On the computational tractability of statistical estimation on amenable graphs

Ahmed El Alaoui[1] · Andrea Montanari[2]

Abstract

We consider the problem of estimating a vector of discrete variables $\theta = (\theta_1, \cdots, \theta_n)$, based on noisy observations Y_{uv} of the pairs (θ_u, θ_v) on the edges of a graph $G = ([n], E)$. This setting comprises a broad family of statistical estimation problems, including group synchronization on graphs, community detection, and low-rank matrix estimation. A large body of theoretical work has established sharp thresholds for weak and exact recovery, and sharp characterizations of the optimal reconstruction accuracy in such models, focusing however on the special case of Erdös–Rényi-type random graphs. An important finding of this line of work is the ubiquity of an information-computation gap. Namely, for many models of interest, a large gap is found between the optimal accuracy achievable by any statistical method, and the optimal accuracy achieved by known polynomial-time algorithms. Moreover, it is expected in many situations that this gap is robust to small amounts of additional side information revealed about the θ_i's. How does the structure of the graph G affect this picture? Is the information-computation gap a general phenomenon or does it only apply to specific families of graphs? We prove that the picture is dramatically different for graph sequences converging to amenable graphs (including, for instance, d-dimensional grids). We consider a model in which an arbitrarily small fraction of the vertex labels is revealed, and show that a linear-time local algorithm can achieve reconstruction accuracy that is arbitrarily close to the information-theoretic optimum. We contrast this to the case of random graphs. Indeed, focusing on group synchro-

✉ Ahmed El Alaoui
ae333@cornell.edu

Andrea Montanari
montanar@stanford.edu

[1] Cornell University, 1176 Comstock Hall, Ithaca, NY 14850, USA

[2] Stanford University, 390 Serra Mall, Stanford, CA 94305, USA

nization on random regular graphs, we prove that local algorithms are unable to have non-trivial performance below the so-called Kesten–Stigum threshold, even when a small amount of side information is revealed.

Keywords Statistical estimation · Local algorithms · Processes on graphs · Information-computation tradeoffs.

Mathematics Subject Classification 62F15 · 82B20 · 82B44

1 Introduction

Classical statistics focuses on problems in which a small number of parameters needs to be estimated from data. As a consequence, it is mostly unconcerned with computational complexity considerations. Fundamental limits to statistical estimation are proven on the basis of information-theoretic considerations. On the contrary, in modern high-dimensional applications, it is not uncommon to come across statistical models that require estimating simultaneously thousands or even millions of parameters. In this setting, a large gap is often observed between information-theoretic limits and what is achieved by the best known polynomial-time algorithms. Indeed, it is expected that no polynomial-time algorithm can achieve optimal statistical performance in general. In specific classes of models, a precise information-computation gap has been conjectured on the basis of current knowledge (see, e.g., [14,18,21,38,43,45] and references therein).

As explained below, most of our understanding of this information-computation gap was developed by analyzing probabilistic models with a high degree of exchangeability. This suggests a natural question: *Is the same gap present in models with other type of structures?*

Statistical estimation on graphs provides a rich and interesting setting to study this question. Let $G_n = (V_n, E_n)$ be a graph on n vertices, $V_n = [n]$. Edges are assumed to be directed in an arbitrary way, i.e., they are ordered pairs $(u, v) \in V_n \times V_n$. We associate to the vertices $u \in V_n$ random variables $\boldsymbol{\theta} = (\theta_u)_{u \in V_n} \sim_{\text{iid}} Unif(\mathcal{X})$, uniformly distributed on a finite alphabet $\mathcal{X}$. For each edge $(u, v) \in E_n$, we observe $Y_{uv} \in \mathcal{Y}$, where $\mathcal{Y}$ is also a finite alphabet. The observations are conditionally independent with $Y_{uv} | \boldsymbol{\theta} \sim Q(\cdot | \theta_u, \theta_v)$, where Q is a probability kernel from $\mathcal{X} \times \mathcal{X}$ to $\mathcal{Y}$. Given the edge observations Y (and, possibly, additional side information, see below), the purpose is to estimate the vertex assignment $\boldsymbol{\theta}$.

This model is general enough to include a broad variety of examples studied in the literature, including group synchronization, community detection, low-rank matrix estimation, and so on. As an example consider the $\mathbb{Z}_q$-*synchronization* problem (further examples are presented in Sect. 3.1). The unknown variables $(\theta_u)_{u \in V_n}$ are i.i.d. uniform in $\mathcal{X} = \mathbb{Z}_q = \{0, \cdots, q-1\}$, which we identify with the cyclic group $\mathbb{Z}/q\mathbb{Z}$ with additive structure. Observations are noisy measurements of the difference between θ_u and θ_v for each edge $(u, v) \in E_n$:

$$Y_{uv} = \begin{cases} \theta_u - \theta_v \ (\text{mod } q) & \text{with probability } 1 - p, \\ w_{uv} & \text{with probability } p, \end{cases} \tag{1.1}$$

where $(w_{uv})_{(u,v) \in E_n}$ is a collection of independent random variables $w_{uv} \sim Unif(\mathcal{X})$, independent of $(\theta_u)_{u \in V_n}$.

In addition to the observations Y, we consider independent observations $(\xi_u^{(\varepsilon)})_{u \in V_n}$ on the vertices of G_n:

$$\xi_u^{(\varepsilon)} = \begin{cases} \theta_u & \text{with probability } \varepsilon, \\ \star & \text{with probability } 1 - \varepsilon, \end{cases} \tag{1.2}$$

where $\star$ is a symbol not belonging to $\mathcal{X}$, so that with probability ε the value of θ_u is directly observed. We will write $\mathcal{X}_\star := \mathcal{X} \cup \{\star\}$. Following the information theory literature, we refer to this noise model as the Binary Erasure Channel, and denote it by BEC($\bar{\varepsilon}$). (It is customary to parametrize the BEC by its erasure probability $\bar{\varepsilon} = 1 - \varepsilon$.) The parameter ε will be considered very small (eventually going to zero as n becomes large). The purpose of this side information is to break the occasional group symmetry (sign symmetry or cyclic shifts in the case of $\mathbb{Z}_q$) that would otherwise be preserved by the observations Y.

We consider two metrics for the estimation accuracy. In our first definition, the goal is to estimate the $n \times n$ rank-one matrix X_f whose entries are

$$(X_f)_{u,v} := f(\theta_u) f(\theta_v), \quad u, v \in V_n, \tag{1.3}$$

where $f : \mathcal{X} \to \mathbb{R}$ is a given real-valued function. For instance by setting $f(\theta) = \mathbf{1}_{\theta = x}$ and then considering all values of $x \in \mathcal{X}$, this allows to estimate whether $\theta_u = \theta_v$ for each pair of vertices $u, v \in V_n$. An estimator is a map $\widehat{X} : \mathcal{Y}^{E_n} \times \mathcal{X}_\star^{V_n} \to \mathbb{R}^{n \times n}$, i.e., a function of the observations Y and the side information $\xi^{(\varepsilon)}$. We evaluate its risk under the square loss

$$\mathcal{R}_n(\widehat{X}; f) := \frac{1}{n^2} \mathbb{E}\left[\left\|X_f - \widehat{X}\right\|_F^2\right]. \tag{1.4}$$

We denote by $\mathcal{R}_n^{\text{Bayes}}(f)$ the minimal achievable error, i.e., the one achieved by the posterior expectation

$$\widehat{X}^{\text{Bayes}} := \left(\mathbb{E}\left[f(\theta_u) f(\theta_v) | Y_{G_n}^{(\varepsilon)}\right]\right)_{u,v \in V_n}. \tag{1.5}$$

(We have made use of the following notation: for a graph $G = (V, E)$, we denote by $Y_G^{(\varepsilon)}$ the union of the vertex and edge observations over G: $Y_G^{(\varepsilon)} = \{Y_{uv} : (u, v) \in E, \xi_u^{(\varepsilon)} : u \in V\}$.) Our second metric for estimation accuracy is the 'overlap', and will be introduced in Sect. 4, see Eq. (4.7).

Statistical estimation on graphs has motivated substantial amount of work. In this context, the first example of a statistical model with a large information-computation

gap is probably the planted clique problem [8,34]. This can be recast in the general framework described above, with G_n the complete graph over n vertices (see Sect. 3.1). Despite more than a quarter century of research, and the study of increasingly powerful classes of algorithms [13,24,27], no known polynomial-time algorithm comes close to saturate the information-theoretic limits for this problem.

In recent years, a much more refined picture of the information-computation gap has emerged, mainly through the careful analysis of a variety of models on sparse random graphs (as well as models on dense graphs in a different noise regime than the hidden clique model). We refer to Sect. 2 for a brief summary of this vast literature. In most of these models an information-computation gap is observed, and has been precisely delineated. This gap is generally conjectured to remain unchanged if a small amount of side information is revealed,[1] as in Eq. (1.2). As mentioned above, most of the theoretical work has focused however on random graphs (Erdös–Rényi random graphs, random regular graphs and their relatives). This motivates the following key question:

Does an information-computation gap exist for statistical estimation on other types of graphs?

In this paper, we consider the case of graph sequences that converge locally to amenable graphs. Roughly, these are graphs for which the boundary of large sets of vertices is negligible compared to their volume. We refer to Sect. 3 for a reminder on the relevant definitions. Our results are already interesting for the simplest example of such graphs, namely large boxes $[1, L] \times \cdots \times [1, L]$ in the d-dimensional grid $\mathbb{Z}^d$ (with $L = n^{1/d}$).

Our main finding is that no information-computation gap exists for such graphs (as long as the gap is defined in terms of polynomial- versus non-polynomial time algorithms). A specific formalization of this finding is given below, and proved in Sect. 4.

Theorem A *Let $f : \mathcal{X} \to \mathbb{R}$ be a function with $\mathbb{E}[f(\theta)] = 0$ for $\theta \sim Unif(\mathcal{X})$. Let $G_n = (V_n, E_n)$ be a sequence of finite graphs (with $|V_n| = n$) converging locally-weakly to a random rooted graph (G, o) which is infinite, locally-finite, almost surely anchored-amenable and tame. Then for each $l \in \mathbb{N}$ there exists an estimator $\widehat{X}^{(l)} : \mathcal{Y}^{E_n} \times \mathcal{X}_\star^{V_n} \to \mathbb{R}^{n \times n}$, with runtime $\mathcal{O}(n^2)$ (for ℓ fixed), such that the following holds. For almost every $\varepsilon > 0$, we have*

$$\lim_{l \to \infty} \lim_{n \to \infty} \left\{ \mathcal{R}_n(\widehat{X}^{(l)}; f) - \mathcal{R}_n^{Bayes}(f) \right\} = 0.$$

The notions of local-weak convergence, anchored-amenability and tameness will be defined in Sect. 3.

More in detail, we present the following contributions:

---

[1] The careful reader will notice that this statement does not apply to the planted clique problem. If the label of εn random vertices is revealed (i.e., whether or not they belong to the clique), then it is easy to find planted cliques of size $k \gg (1/\varepsilon) \log n$, i.e., far below the best known polynomial algorithms for $\varepsilon = 0$. This behavior is however related to the fact that, in the planted clique problem, the underlying graph is the complete graph. Hence a small amount of vertex side information can have large impact even on local algorithms.

No information-computation gap on amenable graphs   Theorem A provides a concrete formalization of the general statement that statistically optimal estimation can be performed using polynomial time algorithms on (asymptotically) amenable graphs. In fact, we will prove that this follows from a more fundamental result, establishing that the vertex marginals of the posterior $\mathbb{P}\big(\boldsymbol{\theta} \,|Y_{G_n}^{(\varepsilon)}\big)$ can be computed to arbitrary accuracy in polynomial time, for almost all values of ε, on asymptotically amenable graphs, cf. Sect. 4.

   Note that approximating the Bayes estimator $\widehat{X}^{\text{Bayes}}$, Eq. (1.5), requires to approximate the joint distribution of pairs of well separated vertices. However, we will use a decoupling argument to reduce ourselves to the case of vertex marginals.

Local algorithms.   Our proof that vertex marginals can be computed efficiently follows from an even stronger, and somewhat surprising fact (as above, holding for almost all ε). The marginal at a vertex v can be well approximated by computing the marginal with respect to the posterior given observations in a large constant-size ball centered at v. In other words, the marginal can be approximated by a local algorithm. The reason for this phenomenon can be explained in information theoretic terms. We will prove that the average conditional mutual information between a random vertex in a region $S \subseteq V$, and the boundary of S, $I\big(\theta_v; \theta_{\partial S}|Y_S^{(\varepsilon)}\big)$ is upper bounded by $|\partial S|/|S|$. Hence, for amenable graphs, the effect of the boundary information is generally negligible.

Robust information-computation gap on random regular graphs   We provide a counter-example, by showing that the conclusions at the previous points do not hold for random regular graphs, converging locally to k-regular trees, which are non-amenable. As mentioned above, several cases of statistical estimation problems have been observed to present an information-computation gap, when the underlying graph is random. While this gap is often expected to be robust to side information about the vertices (see Sect. 2), we are not aware of any result that explicitly establishes robustness—in the setting of the present paper. We consider the $\mathbb{Z}_q$-synchronization problem on random k-regular graphs. We prove that, for a large range of the model parameters and all ε small enough: (i) There exists a statistical estimator that achieves non-trivial reconstruction accuracy uniformly as $\varepsilon \to 0$; (ii) Local algorithms can only achieve accuracy that vanishes as $\varepsilon \to 0$.

2 Related literature

As mentioned in the introduction, large information-computation gaps were observed in a number of statistical estimation problems, when the underlying structure is a random graph, the complete graph, or close relatives. An incomplete list includes community detection in the stochastic block model [1,21,41,46], high-dimensional linear regression and generalized linear models [14,18], low-rank matrix estimation and sparse principal component analysis [9,16,36,38,40], tensor principal component analysis [29,31,45], tensor decomposition, and so on.

In many of these models, two types of results are established. On one hand an 'information-theoretic' analysis allows to characterize the optimal statistical accuracy that is achieved by an ideal estimator. On the other, specific classes of polynomial-time algorithms are analyzed. A gap is observed between the accuracy achieved by known polynomial-time algorithms and the optimal statistical accuracy. Sometimes the resulting statistical estimation limits are stated in terms of specific goals such as 'weak recovery' (the goal in this case is to outperform the trivial estimator) or 'exact recovery' (the goal is to reconstruct exactly θ, with high probability). In the present paper we consider the general goal of estimation with certain expected accuracy, or risk.

The most frequently analyzed classes of algorithms have been spectral methods, local algorithms, and convex relaxations in the sum-of-squares hierarchy. A remarkable dichotomy has emerged from these works. Roughly speaking, in all the examples we know of, either highly sophisticated semidefinite programming hierarchies fail, or simple combinations of spectral methods and local algorithms succeed.

In the context of problems on random graph converging locally to trees with bounded average degree (in particular, the stochastic block model, the censored block model, the group synchronization problem for $\mathbb{Z}_q$), this observed dichotomy can be formalized rather precisely. Consider belief propagation with a small amount ε of side information, run in parallel for ℓ iterations, and denote by $\mathcal{R}_n(BP^{(\ell)}; \varepsilon)$ the corresponding error. This is well known to be the optimal error achieved by any local algorithm. Let[2]

$$\mathcal{R}^{\mathrm{loc}}_\infty(0+) \equiv \lim_{\varepsilon \to 0+} \mathcal{R}^{\mathrm{loc}}_\infty(\varepsilon), \quad \mathcal{R}^{\mathrm{loc}}_\infty(\varepsilon) \equiv \lim_{\ell \to \infty} \lim_{n \to \infty} \mathcal{R}_n(BP^{(\ell)}; \varepsilon). \qquad (2.1)$$

(The order of limits is crucial here.) The results of [1,5,12,19,21,23,41,46] indicate that, in a number of cases, the error $\mathcal{R}^{\mathrm{loc}}_\infty(0+)$ can be achieved in polynomial time *even without side information* (i.e. for $\varepsilon = 0$). Further, this can be strictly larger than the optimal statistical error (also for $\varepsilon = 0$).

A number of works establish that in several cases, broad classes of algorithms (sum-of-squares relaxations, low-degree testing) cannot outperform $\mathcal{R}^{\mathrm{loc}}_\infty(0+)$ [26,29,30,32,37]. These finding suggest that the gap between optimal accuracy and $\mathcal{R}^{\mathrm{loc}}_\infty(0+)$ is a fundamental one. Also, as a consequence of the very definition of $\mathcal{R}^{\mathrm{loc}}_\infty(0+)$, this gap is robust to a small amount of side information, for problems on random graphs. Notice that while the analysis of sum-of-squares relaxations and low-degree testing typically assume $\varepsilon = 0$, they concern algorithms that are robust to small perturbations of the problem.

Unfortunately, none of earlier works fits precisely with the setting of our paper. This motivates our analysis of $\mathbb{Z}_q$-synchronization on random regular graphs (Sect. 5). We establish that a large gap exists between $\mathcal{R}^{\mathrm{loc}}_\infty(\varepsilon)$ and the optimal statistical accuracy for all ε small enough.

---

[2] In the physics literature, see e.g. [21], the BP accuracy is not formally defined in this way, but rather by considering an initialization that is slightly biased towards the true parameters vector. The two definitions are expected to coincide generically.

Let us stress that our main focus is statistical estimation on amenable graphs. Versions of this problem have been studied in a few recent papers [2–4,50,51]. In particular, [4] proved the existence of a weak recovery threshold for $\mathbb{Z}_q$-synchronization[3] on grids in $d \geq 3$ dimensions. However, in contrast with random graphs, no explicit characterization exists (or is likely to exist) for the optimal statistical accuracy nor, in general, for the location of weak recovery thresholds. This poses a clear challenge to us: we want to prove that the optimal statistical accuracy can be achieved by polynomial time algorithms, but *we do not have an explicit characterization for the target accuracy*. Indeed, our proof will be purely conceptual.

Let us finally mention that it is well understood that certain algorithmic tasks are easy on graphs that can be embedded well in $\mathbb{R}^d$ (e.g., on grids). For instance, approximate optimization of a function that decomposes as a sum of edge terms over a grid is easy, by partitioning the grid into large boxes. Unfortunately, these ideas do not have direct implications on the questions addressed in this paper. Even if we can find an approximate-maximum likelihood assignment of the unknown variables θ_i, this is not guaranteed to have any good statistical properties, let alone achieve optimal estimation error. Inference and estimation do not reduce to optimization.

3 Background

3.1 Further examples

It is interesting to check that the framework defined in the introduction is broad enough to encompass a variety of models of interest.

Spiked Wigner and Wishart models Low-rank plus noise models are ubiquitous in statistics and signal processing [35], and can be recast in the language of the present paper. As an example, consider the case of a signal vector $\boldsymbol{\theta} \in \mathbb{R}^n$, with i.i.d. components, and assume we observe the rank-one-plus-noise matrix $Y = \boldsymbol{\theta}\boldsymbol{\theta}^\top + \sigma_n \boldsymbol{W}$. Here W is a noise matrix, with—for instance—$W_{uv} \sim \mathcal{N}(0, 1)$ and σ_n controls the noise level.

We take G_n to be the complete graph, and $(\theta_u)_{u \in V_n}$ be i.i.d. random variables[4] from a distribution P_θ on $\mathbb{R}$. Observations on the edges are given by

$$Y_{uv} \sim Q(\cdot \,|\theta_u, \theta_v) = \mathcal{N}(\theta_u\theta_v; \sigma_n^2), \tag{3.1}$$

where $\mathcal{N}(\mu, \sigma^2)$ denotes the Gaussian distribution.

This example can be easily generalized. For instance, higher rank models can be produced by taking $\theta_u \in \mathbb{R}^r$, $r \geq 1$ fixed. Rectangular (non-symmetric) random

---

[3] For grids in $d = 2$ dimension, the situation is expected to be similar, although the proof is more complicated. Indeed, [4] proves that a threshold exists in the case $q = 2$, and indeed the same is expected to hold for $q \geq 3$ as well. For $d = 1$ no non-trivial threshold exists and weak recovery is always impossible.

[4] Unlike for the model described in the introduction, the variables θ_u's typically take any value in $\mathbb{R}$, and their distribution is non-uniform. However, it is easy to reduce from one case to the other. For instance, we can let $Y_{uv} = \mathcal{N}(h(\theta_u)h(\theta_v), \sigma_n^2)$. We can choose the nonlinear function $h : \mathbb{R} \to \mathbb{R}$ so that $h(\theta_v) \sim P_0$ when $\theta_v \sim Unif([0, 1])$.

matrices of dimensions $n_1 \times n_2$, can also be produced by setting $n = n_1 + n_2$. In this case $\theta_v = (\zeta_v, b_v)$ where $\zeta_u \in \mathbb{R}^r$ and $b_v \in \{1, 2\}$ depending whether v belongs to the first n_1 vertices (left factor) or the last n_2 ones (right factor).

