leq \hat{\epsilon}_k^{(O)}(H_n(W)) + O_P\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)}\right).$$

This holds conditioned on an event E of probability $O(\lambda/n)$. To bound the contribution of E to the overall error, we bound $\frac{1}{\rho(G)}\|\hat{B}\|_2 \leq \frac{1}{\rho(G)}\|\hat{B}\|_\infty \leq \lambda$ and $\|H_n(W)\|_2 \leq \|H_n(W)\|_\infty \leq \lambda$, giving an error contribution of $O_P(\lambda^2/n) = O_P(\sqrt[4]{\lambda^2/\rho n})$ which we can absorb into the error already present.

To prove the almost sure statement, we use that $\epsilon_n(W) \to 0$ almost surely, which by Lemma 12 (part 2) from Appendix C implies that $\rho(Q)/\rho \to 1$ almost surely. Since the error probability in (13) is exponentially small, we can use the Borel-Cantelli Lemma to obtain the a.s. statement. □

4.3 Estimation of the graphon W

To deduce Theorem 2 from Theorem 4, we will bound $\hat{\epsilon}_k^{(O)}(H_n(W))$ in terms of $\epsilon_k^{(O)}(W)$, and $\delta_2(\hat{B}/\rho(G), W)$ in terms of $\hat{\delta}_2(\hat{B}/\rho(G), H_n(W))$. We will show that the leading error in both cases is an additive error of $\epsilon_n(W)$. To do this, we need two lemmas.

Lemma 6. *Fix n and $k \leq n$.*

(i) For each equipartition $\pi : [n] \to [k]$ and each permutation $\sigma : [n] \to [n]$, $\pi \circ \sigma$ is an equipartition.

(ii) For all equipartitions $\pi, \pi' : [n] \to [k]$ there exists a permutation $\sigma : [n] \to [n]$ such that $\pi' = \pi \circ \sigma$.

Proof. Any equipartition must have exactly L_- classes of size $\lfloor n/k \rfloor$ and L_+ classes of size $\lceil n/k \rceil$, where $L_\pm$ are determined by the equations $L_- + L_+ = k$, $L_-\lfloor n/k \rfloor + L_+\lceil n/k \rceil = n$; and any partition with these properties is an equipartition. The statement follows. □

To state the next lemma, we define the *standard equipartition* π of $[n]$ into k classes to be the partition into the classes $I_i = \{n_{i-1} + 1, \ldots, n_i\}$, $i \in [k]$, where $n_i = \lfloor in/k \rfloor$. Note that $n_0 = 0$, $n_k = n$, and $\lfloor n/k \rfloor \leq |n_{i+1} - n_i| \leq \lceil n/k \rceil$.

Lemma 7. *Let B be a symmetric $k \times k$ matrix with nonnegative entries, and let π be the standard equipartition of $[n]$ into k classes. Then*

$$\|W[B] - W[B_\pi]\|_2 \leq \sqrt{\frac{4k}{n}}\|B\|_2.$$

Proof. Let $I_1, \ldots, I_k$ be adjacent intervals of length $1/k$, and for $i = 1, \ldots, k$, let J_i be the set of point in $x \in I_i$ such that $x \leq n_i/n$, and let $\Delta_i = I_i \setminus J_i$. Then $(W[B] - W[B_\pi])(x, y) = 0$ unless (x, y) lies in one of the $3k^2$ sets $R_{ij}^{(1)} = \Delta_i \times \Delta_j$, $R_{ij}^{(2)} = I_i \setminus \Delta_i \times \Delta_j$ or $R_{ij}^{(3)} = \Delta_i \times I_j \setminus \Delta_j$, $i, j \in [k]$. Taking, e.g., $(x, y) \in R_{ij}^{(1)}$ we have that $|(W[B] - W[B_\pi])(x, y)|^2 = |B_{ij} - B_{i+1,j+1}|^2 \leq B_{ij}^2 + B_{i+1,j+1}^2$ (note that the set Δ_k is empty, so that here we only have to consider $i, j \leq k - 1$). In a similar

way, the difference in $R^{(2)}_{ij}$ is bounded by $B^2_{ij} + B^2_{i+1,j}$, and the difference in $R^{(3)}_{ij}$ is bounded by $B^2_{ij} + B^2_{i,j+1}$. The total contribution of all these sets can then be bounded by

$$\sum_{ij} B^2_{ij}(\frac{2}{nk} - \frac{1}{n^2}) + \sum_{ij} B^2_{i,j+1}\frac{1}{nk} + \sum_{ij} B^2_{i+1,j}\frac{1}{nk} + \sum_{ij} B^2_{i+1,j+1}\frac{1}{n^2} \leq \frac{4}{nk}\sum_{ij} B^2_{ij} = \frac{4k}{n}\|B\|_2^2.$$

□

Proof of Theorem 2. We start by bounding $\delta_2(\hat{B}/\rho(G), W)$. Let $\pi : [n] \to [k]$ be a standard equipartition, and let $(I_1, \dots, I_n)$ be a partition of $[0,1]$ into adjacent intervals of lengths $1/n$. By the triangle inequality, the fact that the set of measure preserving bijections $\pi : [0,1] \to [0,1]$ contains all bijections which just permute the intervals $I_1, \dots, I_n$ and Lemma 6

$$\begin{aligned}\delta_2\left(\frac{1}{\rho(G)}\hat{B}, W\right) &\leq \frac{1}{\rho(G)}\delta_2(W[\hat{B}], W[\hat{B}_\pi]) + \delta_2\left(\frac{1}{\rho(G)}W[\hat{B}], W[H_n(W)]\right) + \delta_2(W[H_n(W)], W) \\ &\leq \frac{1}{\rho(G)}\|W[\hat{B}] - W[\hat{B}_\pi]\|_2 + \hat{\delta}_2\left(\frac{1}{\rho(G)}\hat{B}, H_n(W)\right) + \hat{\delta}_2(H_n(W), W).\end{aligned}$$

The third term is equal to $\epsilon_n(W)$. To bound the first term, we first condition on the event (11), and then use (9) together with Lemma 7 to conclude that conditioned on (11), with probability at least $1 - 4e^{-n}$,

$$\frac{1}{\rho(G)}\delta_2(W[\hat{B}], W[\hat{B}_\pi]) \leq O\left(\frac{\rho}{\rho(G)}\sqrt{\frac{\lambda k}{n}}\right) = O\left(\sqrt[4]{\lambda^2\frac{k^2}{\rho n^2}}\right).$$

In view of Lemma 12 from Appendix C, the probability that this bound does not hold is bounded by $4e^{-n} + O(\lambda/n) = O(\lambda/n)$, so in view of the fact that $\hat{B} \in \mathcal{B}_\mu$, which shows that $\|\hat{B}\|_2/\rho(G) \leq \lambda$, we see that the contribution of the failure event is again bounded by $O_P(\lambda^2/n) = O\left(\lambda^2\frac{k}{\rho n}\right) = O\left(\sqrt[4]{\lambda^2\frac{k}{\rho n}}\right)$. All together, this proves that

$$\delta_2\left(\frac{1}{\rho(G)}\hat{B}, W\right) \leq \hat{\delta}_2\left(\frac{1}{\rho(G)}\hat{B}, H_n(W)\right) + \epsilon_n(W) + O_P\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)}\right). \tag{14}$$

The corresponding a.s. bound follows again from the fact that $\rho(Q) \to \rho$ a.s., and the fact that all other failure probabilities are exponentially small.

Next fix B such that it is a minimizer in (5). That implies that B is obtained from W by averaging over a partition of W into k classes, which in particular implies that $\|B\|_2 \leq \|W\|_2 \leq \sqrt{\|W\|_\infty\|W\|_1} \leq \sqrt{\lambda}$. Together with Lemma 7 this implies that there is an equipartition $\pi : [n] \to [k]$ such that

$$\begin{aligned}\epsilon^{(O)}_k(W) &\geq \|W - W[B]\|_2 \\ &\geq \|W - W[B_\pi]\|_2 - \sqrt{\frac{4k}{\lambda}n} \\ &\geq \hat{\delta}_2(B_\pi, W) - \sqrt{\frac{4k}{\lambda}n}\end{aligned}$$

Using Lemma 6 to express $\hat{\delta}_2(B, H_n(W))$ as a minimum over permutations $\sigma : [n] \to [n]$, we then bound

$$\begin{aligned}\hat{\epsilon}_k^{(O)}(W) \leq \hat{\delta}_2(B, H_n(W)) &= \min_\sigma \|B_\pi - [H_n(W)]^\sigma\|_2 \\ &\leq \|W[B_\pi] - W\|_2 + \hat{\delta}_2(H_n(W), W) \\ &\leq \epsilon_k^{(O)}(W) + \epsilon_n(W) + \sqrt{\frac{4k}{\lambda}}n,\end{aligned}$$

where in the first line we use $[H_n(W)]^\sigma$ to denote the matrix with entries $[H_n(W)]_{\sigma(x),\sigma(y)}$. Together with (14) this completes the proof of the theorem. □

5 Analysis of the Private Algorithm

In this section we prove consistency of the private algorithms. Our analysis relies on some basic results on differentially private algorithms from previous work, which are collected in Appendix B.

Compared to the analysis of the non-private algorithms, we need to control several additional error sources which were not present for the nonprivate algorithm. In particular, we will have to control the error between $\hat{\rho}$ and $\rho(G)$, the fact that the algorithm (approximately) maximizes $\widehat{\text{score}}(B; G)$ instead of $Score(B; Q)$, and the error introduced by the exponential sampling error. The necessary bounds are given by the following lemma. To state it, we denote the maximal degree in G by $d_{\max}(G)$.

Lemma 8. *Let $(\hat{\rho}, \hat{B})$ be the output of the randomized Algorithm 1. Then the following properties hold with probability at least $1 - 2e^{-n\rho\epsilon/16}$ with respect to the coin flips of the algorithm:*
1) $|\rho(G) - \hat{\rho}| \leq \rho/4$.
2) If $d_{max}(G) \leq \lambda\rho/4$ and $\rho(G) \geq \rho/2$, then

$$Score(\hat{B}; G) \geq \max_{B \in \mathcal{B}_\mu} Score(B; G) - \frac{16\lambda^2\hat{\rho}^2(k^2+1)\log n}{n\epsilon}.$$

Proof. Observing that $\Pr\{|\text{Lap}(4/n\epsilon)| \geq x\} = \exp(-xn\epsilon/4)$, we get that

$$\Pr\left(|\rho(G) - \hat{\rho}| \geq \delta\rho\right) = e^{-\delta n\rho\epsilon/4}, \tag{15}$$

which immediately gives (1).

To prove (2), we first use (1) and the assumptions on $\rho(G)$ and $d_{\max}(G)$ to bound

$$\lambda\hat{\rho} \geq \lambda(\rho(G) - \rho/4) \geq \lambda\rho/4 \geq d_{\max}(G).$$

This implies that the extended score is equal to the original score.

We conclude the proof by using Lemma 11 to show that with probability at least $1 - e^{-n} \geq 1 - e^{-n\rho\epsilon/16}$, the exponential mechanism returns a matrix $\hat{B}$ such that

$$Score(\hat{B}; G) \geq \max_{B \in \mathcal{B}_\mu} Score(B; G) - \frac{4\Delta\log(|\mathcal{B}_\mu|)}{\epsilon}.$$

where $\Delta = \Delta = \frac{4d\mu}{n^2} = \frac{4\lambda^2\hat{\rho}^2}{n}$. Bounding $|\mathcal{B}_\mu| \leq n^{k^2}$, this completes the proof of the lemma. □

Theorem 1 will follow from the following theorem in the same way as Theorem 2 followed from Theorem 4.

Theorem 5. *Under the assumptions of Theorem 1,*

$$\hat{\delta}_2\Big(\frac{1}{\hat{\rho}}\hat{B}, H_n(W)\Big) \le \hat{\epsilon}_k^{(O)}(H_n(W)) + O_P\left(\sqrt[4]{\frac{\lambda^2 \log k}{\rho n}} + \lambda\sqrt{\frac{k^2 \log n}{n\epsilon}} + \frac{\lambda}{n\rho\epsilon}\right). \tag{16}$$

Moreover, if we replace the assumption $n\rho \ge 6\log n$ in Theorem 1 by the stronger assumption $n\rho\epsilon/\log n \to \infty$, then a.s. as $n \to \infty$,

$$\hat{\delta}_2\left(\frac{1}{\hat{\rho}}\hat{B}, H_n(W)\right) \le \hat{\epsilon}_k^{(O)}(H_n(W)) + O\left(\sqrt[4]{\frac{\lambda^2 \log k}{\rho n}} + \lambda\sqrt{\frac{k^2 \log n}{n\epsilon}} + \frac{\sqrt{\lambda}}{n\rho\epsilon}\right) + o(1).$$

Proof of Theorem 5. With probability at least $1 - e^{-n\rho\epsilon/16}$, we may assume that the output of the private algorithm obeys the conclusions of Lemma 8. With a decrement in probability of at most $P_n = O(\Lambda/n)$, we then have that

$$\frac{\rho}{2} \le \rho(G) \le 2\rho, \qquad \frac{\rho}{2} \le \rho(Q) \le 2\rho \qquad \text{and} \qquad \frac{\rho}{4} \le \hat{\rho} \le 3\rho. \tag{17}$$

Next use the assumption $\rho n \ge 6\log n$, the fact that $1 \le \Lambda \le \lambda/8$, and Lemma 13 from Appendix C with $\beta = \lambda/(8\Lambda)$ to show that at a decrement in probability of at most $e^{\log n - \frac{1}{24}\lambda\rho n} \le e^{-\frac{1}{48}\lambda\rho n}$, the maximal degree in G is at most $(\Lambda + \frac{\lambda}{8})\rho \le \frac{\lambda\rho}{4}$. Lemma 8 then allows us to use Proposition 1 with

$$\nu = \sqrt{\frac{16\lambda^2\hat{\rho}^2(k^2+1)\log n}{n\epsilon}} = O\left(\lambda\rho\sqrt{\frac{k^2 \log n}{n\epsilon}}\right).$$

This introduces an additional error term $\frac{\nu}{\rho} = O\left(\lambda\sqrt{\frac{k^2\log n}{n\epsilon}}\right)$ into the bound (12) and an extra error term of order $O\left(\lambda\rho\sqrt{\frac{k^2\log n}{n\epsilon}}\right)$ in the upper bound (9), leading to the estimate that, with probability at least $1 - 4e^{-n} - P_n - e^{-\frac{1}{48}\lambda\rho n} - e^{-n\rho\epsilon/16} = 1 - O(\Lambda/n) - e^{-\Omega(n\rho\epsilon)}$,

$$\hat{\delta}_2\left(\frac{1}{\rho}\hat{B}, H_n(W)\right) \le \min_{B\in\mathcal{B}_\mu} \hat{\delta}_2\left(\frac{1}{\rho}B, H_n(W)\right) + O\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)} + \lambda\sqrt{\frac{k^2\log n}{n\epsilon}}\right)$$

and

$$\|\hat{B}\|_2 = O\left(\sqrt{\lambda}\rho + \lambda\rho\sqrt{\frac{k^2\log n}{n\epsilon}}\right).$$

From here on we proceed as in the proof of (13), except that we now move from the minimizer B' for $\hat{\epsilon}_k^{(O)}(H_n(W))$ to a matrix $B \in \mathcal{B}_\mu$ by rounding the entries of $\hat{\rho}B'$ down to the nearest multiple of $1/n$. Instead of (13), we now obtain the bound

$$\hat{\delta}_2\left(\frac{1}{\hat{\rho}}\hat{B}, H_n(W)\right) \le \hat{\epsilon}_k^{(O)}(H_n(W)) + O\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)} + \lambda\sqrt{\frac{k^2\log n}{n\epsilon}}\right) + O\left(\sqrt{\lambda}\Big|\frac{\hat{\rho}}{\rho} - 1\Big|\right), \tag{18}$$

a bound which is valid with probability at least $1 - 4e^{-n} - P_n - e^{-\frac{1}{48}\lambda\rho n} - e^{-n\rho\epsilon/16}$. Now the fact that $|\mathrm{Lap}(4/n\epsilon)| = O_P(\frac{1}{n\epsilon})$ and $\rho(G) = \rho(1 + O_P(\sqrt{\Lambda/n}))$ implies that

$$\left|1 - \frac{\hat{\rho}}{\rho}\right|\sqrt{\lambda} = O_P\left(\frac{\lambda}{\sqrt{n}}\right) + O_P\left(\frac{\sqrt{\lambda}}{n\rho\epsilon}\right) = O_P\left(\sqrt[4]{\frac{\lambda^2}{n}}\right) + O_P\left(\frac{\sqrt{\lambda}}{n\rho\epsilon}\right).$$