Community detection The stochastic block model is a popular model for community detection in networks. The model is parametrized by a symmetric 'connectivity' matrix $(c_{rs})_{1 \le r, s \le q}$, whereby $c_{r,s} \in [0, 1]$ is the expected edge density between vertices in communities r and s. (For the sake of simplicity, we consider here the 'balanced' case in which the q communities have all equal expected size.) Each vertex $v \in V_n$ is assigned a label $\theta_v \in [q]$ independently and uniformly at random. Conditional on $\boldsymbol{\theta}$, we generate a graph $\tilde{G}_n = (V_n, \tilde{E}_n)$ by connecting vertices u, v independently with probability $\mathbb{P}((u, v) \in \tilde{E}_n | \boldsymbol{\theta}) = c_{\theta_u, \theta_v}$.

We can encode this model in our general framework as follows. The graph G_n is the complete graph, and observe $Y_{uv} \in \{0, 1\}$ on every edge, where $Q(Y_{uv} = 1 | \theta_u = r, \theta_v = s) = c_{r,s}$. The connection with the standard description is given by the correspondence $\{Y_{uv} = 1\} \Leftrightarrow \{(u, v) \in \tilde{E}_{uv}\}$. The same encoding can be used for the planted clique problem.

Let us note that although the above models are special cases of our framework, we will focus in the rest of the paper onto graphs whose local-weak limit (to be defined shortly) is locally finite. This rules out graphs with diverging typical degree (in particular the complete graph).

3.2 Local-weak convergence and amenability

For the reader's convenience, we collect here some relevant graph-theoretic definitions, referring to [7,15,39] for more details. In this paper, all graphs have a finite or countably infinite vertex set, are connected, and are *locally finite*; i.e., all vertices have finite degree. A *rooted* graph (G, o) is a graph G together with a choice of a vertex $o \in V(G)$, called the *root* of G. We say that two rooted graphs (G, o) and (G', o') are isomorphic—and we write $(G, o) \equiv (G', o')$—if there exists an edge-preserving and root-preserving bijective map $\phi : V(G) \to V(G')$, i.e., $(u, v) \in E(G) \Leftrightarrow (\phi(u), \phi(v)) \in E(G')$, and $\phi(o) = o'$. For an integer $l \ge 0$, define $[G, o]_l$ to be the rooted subgraph spanned by a ball of radius l around the root o on G: this is the rooted graph $((V_l, E_l), o)$ where $V_l = B_G(o, l) := \{u \in V(G) : d_G(o, u) \le l\}$, and $E_l = \{(u, v) \in E : u, v \in V_l\}$. Here, d_G is the graph distance in G.

Definition 1 A sequence of rooted graphs $(G_n, o_n)_{n \ge 1}$ is said to *converge locally* to a rooted graph (G, o), and we write $(G_n, o_n) \xrightarrow{loc.} (G, o)$, if for every radius $l \ge 0$, there exists $n_0 \ge 0$ such that $[G_n, o_n]_l \equiv [G, o]_l$ for all $n \ge n_0$.

This notion of convergence endows the set $\mathcal{G}_*$ (of $\equiv$-equivalence classes) of rooted graphs with a metrizable topology, called the topology of local, or Benjamini–Schramm, convergence [15]. This gives $\mathcal{G}_*$ the structure of a complete separable metric space. Now we can define $\mathscr{P}(\mathcal{G}_*)$, the space of probability measures on $\mathcal{G}_*$ when endowed with its Borel σ-algebra. Then we endow $\mathscr{P}(\mathcal{G}_*)$ with the usual topology of weak convergence.

From a *finite* deterministic graph G, we can construct a random rooted graph (G, o) by choosing the root o uniformly at random from $V(G)$. We denote the law of this random rooted graph by $\rho_G \in \mathcal{P}(\mathcal{G}_*)$.

Definition 2 A sequence of finite graphs $(G_n)_{n \geq 1}$ is said to converge locally-weakly to a random rooted graph (G, o) if the sequence of probability measures $(\rho_{G_n})_{n \geq 1}$ converges weakly to a probability measure $\rho \in \mathcal{P}(\mathcal{G}_*)$, which is the law of (G, o).

In other words, the definition requires that given a fixed finite connected rooted graph (H, o') and a fixed radius l, the probability $\mathbb{P}([G_n, o_n]_l \equiv (H, o'))$ converges to $\mathbb{P}([G, o]_l \equiv (H, o'))$ as $n \to \infty$.

Probability measures $\rho \in \mathcal{P}(\mathcal{G}_*)$ that are local-weak limits of sequences of *finite* graphs as per Definition 2 (such measures are called *sofic* in the literature) inherit a important stationarity property which roughly expresses the intuition that the random graph G should "look the same" when viewed from any of its vertices. A formal definition takes the form of a *mass-transport principle* termed *unimodularity* [7]: Similarly to $\mathcal{G}_*$, we define $\mathcal{G}_{**}$ the space of $\equiv$-equivalence classes of *doubly*-rooted graphs (G, o, o') where the isomorphy relation $\equiv$ and local convergence as per Definition 1 are both extended in the natural way.

Definition 3 A measure $\rho \in \mathcal{P}(\mathcal{G}_*)$ is unimodular if for every Borel function $f : \mathcal{G}_{**} \to \mathbb{R}_+$,

$$
\mathbb{E}_\rho \Big[\sum_{u \in V(G)} f(G, o, u) \Big] = \mathbb{E}_\rho \Big[\sum_{u \in V(G)} f(G, u, o) \Big],
$$

when $(G, o) \sim \rho$.

It is clear that if G is finite then ρ_G is unimodular, since the root is chosen uniformly at random. Furthermore, the property of unimodularity is closed in the topology of local-weak convergence [7], hence *all local-weak limits of sequences of finite graphs are unimodular*.

Next, we define the key concept of *anchored-amenability*.

Definition 4 An infinite rooted graph (G, o) where $G = (V, E)$ is said to be anchored-amenable if its Cheeger constant anchored at o is zero:

$$
\inf \Big\{ |\partial S| / |S| : S \subset V \text{ finite}, o \in S \Big\} = 0.
$$

Here, $\partial S = \{u \in S : \exists v \notin S, (u, v) \in E\}$ is the vertex-boundary of the set $S \subseteq V$.

We will informally use the phrase 'asymptotically amenable' to refer to graph sequences that converge locally-weakly to almost surely anchored-amenable graphs.

Observe that if G is vertex-transitive, the above statement does not depend on the root o, and anchored-amenability reduces to the more classical notion of amenability of (non-rooted) graphs. For instance, the Euclidean lattice $\mathbb{Z}^d$ is amenable, the k-regular tree is not (both graphs being transitive).

Observe that if $(G, o) \sim \rho$ is almost surely anchored-amenable, there exists a sequence of finite sets $S_k \subset V$ such that $o \in S_k$ and which 'witnesses' the amenability of G: $|\partial S_k|/|S_k| \longrightarrow 0$ as $k \to \infty$. Moreover, this random sequence can be chosen in a *measurable* way as a function of the rooted graph (G, o). Indeed, we can for instance label the vertices of G by $\mathbb{N}$, the root being labelled by 0, and for every $k \geq 1$, choose the first finite set $S_k \subseteq V(G)$ (among countably many) in the lexicographic ordering such that $|\partial S_k|/|S_k| \leq 2^{-k}$ and $o \in S_k$. For clarity we make this dependence explicit: $S_k = S_k(G, o)$. We require a technical condition regarding such sets S_k.

Definition 5 We say that $\rho \in \mathscr{P}(\mathcal{G}_*)$ is *tame* if it is supported on anchored-amenable rooted graphs, and there exists a sequence $\{S_k\}_{k \geq 1}$ of sets that witnesses anchored-amenability (i.e., such that $S_k(G, o)$ is a measurable function of (G, o), and $|\partial S_k|/|S_k| \to 0$ almost surely) such that the following holds. For every $\eta > 0$ there exists $\delta > 0$ such that

$$\limsup_{k \to \infty} \rho\left((G, o) \ : \ \sum_{u \in V(G)} \frac{\mathbf{1}_{o \in S_k(G, u)}}{|S_k(G, u)|} \leq \delta\right) \leq \eta. \tag{3.2}$$

By extension, we say that the random rooted graph (G, o) is tame if its law ρ is tame.

Intuitively, tameness is satisfied when the size of the neighborhoods $S_k(G, u)$ of each vertex u around the root is comparable with $S_k(G, o)$. To discuss it further, it is useful to introduce the random variables

$$\alpha_k(G, o) := \sum_{u \in V(G)} \frac{\mathbf{1}_{o \in S_k(G, u)}}{|S_k(G, u)|} . \tag{3.3}$$

The tameness condition requires a uniform upper bound on the lower tail of $(\alpha_k(G, o))_{k \geq 1}$. An equivalent way to express this condition is to say that the sequence of random variables $(1/\alpha_k(G, o))_{k \geq 1}$ is *tight* when $(G, o) \sim \rho$.

Note that $\mathbb{E}_\rho[\alpha_k(G, o)] = 1$ whenever ρ is unimodular. Indeed, by a direct application of the mass-transport principle (for the function $f(G, o, u) = \frac{\mathbf{1}_{o \in S_k(G, u)}}{|S_k(G, u)|}$)

$$\mathbb{E}_\rho[\alpha_k(G, o)] = \mathbb{E}_\rho\left[\sum_{u \in V(G)} \mathbf{1}_{o \in S_k(G, u)} \frac{1}{|S_k(G, u)|}\right]$$

$$= \mathbb{E}_\rho\left[\sum_{u \in V(G)} \mathbf{1}_{u \in S_k(G, o)} \frac{1}{|S_k(G, o)|}\right] = \mathbb{E}_\rho\left[\frac{|S_k(G, o)|}{|S_k(G, o)|}\right] = 1 .$$

Moreover, tameness is satisfied if ρ is supported on vertex-transitive graphs, and in this case $\alpha_k(G, o) = 1$ almost surely. Indeed, assume ρ is supported on a single vertex-transitive graph. Then ρ is unimodular whence $\mathbb{E}_\rho[\alpha_k(G, o)] = 1$, but $\alpha_k(G, o)$ is non-random and therefore $\alpha_k(G, o) = 1$. In the general case where ρ is not an atom, since $\alpha_k(G, o) = 1$ almost surely conditional on (G, o), we have $\alpha_k(G, o) = 1$ almost

surely unconditionally as well. We next provide two examples of non-transitive graphs which are anchored-amenable and tame.

Example 1 (Percolation clusters) Consider the d-dimensional grid $\mathbb{L}^d$, i.e., $V(\mathbb{L}^d) = \mathbb{Z}^d$, and edges connect vertices at distance one $E(\mathbb{L}^d) = \{(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{Z}^d : \|\boldsymbol{x} - \boldsymbol{y}\| = 1\}$. Remove edges independently with probability $1 - p$ and let $G = G_p(o)$ be the connected component of the origin $o = \mathbf{0}$. We consider $p > p_c$, the percolation threshold on $\mathbb{L}^d$ so that G is infinite with positive probability, and condition on the event that G is indeed infinite. In this case we can take $S_k(G, \boldsymbol{x})$ to be the subset of vertices contained in the ℓ_∞ ball of radius ℓ_k around $\boldsymbol{x}$: $S_k(G, \boldsymbol{x}) = \{\boldsymbol{y} \in V(G) : \|\boldsymbol{y} - \boldsymbol{x}\|_\infty \leq \ell_k\}$, for a deterministic sequence of radii $\ell_k \uparrow \infty$. A classical result of Newman and Schulman [48] implies $|S_k(G, o)|/\ell_k^d \to c_0$ almost surely for some non-random constant $c_0 > 0$. Further $\partial S_k \subseteq \{\boldsymbol{x} \in \mathbb{Z}^d : \ell_k - 1 \leq \|\boldsymbol{x}\|_\infty \leq \ell_k\}$ whence $|\partial S_k| \leq c_1 \ell_k^{d-1}$. Hence, there exists a random $k_0 < \infty$ such that almost surely $|\partial S_k|/|S_k| \leq (2c_1/c_0)\,\ell_k^{-1} \to 0$, for all $k \geq k_0$.

Further, $|S_k(G, o)| \leq c_2 \ell_k^d$, and $o \in S_k(G, \boldsymbol{x})$ if and only if $\boldsymbol{x} \in S_k(G, o)$. Therefore

$$
\begin{aligned}
\alpha_k(G, o) &\geq \sum_{\boldsymbol{x} \in V(G)} \mathbf{1}_{o \in S_k(G, \boldsymbol{x})} \frac{1}{c_2 \ell_k^d} \\
&\geq \frac{1}{c_2 \ell_k^d} \sum_{\boldsymbol{x} \in V(G)} \mathbf{1}_{\boldsymbol{x} \in S_k(G, o)} = \frac{1}{c_2 \ell_k^d} |S_k(G, o)|.
\end{aligned}
$$

Therefore $\liminf_{k \to \infty} \alpha_k(G, o) \geq c_0/c_2 > 0$ a.s., whence $\rho\big(\alpha_k(G, o) \leq \delta\big) \to 0$ for all $\delta < c_0/c_2$.

Example 2 (Random geometric graph) In this case the vertices are the points of a Poisson point process on $\mathbb{R}^d$ with constant intensity γ. Any two vertices $\boldsymbol{x}, \boldsymbol{y}$ are connected by an edge if and only if $\|\boldsymbol{x} - \boldsymbol{y}\|_2 \leq r$ for a fixed radius $r > 0$. We choose the root $o \in V(G)$ as the closest vertex to the origin $\mathbf{0}$ and let G be the connected component of $o \in V$. This graph is infinite with positive probability provided γ is larger than the percolation threshold γ_c for this model [49].

The calculations for Bernoulli bond percolation on $\mathbb{Z}^d$ can be applied almost verbatim to the random geometric graph. In particular, letting $S_k(G, o) = \{\boldsymbol{x} \in V(G) : \|\boldsymbol{x}\|_\infty \leq \ell_k\}$ witnesses anchored-amenability and satisfies the tameness assumption.

4 Results for asymptotically amenable graphs

Recall that $Y_G^{(\varepsilon)}$ refers to the union of the vertex- and edge-observations over G: $Y_G^{(\varepsilon)} = \{Y_{uv} : (u, v) \in E, \xi_u^{(\varepsilon)} : u \in V\}$. A natural way to construct an estimator $\hat{\boldsymbol{\theta}}$ is to first estimate the posterior marginals of $\boldsymbol{\theta}$ given $Y_{G_n}^{(\varepsilon)}$ at every vertex:

$$
\mu_{G_n, u}(x) := \mathbb{P}\left(\theta_u = x \big| Y_{G_n}^{(\varepsilon)}\right), \quad \text{for } u \in V_n \text{ and } x \in \mathcal{X}. \tag{4.1}
$$

Letting $(\widehat{\mu}_u)_{u \in G_n}$ be such estimates of the posterior marginals, we can construct $\hat{\boldsymbol{\theta}}$, for instance, by independently sampling from the marginals: $\hat{\theta}_u \sim_{\mathrm{ind}} \widehat{\mu}_u$, for all $u \in V_n$.

Of course, computing the exact posterior probabilities $\mu_{G_n,u}(x)$ is in general intractable. As a tractable alternative, we can compute a *local* version of the vertex marginals by using only observations in a ball of radius l around each vertex. For $u \in V_n$ and $x \in \mathcal{X}$, let

$$\widehat{\mu}_{G_n,u,l}(x) := \mathbb{P}\Big(\theta_u = x \,\big|\, Y^{(\varepsilon)}_{B_{G_n}(u,l)}\Big). \tag{4.2}$$

(Recall that $B_{G_n}(u,l) = \{v \in V : d_{G_n}(u,v) \le l\}$ denotes the set of vertices within graph distance l form u in G_n.) The local marginals $\widehat{\mu}_{G_n,u,l}(x)$ can be computed with complexity at most $|\mathcal{X}|^{|B_{G_n}(u,l)|}$ per vertex. The complexity of estimating all the vertex marginals is linear or nearly linear, under additional assumptions. In particular:

- If G_n has degree bounded by k_{max} independently of n, then $|B_{G_n}(u,l)| \le k_{\mathrm{max}}^{l+1}$.
- If G_n converges to a locally finite unimodular graph, then

$$\lim_{M \to \infty} \lim_{n \to \infty} \mathbb{P}(|B_{G_n}(o_n, l)| \ge M) = 0.$$

In other words, for each δ, there exists $M(\delta)$ such that, for all n large enough, all but a fraction δ of the vertices u have neighborhood of size bounded by $M(\delta)$. Hence $\mu_{G_n,u}(x)$ can be estimated for all but a fraction δ of the vertices in linear time.

Notice that we can safely neglect $o(n)$ atypical vertices for our purposes. For instance, the matrix estimation risk (1.4) is bounded away in the present setting (unless the channel Q is noiseless), and therefore ignoring $o(n)$ vertices has a negligible impact on the asymptotic risk.

Do the local estimates $\widehat{\mu}_{G_n,u,l}$ provide good approximations of the actual marginals $\mu_{G_n,u}$? Our first result shows that this is the case for asymptotically amenable graphs, for almost all $\varepsilon > 0$, and on average over vertices in G_n. This is a consequence of the stronger fact that the state of a vertex becomes independent of a receding boundary around it in the limit of a large graph.

We state this result in terms of the mutual information. To this end let us briefly recall some notions form information theory. Given a discrete random variable (or random vector) X, we denote by $H(X)$ the Shannon entropy of the law of X, namely—with a slight abuse of notation—$H(X) = H(P_X) = -\sum_x P_X(x) \log P_X(x)$. The conditional entropy is defined by $H(X|Y) = H(X,Y) - H(Y)$, and the mutual information by $I(X;Y) = H(X) - H(X|Y) = H(Y) - H(Y|X) = d_{\mathrm{KL}}(\mathbb{P}(X, Y \in \cdot), \mathbb{P}(X \in \cdot) \times \mathbb{P}(Y \in \cdot))$, where d_{KL} is the Kulback-Liebler divergence, or relative entropy. Similarly, the conditional mutual information is defined as $I(X;Y|Z) = H(X|Z) - H(X|Y,Z) = \mathbb{E}_Z d_{\mathrm{KL}}(\mathbb{P}(X, Y \in \cdot|Z), \mathbb{P}(X \in \cdot|Z) \times \mathbb{P}(Y \in \cdot|Z))$.

Theorem B *Let $G_n = (V_n, E_n)$ be a sequence of finite graphs (with $|V_n| = n$) that converges locally-weakly to random rooted graph $(G, o) \sim \rho$ which is almost surely anchored-amenable and tame. Then for almost every $\varepsilon > 0$,*

$$\lim_{l\to\infty} \lim_{n\to\infty} \frac{1}{n} \sum_{u\in V_n} I\left(\theta_u; \theta_{\partial B_{G_n}(u,\ell)} \big| Y^{(\varepsilon)}_{B_{G_n}(u,l)}\right) = 0. \tag{4.3}$$

Consequently,

$$\lim_{l\to\infty} \lim_{n\to\infty} \frac{1}{n} \sum_{u\in V_n} \mathbb{E}\left[d_{\mathrm{TV}}(\widehat{\mu}_{G_n,u,l}, \mu_{G_n,u})\right] = 0. \tag{4.4}$$

The proof to this theorem is presented in Sect. 6.

Remark 1 While Theorem B is stated in non-quantitative form, the main technical lemma in its proof, Lemma 2, is entirely non-asymptotic. Its statement implies that, for most $\varepsilon' \in [0, \varepsilon]$, and most vertices $u \in S$, we have $I\left(\theta_u; \theta_{\partial S} \big| Y^{(\varepsilon')}_S\right) \leq C|\partial S|/(\varepsilon|S|)$.

While this lemma allows to control the dependence upon ε, determining the tradeoff between ε and n would requite stronger asusmptions about the vanishing of the ratio $|\partial S|/|S|$, i.e. the rate at which the infimum in Definition 4 is achieved by bounded regions.

Note that Theorem B is not sufficient to establish Theorem A about the optimality of polynomial-time algorithms to estimate the pairwise correlations $(X_f)_{u,v} = f(\theta_u) f(\theta_v)$. Indeed, the latter requires to approximate the joint distribution of θ_u, θ_v for $u, v \in V_n$ two arbitrary vertices. In order to achieve this goal, we define a *decoupled* estimator:

$$\widehat{X}^{(\mathrm{dec})}_{uv} := \mathbb{E}\left[f(\hat{\theta}_u) \big| Y^{(\varepsilon)}_{G_n}\right] \cdot \mathbb{E}\left[f(\hat{\theta}_v) \big| Y^{(\varepsilon)}_{G_n}\right]$$
$$= \left(\sum_{x\in\mathcal{X}} \mu_{G_n,u}(x) f(x)\right) \cdot \left(\sum_{x\in\mathcal{X}} \mu_{G_n,v}(x) f(x)\right), \quad u, v \in V_n. \tag{4.5}$$

Note that $\widehat{X}^{(\mathrm{dec})}$ may *a priori* have suboptimal accuracy. This is however not the case for almost all ε.

Proposition 1 *Let $\widehat{X}^{(\mathrm{dec})} \in \mathbb{R}^{n\times n}$ be defined as per Eq. (4.5). Then for almost every $\varepsilon > 0$,*

$$\lim_{n\to\infty} \left\{\mathcal{R}_n(\widehat{X}^{(\mathrm{dec})}; f) - \mathcal{R}^{\mathrm{Bayes}}_n(f)\right\} = 0.$$

The proof of the above proposition can be found in Sect. 10.

Given Theorem B and Proposition 1, it is natural to consider the following low complexity version of $\widehat{X}^{(\mathrm{dec})}$:

$$\widehat{X}^{(l)}_{uv} := \left(\sum_{x\in\mathcal{X}} \widehat{\mu}_{G_n,u,l}(x) f(x)\right) \cdot \left(\sum_{x\in\mathcal{X}} \widehat{\mu}_{G_n,v,l}(x) f(x)\right). \tag{4.6}$$

Since we can compute $\widehat{\mu}_{G_n,u,l}(x)$ for all but $o(n)$ vertices in time $\mathcal{O}(1)$, the overall complexity of $\widehat{X}^{(l)}$ is $\mathcal{O}(n^2)$. (Setting $\widehat{X}^{(l)}_{uv} = 0$ for a sublinear fraction of vertices produces a negligible error.) We can now prove Theorem A.

Proof of Theorem A Since Proposition 1 yields $\mathcal{R}_n(\widehat{\boldsymbol{X}}^{(\mathrm{dec})}; f) - \mathcal{R}_n^{\mathrm{Bayes}}(f) \to 0$ for almost all $\varepsilon > 0$, we only need to compare the risks of $\widehat{\boldsymbol{X}}^{(l)}$ and $\widehat{\boldsymbol{X}}^{(\mathrm{dec})}$. We have

$$\mathcal{R}_n(\widehat{\boldsymbol{X}}^{(l)}; f) - \mathcal{R}_n(\widehat{\boldsymbol{X}}^{(\mathrm{dec})}; f) = -\frac{2}{n^2} \mathbb{E}\langle \widehat{\boldsymbol{X}}^{(l)} - \widehat{\boldsymbol{X}}^{(\mathrm{dec})}, \boldsymbol{X}_f\rangle$$
$$+ \frac{1}{n^2}\left(\mathbb{E}\|\widehat{\boldsymbol{X}}^{(l)}\|_F^2 - \mathbb{E}\|\widehat{\boldsymbol{X}}^{(\mathrm{dec})}\|_F^2 \right).$$

We have

$$\mathbb{E}\langle \widehat{\boldsymbol{X}}^{(l)} - \widehat{\boldsymbol{X}}^{(\mathrm{dec})}, \boldsymbol{X}_f\rangle = \sum_{u,v\in V_n} \mathbb{E}\Big[\Big(\mathbb{E}[f(\theta_u)|Y_{B_{G_n}(u,l)}^{(\varepsilon)}]\mathbb{E}[f(\theta_v)|Y_{B_{G_n}(v,l)}^{(\varepsilon)}]$$
$$- \mathbb{E}[f(\theta_u)|Y_{G_n}^{(\varepsilon)}]\mathbb{E}[f(\theta_v)|Y_{G_n}^{(\varepsilon)}] \Big) f(\theta_u)f(\theta_v) \Big].$$

By consecutive triangle inequalities, this is bounded in absolute value by

$$\|f\|_\infty^2 \sum_{u,v\in V_n} \mathbb{E}\Big[\Big| \mathbb{E}[f(\theta_u)|Y_{B_{G_n}(u,l)}^{(\varepsilon)}]\mathbb{E}[f(\theta_v)|Y_{B_{G_n}(v,l)}^{(\varepsilon)}] - \mathbb{E}[f(\theta_u)|Y_{G_n}^{(\varepsilon)}]\mathbb{E}[f(\theta_v)|Y_{G_n}^{(\varepsilon)}] \Big| \Big]$$
$$\leq 2n\|f\|_\infty^3 \sum_{u\in V_n} \mathbb{E}\Big[\Big| \mathbb{E}[f(\theta_u)|Y_{B_{G_n}(u,l)}^{(\varepsilon)}] - \mathbb{E}[f(\theta_u)|Y_{G_n}^{(\varepsilon)}] \Big| \Big]$$
$$\leq 2n\|f\|_\infty^4 \sum_{u\in V_n} \sum_{x\in\mathcal{X}} \mathbb{E}\big[|\widehat{\mu}_{G_n,u,l}(x) - \mu_{G_n,u}(x)| \big]$$
$$= 4n\|f\|_\infty^4 \sum_{u\in V_n} \mathbb{E}\big[d_{\mathrm{TV}}(\widehat{\mu}_{G_n,u,l}, \mu_{G_n,u}) \big].$$

Here, $\|f\|_\infty$ denotes the supremum norm of f.

On the other hand, and following a similar strategy,

$$\mathbb{E}\|\widehat{\boldsymbol{X}}^{(l)}\|_F^2 - \mathbb{E}\|\widehat{\boldsymbol{X}}^{(\mathrm{dec})}\|_F^2 \leq 2n\|f\|_\infty^3 \sum_{u\in V_n} \mathbb{E}\Big[\Big| \mathbb{E}[f(\theta_u)|Y_{B_{G_n}(u,l)}^{(\varepsilon)}] - \mathbb{E}[f(\theta_u)|Y_{G_n}^{(\varepsilon)}] \Big| \Big]$$
$$\leq 4n\|f\|_\infty^4 \sum_{u\in V_n} \mathbb{E}\big[d_{\mathrm{TV}}(\widehat{\mu}_{G_n,u,l}, \mu_{G_n,u}) \big].$$

Invoking Theorem B concludes the proof.   □

We briefly mention that Theorem B allows to control other metrics for the estimation errors beyond $\mathcal{R}_n(\widehat{\boldsymbol{X}}; f)$; for example the 'overlap' (or 'agreement') metric that applies to estimators $\hat{\boldsymbol{\theta}} : \mathcal{Y}^{E_n} \times \mathcal{X}_\star^{V_n} \to \mathcal{X}^{V_n}$ which assign labels to vertices. We define

$$\mathtt{overlap}(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}) := \max_{\sigma\in\mathcal{S}_q} \frac{1}{|V_n|} \sum_{u\in V_n} \mathbf{1}\{\hat{\theta}_u = \sigma(\theta_u)\}, \tag{4.7}$$

where $\mathcal{S}_q$ is the set of permutations on $\mathcal{X}$, with $q = |\mathcal{X}|$. The overlap between a sample from the local marginals and $\boldsymbol{\theta}$ can be lower-bounded in a nontrivial way:

Corollary 1 *For each let $l \geq 1$, let $\hat{\boldsymbol{\theta}}^{(l)} = (\hat{\theta}_u^{(l)})_{u \in V_n}$ where $\hat{\theta}_u^{(l)} \sim \widehat{\mu}_{G_n,u,l}$ independently for all $u \in V_n$. Then for almost every $\varepsilon > 0$,*

$$\liminf_{l \to \infty} \lim_{n \to \infty} \mathbb{E}\big[\texttt{overlap}(\hat{\boldsymbol{\theta}}^{(l)}, \boldsymbol{\theta})\big] \geq \sum_{x \in \mathcal{X}} \mathbb{E}\big[\mu_{G,o}(x)^2\big].$$

In the above, $\mu_{G,o}(x) = \mathbb{P}(\theta_o = x | (G, o), Y_G^{(\varepsilon)})$ where an observation model $(\boldsymbol{\theta}, Y, \xi^{(\varepsilon)})$ is defined on the infinite random graph (G, o). We refer to Sect. 6.2 for the precise definition. The proof of the above corollary can be found in Sect. 10.

Notice that, while Theorem A guarantees that there exists a polynomial-time algorithm that achieves the Bayes risk, Corollary 1 only provides a lower bound. The challenge is that the definition of overlap in Eq. (4.7) includes a maximization over permutations σ, which is not needed for the definition of matrix estimation risk Theorem A. The lower bound of Corollary 1 corresponds to a specific choice of such permutation.

Remark 2 Theorem A and Corollary 1 have algorithmic implications also for the case in which no vertex information is available. A possible approach is as follows:

1. Fix a small $\varepsilon > 0$ let $m \sim \text{Binom}(n, \varepsilon)$. If $m > 2n\varepsilon$, return a fixed vector. Otherwise, select m vertices $U \subseteq V$ uniformly at random.
2. Enumerate all $|\mathcal{X}|^m$ assignments of variables of those vertices $\boldsymbol{\theta}'_U = (\theta'_u)_{u \in U}$.
3. For each such assignment, run the local algorithm for estimating $\boldsymbol{\theta}$, using $\boldsymbol{\theta}'_U$ as the vertex observations, i.e. setting $\xi_u^{(\varepsilon)} = \theta_u$ for $u \in U$ and $\xi_u^{(\varepsilon}) = \star$ otherwise.

This algorithm has complexity of order $\exp\{Cn\varepsilon\}$ and returns a list of at most $N = |\mathcal{X}|^{2n\varepsilon}$ estimates $(\widehat{\boldsymbol{X}}^{(i)})_{i \leq m}$, one of which achieves the guarantees of Theorem A. Since ε can be taken arbitrarily small, this yields a sub-exponential algorithm returning a list of sub-exponential size. This is in general not achieved on non-amenable graphs.

It is an interesting open question whether, given such a list, it is possible to select efficiently a single estimate.

5 Results for random regular graphs

The assumption of anchored-amenability is crucial in the proofs of Theorems A and B. While we do not know whether a weaker condition is sufficient, we show that these results do not hold for at least one non-amenable case, namely, when G_n is a random k-regular graph with constant degree k. For the case of $\mathbb{Z}_q$-synchronization we show that in a certain regime of signal-to-noise ratio (SNR), the local estimates of vertex marginals provide no information about the hidden assignment $\boldsymbol{\theta}$, while in the same regime, it is information-theoretically possible to estimate $\boldsymbol{\theta}$ non-trivially.

As mentioned in the introduction, an information-computation gap has been observed in several statistical models. However, none of the rigorous results in the literature matches the setting of Theorems A and B. To the best of our knowledge, the closest example is the case of the stochastic block model with q communities on

sparse random graphs (see [1] for a comprehensive survey and references therein). As explained in Sect. 3.1, this example fits our framework, although with G_n being the complete graph. In particular, G_n does not converge to a locally finite graphs. In contrast, the example treated in this section satisfies all the assumptions of Theorems A and B except amenability (and tameness).

Notice indeed that random k-regular tree converge locally to the infinite k-regular rooted tree T_k. The latter is the prototypical non-amenable graph. Indeed, for any connected subset of vertices $S \subseteq V(T_k)$, it is easy to see that $|\partial S| = (k-2)|S| + 2$.

Proofs for this section are deferred to Appendices 7 and 8.

5.1 Information-theoretic reconstruction: an exhaustive search algorithm

Given a graph $G = (V, E)$ on n vertices, $\boldsymbol{\theta} \in \mathcal{X}^V$ and $Y \in \mathcal{Y}^E$, we define the edge empirical distribution

$$\hat{v}_{\boldsymbol{\theta},Y}^G := \frac{1}{|E|} \sum_{(u,v) \in E} \delta_{(\theta_u, \theta_v, Y_{uv})} . \tag{5.1}$$

This is a probability distribution on $\mathcal{X} \times \mathcal{X} \times \mathcal{Y}$: $\hat{v}_{\boldsymbol{\theta},Y}^G \in \mathscr{P}(\mathcal{X} \times \mathcal{X} \times \mathcal{Y})$. (Recall that $\mathscr{P}(S)$ denotes the simplex of probability distributions over the set S.) Define $\overline{v} \in \mathscr{P}(\mathcal{X})$ to be the uniform distribution on $\mathcal{X}$ and $\overline{v}_e \in \mathscr{P}(\mathcal{X} \times \mathcal{X} \times \mathcal{Y})$ via

$$\overline{v}_e(\theta_1, \theta_2, y_{12}) = \overline{v}(\theta_1) \, \overline{v}(\theta_2) \, Q(y_{12}|\theta_1, \theta_2) .$$

We then define the set of 'typical' assignments of node variables by

$$\Theta(\eta; G, Y) := \left\{ \boldsymbol{\theta} \in \mathcal{X}^V : \ d_{\mathrm{TV}}(\hat{v}_{\boldsymbol{\theta},Y}^G, \overline{v}_e) \leq \eta \right\} .$$

We then consider the reconstruction algorithm that outputs a typical configuration

$$\hat{\boldsymbol{\theta}}(G, Y) \in \Theta(\eta_n; G, Y), \qquad \eta_n := \frac{\log n}{\sqrt{n}} . \tag{5.2}$$

If $\Theta(\eta_n; G, Y)$ is empty, we define $\hat{\boldsymbol{\theta}}(G, Y)$ arbitrarily (for instance $\hat{\boldsymbol{\theta}}(G, Y) = \boldsymbol{\theta}_*$ for a fixed reference configuration $\boldsymbol{\theta}_* \in \mathcal{X}^V$). If $\Theta(\eta_n; G, Y)$ contains more than one element, then $\hat{\boldsymbol{\theta}}(G, Y)$ selects one arbitrarily, e.g., the first one in lexicographic order. We note that we do not use the additional side information $\xi_{G_n}^{(\varepsilon)}$ in this procedure (we can equally assume that $\varepsilon = 0$ for the remainder of Sect. 5.1.) In fact our proofs apply to any algorithm that satisfy condition (5.2) with high probability. As discussed below (see Remark 3) this condition is also satisfied by the randomized estimator $\hat{\boldsymbol{\theta}} \sim \mathbb{P}(\cdot | Y_{G_n})$ that samples from the posterior.

It is immediate to show that the typical set is non-empty with high probability. (Throughout this section, we use $\boldsymbol{\theta}_0$ for the ground truth, in order to distinguish it from a generic vector $\boldsymbol{\theta} \in \mathcal{X}^n$.)

Lemma 1 *Let G_n be a random k-regular graph on n vertices, and let (θ_0, Y) be distributed according to the random observation model described in the Introduction. Then, there exists $c_0 = c_0(|\mathcal{X}|, |\mathcal{Y}|) > 0$ such that*

$$\mathbb{P}\Big(\theta_0 \in \Theta(\eta_n; G_n, Y)\Big) \geq 1 - c_0^{-1} \exp\big\{ - c_0(\log n)^2\big\}\,.$$

Remark 3 As mentioned above, one might consider a randomized estimator $\hat{\theta}$ that outputs a sample from the posterior: $\hat{\theta} \sim \mathbb{P}(\,\cdot\,|Y_{G_n})$. Note that this satisfies the condition $\hat{\theta} \in \Theta(\eta_n, G_n, Y)$ (cf. Eq. (5.2)) with the same probability $1 - c_0^{-1} \exp\{-c_0(\log n)^2\}$. Indeed this follows simply by noting that, with this definition, the pair $(\hat{\theta}, Y)$ is distributed as (θ_0, Y). Therefore all the results to follow apply to this randomized estimator as well.

Given two assignments $\theta_0, \theta \in \mathcal{X}^V$, we define their joint empirical vertex distribution as

$$\hat{\omega}_{\theta_0, \theta} := \frac{1}{|V|} \sum_{u \in V} \delta_{\theta_{0,u}, \theta_u}\,. \tag{5.3}$$

This is a probability distribution on $\mathcal{X} \times \mathcal{X}$: $\hat{\omega}_{\theta_0, \theta} \in \mathscr{P}(\mathcal{X} \times \mathcal{X})$.

Theorem C *Assume there exists $c_M > 0$ such that $c_M \leq Q(y|x_1, x_2)$ for all $x_1, x_2 \in \mathcal{X}$, $y \in \mathcal{Y}$, and let (θ_1, θ_2, Y) have joint distribution $\overline{v}_e$ (recall that $\overline{v}_e(x_1, x_2, y) = \overline{v}(x_1)\overline{v}(x_2) \times Q(y|x_1, x_2)$ where $\overline{v}$ is the uniform distribution over $\mathcal{X}$). If*

$$\frac{k}{2} I(\theta_1, \theta_2; Y) \geq H(\theta_1) + \delta\,, \tag{5.4}$$

for some $\delta > 0$, then there exists $\delta' = \delta'(\delta, c_M) > 0$ and a constant $c_0 > 0$ such that

$$\mathbb{P}\Big(d_{\mathrm{TV}}(\hat{\omega}_{\hat{\theta}, \theta_0}, \overline{v} \times \overline{v}) \geq \delta'\Big) \geq 1 - c_0^{-1} \exp\{-c_0(\log n)^2\}\,.$$

The proof of this theorem relies on a truncated first moment method, where we count the expected number of typical assignments $\theta \in \Theta(\eta_n; G_n, Y)$ having a given value of the empirical overlap distribution $\hat{\omega}_{\theta, \theta_0}$, conditioned on certain typicality constraints on the instance (G_n, θ_0, Y). The full argument is deferred to Sect. 7.2. This type of approach is standard in information theory [28]. Abbé and Sandon used it to analyze the stochastic block model, and called it "typicality sampling" see [5]. We also refer to [6,22] for similar calculations in somewhat simpler contexts.

Remark 4 For the sake of comparison, Sect. 9 presents an analogous of Theorem C for a version of the stochastic block model with on regular graphs. The model we study is essentially equivalent to the *equitable* stochastic block model of [10] (see also [11]) with the only difference that we allow for $O(1)$ self-loops and multi-edges, which appear with probability bounded away from 0.

For this version of the stochastic block model, the sufficient condition for weak recovery in Eq. (5.4) is replaced by a similar one which we briefly spell out here. Let $\pi \in \Delta^{q-1}$ be the distribution on q colors on the n vertices of the graph and $M \in \mathbb{R}_+^{q \times q}$ be the Markov transition matrix dictating how colors propagate across edges. Then reconstruction is possible if $\frac{k}{2} I(\theta_1; \theta_2) > H(\theta_1)$ where $\theta_1 \sim \pi$ and $\theta_2 \mid \theta_1 \sim M(\theta_1, \cdot)$. This condition reduces to the upper bound in [12, Theorem 1] for the special symmetric case considered in that paper. We refer to Sect. 9 for more details.

The next corollary applies the result of Theorem C to $\mathbb{Z}_q$-synchronization.

Corollary 2 *Consider the $\mathbb{Z}_q$-synchronization problem. If*

$$k > k_*(p; q) := \frac{2 \log q}{\left(1 - p + \frac{p}{q}\right) \log \left(p + q(1 - p)\right) + \left(1 - \frac{1}{q}\right) p \log p},$$

then there exists $\delta, c_0 > 0$ depending on k, p, q such that, with probability at least $1 - c_0^{-1} \exp\{-c_0 (\log n)^2\}$, $d_{\mathrm{TV}}(\hat{\omega}_{\hat{\theta}, \theta_0}, \overline{\nu} \times \overline{\nu}) \geq \delta$.
Furthermore, as $p \to 1$, we have

$$k_*(p; q) = \frac{4 \log q}{(q - 1)(1 - p)^2} + \mathcal{O}\left((1 - p)^{-1}\right).$$

This corollary follows from Theorem C simply by computing $I(\theta_1, \theta_2; Y)$ in the case of $\mathbb{Z}_q$-synchronization. We omit the details. Finally, we deduce from Theorem C the possibility of weak recovery.

Corollary 3 *Under the assumptions of Theorem C, if $\frac{k}{2} I(\theta_1, \theta_2; Y) \geq H(\theta_1) + \varepsilon$, then there exists a constant $\delta = \delta(\varepsilon) > 0$ such that*

$$\liminf_{n \to \infty} \mathbb{E}[\mathrm{overlap}(\hat{\theta}, \theta_0)] \geq \frac{1}{q} + \delta. \tag{5.5}$$

Moreover, there exists a function $f : \mathcal{X} \mapsto \mathbb{R}$ with zero mean, unit variance, and a constant $\delta = \delta(\varepsilon, |\mathcal{X}|, c_M) > 0$ such that

$$\limsup_{n \to \infty} \mathcal{R}_n^{\mathrm{Bayes}}(f) \leq 1 - \delta. \tag{5.6}$$

In particular, the conclusions (5.5) and (5.6) hold in the $\mathbb{Z}_q$-synchronization model if $k > k_*(p; q)$.