Combining this with (18), we obtain that with probability at least $1-4e^{-n}-P_n-e^{-\frac{1}{48}\lambda\rho n}-e^{-n\rho\epsilon/16}$,

$$\hat{\delta}_2\left(\frac{1}{\hat{\rho}}\hat{B}, H_n(W)\right) \leq \hat{\epsilon}_k^{(O)}(H_n(W)) + O_P\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)} + \lambda\sqrt{\frac{k^2\log n}{n\epsilon}} + \frac{\sqrt{\lambda}}{n\rho\epsilon}\right). \tag{19}$$

The contribution of the failure event can now be bounded by

$$O_P\left(\lambda\left(e^{-n} + \frac{\lambda}{n} + e^{-n\rho\lambda/48} + e^{-n\rho\epsilon/16}\right)\right) = O_P\left(\frac{\lambda^2}{n} + \frac{\lambda}{n\rho\epsilon}\right) \tag{20}$$

To complete the proof of the bound in probability, we have to add the error terms from (19) and (20). We can simplify the resulting expression somewhat by first noting that the left hand side of Eq. (19) is of order at most λ, which shows that for the bound on $\hat{\delta}_2$ not to be vacuous, we need $\frac{k^2}{n} \leq \frac{k^2\log n}{n\epsilon} \leq 1$. We can therefore drop the term $\frac{\lambda^2 k^2}{\rho n^2}$ inside the fourth root of (19). Furthermore, by the assumption of the Theorem, $\lambda \leq \sqrt{n}$, which shows that the first term in (20) is $O(\sqrt[4]{\lambda^2/n})$ and can hence be absorbed into the error terms in (19). This gives us the main theorem statement.

To prove bounds which hold a.s., we note that for $n\rho\epsilon/\log n \to \infty$, the error probability in (15) is summable (that is, the probability of error p_n satisfies $\sum_{n=1}^{\infty} p_n < \infty$) for all $\delta > 0$, which together with our previous results implies that $\hat{\rho}/\rho \to 1$ with probability one. Since the probability of failure for all other events necessary for (19) to hold is summable as well, we get that a.s.,

$$\hat{\delta}_2\left(\frac{1}{\hat{\rho}}\hat{B}, H_n(W)\right) \leq \hat{\epsilon}_k^{(O)}(H_n(W)) + O\left(\sqrt[4]{\frac{\lambda^2\log k}{\rho n}} + \lambda\sqrt{\frac{k^2\log n}{n\epsilon}} + \frac{\sqrt{\lambda}}{n\rho\epsilon}\right) + o(1),$$

where again $o(1)$ is a term which goes to zero with probability one as $n \to \infty$. □

Proof of Theorem 1. The proof of Theorem 1 follows from Theorem 5 in essentially same way as Theorem 2 followed from Theorem 4. The only modification needed is that we now have to bound $\frac{1}{\hat{\rho}}\delta_2(W[\hat{B}], W[\hat{B}_\pi])$ instead of $\frac{1}{\rho(G)}\delta_2(W[\hat{B}], W[\hat{B}_\pi])$. But this is even easier, since here we won't need to distinguish several cases. Instead, we just use that $\hat{B} \in \mathcal{B}_\mu$ implies $\|\hat{B}\|_\infty \leq \lambda\hat{\rho}$. With the help of Lemma 7, we then bound this error term by $\lambda\sqrt{\frac{4k}{n}}$, a term which can be incorporated into the error term $O\left(\lambda\sqrt{\frac{k^2\log n}{n\epsilon}}\right)$. □

References

[1] E. Abbe and C. Sandon. Recovering communities in the general stochastic block model without knowing the parameters. arXiv:1503.00609, 2015.

[2] E. Abbe and C. Sandon. Recovering communities in the general stochastic block model without knowing the parameters. Manuscript, 2015.

[3] E. Abbe, A. S. Bandeira, and G. Hall. Exact recovery in the stochastic block model. arXiv:1405.3267, 2014.

[4] E. M. Airoldi, T. Costa, and S. Chan. A non-parametric perspective on network analysis: Theory and consistent estimation. In *Advances in Neural Information Processing Systems (NIPS)*, volume 26, pages 692–700, 2013.

[5] D. Aldous. Representations for partially exchangeable arrays of random variables. *J. Multivar. Anal.*, 11:581–598, 1981.

[6] L. Backstrom, C. Dwork, and J. Kleinberg. Wherefore art thou r3579x? anonymized social networks, hidden patterns, and structural steganography. In *Proc. 16th Intl. World Wide Web Conference*, pages 181–190, 2007.

[7] P. J. Bickel and A. Chen. A nonparametric view of network models and newman-girvan and other modularities. *Proceedings of the National Academy of Sciences of the United States of America*, 106: 21068–21073, 2009.

[8] P. J. Bickel, A. Chen, and E. Levina. The method of moments and degree distributions for network models. *Annals of Statistics*, 39(5):2280–2301, 2011.

[9] J. Blocki, A. Blum, A. Datta, and O. Sheffet. The johnson-lindenstrauss transform itself preserves differential privacy. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 410–419. IEEE Computer Society, 2012. ISBN 978-1-4673-4383-1. doi: 10.1109/FOCS.2012.67. URL `http://dx.doi.org/10.1109/FOCS.2012.67`.

[10] J. Blocki, A. Blum, A. Datta, and O. Sheffet. Differentially private data analysis of social networks via restricted sensitivity. In R. D. Kleinberg, editor, *ITCS*, pages 87–96. ACM, 2013.

[11] B. Bollobas and O. Riordan. Metrics for sparse graphs. In *Surveys in combinatorics 2009 (eds. S. Huczynska, J. D. Mitchell, and C. M. Roney-Dougal)*, pages 211–287. London Math. Soc. Lecture Note Ser. **365**, Cambridge University Press, 2009.

[12] B. Bollobas, S. Janson, and O. Riordan. The phase transition in inhomogeneous random graphs. *Random Struct. Algorithms*, 31:3–122, 2007.

[13] C. Borgs, J. T. Chayes, L. Lovász, V. Sós, and K. Vesztergombi. Counting graph homomorphisms. In *Topics in Discrete Mathematics (eds. M. Klazar, J. Kratochvil, M. Loebl, J. Matousek, R. Thomas, P. Valtr)*, pages 315–371. Springer, 2006.

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

[15] C. Borgs, J. T. Chayes, and L. Lovász. Moments of two-variable functions and the uniqueness of graph limits. *Geometric And Functional Analysis*, 19(6):1597–1619, 2010.

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

[17] C. Borgs, J. T. Chayes, H. Cohn, and Y. Zhao. An L^p theory of sparse graph convergence I: limits, sparse random graph models, and power law distributions. *arXiv:1401.2906*, 2014.

[18] C. Borgs, J. T. Chayes, H. Cohn, and Y. Zhao. An L^p theory of sparse graph convergence II: LD convergence, quotients, and right convergence. *arXiv:1408.0744*, 2014.

[19] F. Caron and E. Fox. Sparse graphs using exchangeable random measures. *arXiv:1401.1137*, 2015.

[20] S. H. Chan and E. M. Airoldi. A consistent histogram estimator for exchangeable graph models. *Journal of Machine Learning Research Workshop and Conference Proceedings*, 32:208–216, 2014.

[21] S. Chatterjee. Matrix estimation by universal singular value thresholding. *Annals of Statistics*, 43(1): 177–214, 2015.

[22] S. Chen and S. Zhou. Recursive mechanism: towards node differential privacy and unrestricted joins. In Ross et al. [57], pages 653–664.

[23] D. S. Choi, P. J. Wolfe, and E. M. Airoldi. Stochastic blockmodels with a growing number of classes. *Biometrika*, 99:273–284, 2012.

[24] P. Diaconis and S. Janson. Graph limits and exchangeable random graphs. *Rendiconti di Matematica*, 28:33—61, 2008.

[25] I. Dinur and K. Nissim. Revealing information while preserving privacy. In *PODS*, pages 202–210. ACM, 2003.

[26] C. Dwork and J. Lei. Differential privacy and robust statistics. In M. Mitzenmacher, editor, *STOC*, pages 371–380. ACM, 2009. ISBN 978-1-60558-506-2.

[27] C. Dwork and A. Roth. *The Algorithmic Foundations of Differential Privacy*. Foundations and Trends in Theoretical Comptuer Science. Now Publishers Inc., 2014. URL `http://www.cis.upenn.edu/~aaroth/privacybook.html`.

[28] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In S. Halevi and T. Rabin, editors, *TCC*, volume 3876, pages 265–284, 2006.

[29] A. Frieze and R. Kannan. Quick approximation to matrices and applications. *Combinatorica*, 19: 175–220, 1999.

[30] C. Gao, Y. Lu, and H. H. Zhou. Rate-optimal graphon estimation. *arXiv:1410.5837*, 2014.

[31] A. Gupta, A. Roth, and J. Ullman. Iterative constructions and private data release. In *TCC*, 2012.

[32] M. Hay, C. Li, G. Miklau, and D. Jensen. Accurate estimation of the degree distribution of private networks. In W. W. 0010, H. Kargupta, S. Ranka, P. S. Yu, and X. Wu, editors, *Int. Conf. Data Mining (ICDM)*, pages 169–178. IEEE Computer Society, 2009. ISBN 978-0-7695-3895-2.

[33] M. Hay, V. Rastogi, G. Miklau, and D. Suciu. Boosting the Accuracy of Differentially Private Histograms Through Consistency. *PVLDB*, 3(1):1021–1032, 2010.

[34] P. D. Hoff, A. E. Raftery, and M. S. Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460):1090–1098, 2002.

[35] P. Holland, K. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Soc Netw*, 5:109–137, 1983.

[36] D. Hoover. Relations on probability spaces and arrays of random variables. *Preprint, Institute for Advanced Study, Princeton, NJ*, 1979.

[37] V. Karwa and A. Slavkovic. Inference using noisy degrees: Differentially private -model and synthetic graphs. *stat.ME*, arXiv:1205.4697v3 [stat.ME], 2014.

[38] V. Karwa and A. B. Slavkovic. Differentially private graphical degree sequences and synthetic graphs. In J. Domingo-Ferrer and I. Tinnirello, editors, *Privacy in Statistical Databases*, volume 7556 of *Lecture Notes in Computer Science*, pages 273–285. Springer, 2012. ISBN 978-3-642-33626-3.

[39] V. Karwa, S. Raskhodnikova, A. Smith, and G. Yaroslavtsev. Private analysis of graph structure. *PVLDB*, 4(11):1146–1157, 2011.

[40] V. Karwa, A. B. Slavkovic, and P. N. Krivitsky. Differentially private exponential random graphs. In J. Domingo-Ferrer, editor, *Privacy in Statistical Databases - UNESCO Chair in Data Privacy, International Conference, PSD 2014, Ibiza, Spain, September 17-19, 2014. Proceedings*, volume 8744 of *Lecture Notes in Computer Science*, pages 143–155. Springer, 2014. ISBN 978-3-319-11256-5. doi: 10.1007/978-3-319-11257-2_12. URL `http://dx.doi.org/10.1007/978-3-319-11257-2_12`.

[41] S. P. Kasiviswanathan, K. Nissim, S. Raskhodnikova, and A. Smith. Analyzing graphs with node-differential privacy. In *Theory of Cryptography Conference (TCC)*, pages 457–476, 2013.

[42] M. Kirszbraun. Über die zusammenziehende und lipschitzsche transformationen. *Fundamenta Mathematicae*, 22(1):77–108, 1934. URL `http://eudml.org/doc/212681`.

[43] P. Latouche and S. Robin. Bayesian model averaging of stochastic block models to estimate the graphon function and motif frequencies in a w-graph model. *ArXiv:1310.6150*, 2013.

[44] B.-R. Lin and D. Kifer. Information preservation in statistical privacy and Bayesian estimation of unattributed histograms. In Ross et al. [57], pages 677–688.

[45] J. R. Lloyd, P. Orbanz, Z. Ghahramani, and D. M. Roy. Random function priors for exchangeable arrays with applications to graphs and relational data. In *Advances in Neural Information Processing Systems (NIPS)*, volume 25, pages 1007–1015, 2012.

[46] L. Lovász and B. Szegedy. Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B*, 96:933–957, 2006.

[47] W. Lu and G. Miklau. Exponential random graph estimation under differential privacy. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 921–930. ACM, 2014.

[48] E. J. McShane. Extension of range of functions. *Bull. Amer. Math. Soc.*, 40(12):837–842, 1934.

[49] F. McSherry and I. Mironov. Differentially private recommender systems: building privacy into the net. In *Symp. Knowledge Discovery and Datamining (KDD)*, pages 627–636. ACM New York, NY, USA, 2009.

[50] F. McSherry and K. Talwar. Mechanism design via differential privacy. In *FOCS*, pages 94–103. IEEE, 2007.

[51] D. J. Mir and R. N. Wright. A differentially private estimator for the stochastic kronecker graph model. In D. Srivastava and I. Ari, editors, *EDBT/ICDT Workshops*, pages 167–176. ACM, 2012. ISBN 978-1-4503-1143-4.

[52] A. Narayanan and V. Shmatikov. De-anonymizing social networks. In *IEEE Symp. Security and Privacy*, pages 173–187, 2009.

[53] K. Nissim, S. Raskhodnikova, and A. Smith. Smooth sensitivity and sampling in private data analysis. In *Symp. Theory of Computing (STOC)*, pages 75–84. ACM, 2007. Full paper: `http://www.cse.psu.edu/~asmith/pubs/NRS07`.

[54] S. Raskhodnikova and A. Smith. High-dimensional lipschitz extensions and node-private analysis of network data. *arXiv:1504.07912*, 2015.

[55] V. Rastogi, M. Hay, G. Miklau, and D. Suciu. Relationship privacy: output perturbation for queries with joins. In *Symp. Principles of Database Systems (PODS)*, pages 107–116, 2009.

[56] K. Rohe, S. Chatterjee, and B. Yu. Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Statist.*, 39(4):1878–1915, 08 2011.

[57] K. A. Ross, D. Srivastava, and D. Papadias, editors. *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2013, New York, NY, USA, June 22-27, 2013*, 2013. ACM.

[58] M. Tang, D. L. Sussman, and C. E. Priebe. Universally consistent vertex classification for latent positions graphs. *Ann. Statist.*, 41(3):1406–1430, 06 2013. doi: 10.1214/13-AOS1112. URL `http://dx.doi.org/10.1214/13-AOS1112`.

[59] P. Wolfe and S. C. Olhede. Nonparametric graphon estimation. *arXiv:1309.5936*, 2013.

[60] Q. Xiao, R. Chen, and K. Tan. Differentially private network data release via structural inference. In S. A. Macskassy, C. Perlich, J. Leskovec, W. Wang, and R. Ghani, editors, *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014*, pages 911–920. ACM, 2014. ISBN 978-1-4503-2956-9. doi: 10.1145/2623330.2623642. URL `http://doi.acm.org/10.1145/2623330.2623642`.

[61] J. J. Yang, Q. Han, and E. M. Airoldi. Nonparametric estimation and testing of exchangeable graph models. In *Proceedings of 17th AISTATS (JMLR: W&CP volume 33)*, 2014.

A Comparison to Nonprivate Bounds from Previous Work

The most relevant previous works are those of Wolfe and Olhede [59], Chatterjee [21] and [30]. We provide comparisons for two types of bounded graphons: (1) k-block graphons, and (2) α-Hölder graphons.

k-block graphons. A k-block graphon is a function on $[0,1]^2$ that is constant on rectangles of the form $I_i \times I_j$, where $I_1, ..., I_k$ form a partition of the interval $[0,1]$.

In this setting, the oracle error $\epsilon_k^{(O)}(W) = 0$ and $\epsilon_n(W) = O_P(\sqrt[4]{k/n})$ (see Appendix D). Our nonprivate estimator then has asymptotic error $\sqrt[4]{\frac{k}{n} + \frac{\log k}{\rho n} + \frac{k^2}{\rho n^2}}$; this is dominated by $\epsilon_n(W)$ as long as $k \leq \rho n$ and $\rho \geq \frac{\log k}{k}$. For constant ϵ, our private estimator has asymptotic error at most $\sqrt[4]{\frac{k}{n} + \frac{\log k}{\rho n} + \frac{k^4 \log^3 n}{n^2}}$. For $k \leq (n/\log^2 n)^{1/3}$ and density $\rho \geq \frac{\log k}{k}$, the error of our estimator is again dominated by the (unavoidable) error of $\epsilon_n(W) = \sqrt[4]{k/n}$.

Several works analyze procedures for estimating the edge-probability matrix Q assuming that it is (exactly) a k-block matrix. In the dense case $\rho = \Omega(1)$, Gao et al. [30, Theorem 1.1] show that the least squares estimator achieves error $\|\frac{1}{\rho}\hat{Q} - \frac{1}{\rho}Q\|_2 = O_P(\frac{k}{n} + \sqrt{\frac{\log k}{n}})$. They also give a matching lower bound, which shows that the MLE is optimal with respect to ℓ_2 estimation of Q. Chatterjee [21, Theorem 2.3] gives a polynomial-time algorithm with higher error $\|\frac{1}{\rho}\hat{Q} - \frac{1}{\rho}Q\|_2 = O_P(\sqrt[4]{\frac{k}{n}})$.