5.2 Performance of the local algorithm

In this section we examine the asymptotics of the local marginals

$$\widehat{\mu}_{G_n, u, l}(x) = \mathbb{P}\left(\theta_u = x \mid Y_{B_{G_n}(u, l)}^{(\varepsilon)}\right),$$

when G_n is a random k-regular graph, in the special case of $\mathbb{Z}_q$-synchronization with side information from BEC($\bar{\varepsilon}$).

We have seen in the previous section that weak recovery is possible (albeit non-efficiently) when $(1-p)^2 k > \frac{4 \log q}{q-1} + \mathcal{O}(1-p)$ even in the absence of side information (Corollary 3). We show on the other hand that the local marginals are approximately uniform if $(1 - p)^2(k - 1) < 1$. The latter condition is known as the *Kesten-Stigum threshold* for the problem of robust reconstruction on the tree [33].

Theorem D *Consider $\mathbb{Z}_q$-synchronization with side information from BEC($\bar{\varepsilon}$) on a random k-regular graph G_n. There exist constants $c = c(k, p, q)$ and $C = C(k, p, q)$ such that the following holds. If $(1 - p)^2(k - 1) < 1$ and $\varepsilon \le c$ then*

$$\limsup_{l \to \infty} \limsup_{n \to \infty} \frac{1}{n} \sum_{u \in V_n} \mathbb{E}\left[d_{\text{TV}}(\widehat{\mu}_{G_n, u, l}, \overline{v})^2\right] \le C\varepsilon. \tag{5.7}$$

The above theorem implies that all estimators $(\hat{\boldsymbol{\theta}}^{(l)})_{l \ge 1}$ where $\hat{\theta}_u^{(l)} \sim \widehat{\mu}_{u,l}$ independently for all $u \in V_n$, have almost trivial performance. Recall the definition of the matrix $\widehat{X}^{(l)}$:

$$\widehat{X}_{uv}^{(l)} = \mathbb{E}\left[f(\theta_u)\big|Y_{B_{G_n}(u,l)}^{(\varepsilon)}\right] \cdot \mathbb{E}\left[f(\theta_v)\big|Y_{B_{G_n}(v,l)}^{(\varepsilon)}\right], \quad u, v \in V_n.$$

Corollary 4 *In the setting of Theorem D, if $(1 - p)^2(k - 1) < 1$, then there exists constants $c_1, c_2 > 0$ depending on k and q such that*

$$\limsup_{l \to \infty} \limsup_{n \to \infty} \mathbb{E}\left[\texttt{overlap}(\hat{\boldsymbol{\theta}}^{(l)}, \boldsymbol{\theta}_0)\right] \le \frac{1}{q} + c_1\sqrt{\varepsilon}.$$

Moreover, for all $f : \mathbb{Z}_q \mapsto \mathbb{R}$ with zero mean and unit variance,

$$\liminf_{l \to \infty} \liminf_{n \to \infty} \mathcal{R}_n\left(\widehat{X}^{(l)}; f\right) \ge 1 - c_2 \|f\|_\infty^2 \varepsilon.$$

Remark 5 The above implies that no local algorithm can estimate X_f with non-trivial accuracy. Indeed, the estimator of $f(\theta_u)f(\theta_v)$ of minimal risk based on the information contained is the balls of radius l centered around u and v respectively is $\mathbb{E}\left[f(\theta_u)f(\theta_v)|Y_{B_{G_n}(u,l)\cup B_{G_n}(v,l)}^{(\varepsilon)}\right]$. The latter quantity is equal to $\widehat{X}_{uv}^{(l)}$ if the two balls are disjoint, which is the case for $1 - o_n(1)$ fraction of pairs of vertices (u, v) when l is held constant.

The proof of Theorem D is deferred to Sect. 8.1, but we give here an outline. We use local-weak convergence to first lift the problem to the infinite k-regular tree, in which the study of the local marginals reduces to the study of a certain distributional recursion. Then we prove that below the Kesten–Stigum threshold, the uniform distribution $\overline{v}$ is a stable fixed point of this recursion. The argument proceeds as follows. Let o be the

root of infinite $(k-1)$-ary tree T and denote by $T(l)$ the subtree consisting of the first l generations of T rooted at o. Now let $\mu_{o,l}(x) := \mathbb{P}\big(\theta_o = x | Y^{(\varepsilon)}_{T_k(l)}\big)$ for all $x \in \mathbb{Z}_q$ and consider the sequence $z_l := \mathbb{E}[\mu_{o,l}(\theta_o)|\xi_o = \star] - \frac{1}{q}$ which measures the deviation from uniformity of the local marginal at the root. We use the recursive structure of the tree to show that for ε small enough and $\kappa = (1-p)^2(k-1)$, the sequence $(z_l)_{l\geq 0}$ satisfies the approximate recursion

$$\Big| z_{l+1} - (1-\varepsilon)\kappa z_l - \varepsilon\kappa \frac{q-1}{q} \Big| \leq C(q)\kappa^2\big(z_l^2 + \varepsilon^2\big), \tag{5.8}$$

where $C(q)$ is constant depending only on q. Since $z_0 = 0$, this implies that if $\kappa < 1$ then the sequence stays within an interval of size $C'(q,\kappa)\varepsilon$ around the origin. This, in turn, can be converted to the claim of Theorem D. The analysis of this recursion originates in the study of the robust reconstruction problem on the tree. In this problem, a spin at the root (an $\mathcal{X}$-valued r.v.) is broadcast through noisy channels along edges of the tree. The statistician observes a noisy realization of this process on the leaves of $T(l)$ for large l, and is tasked with inferring the value at the root (see e.g., [25,33,47]). In particular, our analysis builds on ideas from [42,52].

6 Proof of Theorem B

6.1 Proof of Eq. (4.4)

We first deduce the statement about the average total distance between $\widehat{\mu}_{G_n,u,l}$ and $\mu_{G_n,u}$, Eq. (4.4) from the statement about the mutual information, Eq. (4.3). Since

$$\mathbb{P}\big(\theta_u = x | Y^{(\varepsilon)}_{G_n}\big) = \mathbb{E}\Big[\mathbb{P}\big(\theta_u = x | \theta_{\partial B_{G_n}(u,l)}, Y^{(\varepsilon)}_{G_n}\big)|Y^{(\varepsilon)}_{G_n}\Big]$$
$$= \mathbb{E}\Big[\mathbb{P}\big(\theta_u = x | \theta_{\partial B_{G_n}(u,l)}, Y^{(\varepsilon)}_{B_{G_n}(u,l)}\big)|Y^{(\varepsilon)}_{G_n}\Big],$$

where the last line follows from the spatial Markov property satisfied by the model (i.e., conditionally on the boundary variables $\theta_{\partial B_{G_n}(u,l)}$, θ_u is independent of $Y^{(\varepsilon)}_{vw}$ for all $(v,w) \in E$ such that either v or w is outside the ball $B_{G_n}(u,l)$), we have by the triangle inequality (letting $S = B_{G_n}(u,l)$ for ease of notation),

$$\mathbb{E}d_{\mathrm{TV}}\big(\mu_{G_n,u}, \widehat{\mu}_{G_n,u,l}\big) = \frac{1}{2}\mathbb{E}\sum_{x\in\mathcal{X}}\Big|\mathbb{P}\big(\theta_u = x|Y^{(\varepsilon)}_{G_n}\big) - \mathbb{P}\big(\theta_u = x|Y^{(\varepsilon)}_S\big)\Big|$$

$$\leq \frac{1}{2}\mathbb{E}\sum_{x\in\mathcal{X}}\Big|\mathbb{P}\big(\theta_u = x|\theta_{\partial S}, Y^{(\varepsilon)}_S\big) - \mathbb{P}\big(\theta_u = x|Y^{(\varepsilon)}_S\big)\Big|$$

$$= \frac{1}{2}\mathbb{E}\sum_{x\in\mathcal{X}, \sigma\in\mathcal{X}^{\partial S}}\mathbb{P}\big(\theta_{\partial S} = \sigma|Y^{(\varepsilon)}_S\big)\cdot\Big|\mathbb{P}\big(\theta_u = x|\theta_{\partial S} = \sigma, Y^{(\varepsilon)}_S\big) - \mathbb{P}\big(\theta_u = x|Y^{(\varepsilon)}_S\big)\Big|$$

$$= \mathbb{E}d_{\mathrm{TV}}\Big(\mathbb{P}\big((\theta_u, \theta_{\partial S}) \in \cdot|Y^{(\varepsilon)}_S\big), \mathbb{P}\big(\theta_u \in \cdot|Y^{(\varepsilon)}_S\big) \times \mathbb{P}\big(\theta_{\partial S} \in \cdot|Y^{(\varepsilon)}_S\big)\Big)$$

$$\leq \sqrt{I\big(\theta_u; \theta_{\partial S}|Y_S^{(\varepsilon)}\big)/2},$$

where the last line follows from Pinsker's inequality. Therefore Eq. (4.4) follows from Eq. (4.3) via Jensen's inequality.

6.2 Proof of Eq. (4.3)

Now we prove Eq. (4.3). We define an observation model $(\boldsymbol{\theta}, Y, \xi^{(\varepsilon)})$ on the infinite random graph G exactly as for the finite graphs G_n. We then let

$$\mu_{G,o}(x) := \mathbb{P}\Big(\theta_o = x \big| (G, o), Y_G^{(\varepsilon)}\Big),$$

where we condition on the realization of the rooted graph and on σ-algebra generated by the sequence of random variables $(Y_{B_G(o,l)}^{(\varepsilon)})_{l \geq 0}$. Equivalently, we can also define $\mu_{G,o}(x)$ as the almost-sure limit of the sequence $(\mathbb{P}(\theta_o = x|(G,o), Y_{B_G(o,l)}^{(\varepsilon)}))_{l \geq 0}$, where convergence is guaranteed by Lévy's upward theorem.

Let $l \geq 1$. The functions $f, g : \mathcal{G}_* \to [0, \log|\mathcal{X}|]$ defined respectively by

$$f(G, o) = H(\theta_o|Y_{[G,o]_l}^{(\varepsilon)}) \quad \text{and} \quad g(G, o) = H(\theta_o|\theta_{\partial[G,o]_l}, Y_{[g,o]_l}^{(\varepsilon)})$$

are continuous in the topology of local convergence. Indeed for $(G_n, o_n) \xrightarrow{loc.} (G, o)$, let $n_0 \geq 1$ such that $[G_n, o_n]_l \equiv [G, o]_l$ for all $n \geq n_0$. Then $f(G_n, o_n) = f(G, o)$ for all $n \geq n_0$, and similarly for g. Since the entropy is bounded between 0 and $\log|\mathcal{X}|$, we obtain by local-weak convergence under uniform rooting that

$$\frac{1}{n} \sum_{u \in V_n} I\big(\theta_u; \theta_{\partial B_{G_n}(u,l)}|Y_{B_{G_n}(u,l)}^{(\varepsilon)}\big)$$

$$= \frac{1}{n} \sum_{u \in V_n} H(\theta_u|Y_{B_{G_n}(u,l)}^{(\varepsilon)}) - H(\theta_u|\theta_{\partial B_{G_n}(u,l)}, Y_{B_{G_n}(u,l)}^{(\varepsilon)})$$

$$= \mathbb{E}_{\rho_{G_n}}\big[f(G_n, o_n)\big] - \mathbb{E}_{\rho_{G_n}}\big[g(G_n, o_n)\big]$$

$$\xrightarrow{n \to \infty} \mathbb{E}_\rho\big[f(G, o)\big] - \mathbb{E}_\rho\big[g(G, o)\big]$$

$$= \mathbb{E}_\rho\Big[I\big(\theta_o; \theta_{\partial B_G(o,l)}|Y_{B_G(o,l)}^{(\varepsilon)}\big)\Big].$$

Furthermore, since conditioning reduces the entropy, we have for all $o \in S_1 \subseteq S_2 \subseteq V(G)$,

$$I\big(\theta_o; \theta_{\partial S_2}|Y_{S_2}^{(\varepsilon)}\big) \leq I\big(\theta_o; \theta_{\partial S_1}|Y_{S_1}^{(\varepsilon)}\big). \tag{6.1}$$

Therefore the map $l \to I\big(\theta_o; \theta_{\partial B_G(o,l)}|Y_{B_G(o,l)}^{(\varepsilon)}\big)$ is non-increasing. Let

$$I_\star^{(\varepsilon)}(G, o) := \lim_{l \to \infty} I\big(\theta_o; \theta_{\partial B_G(o,l)}|Y_{B_G(o,l)}^{(\varepsilon)}\big). \tag{6.2}$$

It remains to prove that $\mathbb{E}_\rho[I_\star^{(\varepsilon)}(G, o)] = 0$ for almost all $\varepsilon > 0$. We use the following generic lemma which allows to control the dependency under the posterior between variables θ_u associated to vertices in the interior of a set $S \subset V(G)$ and variables $\theta_{\partial S}$ associated to the boundary of this set. Our first lemma bounds the mutual information between θ_u and $\theta_{\partial S}$. This result is inspired by Lemma 3.1 in [44].

Lemma 2 *Let G be a graph, and $S \subset V(G)$ finite and non-empty. For all $\varepsilon \geq 0$, we have*

$$\sum_{u \in S} \int_0^\varepsilon I\left(\theta_u; \theta_{\partial S} \middle| Y_S^{(\varepsilon')}\right) d\varepsilon' \leq \log |\mathcal{X}| \cdot |\partial S|.$$

Proof The argument relies on differentiating the conditional Shannon entropy of $\theta_{\partial S}$ given $Y_S^{(\varepsilon)}$ with respect to ε. Let us first replace the single parameter ε (the probability of non-erasure) by a set of parameters $\underline{\varepsilon} = (\varepsilon_u)_{u \in S}$: for each vertex u, θ_u is revealed with probability ε_u. We also replace the notation $Y_S^{(\varepsilon)}$ by (Y, ξ), omitting an explicit reference to $\underline{\varepsilon}$ and to the ball S. We finally denote $\xi^{\backslash(u)} = \{\xi_v : v \in S, v \neq u\}$ with ξ_u removed. We have

$$H(\theta_{\partial S}|Y, \xi) = \varepsilon_u H\left(\theta_{\partial S}|Y, \xi^{\backslash(u)}, \theta_u\right) + (1 - \varepsilon_u) H\left(\theta_{\partial S}|Y, \xi^{\backslash(u)}\right).$$

Taking a derivative w.r.t. ε_u yields:

$$\frac{d}{d\varepsilon_u} H(\theta_{\partial S}|Y, Z) = H\left(\theta_{\partial S}|Y, \xi^{\backslash(u)}, \theta_u\right) - H\left(\theta_{\partial S}|Y, \xi^{\backslash(u)}\right)$$
$$= -I\left(\theta_u; \theta_{\partial S}|Y, \xi^{\backslash(u)}\right),$$

where the latter is the conditional mutual information of θ_u and $\theta_{\partial S}$ given $(Y, \xi^{\backslash(u)})$. Now we set $\varepsilon_u = \varepsilon$ for all $u \in S$. We obtain

$$\frac{d}{d\varepsilon} H(\theta_{\partial S}|Y, \xi) = -\sum_{u \in S} I\left(\theta_u; \theta_{\partial S}|Y, \xi^{\backslash(u)}\right).$$

We now integrate w.r.t. ε:

$$\int_0^\varepsilon \sum_{u \in S} I\left(\theta_u; \theta_{\partial S}|Y, \xi^{\backslash(u)}\right) d\varepsilon' = H(\theta_{\partial S}|Y, \xi^{(\varepsilon=0)}) - H(\theta_{\partial S}|Y, \xi^{(\varepsilon)})$$
$$\leq H(\theta_{\partial S}|Y, \xi^{(\varepsilon=0)})$$
$$\leq H(\theta_{\partial S})$$
$$\leq \sum_{u \in \partial S} H(\theta_u)$$
$$= \log |\mathcal{X}| \cdot |\partial S|.$$

The second line is by positivity of entropy, the third line follows from the fact that conditioning reduces the entropy, the fourth line is by sub-additivity, and the last line is since θ_u is marginally uniform on $\mathcal{X}$. Now we finish the proof by observing that $I\big(\theta_u; \theta_{\partial S}|Y, \xi\big) = I\big(\theta_u; \theta_{\partial S}|Y, \xi^{\setminus(u)}, \xi_u\big) \leq I\big(\theta_u; \theta_{\partial S}|Y, \xi^{\setminus(u)}\big)$ because the left-hand side vanishes whenever $\xi_u \neq \star$. □

Now we are in a position to prove Eq. (4.3).

Assume (G, o) is almost surely anchored-amenable and tame. Let $(S_k = S_k(G, o))_{k \geq 1}$ be the sequence of finite measurable random subsets of $V(G)$ satisfying the conditions of Definition 5 (recall in particular that $o \in S_k$.) We use Lemma 2 with this choice of sequence $(S_k)_{k \geq 1}$, and then average over the realization of the rooted graph $(G, o) \sim \rho$:

$$
\mathbb{E}_\rho \left[\frac{1}{|S_k(G,o)|} \sum_{u \in S_k(G,o)} \int_0^\varepsilon I\Big(\theta_u; \theta_{\partial S_k(G,o)}\big|Y^{(\varepsilon')}_{S_k(G,o)}\Big) \mathrm{d}\varepsilon' \right]
$$

$$
\leq \mathbb{E}_\rho \left[\sqrt{\log |\mathcal{X}| \cdot |\partial S_k(G,o)|/|S_k(G,o)|} \right]
$$

$$
\leq \sqrt{\log |\mathcal{X}|}\, \Delta_k,
$$

where $\Delta_k \to 0$ by an application of dominated convergence (since $|\partial S_k(G,o)|/|S_k(G,o)| \to 0$ almost surely by assumption). Now, for a given k, let $L \geq \mathrm{diam}(S_k(G,o))$ be such that $S_k(G,o) \subseteq B_G(u, L)$ for all $u \in S_k(G,o)$. By the inequality (6.1) we have for all $u \in S_k(G,o)$,

$$
I\Big(\theta_u; \theta_{\partial B_G(u,L)}\big|Y^{(\varepsilon')}_{B_G(u,L)}\Big) \leq I\Big(\theta_u; \theta_{\partial S_k(G,o)}\big|Y^{(\varepsilon')}_{S_k(G,o)}\Big),
$$

so we obtain

$$
\mathbb{E}_\rho \left[\frac{1}{|S_k(G,o)|} \sum_{u \in S_k(G,o)} \int_0^\varepsilon I\Big(\theta_u; \theta_{\partial B_G(u,L)}\big|Y^{(\varepsilon')}_{B_G(u,L)}\Big) \mathrm{d}\varepsilon' \right] \leq \sqrt{\log |\mathcal{X}|}\, \Delta_k. \quad (6.3)
$$

Now we let $f_k : \mathcal{G}_{**} \mapsto \mathbb{R}_+$ defined by

$$
f_k(G, o, u) := \frac{1}{|S_k(G,o)|} \mathbf{1}_{u \in S_k(G,o)} \int_0^\varepsilon I\Big(\theta_u; \theta_{\partial B_G(u,L)}\big|Y^{(\varepsilon')}_{B_G(u,L)}\Big) \mathrm{d}\varepsilon'.
$$

With this notation, expression (6.3) is equal to $\mathbb{E}_\rho\big[\sum_{u \in V(G)} f_k(G, o, u)\big]$. By unimodularity of ρ, this is also equal to

$$
\mathbb{E}_\rho \left[\sum_{u \in V(G)} f_k(G, u, o) \right]
$$

$$
= \mathbb{E}_\rho \left[\Big(\sum_{u \in V(G)} \frac{\mathbf{1}_{o \in S_k(G,u)}}{|S_k(G,u)|} \Big) \int_0^\varepsilon I\Big(\theta_o; \theta_{\partial B_G(o,L)}\big|Y^{(\varepsilon')}_{B_G(o,L)}\Big) \mathrm{d}\varepsilon' \right]. \quad (6.4)
$$

Taking the limit $L \to \infty$ and using Eq. (6.3) and Eq. (6.4) we obtain

$$\mathbb{E}_\rho \left[\alpha_k(G, o) \cdot \int_0^\varepsilon I_\star^{(\varepsilon')}(G, o) d\varepsilon' \right] \leq \sqrt{\log |\mathcal{X}|} \, \Delta_k, \qquad (6.5)$$

where $\alpha_k(G, o) = \sum_{u \in V(G)} \mathbf{1}_{o \in S_k(G, u)} |S_k(G, u)|^{-1}$ and $I_\star$ is defined in (6.2). Now we use the tameness assumption: the sequence $(1/\alpha_k(G, o))_k$ is tight. Let $Z(G, o) = \int_0^\varepsilon I_\star^{(\varepsilon')}(G, o) d\varepsilon'$ For $\eta > 0$ let $\delta > 0$ such that

$$\limsup_{k \to \infty} \rho(\alpha_k(G, o) \leq \delta) \leq \eta.$$

Since all involved quantities are nonnegative, we have

$$\mathbb{E}_\rho \big[\alpha_k(G, o) \cdot Z(G, o) \big] \geq \delta \mathbb{E}_\rho \big[\mathbf{1}_{\alpha_k(G, o) > \delta} \cdot Z(G, o) \big]$$
$$= \delta \big(\mathbb{E}_\rho \big[Z(G, o) \big] - \mathbb{E}_\rho \big[\mathbf{1}_{\alpha_k(G, o) \leq \delta} \cdot Z(G, o) \big] \big).$$