These bounds apply to estimating the edge-probability matrix Q, but do not apply directly to estimating an underlying block graphon W. Lemma 14 shows that $\frac{1}{\rho}Q$ converges to W in the δ_2 metric at a rate of $O_P(\sqrt[4]{k/n})$. Using either of the algorithms above for estimating W gives a net error rate of $O_P(\sqrt[4]{k/n})$ for δ_2 estimation of W. This is the best known nonprivate rate, and is matched by our nonprivate rate.

In the sparse case, where $\rho \to 0$ as $n \to \infty$, Wolfe and Olhede showed under additional assumptions (roughly, that entries of Q are bounded above and below by multiples of ρ) that the MLE produces an estimate $\hat{Q}$ of Q that satisfies $\|\frac{1}{\rho}\hat{Q} - \frac{1}{\rho}Q\|_2 = O_P\left(\frac{k}{n} \cdot \sqrt{\frac{\log(n)}{\rho}} + \sqrt[4]{\frac{\log^2(1/\rho)\log(k)}{n\rho}}\right)$ [59, Theorem 5.1][1]. Again, one can combine these with Lemma 14 to get a rate of $\sqrt[4]{\frac{k}{n}} + \frac{k}{n}\sqrt{\frac{\log(n)}{\rho}} + \sqrt[4]{\frac{\log^2(1/\rho)\log(k)}{n\rho}}$ for estimating an underlying k-block graphon W. Note that when ρ is small, any of these three terms may dominate the rate.

Hölder-continuous graphons. The known algorithms for estimating continuous graphons proceed by fitting a k-block model to the observed data, and arguing that this model approximates the underlying graphon.

Our results show that if ϵ is constant and W is α-Hölder continuous (Lipschitz continuity corresponds to $\alpha = 1$), then the nonprivate error scales as $\left(\frac{1}{n\sqrt{\rho}}\right)^{\frac{\alpha}{2\alpha+1}} + \sqrt[4]{\frac{\log n}{\rho n}} + n^{-\alpha/2}$ for an appropriate choice of k, while the private error scales as $\left(\frac{\log n}{n}\right)^{\frac{\alpha}{2\alpha+2}} + \sqrt[4]{\frac{\log n}{\rho n}} + n^{-\alpha/2}$ for an appropriate choice of k. See Remark 2 for details.

[1] The guarantee in [59, Theorem 5.1] is given in terms of KL divergence. One can convert to ℓ_2 using the fact that $D(p\|q) = \Theta((q-p)^2/p)$ when $q - p$ is small relative to p.

In the dense case ($\rho = \Omega(1)$), [30] show that one can estimate a α-Hölder continuous graphon by a k-block graphon with error

$$\delta_2(W, \hat{W}_{LS}) = O_P(\tfrac{k}{n} + \sqrt{\tfrac{\log k}{n}} + k^{-\alpha} + n^{-\alpha/2}),$$

with the last term accounting for the difference between estimating Q and W. Setting k to the optimal value of $k = 1/n^{\alpha+1}$ gives a rate which except for the case $\alpha = 1$ is dominated by the term $\epsilon_n(W) = O(n^{-\alpha/2})$. Our nonprivate bound matches this bound for $\alpha < 1/2$, and is worth for $\alpha > 1/2$, while the private one is always worth.

Wolfe and Olhede [59] analyse the MLE in the sparse case, again restricting to k-block models. They show[2]

$$\delta_2(W, \hat{W}_{MLE}) = O_P\left(\tfrac{k}{n} \cdot \sqrt{\tfrac{\log(n)}{\rho}} + \sqrt[4]{\tfrac{\log^2(1/\rho)\log(k)}{n\rho}} + k^{-\alpha} + \tfrac{\sqrt{\log(n/\rho)}}{n^{\alpha/4}}\right) \tag{21}$$

The value of k that maximizes this expression (asymptotically) can be found by setting the first and third terms to be equal; we get $k = \left(n\sqrt{\rho/\log n}\right)^{\frac{1}{\alpha+1}}$ and a resulting bound of

$$\delta_2(W, \hat{W}_{MLE}) = O_P\left(\left(\tfrac{1}{n} \cdot \sqrt{\tfrac{\log n}{\rho}}\right)^{\frac{\alpha}{\alpha+1}} + \sqrt[4]{\tfrac{\log^2(1/\rho)\log(n)}{n\rho}} + \tfrac{\sqrt{\log(n/\rho)}}{n^{\alpha/4}}\right).$$

Next, note that ρ must be $\Omega(\log^3(n)/n)$ for the middle term to be less than 1. This means that the first term is $O(n^{\frac{-\alpha}{2(\alpha+1)}})$. For $\alpha \leq 1$, this is never larger than the third term. We may therefore simplify the bound to $O_P\left(\sqrt[4]{\tfrac{\log^2(1/\rho)\log(n)}{n\rho}} + \tfrac{\sqrt{\log(n/\rho)}}{n^{\alpha/4}}\right)$, as stated in the introduction.

B Background on Differential Privacy

The notion of node-privacy defined in Section 2.3 "composes" well, in the sense that privacy is preserved (albeit with slowly degrading parameters) even when the adversary gets to see the outcome of an adaptively chosen sequence of differentially private algorithms run on the same data set.

Lemma 9 (Composition, post-processing [49, 26])**.** *If an algorithm $\mathcal{A}$ runs t randomized algorithms $\mathcal{A}_1, \ldots, \mathcal{A}_t$, each of which is ϵ-differentially private, and applies an arbitrary (randomized) algorithm g to their results, i.e., $\mathcal{A}(G) = g(\mathcal{A}_1(G), \ldots, \mathcal{A}_t(G))$, then $\mathcal{A}$ is $t\epsilon$-differentially private. This holds even if for each $i > 1$, $\mathcal{A}_i$ is selected adaptively based on $\mathcal{A}_1(G), \ldots, \mathcal{A}_{i-1}(G)$.*

Output Perturbation. One common method for obtaining efficient differentially private algorithms for approximating real-valued functions is based on adding a small amount of random noise to the true answer. A *Laplace* random variable with mean 0 and standard deviation $\sqrt{2}\lambda$ has density $h(z) = \frac{1}{2\lambda}e^{-|z|/\lambda}$. We denote it by $\mathrm{Lap}(\lambda)$.

In the most basic framework for achieving differential privacy, Laplace noise is scaled according to the *global sensitivity* of the desired statistic f. This technique extends directly to graphs as long

[2] We state [59, Theorem 3.1] for the special case where one searches over k-block graphons in which all intervals have size $\Theta(k/n)$ (since allowing nonuniformly sized blocks only makes their bounds worse), the original graphon takes values in a range $[\lambda_a, \lambda_b]$ defined by two constants such that $0 < \lambda_a < \lambda_b$, and k is polynomially smaller than n.

as we measure sensitivity with respect to the metric used in the definition of the corresponding variant of differential privacy. Below, we explain this (standard) framework in terms of node privacy. Let $\mathbb{G}$ denote the set of all graphs.

Definition 2 (Global Sensitivity [28]). *The ℓ_1-global node sensitivity of a function $f : \mathbb{G} \to \mathbb{R}^p$ is:*

$$\Delta f = \max_{G,G' \text{ node neighbors}} \|f(G) - f(G')\|_1 .$$

For example, the edge density of an n-node graph has node sensitivity $2/n$, since adding or deleting a node and its adjacent edges can add or remove at most $n-1$ edges. In contrast, the number of nodes in a graph has node sensitivity 1.

Lemma 10 (Laplace Mechanism [28]). *The algorithm $\mathcal{A}(G) = f(G) + \mathrm{Lap}(\Delta f/\epsilon)^p$ (which adds i.i.d. noise $\mathrm{Lap}(\Delta f/\epsilon)$ to each entry of $f(G)$) is ϵ-node-private.*

Thus, we can release the number of nodes, $v(G)$, in a graph G with noise of expected magnitude $1/\epsilon$ while satisfying node differential privacy. Given a public bound n on $v(G)$, we can release the number of edges, $e(G)$, with additive noise of expected magnitude n/ϵ.

Exponential Mechanism. Sensitivity plays a crucial role in another basic design tool for differentially private algorithms, called the *exponential mechanism*.