Since $Z(G, o) \leq \varepsilon \log |\mathcal{X}|$ a.s., we obtain

$$\mathbb{E}_\rho \big[Z(G, o) \big] \leq \delta^{-1} \mathbb{E}_\rho \big[\alpha_k(G, o) Z(G, o) \big] + \varepsilon \log |\mathcal{X}| \cdot \rho(\alpha_k(G, o) \leq \delta)$$
$$\leq \delta^{-1} \sqrt{\log |\mathcal{X}|} \, \Delta_k + \varepsilon \log |\mathcal{X}| \cdot \rho(\alpha_k(G, o) \leq \delta),$$

where we have used (6.5) to obtain the last display. Letting $k \to \infty$ and then $\eta \to 0$ we obtain for all $\varepsilon \geq 0$,

$$\mathbb{E}_\rho \left[\int_0^\varepsilon I_\star^{(\varepsilon')}(G, o) d\varepsilon' \right] = 0,$$

and this concludes the proof. $\qquad\square$

7 Information-theoretic reconstruction on random graphs: Proofs

7.1 Proof of Lemma 1

This is a consequence of McDiarmid's bounded differences inequality. For $(u, v) \in E$ and $(x_1, x_2, y_{12}) \in \mathcal{X} \times \mathcal{X} \times \mathcal{Y}$, we let $X_{uv}(x_1, x_2, y_{12}) = \mathbf{1}\{\theta_{0,u} = x_1, \theta_{0,v} = x_2, Y_{uv} = y_{12}\}$, and let $Z(x_1, x_2, y_{12}) = \frac{1}{|E|} \sum_{(u,v) \in E} (X_{uv}(x_1, x_2, y_{12}) - \mathbb{E}[X_{uv}(x_1, x_2, y_{12})])$. Since

$$d_{\mathrm{TV}}(\hat{v}_{\theta_0, Y}^{G_n}, \overline{v}_e) = \frac{1}{2} \sum_{x_1, x_2, y_{12}} \big| Z(x_1, x_2, y_{12}) \big|,$$

we have

$$\mathbb{P}\big(d_{\mathrm{TV}}(\hat{v}_{\boldsymbol{\theta}_0, Y}^{G_n}, \overline{v}_{\mathrm{e}}) \geq \eta\big) \leq \sum_{x_1, x_2, y_{12}} \mathbb{P}\Big(\big|Z(x_1, x_2, y_{12})\big| \geq \frac{2\eta}{|\mathcal{X}|^2 |\mathcal{Y}|}\Big).$$

We associate to each edge $(i, j) \in E$ an independent random variable $U_{ij} \sim Unif([0, 1])$. We can then construct a function $f : \mathcal{X} \times \mathcal{X} \times [0, 1] \to \mathcal{Y}$, such that $Q(y_{12}|\theta_1, \theta_2) = \mathbb{P}(f(\theta_1, \theta_2, U_{12}) = y_{12}|\theta_1, \theta_2)$. Hence we can define $Y, \boldsymbol{\theta}_0$ by letting $Y_{uv} = f(\theta_{0,u}, \theta_{0,v}, U_{uv})$ for each $(u, v) \in E$, and we view $Z(x_1, x_2, y_{12})$ as a function of the independent random variables $\{\theta_{0,u}, U_{uv}\}$.

Moreover, if we change the value $\theta_{0,u}$ at vertex u to $\theta'_{0,u}$ and call $Z'(x_1, x_2, y_{12})$ the resulting value of $Z(x_1, x_2, y_{12})$, we have $|Z - Z'| \leq \frac{k}{|E|} = \frac{2}{n}$ (recall that k is the degree of u and $|E| = \frac{nk}{2}$). If we further change U_{uv} to U'_{uv} at an edge $(u, v) \in E$, we have $|Z - Z'| \leq 1/|E|$. Applying the bounded differences inequality conditionally on G_n, we get

$$\mathbb{P}\Big(\big|Z(x_1, x_2, y_{12})\big| \geq \eta'\Big) \leq 2 \exp\bigg\{-\frac{2\eta'^2}{(n(\frac{2}{n})^2 + |E|(\frac{1}{|E|})^2)}\bigg\} \leq 2e^{-n\frac{\eta'^2}{3}}.$$

Now we let $\eta' = 2\frac{\eta}{|\mathcal{X}|^2 |\mathcal{Y}|}$ and $\eta = \frac{\log n}{\sqrt{n}}$.

7.2 Proof of Theorem C: a truncated first moment method

Instead of working directly with the ensemble of random regular graphs, we will use the *configuration model* [17] for our moment computations. Let kn be even and let $\mathcal{M}_{nk}$ be the set of perfect matchings on nk vertices. For $\mathfrak{m} \in \mathcal{M}_{nk}$ we define the multi-graph $G(\mathfrak{m})$ on n vertices where a vertex $i' \in [nk]$ in $\mathfrak{m}$ is sent to a vertex i in $G(\mathfrak{m})$ through the mapping $i' \mapsto i = i' \ (mod) \ n$. The resulting multi-graph may contain multiple edges and self-loops. The configuration model is the probability measure $\mathbb{P}_{n,k}^{\mathrm{cm}}$ on multi-graphs induced by the uniform measure on perfect matchings through the above mapping. The measure $\mathbb{P}_{n,k}^{\mathrm{cm}}$ conditioned on the multi-graph $G(\mathfrak{m})$ being *simple* (i.e., not having self-loops nor multiple edges) is the uniform measure on k-regular graphs $\mathbb{P}_{n,k}^{\mathrm{reg}}$. The probability that $G(\mathfrak{m})$ is simple under $\mathbb{P}_{n,k}^{\mathrm{cm}}$ is $(1 - \mathcal{O}(k^3/n))e^{(1-k^2)/4}$ for large n by a formula of McKay and Wormald [53]. Therefore, for any event A and sequence $\varepsilon_n \to 0$, $\mathbb{P}_{n,k}^{\mathrm{cm}}(A) \geq 1 - \varepsilon_n$ implies $\mathbb{P}_{n,k}^{\mathrm{reg}}(A) \geq 1 - c(k)\varepsilon_n$ with $c(k) > 0$ depending only on k.

Let $G_n = (V_n, E_n)$ be from the configuration model with $V_n = [n]$. We will assume edges to be directed, and the direction to be chosen uniformly at random. Assigning directions is useful for book-keeping when counting tuples $(G, \boldsymbol{\theta}, Y)$. The number of graphs $G_n = (V_n, E_n)$, (counting edge directions) is

$$N_{n,k} = \frac{(nk)!}{(nk/2)!} = \exp\bigg\{\frac{nk}{2} \log\Big(\frac{2nk}{e}\Big) + \mathcal{O}(1)\bigg\}. \tag{7.1}$$

Indeed, $N_{n,k}$ is the number of ordered pairings of the nk half-edges. Such a pairing can be constructed by ordering the nk half-edges (which can be done in $(nk)!$ possible ways), and then pairing consecutive half edges following this ordering. Each pairing can arise in $(nk/2)!$ possible ways.

We next state a standard counting lemma that will be useful in what follows. Given finite alphabets $\overline{\mathcal{X}}$, $\overline{\mathcal{Y}}$, and integers n, k with $nk = 2m$ even, let $\mathscr{P}_k(\overline{\mathcal{X}} \times \overline{\mathcal{X}} \times \overline{\mathcal{Y}}) \subseteq \mathscr{P}(\overline{\mathcal{X}} \times \overline{\mathcal{X}} \times \overline{\mathcal{Y}})$ be the subset of probability distributions $\nu \in \mathscr{P}(\overline{\mathcal{X}} \times \overline{\mathcal{X}} \times \overline{\mathcal{Y}})$ such that $\nu(x_1, x_2, y) \in \mathbb{N}/m$ for all $x_1, x_2 \in \overline{\mathcal{X}}$, $y \in \overline{\mathcal{Y}}$, and $\sum_{\tilde{x}, y}(\nu(x, \tilde{x}, y) + \nu(\tilde{x}, x, y)) \in \mathbb{N}/n$ for all $x \in \overline{\mathcal{X}}$.

Given $\nu \in \mathscr{P}(\overline{\mathcal{X}} \times \overline{\mathcal{X}} \times \overline{\mathcal{Y}})$, we let $\pi_1\nu(x) \equiv \sum_{\tilde{x} \in \overline{\mathcal{X}}, y \in \overline{\mathcal{Y}}} \nu(x, \tilde{x}, y)$, $\pi_2\nu(x) \equiv \sum_{\tilde{x} \in \overline{\mathcal{X}}, y \in \overline{\mathcal{Y}}} \nu(\tilde{x}, x, y)$. We further let $\pi_{12}\nu(x_1, x_2) = \sum_{y \in \overline{\mathcal{Y}}} \nu(x_1, x_2, y)$.

Recall that Shannon entropy of a probability distribution p on the finite set $\mathcal{X}$ is $H(p) = -\sum_{x \in S} p(x) \log p(x)$, and the joint empirical edge distribution of $(\boldsymbol{\theta}, Y)$ on a graph G is

$$\hat{\nu}_{\boldsymbol{\theta}, Y}^G = \frac{1}{|E|} \sum_{(u,v) \in E} \delta_{(\theta_u, \theta_v, Y_{uv})} \in \mathscr{P}(\mathcal{X} \times \mathcal{X} \times \mathcal{Y}).$$

Lemma 3 *For such ν, let $N_{n,k}(\nu)$ be the number of triples $(G, \boldsymbol{\theta}, Y)$ where $G = (V = [n], E)$ is a graph from the configuration model, $\boldsymbol{\theta} \in \overline{\mathcal{X}}^V$, $Y \in \overline{\mathcal{Y}}^E$, with edge empirical distribution equal to ν. Let $\nu_v \equiv (\pi_1\nu + \pi_2\nu)/2$. Then*

$$N_{n,k}(\nu) \leq (e\sqrt{m})^{|\overline{\mathcal{X}}|^2} \cdot \exp\left\{ nA(\nu) + \frac{nk}{2} \log\left(\frac{2nk}{e}\right) \right\},$$

$$A(\nu) \equiv \frac{k}{2} H(\nu) - (k-1)H(\nu_v).$$

Proof Recall that $m = nk/2$ is the number of edges in G. Note that $m\pi_1\nu(x)$ is the number of edges (u, v) such that $\theta_u = x$, and $m\pi_2\nu(x)$ is the number of edges (u, v) such that $\theta_v = x$. Therefore $m(\pi_1\nu(x) + \pi_2\nu(x))/k = n(\pi_1\nu(x) + \pi_2\nu(x))/2$ is the number of vertices u such that $\theta_u = x$. Further $m\pi_{12}\nu(x_1, x_2)$ is the number of edges (u, v) such that $\theta_u = x_1$ and $\theta_v = x_2$.

Given a non-negative integer vector $(b(x))_{x \in S}$ with $b_{\text{sum}} \equiv \sum_{x \in S} b(x)$, we denote the corresponding multinomial coefficient by

$$\binom{b_{\text{sum}}}{b(\cdot)} \equiv \frac{b_{\text{sum}}!}{\prod_{x \in S} b(x)!}.$$

We then obtain the following exact counting formula (where $\nu_v(x) \equiv (\pi_1\nu(x) + \pi_2\nu(x))/2$ and $\nu_{12} = \pi_{12}\nu$):

$$N_{n,k}(\nu) = \binom{n}{n\nu_v(\cdot)} \prod_{x \in \mathcal{X}} \frac{[nk\nu_v(x)]!}{\prod_{\tilde{x} \in \overline{\mathcal{X}}}[m\nu_{12}(x, \tilde{x})]! \prod_{\tilde{x} \in \overline{\mathcal{X}}}[m\nu_{12}(\tilde{x}, x)]!}$$

$$\times \prod_{x,\tilde{x}\in\overline{\mathcal{X}}} [m v_{12}(x,\tilde{x})]! \prod_{x,\tilde{x}\in\overline{\mathcal{X}}} \binom{m v_{12}(x,\tilde{x})}{m v(x,\tilde{x},\,\cdot\,)}.$$

The first factor account for the number of ways of choosing $\boldsymbol{\theta}$. The second corresponds to the ways of giving a matching type to half-edges. The third factor counts the number of ways of matching half-edges, and the last one the number of ways of assigning labels in $\overline{\mathcal{Y}}$ to edges.

Equation (7.2) follows by using the following elementary bounds (that hold for any $N \in \mathbb{N}$ and any $p \in \mathscr{P}(S)$ [20]):

$$N! \le e\sqrt{N}\left(\frac{N}{e}\right)^N, \qquad \binom{N}{Np(\cdot)} \le e^{N H(p)}. \tag{7.2}$$

$\square$

Now recall the joint empirical distribution of two assignments $\boldsymbol{\theta}_0, \boldsymbol{\theta} \in \mathcal{X}^V$:

$$\hat{\omega}_{\boldsymbol{\theta}_0,\boldsymbol{\theta}} = \frac{1}{|V|} \sum_{u\in V} \delta_{\theta_{0,u},\theta_u} \in \mathscr{P}(\mathcal{X}\times\mathcal{X}).$$

Further, let $\overline{v}_e(x_1, x_2, y) = \overline{v}(x_1)\,\overline{v}(x_2)\,Q(y|x_1, x_2)$, $\overline{v}$ being the uniform distribution on $\mathcal{X}$, and

$$\Theta(\eta; G, Y) = \left\{\boldsymbol{\theta} \in \mathcal{X}^V : \ d_{\mathrm{TV}}(\hat{v}_{\boldsymbol{\theta},Y}^G, \overline{v}_e) \le \eta\right\}.$$

Given a graph G, a true assignment $\boldsymbol{\theta}_0$, observations Y, and a closed set $\mathcal{S} \subseteq \mathscr{P}(\mathcal{X}\times\mathcal{X})$ we define

$$Z(\mathcal{S}; G, \boldsymbol{\theta}_0, Y) = \left|\left\{\boldsymbol{\theta} \in \Theta(\eta_n; G, Y) : \ \hat{\omega}_{\boldsymbol{\theta}_0,\boldsymbol{\theta}} \in \mathcal{S}\right\}\right|, \tag{7.3}$$

where $\eta_n = (\log n)/\sqrt{n}$. We denote by $\mathcal{G}_n$ the set of instances, i.e., triples $(G_n, \boldsymbol{\theta}_0, Y)$, where G_n is a graph over n vertices, $\boldsymbol{\theta}_0 \in \mathcal{X}^{V_n}$ and $Y \in \mathcal{Y}^{E_n}$.

Lemma 4 *Assume there exists $c_M > 0$ such that $c_M^{-1} \le Q(y|x_1, x_2) \le c_M$ for all $x_1, x_2 \in \mathcal{X}$, $y \in \mathcal{Y}$. Define the map $S : \mathscr{P}(\mathcal{X}^2 \times \mathcal{X}^2 \times \mathcal{Y}) \to \mathbb{R}$ by*

$$S(\Omega) \equiv \frac{k}{2}H(\Omega) - (k-1)H\big((\pi_1\Omega + \pi_2\Omega)/2\big) - \frac{k}{2}H(\overline{v}_e) + (k-1)H(\overline{v}). \tag{7.4}$$

(Here π_1, π_2 are defined as in Lemma 3, with $\overline{\mathcal{X}} = \mathcal{X}^2$, and H denotes the Shannon entropy.) Further define $S_* : \mathscr{P}(\mathcal{X}\times\mathcal{X}) \to \mathbb{R}$ by

$$S_*(\omega) \equiv \ \max \quad S(\Omega),$$
$$\text{subj. to} \quad (\pi_1\Omega + \pi_2\Omega)/2 = \omega,$$

$$\sum_{\tilde{x}_1, \tilde{x}_2 \in \mathcal{X}} \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2, y) = \bar{v}_{\mathrm{e}}(x_1, x_2, y),$$

$$\sum_{x_1, x_2 \in \mathcal{X}} \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2, y) = \bar{v}_{\mathrm{e}}(\tilde{x}_1, \tilde{x}_2, y). \qquad (7.5)$$

There is a set $\mathcal{G}_n^* \subseteq \mathcal{G}_n$ of 'good' instances such that the following happens. For $\mathcal{S} \subseteq \mathscr{P}(\mathcal{X} \times \mathcal{X})$ a closed set, we have

$$\mathbb{P}\big((G_n, \boldsymbol{\theta}_0, Y) \in \mathcal{G}_n^*\big) \geq 1 - c_0^{-1} \exp\left\{-c_0 (\log n)^2\right\},$$

$$\mathbb{E}\left[Z(\mathcal{S}; G_n, \boldsymbol{\theta}_0, Y)\, \mathbf{1}_{(G_n, \boldsymbol{\theta}_0, Y) \in \mathcal{G}_n^*}\right] \leq \exp\left\{n \sup_{\omega \in \mathcal{S}} S_*(\omega) + C\sqrt{n} \log n\right\}. \qquad (7.6)$$

Proof Given a tuple $(G, \boldsymbol{\theta}_0, \boldsymbol{\theta}, Y)$, where $G = (V, E)$ is a graph, $\boldsymbol{\theta}, \boldsymbol{\theta}_0 \in \mathcal{X}^V, Y \in \mathcal{Y}^E$, we define its joint edge empirical distribution $\widehat{\Omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y}^G \in \mathscr{P}(\mathcal{X} \times \mathcal{X} \times \mathcal{X} \times \mathcal{X} \times \mathcal{Y})$ as

$$\widehat{\Omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y}^G = \frac{1}{|E|} \sum_{(u,v) \in E} \delta_{\theta_{0,u}, \theta_u, \theta_{0,v}, \theta_v, Y_{uv}}. \qquad (7.7)$$

In other words $\widehat{\Omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y}^G(x_1, \tilde{x}_1, x_2, \tilde{x}_2, y)$ is the probability that, sampling an edge $(u, v) \in E$ uniformly at random, we have $\theta_{0,u} = x_1, \theta_u = \tilde{x}_1, \theta_{0,v} = x_2, \theta_v = \tilde{x}_2$, $Y_{uv} = y$. Let $\mathscr{P}_{nk}(\mathcal{X}^4 \times \mathcal{Y}) \subseteq \mathscr{P}(\mathcal{X}^4 \times \mathcal{Y})$ be the subset of probability distributions with entries that are integer multiples of $1/|E| = 2/(nk)$. For $\Omega \in \mathscr{P}_{nk}(\mathcal{X}^4 \times \mathcal{Y})$, we let $N_{n,k}(\Omega)$ denote the number of tuples with edge empirical distribution equal to Ω:

$$N_{n,k}(\Omega) = \left|\left\{(G, \boldsymbol{\theta}_0, \boldsymbol{\theta}, Y): \ \widehat{\Omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y}^G = \Omega\right\}\right|. \qquad (7.8)$$

Notice that setting $\overline{\mathcal{X}} = \mathcal{X} \times \mathcal{X}$, we can view $(\boldsymbol{\theta}_0, \boldsymbol{\theta})$ as a vector in $\overline{\mathcal{X}}^V$ and Ω as a probability distribution in $\mathscr{P}(\overline{\mathcal{X}} \times \overline{\mathcal{X}} \times \mathcal{Y})$. Applying Eq. (7.1) and Lemma 3, we get

$$\frac{N_{n,k}(\Omega)}{N_{n,k}} \leq C\, e^{nA(\Omega) + C \log(n)}. \qquad (7.9)$$

We define

$$\mathcal{G}_n^* \equiv \left\{(G, \boldsymbol{\theta}_0, Y) \in \mathcal{G}_n: \ d_{\mathrm{TV}}(\hat{v}_{\boldsymbol{\theta}_0, Y}^G, \bar{v}_{\mathrm{e}}) \leq \eta_n\right\}.$$

Then Eq. (7.6) follows immediately from Lemma 1.

We also define $B_{\mathrm{TV}}(\bar{v}_{\mathrm{e}}; \eta_n) \equiv \{v \in \mathscr{P}(\mathcal{X} \times \mathcal{X} \times \mathcal{Y}): \ d_{\mathrm{TV}}(v, \bar{v}_{\mathrm{e}}) \leq \eta_n\}$. With this notation

$$Z(\mathcal{S}; G, \boldsymbol{\theta}_0, Y) = \sum_{\boldsymbol{\theta} \in \mathcal{X}^V} \mathbf{1}_{\hat{v}_{\boldsymbol{\theta}, Y}^G \in B_{\mathrm{TV}}(\bar{v}_{\mathrm{e}}; \eta_n)} \mathbf{1}_{\hat{\omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}} \in \mathcal{S}},$$

and therefore, using Eq. (7.9),

$$
\mathbb{E}\left[Z(\mathcal{S}; G_n, \boldsymbol{\theta}_0, Y)\,\mathbf{1}_{(G_n, \boldsymbol{\theta}_0, Y)\in\mathcal{G}_n^*}\right]
$$

$$
= \frac{1}{N_{n,k}|\mathcal{X}|^n} \sum_G \sum_{Y\in\mathcal{Y}^E} \sum_{\boldsymbol{\theta}_0, \boldsymbol{\theta}\in\mathcal{X}^V} \prod_{(u,v)\in E} Q(Y_{uv}|\theta_{0u}, \theta_{0v})\mathbf{1}_{\hat{v}^G_{\boldsymbol{\theta}_0, Y}\in B_{\mathrm{TV}}(\overline{v}_e; \eta_n)}
$$

$$
\cdot\,\mathbf{1}_{\hat{v}^G_{\boldsymbol{\theta}, Y}\in B_{\mathrm{TV}}(\overline{v}_e; \eta_n)}\,\mathbf{1}_{\hat{\omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}}\in\mathcal{S}}. \quad (7.10)
$$

Recall the definition of $\widehat{\Omega}^G_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y}$ from Eq. (7.7). We observe that this empirical measure has the following marginals:

$$
\frac{1}{2}(\pi_1\widehat{\Omega}^G_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y} + \pi_2\widehat{\Omega}^G_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y}) = \hat{\omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}},
$$

$$
\sum_{\tilde{x}_1, \tilde{x}_2\in\mathcal{X}} \widehat{\Omega}^G_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y}(x_1, \tilde{x}_1, x_2, \tilde{x}_2, y) = \hat{v}^G_{\boldsymbol{\theta}_0, Y}(x_1, x_2, y),
$$

$$
\text{and } \sum_{x_1, x_2\in\mathcal{X}} \widehat{\Omega}^G_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y}(x_1, \tilde{x}_1, x_2, \tilde{x}_2, y) = \hat{v}^G_{\boldsymbol{\theta}, Y}(\tilde{x}_1, \tilde{x}_2, y).
$$

Moreover, if Q does not vanish, we have

$$
\prod_{(u,v)\in E} Q(Y_{uv}|\theta_{0u}, \theta_{0v}) = \exp\left\{\sum_{(u,v)\in E} \log Q(Y_{uv}|\theta_{0u}, \theta_{0v})\right\}
$$

$$
= \exp\left\{|E|\int \log Q(y|x_1, x_2)\mathrm{d}\widehat{\Omega}^G_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y}(x_1, \tilde{x}_1, x_2, \tilde{x}_2, y)\right\}
$$

$$
=: F\big(\widehat{\Omega}^G_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y}\big).
$$

Therefore the summand in the formula (7.10) depends only in the empirical edge distribution $\widehat{\Omega}^G_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y}$ of the instance $(G, \boldsymbol{\theta}_0, \boldsymbol{\theta}, Y)$. Now let $\mathcal{Q}(\eta_n) \subseteq \mathcal{P}(\mathcal{X}^4 \times \mathcal{Y})$ be the set of $\Omega \in \mathcal{P}(\mathcal{X}^4 \times \mathcal{Y})$ satisfying the constraints

$$
\begin{cases}
\frac{1}{2}(\pi_1\Omega + \pi_2\Omega) \in \mathcal{S}, \\
\left(\sum_{\tilde{x}_1, \tilde{x}_2\in\mathcal{X}} \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2, y)\right)_{x_1, x_2, y} \in B_{\mathrm{TV}}(\overline{v}_e; \eta_n), \\
\left(\sum_{x_1, x_2\in\mathcal{X}} \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2, y)\right)_{\tilde{x}_1, \tilde{x}_2, y} \in B_{\mathrm{TV}}(\overline{v}_e; \eta_n).
\end{cases} \quad (7.11)
$$

We have

$$
\mathbb{E}\left[Z(\mathcal{S}; G_n, \boldsymbol{\theta}_0, Y)\,\mathbf{1}_{(G_n, \boldsymbol{\theta}_0, Y)\in\mathcal{G}_n^*}\right]
$$

$$
= \frac{1}{N_{n,k}|\mathcal{X}|^n} \sum_{\Omega\in\mathcal{Q}(\eta_n)\cap\mathcal{P}_{nk}(\mathcal{X}^4\times\mathcal{Y})} F(\Omega) \sum_{(G, \boldsymbol{\theta}_0, \boldsymbol{\theta}, Y)\in\mathcal{G}_n} \mathbf{1}\{\widehat{\Omega}^G_{\boldsymbol{\theta}_0, \boldsymbol{\theta}, Y} = \Omega\}
$$

$$
= \frac{1}{N_{n,k}|\mathcal{X}|^n} \sum_{\Omega\in\mathcal{Q}(\eta_n)\cap\mathcal{P}_{nk}(\mathcal{X}^4\times\mathcal{Y})} F(\Omega)\, N_{n,k}(\Omega)
$$

$$\leq \frac{Cn^C}{|\mathcal{X}|^n} \sum_{\Omega \in \mathcal{Q}(\eta_n) \cap \mathscr{P}_{nk}(\mathcal{X}^4 \times \mathcal{Y})} F(\Omega)\, e^{nA(\Omega)}\,.$$

We applied the bound (7.9) in the last line above. Due to the second constraint in (7.11), we can upper bound $F(\Omega)$ as follows

$$F(\Omega) \leq \exp\left\{\frac{nk}{2} \int \log Q(y|x_1, x_2)\mathrm{d}\bar{\nu}_{\mathrm{e}} + Cn\eta_n\right\}$$
$$= \exp\left\{\frac{nk}{2}\left(- H(\bar{\nu}_{\mathrm{e}}) + 2H(\bar{\nu})\right) + Cn\eta_n\right\}.$$

Therefore, letting

$$S(\Omega) = A(\Omega) - \frac{k}{2}H(\bar{\nu}_{\mathrm{e}}) + (k - 1)H(\bar{\nu})$$
$$= \frac{k}{2}H(\Omega) - (k - 1)H\left((\pi_1\Omega + \pi_2\Omega)/2\right) - \frac{k}{2}H(\bar{\nu}_{\mathrm{e}}) + (k - 1)H(\bar{\nu}),$$

we arrive at

$$\mathbb{E}\left[Z(\mathcal{S}; G_n, \boldsymbol{\theta}_0, Y)\,\mathbf{1}_{(G_n, \boldsymbol{\theta}_0, Y) \in \mathcal{G}_n^*}\right] \leq C\, n^C \sum_{\Omega \in \mathcal{Q}(\eta_n) \cap \mathscr{P}_{nk}(\mathcal{X}^4 \times \mathcal{Y})} \exp\left\{nS(\Omega) + Cn\eta_n\right\}$$
$$\leq C\, n^C\, |\mathscr{P}_{nk}(\mathcal{X}^4 \times \mathcal{Y})| \exp\left\{n \sup_{\omega \in \mathcal{S}} S_*(\omega) + C\sqrt{n} \log n\right\}$$
$$\leq C\, n^C \exp\left\{n \sup_{\omega \in \mathcal{S}} S_*(\omega) + C\sqrt{n} \log n\right\},$$

which implies the claim. $\square$

The next result provides a sufficient condition for weak recovery using the estimator $\hat{\boldsymbol{\theta}}$ satisfying Eq. (5.2); this is a more general version of Theorem C.

Theorem E *Assume there exists $c_M > 0$ such that $c_M \leq Q(y|x_1, x_2)$ for all $x_1, x_2 \in \mathcal{X}$, $y \in \mathcal{Y}$. Assume $S_*(\bar{\nu} \times \bar{\nu}) < -\varepsilon < 0$. Then there exists $\delta = \delta(\varepsilon, c_M) > 0$ such that, with probability at least $1 - c_0^{-1} \exp\{-c_0(\log n)^2\}$, the following happens*

$$d_{\mathrm{TV}}(\hat{\omega}_{\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0}, \bar{\nu} \times \bar{\nu}) \geq \delta\,. \tag{7.12}$$

Proof Recall that $B_{\mathrm{TV}}(\bar{\nu} \times \bar{\nu}; \delta)$ denotes the set of probability distributions $\omega \in \mathscr{P}(\mathcal{X} \times \mathcal{X})$ such that $d_{\mathrm{TV}}(\omega, \bar{\nu} \times \bar{\nu}) \leq \delta$. We claim that, under the stated assumptions there exists $\delta, c_1 > 0$ such that, setting $\mathcal{S}_\delta = B_{\mathrm{TV}}(\bar{\nu} \times \bar{\nu}; \delta)$, and $\mathcal{G}_*$ as in Lemma 4, we have

$$\mathbb{E}\left[Z(\mathcal{S}_\delta; G_n, \boldsymbol{\theta}_0, Y)\,\mathbf{1}_{(G_n, \boldsymbol{\theta}_0, Y) \in \mathcal{G}_n^*}\right] \leq e^{-c_1 n}\,. \tag{7.13}$$

Hence applying Lemma 4, it follows that with probability at least $1 - c_0^{-1} e^{-c_0 (\log n)^2}$ (possibly adjusting the constant c_0), $Z(\mathcal{S}_\delta; G_n, \boldsymbol{\theta}_0, Y) = 0$. Hence $\hat{\omega}_{\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0} \notin \mathcal{S}_\delta$ by construction of $\hat{\boldsymbol{\theta}}$, and therefore the claim follows.

We are left with the task of proving Eq. (7.13), which will follow from Lemma 4 and $S_*(\overline{\nu} \times \overline{\nu}) < -\varepsilon$ once we show that $\omega \mapsto S_*(\omega)$ is continuous. For $\omega \in \mathscr{P}(\mathcal{X} \times \mathcal{X})$, define $K(\omega) \subseteq \mathscr{P}(\mathcal{X}^2 \times \mathcal{X}^2 \times \mathcal{Y})$ the set of constraints of the optimization problem in Eq. (7.5), i.e.,

$$
\begin{aligned}
K(\omega) = \Big\{ \Omega \in \mathscr{P}(\mathcal{X}^2 \times \mathcal{X}^2 \times \mathcal{Y}) \; : \; (\pi_1 \Omega + \pi_2 \Omega)/2 = \omega, \\
\sum_{\tilde{x}_1, \tilde{x}_2 \in \mathcal{X}} \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2, y) = \overline{\nu}_{\mathrm{e}}(x_1, x_2, y), \\
\sum_{x_1, x_2 \in \mathcal{X}} \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2, y) = \overline{\nu}_{\mathrm{e}}(\tilde{x}_1, \tilde{x}_2, y) \Big\}.
\end{aligned} \tag{7.14}
$$

Since $\Omega \mapsto S(\Omega)$ is a uniformly continuous function (the entropy function is Hölder), it suffices to show that the map $\omega \mapsto K(\omega)$ is continuous in the Gromov–Hausdorff metric: Fix $r > 0$ and $\omega \in \mathscr{P}(\mathcal{X} \times \mathcal{X})$. We will show that there exists a $\delta > 0$ such that for all $\omega' \in B_{\mathrm{TV}}(\omega, \delta)$ we have $K(\omega) \in [K(\omega')]_r$ and $K(\omega') \in [K(\omega)]_r$. Here, $[A]_r$ denotes the r-enlargement of the set A in d_{TV} distance: $[A]_r = \cup_{x \in A} B_{\mathrm{TV}}(x, r)$. The smallest r such that $B \subseteq [A]_r$ and $A \subseteq [B]_r$ is the Gromov-Hausdorff distance between A and B.

Let $\Omega \in K(\omega)$. For $\omega' \in \mathscr{P}(\mathcal{X} \times \mathcal{X})$ let $\Omega' = \arg\min d_{\mathrm{KL}}(\Omega'' || \Omega)$ s.t. $\Omega'' \in K(\omega')$. We want to show that Ω' is close to Ω when ω' is close to ω. It is possible to write an expression for Ω' using Lagrange multipliers:

$$
\begin{aligned}
\Omega'(x_1, \tilde{x}_1, x_2, \tilde{x}_2) = Z^{-1} \cdot \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2) \\
\cdot \exp\{(\mu(x_1, \tilde{x}_1) + \mu(x_2, \tilde{x}_2))/2 + \lambda_1(x_1, x_2, y) + \lambda_2(\tilde{x}_1, \tilde{x}_2, y)\},
\end{aligned}
$$

where Z is the normalizing constant and $\{\mu(x, x'), \lambda_1(x, x', y), \lambda_2(x, x', y) \; : \; x, x', y \in \mathcal{X}^2 \times \mathcal{Y}\}$ are uniquely determined by the system of equations defining $K(\omega')$, c.f. (7.14), and we have

$$
\begin{aligned}
d_{\mathrm{KL}}(\Omega' || \Omega) = \mathbb{E}_{\Omega'} \Big[\frac{1}{2} (\mu(X_1, \tilde{X}_1) + \mu(X_2, \tilde{X}_2)) + \lambda_1(X_1, X_2, Y) + \lambda_2(\tilde{X}_1, \tilde{X}_2, Y) \Big] \\
- \log \mathbb{E}_{\Omega'} \Big[e^{(\mu(X_1, \tilde{X}_1) + \mu(X_2, \tilde{X}_2))/2 + \lambda_1(X_1, X_2, Y) + \lambda_2(\tilde{X}_1, \tilde{X}_2, Y)} \Big],
\end{aligned}
$$

where $(X_1, \tilde{X}_1, X_2, \tilde{X}_2, Y) \sim \Omega'$. We remark that $\mu, \lambda_1, \lambda_2$ are continuous functions of ω', and that for $\omega' = \omega$ we have $\mu = 0$ and $\lambda_1 = \lambda_2 = 0$ by strict convexity of the KL divergence. Therefore, there exists $\delta > 0$ such that if $d_{\mathrm{TV}}(\omega, \omega') \le \delta$ then $d_{\mathrm{KL}}(\Omega' || \Omega) \le 2r^2$, which by Pinsker's inequality implies $d_{\mathrm{TV}}(\Omega', \Omega) \le r$ and hence $K(\omega) \subseteq [K(\omega')]_r$. Switching the roles of the pairs (ω, Ω) and (ω', Ω') in this construction, we also obtain that $K(\omega') \subseteq [K(\omega)]_r$ when $d_{\mathrm{TV}}(\omega, \omega') \le \delta$ for a possibly smaller value of δ. This implies that the Gromov-Hausdorff distance between

$K(\omega)$ and $K(\omega')$ is smaller than r whenever $d_{\mathrm{TV}}(\omega, \omega') \leq \delta$, showing that the map $\omega \mapsto K(\omega)$ is continuous. $\qquad\square$

The condition $S_*(\bar{v} \times \bar{v}) < -\varepsilon$ might be hard to verify in practice because it requires solving the optimization problem (7.5). We provide a simpler sufficient condition, which is the content of Theorem C:

Lemma 5 *Let* $(\theta_1, \theta_2, Y) \sim \bar{v}_e$, *with* $\bar{v}_e(x_1, x_2, y) = \bar{v}(x_1)\bar{v}(x_2)Q(y|x_1, x_2)$. *We have* $S_*(\bar{v} \times \bar{v}) \leq -\frac{k}{2}I(\theta_1, \theta_2; Y) + H(\theta_1)$.

Proof Let $\Omega_* \in \mathscr{P}(\mathcal{X}^2 \times \mathcal{X}^2 \times \mathcal{Y})$ be any distribution achieving the maximum in (7.5) for $\omega = \bar{v} \times \bar{v}$, and let $(X_1, \tilde{X}_1, X_2, \tilde{X}_2, Y)$ have distribution Ω_*. Note that $(X_1, X_2, Y) \sim \bar{v}_e$, $(\tilde{X}_1, \tilde{X}_2, Y) \sim \bar{v}_e$, $(\pi_1\Omega_* + \pi_2\Omega_*)/2 = \bar{v} \times \bar{v} = \mathrm{Law}(X_1, X_2)$, whence

$$
\begin{aligned}
S_*(\bar{v} \times \bar{v}) &= S(\Omega_*) \\
&= \frac{k}{2}H(X_1, \tilde{X}_1, X_2, \tilde{X}_2, Y) - (k-1)H(X_1, X_2) - \frac{k}{2}H(X_1, X_2, Y) + (k-1)H(X_1) \\
&= \frac{k}{2}H(X_1, \tilde{X}_1, X_2, \tilde{X}_2|Y) + \frac{k}{2}H(Y) - 2(k-1)H(X_1) - \frac{k}{2}H(X_1, X_2, Y) + (k-1)H(X_1) \\
&\overset{(a)}{\leq} \frac{k}{2}H(X_1, X_2|Y) + \frac{k}{2}H(\tilde{X}_1, \tilde{X}_2|Y) + \frac{k}{2}H(Y) - \frac{k}{2}H(X_1, X_2, Y) - (k-1)H(X_1) \\
&= \frac{k}{2}H(X_1, X_2|Y) - (k-1)H(X_1) \\
&= -\frac{k}{2}I(X_1, X_2; Y) + H(X_1).
\end{aligned}
$$

Step (a) follows by sub-additivity of entropy. $\qquad\square$

Hence, if $\frac{k}{2}I(\theta_1, \theta_2; Y) \geq H(\theta_1) + \varepsilon$, then $S_*(\bar{v} \times \bar{v}) < -\varepsilon < 0$, and the claim follows by applying Theorem E.

7.3 Proof of Corollary 3

Let $\mathcal{B}_{q \times q}$ is the set of all $q \times q$ non-negative doubly stochastic matrices (with $q = |\mathcal{X}|$). It holds that

$$
\mathtt{overlap}(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0) = \max_{\sigma \in \mathcal{S}_q} \sum_{x \in \mathcal{X}} \hat{\omega}_{\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0}(x, \sigma(x)) = \max_{\pi \in \mathcal{B}_{q \times q}} \sum_{x, x' \in \mathcal{X}} \pi(x, x')\hat{\omega}_{\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0}(x, x').
$$

$$(7.15)$$

Indeed, since the right-most expression in the above display is a linear program, the objective value is maximized at the extreme points of the polytope $\mathcal{B}_{q \times q}$, which by Birkhoff's theorem are permutation matrices: $\pi(x, y) = \mathbf{1}_{y=\sigma(x)}$ for $\sigma \in \mathcal{S}_q$, hence the equality.

Since $q\hat{\omega}_{\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0} \in \mathcal{B}_{q \times q}$ (we abused notation and identified the joint distribution $\hat{\omega}_{\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0}$ on $\mathcal{X} \times \mathcal{X}$ with a $q \times q$ matrix), we have

$$
\mathtt{overlap}(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0) \geq q \sum_{x, x'} \left(\hat{\omega}_{\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0}(x, x')\right)^2.
$$

Now, on the event $d_{\mathrm{TV}}(\hat{\omega}_{\hat{\theta},\theta_0}, \overline{v} \times \overline{v}) \geq \delta$, we have $\sum_{x,x'} (\hat{\omega}_{\hat{\theta},\theta_0}(x, x'))^2 \geq \frac{1}{q^2} + \frac{\delta^2}{q}$. Hence $\texttt{overlap}(\hat{\theta}, \theta_0) \geq \frac{1}{q} + \delta^2$ on the same event.

Next, we prove the second statement. For two functions $f, g : \mathcal{X} \mapsto \mathbb{R}$, we let $\hat{\omega}_{\hat{\theta},\theta_0}(f, g) := \sum_{x_1,x_2} \hat{\omega}_{\hat{\theta},\theta_0}(x_1, x_2) f(x_1) g(x_2)$. Theorem C implies

$$\mathbb{P}\left(\exists f, g : \mathcal{X} \mapsto \mathbb{R} \text{ s.t. } |\hat{\omega}_{\hat{\theta},\theta_0}(f, g)| \geq \frac{\delta}{q(q-1)}\right) \geq 1 - c_0^{-1} \exp\{-c_0(\log n)^2\}.$$

Indeed, if $d_{\mathrm{TV}}(\hat{\omega}_{\hat{\theta},\theta_0}, \overline{v} \times \overline{v}) \geq \delta$ then there exist $x_1, x_2 \in \mathcal{X}$ such that $|\hat{\omega}_{\hat{\theta},\theta_0}(x_1, x_2) - \frac{1}{q^2}| \geq \frac{\delta}{q^2}$. Now take $f = (\delta_{x_1} - \frac{1}{q})\frac{q}{q-1}$ and $g = (\delta_{x_2} - \frac{1}{q})\frac{q}{q-1}$.