Suppose we are given a collection of Q functions, $q_1, .., q_Q$ from $\mathbb{G}_n$ to $\mathbb{R}$, each with sensitivity at most Δ. The exponential mechanism, due to McSherry and Talwar [50], takes a data set (in our case, a graph G) and aims to output the index i^* of a function in the collection which has nearly maximal value at G, that is, such that $q_i^*(G) \approx \max_i q_i(G)$. The algorithm $\mathcal{A}$ samples an index i such that

$$\Pr(\mathcal{A}(G) = i) \propto \exp\left(\frac{\epsilon}{2\Delta} q_i(G)\right).$$

Lemma 11 (Exponential Mechanism [50], see also [27, Sec. 3.4]). *The algorithm $\mathcal{A}$ is ϵ-differentially private. Moreover, with probability at least $1-\eta$, its output i^* satisfies*

$$q_{i^*}(G) \geq \max_i \left(q_i(G)\right) - \frac{2\Delta \ln(Q/\eta)}{\epsilon} .$$

Lipschitz Extensions. There are cases (and we will encounter them in this paper), where the sensitivity of a function can only be guaranteed to be low if the graph in question has sufficiently low degrees. In this situation, it is useful to consider extensions of these functions from graphs obeying a certain degree bound to those without this restriction.

Definition 3 ($\mathbb{G}_{n,d}$ and vertex extensions). *Let $\mathbb{G}_{n,d}$ denote the set of graphs with degree at most d. Given functions $f : \mathbb{G}_{n,d} \to \mathbb{R}$ and $\hat{f} : \mathbb{G}_n \to \mathbb{R}$, we say $\hat{f}$ is a vertex Lipschitz extension of f from $\mathbb{G}_{n,d}$ to $\mathbb{G}_n$ if $\hat{f}$ agrees with f on $\mathbb{G}_{n,d}$ and $\hat{f}$ has the same node-sensitivity as f, that is*

$$\sup_{\substack{G,G' \in \mathbb{G}_n \\ \text{vertex neighbors}}} |\hat{f}(G) - \hat{f}(G')| = \sup_{\substack{G,G' \in \mathbb{G}_{n,d} \\ \text{vertex neighbors}}} |f(G) - f(G')| .$$

We close this section with the proof of Lemma 2.

Proof of Lemma 2. The existence of $\hat{f}$ follows from a very general result (e.g., [48, 42]), which states that for any metric spaces X and Y such that $Y \subset X$, and any Lipschitz function $f : Y \to \mathbb{R}$, there exists an extension $\hat{f} : X \to \mathbb{R}$ with the same Lipschitz constant. The explicit, efficient construction of extensions for linear functions is due to Kasiviswanathan et al. [41]. The idea is to replace $f(G)$ with the maximum of $f(C)$ where C ranges over weighted subgraphs of G with (weighted) degree at most d. It is the value of the following linear program:

$$\hat{f}(G) = \max f(C) \text{ such that } \begin{cases} C \in [0,1]^{n\times n} \text{ is symmetric, and} \\ C_{i,j} \le A(G)_{i,j} \text{ for all } i,j, \text{ and} \\ \sum_{j\neq i} C_{i,j} \le d \text{ for all } i \in [n]\,. \end{cases}$$

See [41] for the analysis of this program's properties. □

C Auxiliary Bounds on Densities and Degrees

Lemma 12. *Let $W : [0,1]^2 \to [0,\Lambda]$ be a normalized graphon, let $\rho \in (0, \Lambda^{-1}]$, let $Q = H_n(\rho W)$, let $G = G_n(\rho W)$ and assume that ρn is bounded away from zero. Then*

1. *$\mathbb{E}\rho(Q) = \mathbb{E}\rho(G) = \rho$, $Var(\rho(Q)) = O(\rho^2\Lambda/n)$ and $Var(\rho(G)) = O(\rho^2\Lambda/n)$, so in particular*

$$\Pr\{|\rho(G) - \rho| \ge \delta\rho\} = O\Big(\frac{\Lambda}{n\delta^2}\Big) \qquad \text{and} \qquad \Pr\{|\rho(Q) - \rho| \ge \delta\rho\} = O\Big(\frac{\Lambda}{n\delta^2}\Big).$$

   for any $\delta > 0$.

2. *Let $\epsilon_n(W) = \|W - W[Q]\|_2$. Then*

$$(1 - \epsilon_n(W))\frac{n}{n-1} \le \frac{\rho(Q)}{\rho} \le (1 + \epsilon_n(W))\frac{n}{n-1}.$$

3. *Let $\delta \in 0 < \delta < 1$. With probability at least $1 - 2e^{-\frac{1}{6}\delta^2\rho(Q)n^2}$,*

$$1 - \delta \le \frac{\rho(G)}{\rho(Q)} \le 1 + \delta.$$

Proof. 1. Clearly, $\mathbb{E}\rho(Q) = \mathbb{E}\rho(G) = \rho$.

To bound the variance of $\rho(Q)$, we expand $\frac{1}{\rho^2}Var(\rho(Q))$ as a sum of $n^2(n-1)^2/4$ terms of the form $\mathbb{E}[W(x_i,x_j)W(x_k,x_\ell)] - \mathbb{E}[W(x_i,x_j)]\mathbb{E}[W(x_k,x_\ell)]$ with $i < j$ and $k < \ell$. Observing that only those terms contribute for which either $i = k$, $j = \ell$ or $j = k$, and bounding $\mathbb{E}[W(x_i,x_j)W(x_k,x_\ell)] \le \|W\|_2^2 \le \|W\|_\infty\|W\|_1$, we obtain that the variance of $\frac{1}{\rho}\rho(Q)$ is $O(\Lambda/n)$.

To bound the variance of $\rho(G)$, we first condition on $X = (x_1,\dots,x_n)$, and bound

$$\mathbb{E}[\rho^2(G) \mid X] = \rho^2(Q) + \frac{4}{n^2(n-1)^2}\sum_{i<j}\big(Q_{ij} - Q_{ij}^2\big) \le \rho^2(Q) + \frac{2}{n(n-1)}\rho(Q).$$

Taking the expectation over X and using the bound on the variance of $\rho(Q)$, we obtain that

$$Var(\rho(G)) = O\Big(\frac{\rho^2\Lambda}{n} + \frac{\rho}{n^2}\Big) = O\Big(\frac{\rho^2\Lambda}{n}\Big).$$

where in the last step we used the assumption that ρn is bounded away from zero.

2. Note that $\rho(Q) = \frac{n}{n-1}\|Q\|_1$. Next, we use the triangle inequality and the fact that the L_1-norm is bounded by the L_2-norm to see that

$$\Big|\frac{\|Q\|_1}{\rho} - 1\Big| = \Big|\frac{\|Q\|_1}{\rho} - \|W\|_1\Big| \le \|\frac{1}{\rho}W[Q] - W\|_1 \le \|\frac{1}{\rho}W[Q] - W\|_2 = \epsilon_n(W)$$

3. Conditioned on X, $S = \frac{n(n-1)}{2}\rho(G)$ is a sum of Bernouilli random variables with mean $E(S) = \frac{n(n-1)}{2}\rho(Q)$. By the multiplicative Chernov bound from Lemma 17, we have that for all $\beta \le 1$

$$\Pr\Big\{|\rho(G) - \rho(Q)| > \delta\rho(Q) \,\Big|\, X\Big\} \le 2\exp\Big(-\frac{\delta^2}{3}\frac{n(n-1)}{2}\rho(Q)\Big) \le 2\exp\Big(-\frac{\delta^2}{6}n^2\rho(Q)\Big).$$

□

Next we bound the maximal degree in $G = G_n(\rho W)$.

Lemma 13. *Let W be a normalized graphon with $\|W\|_\infty \le \Lambda$, let $0 < \rho\Lambda \le 1$, and let $\beta \ge 1$. Then with probability at least $1 - n\exp(-\beta\Lambda\rho n/3)$, the maximal degree in $G_n(\rho W)$ is bounded by $(1+\beta)\Lambda\rho n$.*

Proof. Note that the degrees in G_n are stochastically dominated by those in an Erdos-Renyi random graph where edges are chosen i.i.e. with probability $\rho\Lambda$. In such a graph, the degree of a given vertex i is a sum of $n-1$ i.i.d Bernouilli random variables, and the standard Chernov bound implies that for all $\beta \ge 1$

$$\Pr\big(d_i \ge n\rho\Lambda(1+\beta)\big) \le exp\Big(-\frac{\beta}{3}n\rho\Lambda\Big).$$

Taking the union bound over all n vertices in $G_n(\rho W)$ this proves the claim. □

D Convergence of the edge-probability matrix for k-block graphons

The sampling error $\epsilon_n(W)$ plays a key role in the statements of Theorems 2 and 1. In some settings such as Hölder-continuous graphons, this sampling error is dominated by the oracle error $\epsilon_k^{(O)}(W)$. For k-block graphons, however, $\epsilon_k^{(O)}(W) = 0$ and $\epsilon_n(W)$ becomes more significant.