On the other hand, letting $\mathcal{F} := \{f = (\delta_x - \frac{1}{q})\frac{q}{q-1}, x \in \mathcal{X}\}$, a union bound implies

$$\mathbb{P}\left(\exists f, g : \mathcal{X} \mapsto \mathbb{R} \text{ s.t. } |\hat{\omega}_{\hat{\theta},\theta_0}(f, g)| \geq \frac{\delta}{q(q-1)}\right)$$
$$\leq q^2 \max_{f,g \in \mathcal{F}} \mathbb{P}\left(|\hat{\omega}_{\hat{\theta},\theta_0}(f, g)| \geq \frac{\delta}{q(q-1)}\right).$$

Therefore, there exists a (deterministic) pair $f, g \in \mathcal{F}$ such that $\mathbb{P}(|\hat{\omega}_{\hat{\theta},\theta_0}(f, g)| \geq \frac{\delta}{q(q-1)}) \geq \frac{1-o_n(1)}{q^2} > c_0 > 0$. By Markov's inequality, this in turn implies that for this specific pair $f, g \in \mathcal{F}$ we have

$$\mathbb{E}[\hat{\omega}_{\hat{\theta},\theta_0}(f, g)^2] \geq c_0 \frac{\delta^2}{(q(q-1))^2} = c(q)\delta^2. \tag{7.16}$$

Now consider estimating the matrix X_f (recall that $(X_f)_{uv} = f(\theta_u) f(\theta_v)$) with the matrix $\widehat{X}^{(\lambda)}$ having entries $\widehat{X}_{uv}^{(\lambda)} = \lambda g(\hat{\theta}_u) g(\hat{\theta}_v)$, with

$$\lambda = n^2 \mathbb{E}[\hat{\omega}_{\hat{\theta},\theta_0}(f, g)^2] / \mathbb{E}[\|\widehat{X}^{(1)}\|_F^2].$$

Since

$$\frac{1}{n^2}\langle\widehat{X}^{(\lambda)}, X_f\rangle = \frac{\lambda}{n^2} \sum_{u,v \in V_n} f(\theta_u) f(\theta_v) g(\hat{\theta}_u) g(\hat{\theta}_v) = \lambda \hat{\omega}_{\hat{\theta},\theta_0}(f, g)^2,$$

the loss $\mathcal{R}_n$ incurred is

$$\mathcal{R}_n(\widehat{X}^{(\lambda)}; f) = \frac{1}{n^2}\mathbb{E}\|X_f\|_F^2 - 2\lambda\mathbb{E}[\hat{\omega}_{\hat{\theta},\theta_0}(f, g)^2] + \frac{\lambda^2}{n^2}\mathbb{E}\|\widehat{X}^{(1)}\|_F^2$$
$$= \frac{1}{n^2}\mathbb{E}\|X_f\|_F^2 - \frac{\mathbb{E}[\hat{\omega}_{\hat{\theta},\theta_0}(f, g)^2]^2}{(\mathbb{E}\|\widehat{X}^{(1)}\|_F^2/n^2)}.$$

We have

$$\mathbb{E}\|X_f\|_F^2 = \sum_{u,v \in V_n} \mathbb{E}[f(\theta_u)^2 f(\theta_v)^2] = \frac{n}{q} \sum_{x \in \mathbb{Z}_q} f(x)^4 + n(n-1).$$

So $\lim \frac{1}{n^2} \mathbb{E}\|X_f\|_F^2 = 1$. Furthermore, since $\|g\|_\infty = 1$,

$$\mathbb{E}\|\widehat{X}^{(1)}\|_F^2 = \sum_{u,v \in V_n} \mathbb{E}[g(\hat{\theta}_u)^2 g(\hat{\theta}_v)^2] \le n^2.$$

Combining these estimates with the lower bound (7.16) implies $\limsup \mathcal{R}_n(\widehat{X}^{(\lambda)}; f) \le 1 - c(q)\delta^2$. Since $\mathcal{R}_n^{\text{Bayes}}(f) \le \mathcal{R}_n(\widehat{X}^{(\lambda)}; f)$ this concludes the proof.

8 Local algorithms on random graphs: Proofs

8.1 Proof of Theorem D

8.1.1 Preliminaries

Let (T_k, o) denote the infinite k-regular tree rooted at o. (Except the root o, every vertex has $k-1$ offsprings.) By expanding the square, we get

$$\mathbb{E}\big[d_{\text{TV}}(\widehat{\mu}_{G_n,u,l}, \overline{v})^2\big] \le \frac{q}{4} \mathbb{E}\big[d_{\ell_2}(\widehat{\mu}_{G_n,u,l}, \overline{v})^2\big] = \frac{q}{4}\left(\sum_{x \in \mathbb{Z}_q} \mathbb{E}\{\widehat{\mu}_{G_n,u,l}(x)^2\} - \frac{1}{q}\right).$$

(Here, d_{ℓ_2} is the ℓ_2 distance in $\mathbb{R}^q$.) Since the graph sequence $(G_n)_{n \ge 1}$ almost surely converges locally-weakly to (a Dirac delta on) (T_k, o), we have

$$\lim_{n \to \infty} \frac{1}{|V_n|} \sum_{u \in V_n} \mathbb{E}\big[d_{\ell_2}(\widehat{\mu}_{G_n,u,l}, \overline{v})^2\big] = \sum_{x \in \mathbb{Z}_q} \mathbb{E}\Big[\mathbb{P}\big(\theta_o = x | \underline{Y}_{B_{T_k}(o,l)}^{(\varepsilon)}\big)^2\Big] - \frac{1}{q}. \quad (8.1)$$

Recall

$$\mu_{o,l}(x) = \mathbb{P}\Big(\theta_o = x | Y_{B_{T_k}(o,l)}^{(\varepsilon)}\Big) \quad \text{for all } x \in \mathbb{Z}_q.$$

Let $Q_x = \text{Law}(\mu_{o,l}|\theta_o = x, \xi_o^{(\varepsilon)} = \star)$ be the conditional law of $\mu_{o,l}$ given the value at the root being x and no information revealed by the side channel. This is a probability distribution on the simplex $\Delta^{q-1} = \mathscr{P}(\mathbb{Z}_q)$: $Q_x \in \mathscr{P}(\Delta^{q-1})$. Furthermore, let $Q = \text{Law}(\mu_{o,l}|\xi_o^{(\varepsilon)} = \star) = \frac{1}{q}\sum_{x \in \mathbb{Z}_q} Q_x$. The following simple lemma from [42] is quite useful.

Lemma 6 *For every $x \in \mathbb{Z}_q$, Q_x has a density w.r.t. Q, and $\frac{\mathrm{d}Q_x}{\mathrm{d}Q}(\mu) = q\mu(x)$ for all $\mu \in \Delta^{q-1}$.*

Proof Let $\psi : \Delta^{q-1} \to \mathbb{R}$ be bounded measurable. We let $Y \equiv \{Y^{(\varepsilon)}_{B_{T_k}(o,l)}\}$. Then

$$
\begin{aligned}
\mathbb{E}\big[\psi(\mu_{o,l})|\theta_o = x, \xi_o^{(\varepsilon)} = \star\big] &= q\mathbb{E}\big[\psi(\mu_{o,l})\mathbf{1}\{\theta_o = x\}|\xi_o^{(\varepsilon)} = \star\big] \\
&= q\mathbb{E}\big[\mathbb{E}\big[\psi(\mu_{o,l})\mathbf{1}\{\theta_o = x\}|Y, \xi_o^{(\varepsilon)} = \star\big]|\xi_o^{(\varepsilon)} = \star\big] \\
&= q\mathbb{E}\big[\psi(\mu_{o,l})\mathbb{E}\big[\mathbf{1}\{\theta_o = x\}|Y, \xi_o^{(\varepsilon)} = \star\big]|\xi_o^{(\varepsilon)} = \star\big] \\
&= q\mathbb{E}\big[\psi(\mu_{o,l})\mu_{o,l}(x)|\xi_o^{(\varepsilon)} = \star\big].
\end{aligned}
$$

Therefore $dQ_x/dQ(\mu) = q\mu(x)$. $\qquad\square$

With the above lemma in hand, the right-hand side in (8.1) can be written as

$$
\lim_{n\to\infty} \frac{1}{|V_n|} \sum_{u\in V_n} \mathbb{E}\big[d_{\ell_2}(\widehat{\mu}_{G_n,u,l}, \overline{\nu})^2\big]
$$

$$
\begin{aligned}
&= \varepsilon \sum_x \mathbb{E}[\mathbf{1}\{x = \theta_o\}|\xi_o \neq \star] + (1-\varepsilon)\sum_x \mathbb{E}[\mu_{o,l}(x)^2|\xi_o = \star] - \frac{1}{q} \\
&= \varepsilon + \frac{1-\varepsilon}{q}\sum_x \mathbb{E}[\mu_{o,l}(x)|\theta_o = x, \xi_o = \star] - \frac{1}{q} \\
&= \varepsilon\frac{q-1}{q} + (1-\varepsilon)\Big(\mathbb{E}[\mu_{o,l}(\theta_o)|\xi_o = \star] - \frac{1}{q}\Big).
\end{aligned}
$$

The first equality follows by conditioning on $\xi_o^{(\varepsilon)}$ as noting that conditional on $\xi_o^{(\varepsilon)} \neq \star$, $\mu_o(x) = \mathbf{1}\{x = \xi_o^{(\varepsilon)}\}$. Lemma 6 was used to obtain the second equality.

In light of the above expression, we will track the evolution of the sequence

$$
\hat{z}_{o,l} := \mathbb{E}[\mu_{o,l}(\theta_o)|\xi_o = \star] - \frac{1}{q}, \quad l \geq 0,
$$

which measures the deviation from uniformity of the local marginal at the root. In order to exploit the recursive structure of the tree, we will need to work at the level of the first offsprings of o. For every offspring u of o, we denote by $T^{\downarrow}(u,l)$ the first l generations of the subtree rooted at u not containing o; this is a $(k-1)$-ary tree. Now, (with a slight notation override) we redefine

$$
\mu_{u,l}(x) := \mathbb{P}\Big(\theta_u = x|Y^{(\varepsilon)}_{T_k^{\downarrow}(u,l)}\Big) \quad \text{for all } x \in \mathbb{Z}_q,
$$

and consider the auxiliary sequence

$$
z_l := \mathbb{E}[\mu_{u,l}(\theta_u)|\xi_u = \star] - \frac{1}{q}, \quad l \geq 0.
$$

Note that the above definition does not depend on u since $\mu_{u,l}(\theta_u)$ have the same distribution for all $u \sim o$. In the next proposition, we relate the two sequences $(\hat{z}_{o,l})_{l\geq 0}$ and $(z_l)_{l\geq 0}$, and establish a recursion for the latter.

Proposition 2 *Let $\kappa = (k-1)(1-p)^2$ and $\hat{\kappa} = k(1-p)^2$. There exists constants $c, C > 0$ depending only on q such that the following holds. If for some $l \geq 1$, $\hat{\kappa}|z_{l-1}| \leq c$ and $\hat{\kappa}\varepsilon \leq c$, then*

$$\left| \hat{z}_{o,l} - \varepsilon\hat{\kappa}\frac{q-1}{q} - (1-\varepsilon)\hat{\kappa}z_{l-1} \right| \leq C\hat{\kappa}^2(z_{l-1}^2 + \varepsilon^2),$$

and $$\left| z_l - \varepsilon\kappa\frac{q-1}{q} - (1-\varepsilon)\kappa z_{l-1} \right| \leq C\kappa^2(z_{l-1}^2 + \varepsilon^2).$$

The proof of this proposition is presented in Sect. 8.1.2 Theorem D follows directly from Proposition 2, as shown in the next Corollary.

Corollary 5 *If $\kappa < 1$ and $\hat{\kappa}\varepsilon < c$ for a constant $c = c(q, \kappa)$ then there exists $L = L(q, \kappa)$ such that $|\hat{z}_{o,l}| \leq L\varepsilon$ for all $l \geq 0$.*

Proof We only need to prove that $|z_l| \leq L\varepsilon$, which we will achieve by induction. Since $z_0 = 0$, let's assume that $|z_l| \leq L\varepsilon$ for a fixed $l \geq 0$. Then we obtain from Proposition 2 that

$$|z_{l+1}| \leq \varepsilon\kappa\frac{q-1}{q} + \kappa L\varepsilon + C\kappa^2(L^2 + 1)\varepsilon^2.$$

It suffices to find an L (independent of ε) such that the above upper bound is smaller than $L\varepsilon$ for all ε. This is equivalent to the quadratic inequality $\kappa\frac{q-1}{q} + C\kappa^2\varepsilon - (1 - \kappa)L + C\kappa^2\varepsilon L^2 \leq 0$. The smallest solution to this inequality is $L_* = \frac{1-\kappa-\sqrt{\Delta}}{2\alpha}$, with $\alpha = C\kappa^2\varepsilon$, and $\Delta = (1 - \kappa)^2 - 4\alpha(\alpha + \frac{q-1}{q}\kappa)$. Latter is non-negative provided that $\varepsilon < c_0(q)(\frac{1}{\kappa} - 1)^2$ for constant some $c_0(q) > 0$. Moreover, for ε small enough we can write $\sqrt{\Delta} = (1 - \kappa)(1 - 2\alpha(\alpha + \frac{q-1}{q}\kappa)/(1 - \alpha)^2) + \mathcal{O}(\varepsilon^2)$, so that $L_* = (\alpha + \frac{q-1}{q}\kappa)/(4(1-\kappa)) + \mathcal{O}(\varepsilon^2)$. Therefore, we can take $L = (C + \frac{q-1}{q}\kappa)/(4(1-\kappa)) + 1$. □

8.1.2 Proof of Proposition 2: analysis of the recursion on the tree

Here, we prove Proposition 2. The two statements can be treated in exactly the same way; the only difference being that the root o has k children, while every other vertex has $k-1$ children. For this reason we only write a detailed proof for the first statement; the second one is obtained merely by replacing k by $k-1$.

Observe that conditional on $\xi_o^{(\varepsilon)} = \star$ the marginal at o is obtained from the marginals at its offsprings $u \sim o$ by a sum-product relation which, in the case of $\mathbb{Z}_q$-synchronization, has the form

$$\mu_{o,l}(x) = \frac{1}{\Sigma_{o,l}} \prod_{u \sim o} \sum_{y \in \mathbb{Z}_q} M_{x,y}(Y_{ou})\mu_{u,l-1}(y)$$

$$= \frac{1}{\Sigma_{o,l}} \prod_{u \sim o} \left(\frac{p}{q} + (1-p)\mu_{u,l-1}(x - Y_{ou}) \right). \tag{8.2}$$

where $\Sigma_{o,l}$ is the normalizing constant, and $M_{x,y}(Y_{ou}) = \mathbb{P}(\theta_o = x | \theta_u = y, Y_{ou}) = \frac{p}{q} + (1-p)\mathbf{1}\{Y_{ou} = x - y\}$ is the Markov transition matrix associated to a 'broadcasting process' on the tree according to the $\mathbb{Z}_q$-synchronization model.

The recursion (8.2) induces a deterministic recursion over probability distributions over the simplex $\Delta^{q-1} = \mathscr{P}(\mathbb{Z}_q)$. Namely, if we define $Q_x^{(l)} := \text{Law}(\mu_{o,l} | \theta_o = x) \in \mathscr{P}(\Delta^{q-1})$, we obtain a recursion that determines $Q_x^{(l)}$ in terms of $Q_x^{(l-1)}$ (notice that, by Lemma 6, once $Q_x^{(l)}$ is given for one value of x, it is determined for the other values as well.) The laws of $\mu_{u,l-1}$ are given by $\text{Law}(\mu_{u,l-1} | \theta_u = x) = Q_x^{(l-1)}$ for all $u \sim o$. Note that this law does not depend on u since $\mu_{u,l-1}$ are i.i.d. given θ_o. Then $Q_x^{(l)}$ can be obtained from $Q_x^{(l-1)}$ as follows:

1. Draw θ_o and θ_u, $\forall u \sim o$ independently and uniformly at random from $\mathbb{Z}_q$.
2. Construct $\{Y_{ou}, u \sim o\}$ according to the $\mathbb{Z}_q$-synchronization model (1.1).
3. Draw $\mu_{u,l-1}$ from $Q_{\theta_u}^{(l-1)}$ independently for each $u \sim o$.
4. Construct a distribution μ according to (8.2).
5. Then, given $\xi_o^{(\varepsilon)} = \star$, $\mu_{o,l}$ has the same law as μ.

We now analyze the map described above. Define

$$Z_o(x) := \prod_{u \sim o} \left(p + (1 - p)q\mu_u(x - Y_{ou}) \right),$$

so $\mu_o(x) = Z_o(x)/\sum_y Z_o(y)$, where we have dropped the indices l for convenience. Following the analysis of [52], we use the identity $\frac{a}{b+c} = \frac{a}{b} - \frac{ac}{b^2} + \frac{c^2}{b^2}\frac{a}{b+c}$ with $a = Z_o(x)$, $b = q$ and $c = \sum_y Z_o(y) - q$ to write

$$\mu_o(x) = \frac{1}{q}Z_o(x) - \frac{1}{q^2}Z_o(x)\left(\sum_y Z_o(y) - q\right) + \left(\frac{1}{q}\sum_y Z_o(y) - 1\right)^2 \frac{Z_o(x)}{\sum_y Z_o(y)}.$$
(8.3)

Next we compute the conditional expectations of $Z_o(y)$ and $Z_o(y)Z_o(y')$ (given $\theta_o = x$ and $\xi_o = \star$) in order to control $\mathbb{E}[\mu_o(\theta_o)|\xi_o = \star]$.