Lemma 14. *If W is a k-block graphon, then $\mathbb{E}(\epsilon_n(W)^2) = O(\Lambda^2\sqrt{k/n})$ and $\mathbb{E}(\epsilon_n(W)) = O(\Lambda\sqrt[4]{k/n})$ where $\Lambda = \|W\|_\infty$.*

Proof. Fix a k-block graphon W, and let $p_1, ..., p_k$ be the lengths of the "blocks", that is the intervals defining the block representation (so that $p_t \ge 0$ and $\sum_t p_t = 1$). Given a sample $x_1, ..., x_n$ of i.i.d. uniform values in $[0,1]$, let $\hat{p}_t$ denote the fraction of the x_i that land in each block t.

Aligning $Q = (W(x_i, x_j))_{i,j\in[n]}$ with W consists of finding a permutation π of $[n]$. This maps each x_i to one of the intervals $I_1, .., I_n$ where $I_\ell = [\frac{\ell-1}{n}, \frac{\ell}{n}]$.

We say x_i is correctly aligned if its interval $I_{\pi(i)}$ is contained in the block in which x_i landed. For each block t, we can ensure that $n\min\{p_t, \hat{p}_t\} - 2$ of the points x_i that landed in t get aligned with t (the -2 term accounts for the fact that up to $1/n$ of the length at each end of the interval

does not line up exactly with one of the intervals I_ℓ). Thus, the number of points x_i that get incorrectly aligned is at most $n\sum_t(\hat{p}_t - \min\{\hat{p}_t, p_t\} + 2/n) = \frac{n}{2}\|p-\hat{p}\|_1 + 2k$.

Each misaligned point contributes at most $2\Lambda^2/n$ to the total squared error $\|W-Q\|_2^2$, so we have

$$\hat{\delta}_2^2(W,Q) \le \Lambda^2(\|p-\hat{p}\|_1 + 4k/n).$$

Each term $|p_t - \hat{p}_t|$ in the ℓ_1 norm on the right-hand side is the deviation of a binomial from it's mean, and has standard deviation $\sqrt{p_t(1-p_t)/n}$. This upper bounds the expected absolute deviation by Jensen's inequality. Thus, $\mathbb{E}(\hat{\delta}_2^2(W,Q)) \le \Lambda^2\sum_t\sqrt{p_t/n} + 4\Lambda^2 k/n$. The sum $\sum_t\sqrt{p_t/n}$ is maximized when $p_t = 1/k$ for all t; it then takes the value $k\sqrt{1/(kn)} = \sqrt{k/n}$. Hence $\mathbb{E}(\hat{\delta}_2^2(W,Q)) \le \Lambda^2\sqrt{k/n} + 4\Lambda^2 k/n \le 5\Lambda^2\sqrt{k/n}$. By Jensen's inequality, $\mathbb{E}(\delta_2(W,Q)) \le \sqrt{5}\Lambda\sqrt[4]{k/n}$, as desired. □

Corollary 1. *For any graphon W, we have $\mathbb{E}(\epsilon_n(W)) \le 2\epsilon_k^{(O)}(W) + O(\sqrt[4]{k/n})$.*

Proof. Fix a matrix W, and let W_P denote the best k-block approximation to W in the L_2 norm (that is, the minimizer of $\epsilon_k^{(O)}(W)$). Given a uniform i.i.d. sample in $x_1, ..., x_n$, let $H_n(W)$ denote the matrix $(W(x_i,x_j))_{i,j\in[n]}$, and let $H_n(W_P)$ denote $(W_P(x_i,x_j))_{i,j\in[n]}$. By the triangle inequality,

$$\epsilon_n(W) = \hat{\delta}_2(W,H_n) \le \underbrace{\|W-W_P\|_2}_{\epsilon_k^{(O)}(W)} + \epsilon_n(W_P) + \|H_n(W) - H_n(W_P)\|_2.$$

Lemma 14 bounds $\epsilon_n(W_P)$ by $\sqrt[4]{k/n}$. It remains to bound the last term. Squaring it, we have $\|H_n(W) - H_n(W_P)\|_2^2 = \frac{2}{n^2}\sum_{i<j}|(W-W_P)(x_i,x_j)|^2$. These terms are not independent, but each individually has expectation $\|W-W_P\|^2 = \epsilon_k^{(O)}(W)^2$. By linearity of expectation, $\mathbb{E}\|H_n(W) - H_n(W_P)\|_2^2 \le \epsilon_k^{(O)}(W)^2$, and hence $\mathbb{E}\|H_n(W) - H_n(W_P)\|_2 \le \epsilon_k^{(O)}(W)$. □

E Bounds for Hölder-Continuous Graphons

In this section we prove Remark 2. Throughout this section we assume that $W : [0,1]^2 \to [0,1]$ is α-Hölder continuous for some $\alpha \in (0,1]$, i.e., we assume that there exists a constant $C < \infty$ such that

$$|W(x,y) - W(x',y')| \le C\Big(|x-x'| + |y-y'|\Big)^\alpha.$$

Lemma 15. *Let $H = H_n(W)$, and assume that the vertices of H are reordered in such a way that $x_1 < x_2 < \ldots, x_n$. Then*

$$\|W - W[H_n]\|_2 = O_P(n^{-\alpha/2}).$$

Proof. We first approximate W in terms of the weighted graph $\tilde{H}$ with weights $(\tilde{H})_{ij} = W(\bar{x}_i, \bar{x}_j)$, where $\bar{x}_i = \frac{i}{n+1}$ is the expectation of x_i. Since $|x - \bar{x}_i| \le \frac{1}{n}$ when $x \in [\frac{i-1}{n}, \frac{i}{n}]$, we can use the Hölder continuity of W to conclude that

$$\|W - W[\tilde{H}]\|_2 \le \|W - W[\tilde{H}]\|_\infty \le C\Big(\frac{2}{n}\Big)^\alpha.$$

To prove the lemma, it is therefore enough to prove that

$$\mathbb{E}\Big[\|W[\tilde{H}] - W[H]\|_2^2\Big] = O(n^{-\alpha}),$$

where $\mathbb{E}[\cdot]$ denotes expectations with respect to the random variables $x_1, \dots, x_n$. Using the Hölder continuity of W together with Jensen's inequality, we bound

$$\begin{aligned}
\mathbb{E}\Big[\|W[\tilde{H}_n] - W[H]\|_2^2\Big] &= \frac{1}{n^2}\sum_{i,j\in[n]} \mathbb{E}\Big[\big(W(\bar{x}_i,\bar{x}_j) - W(x_i,x_j)\big)^2\Big] \\
&\le \frac{C^2}{n^2}\sum_{i,j\in[n]} \mathbb{E}\Big[\big(|\bar{x}_i - x_i| + |\bar{x}_j - x_j|\big)^{2\alpha}\Big] \\
&\le \frac{C^2}{n^2}\sum_{i,j\in[n]} \mathbb{E}\Big[\big(2|\bar{x}_i - x_i|^2 + 2|\bar{x}_j - x_j|^2\big)^{\alpha}\Big] \\
&\le \frac{C^2}{n^2}\sum_{i,j\in[n]} \Big(2\mathbb{E}\Big[|\bar{x}_i - x_i|^2\Big] + 2\mathbb{E}\Big[|\bar{x}_j - x_j|^2\Big]\Big)^{\alpha}
\end{aligned}$$

Using the fact that x_i has expectation $\bar{x}_i = \frac{i}{n+1}$ and variance $\frac{1}{n+2}\bar{x}_i(1-\bar{x}_i) \le \frac{1}{4n}$, we bound the right hand side by $C^2 n^{-\alpha}$, completing the proof. □

Lemma 16. *Let $\mathcal{P}_k$ be a partition of $[0,1]$ into adjacent intervals of lengths $1/k$. Then*

$$\epsilon_k^{(O)}(W) \le \|W - W_{\mathcal{P}_k}\|_\infty \le C\Big(\frac{2}{k}\Big)^{\alpha}.$$

Proof. Let $\mathcal{P}_k = (I_1, \dots, I_k)$. For $(x,y) \in I_i \times I_j$, $W_{\mathcal{P}_k}(x,y)$ is an average over points in $I_i \times I_j$, implying that $|W(x,y) - W_{\mathcal{P}_k}(x,y)| \le C(2/k)^{\alpha}$. □

F Consistency of Multi-way cuts

In this appendix, we prove the following theorem which implies Theorem 3 by the same arguments as those which lead from Theorems 5 and 4 to Theorems 1 and 2.

Theorem 6. *Let $q \ge 2$ be an integer.*

(i) Under the assumptions of Theorem 2,

$$d_\infty^{Haus}(S_q(G), \hat{S}_q(\hat{B}_{nonprivate})) \le 2\hat{\epsilon}_k^{(O)}(H_n(W)) + O_P\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)}\right). \tag{22}$$

(ii) Under the assumptions of Theorem 1,

$$d_\infty^{Haus}(S_q(G), \hat{S}_q(\hat{B}_{private})) \le 2\hat{\epsilon}_k^{(O)}(H_n(W)) + O_p\left(\sqrt[4]{\frac{\lambda^2\log k}{\rho n}} + \lambda\sqrt{\frac{k^2\log n}{n\epsilon}} + \frac{\lambda}{n\rho\epsilon}\right). \tag{23}$$

Before we prove the theorem, we start with a few bounds on the Hausdorff distance of various sets of q-way cuts. First, using the definition of the cut-distance (and the fact that the set of q-way cuts of a graph is invariant under relabelings), it is easy to see (see also [18]) that whenever G and G' are weighted graphs on $[n]$ and $\mathcal{P}$ is a partition of $[n]$, then

$$\|G/\mathcal{P} - G'/\mathcal{P}\|_\infty \le \hat{\delta}_\square\Big(\frac{1}{\|G\|_1}G, \frac{1}{\|G'\|_1}G'\Big)$$

implying that

$$d_\infty^{\mathrm{Haus}}(S_q(G), S_q(G')) \le \hat{\delta}_\square\Big(\frac{1}{\|G\|_1}G, \frac{1}{\|G'\|_1}G'\Big). \tag{24}$$

In a similar way, we have that for two graphons W, W',

$$d_\infty^{\mathrm{Haus}}(\hat{S}_q(W), \hat{S}_q(W')) \le \delta_\square\Big(\frac{1}{\|W\|_1}W, \frac{1}{\|W'\|_1}W'\Big). \tag{25}$$

We will also need to compare the fractional and integer cuts, $\hat{S}_q(G)$ and $S_q(G')$. To do so, one can use a simple rounding argument, as in Theorem 5.4 and its proof from [16]. This gives the bound[3]

$$d_\infty^{\mathrm{Haus}}(S_q(G), \hat{S}_q(G)) \le \frac{5}{\|G\|_1\sqrt{n}}, \tag{26}$$

valid for any weighted graph G with node weights 1 on $[n]$ and maximal edge-weight 1. We also note that for any weighted graph Q,

$$\hat{S}_q(Q) = \hat{S}_q(W[Q]) \tag{27}$$

(see [16, Proposition 5.3] [4]). Finally, we note that $\big|\|W\|_1 - \|W'\|_1\big| \le \|W - W'\|_\square$. In particular,

$$\Big|\|G\|_1 - \|G'\|_1\Big| \le \hat{\delta}_\square(G, G') \quad \text{and} \quad \Big|\|W\|_1 - \|W'\|_1\Big| \le \delta_\square(W, W'). \tag{28}$$

Proof of Theorem 6. Let $Q = H_n(\rho W)$ and $G = G(Q)$. By Lemma 12, we have that $\rho(Q) \in [\rho/2, 2\rho]$ with probability at least $1 - O(\Lambda/n)$. By the assumptions of the two theorems, $\rho n \ge 2\log 2$. We apply [17, Lemma 7.2] to show that $\hat{\delta}_\square(G, Q) = \rho O\Big(\sqrt{\frac{1}{\rho n}}\Big)$ with probability at least $1 - O(\Lambda/n)$. As a consequence, again with probability $1 - O(\Lambda/n)$,

$$\hat{\delta}_\square\Big(\frac{1}{\|Q\|_1}Q, \frac{1}{\|Q\|_1}G\Big) = O\Big(\sqrt{\frac{1}{\rho n}}\Big).$$

By (24) and (28), this implies that with the same probability

$$d_\infty^{\mathrm{Haus}}(S_q(G), S_q(Q)) = O\Big(\sqrt{\frac{1}{\rho n}}\Big).$$

Next we apply (26) to the weighted graph $Q' = \frac{1}{\|Q\|_\infty}Q$. Since $\frac{1}{\|Q'\|_1} = \frac{\|Q\|_\infty}{\|Q\|_1} \le \frac{\Lambda\rho}{\|Q\|_1}$, we conclude that with probability at least $1 - O(\Lambda/n)$,

$$d_\infty^{\mathrm{Haus}}(S_q(Q), \hat{S}_q(Q)) = d_\infty^{\mathrm{Haus}}(S_q(Q), \hat{S}_q(Q')) = O\Big(\sqrt{\frac{1}{\rho n}} + \frac{\Lambda}{\sqrt{n}}\Big).$$

Since $\|F\|_\infty \le 1$ for all $F \in S_q(G)$ and all $F \in \hat{S}_q(Q)$, we can easily absorb the failure event, getting

$$d_\infty^{\mathrm{Haus}}(S_q(G), \hat{S}_q(Q)) = O_P\Big(\sqrt{\frac{1}{\rho n}} + \frac{\Lambda}{\sqrt{n}}\Big). \tag{29}$$

[3] To translate the results from [16] into (26), we need to take into account that in [16], quotients where defined with a normalization of $\frac{1}{n^2}$ instead of $\frac{1}{n^2\|G\|_1}$ (leading to the factor $\frac{1}{\|G\|_1}$ on the right hand side of (26)), and that Hausdorff distances were defined with respect to the L_1-norm (leading to a bound which is better by a factor q than the bounds in [16]).

[4] Note that in [16] the notation for integer and fractional partitions is the reverse of the one used here and in [18].

To complete the proof, we proceed as in the proof of (14) to show that

$$\delta_2\Big(H_n(W), \frac{1}{\rho(G)}\hat{B}\Big) \le \hat{\delta}_2\Big(\frac{1}{\rho(G)}\hat{B}, H_n(W)\Big) + O_P\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)}\right).$$

Combined with the bound from Theorem 2, the fact that the cut-norm is bounded by the L_2-norm, and the fact that $\delta_\square(H_n(W), \frac{1}{\|Q\|_1}Q) \le |1 - \rho(Q)/\rho| = O_P(\Lambda/n)$, we conclude that

$$\delta_\square\Big(\frac{1}{\|Q\|_1}Q, \frac{1}{\rho(G)}W[\hat{B}]\Big) \le \hat{\epsilon}_k^{(O)}(H_n(W)) + O_P\left(\sqrt[4]{\lambda^2\left(\frac{k^2}{n^2\rho} + \frac{\log k}{n\rho}\right)}\right).$$

Combined with (27), (25), (28), and the bound (29), this proves (22). The proof of (23) is essentially identical, except that now we use Theorem 1. □

G Useful Lemmas

Lemma 17 (Multiplicative Chernoff bound)**.** *Let $X_1, X_2, ..., X_n$ be independent random variables taking values in $[0, 1]$, and let $X = \sum_{i=1}^n X_i$. If $\mathbb{E}(X) \le \mu_0$ and $\beta \le 1$, then*

$$\Pr(|X - \mathbb{E}(X)| \ge \beta\mu_0) \le 2\exp(-\beta^2\mu_0/3)\,.$$

For $\beta > 1$, the probability is at most

$$\Pr(|X - \mathbb{E}(X)| \ge \beta\mu_0) \le 2\exp(-\beta\mu_0/3)\,.$$

Proof. Let $\mu = \mathbb{E}(X)$ denotes the exact mean of X (so $\mu \le \mu_0$). The standard multiplicative form of the Chernoff bound states that for $\delta > 0$ (not necessarily less than 1), we have

$$\Pr(|X - \mathbb{E}X| \ge \delta\mu) \le 2\max(e^{-\frac{1}{3}\delta^2\mu}, e^{-\frac{1}{3}\delta\mu})\,.$$

Setting $\delta\mu = \beta\mu_0$ (that is, $\delta = \frac{\beta\mu_0}{\mu}$), the bound above becomes $2\max(e^{-\frac{1}{3}\frac{\beta^2\mu_0^2}{\mu}}, e^{-\frac{1}{3}\beta\mu_0})$. Both of these terms are bounded above by $2\exp(-\beta^2\mu_0/3)$: the first, since $\mu_0 \le \mu$; and the second, since $\beta \le 1$. □