Lemma 7 *Let $\delta_u := \mu_{u,l} - \frac{1}{q}$ for $u \sim o$. For all $x, y, y' \in \mathbb{Z}_q$, we have*

$$\mathbb{E}[Z_o(y)|\theta_o = x, \xi_o = \star] = \left(1 + \varepsilon(1 - p)^2q\left(\mathbf{1}_{y=x} - \frac{1}{q}\right)\right.$$

$$\left. + (1 - \varepsilon)(1 - p)^2q\mathbb{E}[\delta_u(y - x + \theta_u)|\xi_u = \star]\right)^k,$$
(8.4)

and

$$\mathbb{E}[Z_o(y)Z_o(y')|\theta_o = x, \xi_o = \star] = \left(1 + \varepsilon p(1 - p)^2q\left(\left(\mathbf{1}_{y=x} - \frac{1}{q}\right) + \left(\mathbf{1}_{y'=x} - \frac{1}{q}\right)\right)\right.$$

$$+ \varepsilon(1 - p)^2q\left(\mathbf{1}_{y=y'} - \frac{1}{q}\right)$$

$$+ \varepsilon(1-p)^3 q^2 \mathbf{1}_{y=y'} \left(\mathbf{1}_{y=x} - \frac{1}{q} \right)$$

$$+ (1-\varepsilon)(1-p)^2 q \mathbb{E}[\delta_u(y - x + \theta_u)|\xi_u = \star]$$

$$+ (1-\varepsilon)(1-p)^2 q \mathbb{E}[\delta_u(y' - x + \theta_u)|\xi_u = \star]$$

$$+ (1-\varepsilon)p(1-p)^2 q \sum_z \mathbb{E}[\delta_u(y - z)\delta_u(y' - z)|\xi_u = \star]$$

$$\left. + (1-\varepsilon)(1-p)^3 q^2 \mathbb{E}[\delta_u(y - x + \theta_u)\delta_u(y' - x + \theta_u)|\xi_u = \star] \right)^k.$$

$$(8.5)$$

Proof We start with the first identity (8.4). Since the distributions $\{(\mu_u, Y_{ou}) : u \sim o\}$ are conditionally independent given θ_o, we have

$$\mathbb{E}[Z_o(y)|\theta_o = x, \xi_o = \star] = \prod_{u \sim o} \left(p + (1-p)q \mathbb{E}[\mu_u(y - Y_{ou})|\theta_o = x, \xi_o = \star] \right).$$

Moreover,

$$\mathbb{E}[\mu_u(y - Y_{ou})|\theta_o = x, \xi_o = \star] = \varepsilon \mathbb{E}[\mathbf{1}_{y-Y_{ou}=\xi_u}|\theta_o = x, \xi_o = \star, \xi_u \neq \star]$$
$$+ (1-\varepsilon)\mathbb{E}[\mu_u(y - Y_{ou})|\theta_o = x, \xi_o = \star, \xi_u = \star].$$

The first term in the right-hand side is $\mathbb{P}(y - Y_{ou} = \theta_u|\theta_o = x) = (1-p)\mathbf{1}_{x=y} + \frac{p}{q}$. The second term is

$$(1-p)\mathbb{E}[\mu_u(y - x + \theta_u)|\xi_u = \star] + \frac{p}{q} \sum_z \mathbb{E}[\mu_u(y - x + z)|\xi_u = \star]$$

$$= (1-p)\mathbb{E}[\mu_u(y - x + \theta_u)|\xi_u = \star] + \frac{p}{q}.$$

Therefore

$$\mathbb{E}[\mu_u(y - Y_{ou})|\theta_o = x, \xi_o = \star] = \varepsilon \left((1-p)\mathbf{1}_{x=y} + \frac{p}{q} \right)$$

$$+ (1-\varepsilon)\left((1-p)\mathbb{E}[\mu_u(y - x + \theta_u)|\xi_u = \star] + \frac{p}{q} \right)$$

$$= \frac{1}{q} + \varepsilon(1-p)(\mathbf{1}_{x=y} - \frac{1}{q})$$

$$+ (1-\varepsilon)(1-p)\mathbb{E}[\delta_u(y - x + \theta_u)|\xi_u = \star].$$

So we obtain

$$\mathbb{E}[Z_o(y)|\theta_o = x, \xi_o = \star]$$
$$= \prod_{u \sim o} \left(1 + \varepsilon(1-p)^2(q\mathbf{1}_{x=y} - 1) + (1-\varepsilon)(1-p)^2 q \mathbb{E}[\delta_u(y - x + \theta_u)|\xi_u = \star] \right)$$
$$= \left(1 + \varepsilon(1-p)^2(q\mathbf{1}_{x=y} - 1) + (1-\varepsilon)(1-p)^2 q \mathbb{E}[\delta_u(y - x + \theta_u)|\xi_u = \star] \right)^k,$$

where u is an arbitrary offspring since terms participating in the product are all equal. Now we deal with the second identity (8.5):

$$\mathbb{E}[Z_o(y)Z_o(y')|\theta_o = x, \xi_o = \star] = \prod_{u \sim o} \mathbb{E}\Big[\big(p + (1-p)q\mu_u(y - Y_{ou})\big)$$
$$\big(p + (1-p)q\mu_u(y' - Y_{ou})\big)|\theta_o = x, \xi_o = \star\Big]$$
$$= \Big(p^2 + p(1-p)q\big(\mathbb{E}[\mu_u(y - Y_{ou})|\theta_o = x, \xi_o = \star]$$
$$+ \mathbb{E}[\mu_u(y' - Y_{ou})|\theta_o = x, \xi_o = \star]\big)$$
$$+ (1-p)^2 q^2 \mathbb{E}[\mu_u(y - Y_{ou})\mu_u(y' - Y_{ou})|\theta_o = x, \xi_o = \star]\Big)^k.$$

Similarly to a previous computation, we have

$$\mathbb{E}[\mu_u(y - Y_{ou})|\theta_o = x, \xi_o = \star] = \varepsilon\Big((1-p)\mathbf{1}_{x=y} + \frac{p}{q}\Big)$$
$$+ (1-\varepsilon)\Big((1-p)\mathbb{E}[\mu_u(y - x + \theta_u)|\xi_u = \star] + \frac{p}{q}\Big),$$

and

$$\mathbb{E}[\mu_u(y - Y_{ou})\mu_u(y' - Y_{ou})|\theta_o = x, \xi_o = \star] = \varepsilon\mathbb{P}(y - Y_{ou} = y' - Y_{ou} = \theta_u|\theta_o = x)$$
$$+ (1-\varepsilon)\mathbb{E}[\mu_u(y - Y_{ou})\mu_u(y' - Y_{ou})|\theta_o = x, \xi_o = \star, \xi_u = \star]$$
$$= \varepsilon\mathbf{1}_{y=y'}\Big((1-p)\mathbf{1}_{y=x} + \frac{p}{q}\Big)$$
$$+ (1-\varepsilon)\Big((1-p)\mathbb{E}[\mu_u(y - x + \theta_u)\mu_u(y' - x + \theta_u)|\xi_u = \star]$$
$$+ \frac{p}{q}\sum_z \mathbb{E}[\mu_u(y - z)\mu_u(y' - z)|\xi_u = \star]\Big).$$

Combining and rearranging terms we obtain the desired result. □

Now we use the expressions just obtained to produce Taylor estimates for each term in the decomposition (8.3).

Lemma 8 *Let $X = \mathbb{E}[\delta_u(\theta_u)|\xi_u = \star]$ and $\hat{\kappa} = k(1-p)^2$. There exists constants c, C depending only on q such that if $\hat{\kappa}|X| \le c$ and $\hat{\kappa}\varepsilon < c$, then*

$$\Big|\mathbb{E}[Z_o(x)|\theta_o = x, \xi_o = \star] - 1 - \varepsilon\hat{\kappa}(q - 1) - (1-\varepsilon)\hat{\kappa}qX\Big| \le C\hat{\kappa}^2(X^2 + \varepsilon^2),$$

(8.6)

$$\Big|\mathbb{E}\Big[Z_o(x)\Big(\sum_y Z_o(y) - q\Big)\Big|\theta_o = x, \xi_o = \star\Big]\Big| \le C\hat{\kappa}^2(X^2 + \varepsilon^2),$$

(8.7)

$$\text{and } \mathbb{E}\left[\left(\sum_{y\in\mathbb{Z}_q} Z_o(y) - q\right)^2 \middle| \theta_o = x, \xi_o = \star\right] \leq C\hat{\kappa}^2(X^2 + \varepsilon^2).$$

(8.8)

Proof We use $|(1+x)^d - 1 - x| \leq e^c d^2 x^2$ for all x such that $d|x| \leq c$. Applying this to (8.4) yields (8.6). For $k(1-p)^2 q|X| \leq 1/2$ and $k(1-p)^2(q-1)\varepsilon < 1/2$ we have

$$\left|\mathbb{E}[Z_o(x)|\theta_o = x, \xi_o = \star] - 1 - \varepsilon\hat{\kappa}(q-1) - (1-\varepsilon)\hat{\kappa}qX\right|$$
$$\leq e\hat{\kappa}^2\big(\varepsilon(q-1) + (1-\varepsilon)qX\big)^2 \leq 2e\hat{\kappa}^2 q^2\big(\varepsilon^2 + X^2\big).$$

Next, we use (8.5), combined with the fact $\sum_y \delta_u(y) = 0$ to obtain that if $\hat{\kappa}(|X| \vee \varepsilon) \leq c(q)$ for some constant $c(q)$ then

$$\left|\sum_y \mathbb{E}[Z_o(x)Z_o(y)|\theta_o = x, \xi_o = \star] - q - \varepsilon\hat{\kappa}q(q-1) - (1-\varepsilon)\hat{\kappa}q^2 X\right| \leq C(q)k^2\Sigma^2,$$

where Σ gathers all the terms other than 1 in the expression (8.5), and the constant C depends on c. We use the inequality $(\sum_{i=1}^n x_i)^2 \leq n\sum_i x_i^2$ to obtain

$$k^2\Sigma^2 \leq C(q)\hat{\kappa}^2\Big(\varepsilon^2 + \mathbb{E}[\delta_u(\theta_u)|\xi_u = \star]^2$$
$$+ \sum_{y\in\mathbb{Z}_q} \mathbb{E}[\delta_u(y-x+\theta_u)|\xi_u = \star]^2 + \max_{z\in\mathbb{Z}_q} \mathbb{E}[\delta_u(z)^2|\xi_u = \star]^2\Big). \quad (8.9)$$

The last term was obtained by using Cauchy-Schwarz on the term $\sum_z \mathbb{E}[\delta_u(y-z)\delta_u(y'-z)|\xi_u = \star]$ in (8.5) and then replacing sums over y, z by maxima. Now it remains to show that the last two terms in (8.9) are bounded by X^2. Starting with the last term, we have

$$\max_{z\in\mathbb{Z}_q} \mathbb{E}[\delta_u(z)^2|\xi_u = \star]^2 \leq \left(\sum_{z\in\mathbb{Z}_q} \mathbb{E}[\delta_u(z)^2|\xi_u = \star]\right)^2 = \mathbb{E}[\delta_u(\theta_u)|\xi_u = \star]^2 = X^2.$$

As for the remaining term,

Lemma 9 We have $\sum_{y\in\mathbb{Z}_q} \mathbb{E}[\delta_u(y-x+\theta_u)|\xi_u = \star]^2 \leq qX^2$.

This implies $k^2\Sigma^2 \leq C(q)\hat{\kappa}^2(\varepsilon^2 + X^2)$. This, combined with (8.6), allows us to deduce (8.7). Now we treat the last term (8.8):

$$\mathbb{E}\left[\left(\sum_y Z_o(y) - q\right)^2 \middle| \theta_o = x, \xi_o = \star\right] = \sum_{y,y'} \mathbb{E}[Z_o(y)Z_o(y')|\theta_o = x, \xi_o = \star]$$
$$- 2q\sum_y \mathbb{E}[Z_o(y)|\theta_o = x, \xi_o = \star] + q^2.$$

Similarly to our treatment of the quantity Σ, we use expression (8.5) and perform a Taylor expansion to obtain

$$\left| \sum_{y,y'} \mathbb{E}[Z_o(y)Z_o(y')|\theta_o = x, \xi_o = \star] - q^2 \right| \leq C(q)\hat{\kappa}^2(\varepsilon^2 + X^2).$$

Using (8.4) the cross term can be estimated as

$$\left| \sum_{y \in \mathbb{Z}_q} \mathbb{E}[Z_o(y)|\theta_o = x, \xi_o = \star] - q \right| \leq C(q)\hat{\kappa}^2(\varepsilon^2 + X^2).$$

Now we conclude

$$\mathbb{E}\left[\left(\sum_{y \in \mathbb{Z}_q} Z_o(y) - q \right)^2 |\theta_o = x, \xi_o = \star \right] \leq C(q)\hat{\kappa}^2(\varepsilon^2 + X^2).$$

$\square$

Proof (Lemma 9) For $y \in \mathbb{Z}_q$, using Lemma 6 we have $\mathbb{E}[\delta_u(y + \theta_u)|\xi_u = \star] = \frac{1}{q}\sum_{z \in \mathbb{Z}_q} \mathbb{E}[\delta_u(y + z)|\theta_u = z, \xi_u = \star] = \sum_{z \in \mathbb{Z}_q} \mathbb{E}[\delta_u(y + z)\mu_u(z)|\xi_u = \star] = \sum_{z \in \mathbb{Z}_q} \mathbb{E}[\delta_u(y + z)\delta_u(z)|\xi_u = \star]$. The last equality follows from $\sum_z \delta_u(z) = 0$. Then

$$\sum_{y \in \mathbb{Z}_q} \mathbb{E}[\delta_u(y + \theta_u)|\xi_u = \star]^2 = \sum_{y \in \mathbb{Z}_q} \left(\sum_{z \in \mathbb{Z}_q} \mathbb{E}[\delta_u(y + z)\delta_u(z)|\xi_u = \star] \right)^2$$

$$\overset{(a)}{\leq} q \sum_{y,z \in \mathbb{Z}_q} \mathbb{E}[\delta_u(y + z)\delta_u(z)|\xi_u = \star]^2$$

$$\overset{(b)}{\leq} q \sum_{y,z \in \mathbb{Z}_q} \mathbb{E}[\delta_u(y + z)^2|\xi_u = \star]\mathbb{E}[\delta_u(z)^2|\xi_u = \star]$$

$$= q \left(\sum_{y \in \mathbb{Z}_q} \mathbb{E}[\delta_u(y)^2|\xi_u = \star] \right)^2.$$

Inequality (a) follows from $(\sum_{i=1}^n x_i)^2 \leq n \sum_i x_i^2$, and (b) follows from Cauchy-Schwarz. Lastly, we have $\sum_{y \in \mathbb{Z}_q} \mathbb{E}[\delta_u(y)^2|\xi_u = \star] = \mathbb{E}[\delta_u(\theta_u)|\xi_u = \star] = X$. $\square$

Now we plug the estimates of Lemma 8 in (8.3). Using the fact $0 \leq Z_o(x)/\sum_y Z_o(y) \leq 1$, we obtain

$$\left| \hat{z}_{o,l} - \varepsilon\hat{\kappa}\frac{q - 1}{q} - (1 - \varepsilon)\hat{\kappa}z_{l-1} \right| \leq C(q)\hat{\kappa}^2(z_{l-1}^2 + \varepsilon^2).$$

where $C(q)$ is a constant that depends only on q.

8.2 Proof of Corollary 4

We first prove the result concerning the overlap with $\boldsymbol{\theta}_0$. Let $\sigma \in \mathcal{S}_q$ be a fixed permutation. We have

$$\mathbb{P}\big(\hat{\theta}_u^{(l)} = \sigma(\theta_u)\big) - \frac{1}{q} = \sum_{x \in \mathbb{Z}_q} \mathbb{E}\big[\hat{\mu}_{G_n,u,l}(\sigma(x)) \mu_{G_n,u}(x)\big] - \frac{1}{q}$$

$$= \sum_{x \in \mathbb{Z}_q} \mathbb{E}\Big[\big(\hat{\mu}_{G_n,u,l}(\sigma(x)) - \frac{1}{q}\big) \mu_{G_n,u}(x)\Big]$$

$$\leq \sum_{x \in \mathbb{Z}_q} \mathbb{E}\Big[\big|\hat{\mu}_{G_n,u,l}(\sigma(x)) - \frac{1}{q}\big|\Big]$$

$$\leq \sqrt{q} \, \mathbb{E}\big[d_{\ell_2}(\hat{\mu}_{G_n,u,l}, \bar{\nu})^2\big]^{1/2}.$$

The last line follows by Cauchy-Schwarz and then Jensen's inequality. Averaging over $u \in V_n$, applying Jensen's inequality once more, and then using Theorem D yields the first statement.

Next, let $f : \mathbb{Z}_q \mapsto \mathbb{R}$ with $\sum_{x \in \mathbb{Z}_q} f(x) = 0$ and $\frac{1}{q} \sum_{x \in \mathbb{Z}_q} f(x)^2 = 1$. The loss of $\hat{\boldsymbol{X}}^{(l)}$ is

$$\mathcal{R}_n(\hat{\boldsymbol{X}}^{(l)}; f) = \frac{1}{n^2} \mathbb{E}\|\boldsymbol{X}_f\|_F^2 - \frac{2}{n^2} \mathbb{E}\langle \hat{\boldsymbol{X}}^{(l)}, \boldsymbol{X}_f \rangle + \frac{1}{n^2} \mathbb{E}\|\hat{\boldsymbol{X}}^{(l)}\|_F^2$$

$$\geq \frac{1}{n^2} \mathbb{E}\|\boldsymbol{X}_f\|_F^2 - \frac{2}{n^2} \mathbb{E}\langle \hat{\boldsymbol{X}}^{(l)}, \boldsymbol{X}_f \rangle.$$

We have $\mathbb{E}\|\boldsymbol{X}_f\|_F^2 = \sum_{u,v \in V_n} \mathbb{E}[f(\theta_u)^2 f(\theta_v)^2] = \frac{n}{q} \sum_{x \in \mathbb{Z}_q} f(x)^4 + n(n-1)$. So $\lim \frac{1}{n^2} \mathbb{E}\|\boldsymbol{X}_f\|_F^2 = 1$. On the other hand, since $\sum_{x \in \mathbb{Z}_q} f(x) = 0$, we have

$$\mathbb{E}\Big[f(\theta_u)\big|Y_{B_{G_n}(u,l)}^{(\varepsilon)}\Big] = \sum_{x \in \mathbb{Z}_q} \big(\hat{\mu}_{G_n,u,l}(x) - \frac{1}{q}\big) f(x)$$

$$= \sum_{x \in \mathbb{Z}_q} \hat{\delta}_{u,l,G_n}(x) f(x),$$

where $\hat{\delta}_{u,l,G_n}(x) = \hat{\mu}_{G_n,u,l}(x) - \frac{1}{q}$, $x \in \mathbb{Z}_q$. On the other hand we have

$$\mathbb{E}\langle \hat{\boldsymbol{X}}^{(l)}, \boldsymbol{X}_f \rangle = \mathbb{E}\Big[\Big(\sum_{u \in V_n} f(\theta_u) \mathbb{E}\big[f(\theta_u)\big|Y_{B_{G_n}(u,l)}^{(\varepsilon)}\big]\Big)^2\Big]$$

$$= \mathbb{E}\Big[\Big(\sum_{x \in \mathbb{Z}_q} f(x)\Big(\sum_{u \in V_n} \hat{\delta}_{u,l,G_n}(x) f(\theta_u)\Big)\Big)^2\Big].$$

We use Cauchy-Schwarz inequality and the fact $\sum_{x \in \mathbb{Z}_q} f(x)^2 = q$ to obtain

$$
\begin{aligned}
\mathbb{E}\langle \widehat{\boldsymbol{X}}^{(l)}, \boldsymbol{X}_f \rangle &\leq q \mathbb{E}\Big[\sum_{x \in \mathbb{Z}_q} \Big(\sum_{u \in V_n} \widehat{\delta}_{u,l,G_n}(x) f(\theta_u) \Big)^2 \Big] \\
&\leq q \mathbb{E}\Big[\Big(\sum_{u \in V_n} f(\theta_u)^2 \Big) \Big(\sum_{u \in V_n} d_{\ell_2}(\widehat{\mu}_{G_n,u,l}, \overline{v})^2 \Big) \Big] \\
&\leq nq \| f \|_\infty^2 \sum_{u \in V_n} \mathbb{E} d_{\ell_2}(\widehat{\mu}_{G_n,u,l}, \overline{v})^2.
\end{aligned}
$$

We apply Theorem 5.7 to obtain $\limsup_l \limsup_n \frac{1}{n^2} \mathbb{E}\langle \widehat{\boldsymbol{X}}^{(l)}, \boldsymbol{X}_f \rangle \leq C \| f \|_\infty^2 \varepsilon$, and this yields the desired result.

9 Reconstruction for the Stochastic Block Model

Here we define a version of the stochastic block model and sketch the proof of an analogue of Theorem C in this context. Let $\pi \in \Delta^{q-1}$ be a distribution on q colors and $M \in \mathbb{R}_+^{q \times q}$ be a Markov transition matrix, i.e., the rows of M are elements of Δ^{q-1}, and assume π and M satisfy the detailed balance condition

$$
\pi(x_1) M(x_1, x_2) = \pi(x_2) M(x_2, x_1) \quad \forall x_1, x_2 \in [q].
$$

As before we let n, k such that $nk = 2m$. We first draw an assignment $\boldsymbol{\theta}_0 \in [q]^n$ uniformly at random among the $\binom{n}{n\pi(\cdot)}$-many color assignments with proportions given by π. We then draw a graph $G_n = (V_n, E_n)$ with $V_n = [n]$ from the configuration model $\mathbb{P}_{n,k}^{cm}$ as described in Sect. 7.2 conditioned on the following edge-color proportions

$$
\forall x_1, x_2 \in [q], \; \big| \{ (u, v) \in E_n : \theta_{0u} = x_1, \theta_{0v} = x_2 \} \big| = m \, \pi(x_1) M(x_1, x_2). \quad (9.1)
$$

Remark 6 We note that the model is only defined when $n\pi(x)$ and $m\pi(x)M(x, x')$ are integers for every $x, x' \in [q]$, which we assume throughout. Dispensing with these restrictions requires putting error bounds on $d_{TV}(\hat{v}_{\boldsymbol{\theta}_0}^G, \pi M)$ where $\hat{v}_{\boldsymbol{\theta}_0}^G$ the joint empirical distribution of the pair $(G, \boldsymbol{\theta}_0)$ (see below) which would reflect an error term in (9.1), and doing a truncation in the first moment computation (i.e., restricting the computation to a notion of a 'good instance'), exactly in the same way done in Sect. 7. We forgo those details and work with equalities.

We observe the graph G_n and we are asked to infer the hidden coloring $\boldsymbol{\theta}_0$. As before, for any graph $G = (V = [n], E)$ and any pair $(\boldsymbol{\theta}_0, \boldsymbol{\theta}) \in [q]^V \times [q]^V$ we define the joint edge-empirical distribution of $(G, \boldsymbol{\theta})$ and the joint vertex-empirical distribution of $(\boldsymbol{\theta}_0, \boldsymbol{\theta})$ respectively as follows:

$$
\hat{v}_{\boldsymbol{\theta}}^G := \frac{1}{|E|} \sum_{(u,v) \in E} \delta_{\theta_u, \theta_v}, \quad \hat{\omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}} := \frac{1}{|V|} \sum_{u \in V} \delta_{\theta_{0,u}, \theta_u}.
$$

Define $\bar{v}_e(x_1, x_2) = \pi(x_1)M(x_1, x_2)$, and consider the estimator which outputs an arbitrary element of the set of typical assignments

$$\hat{\theta}(G_n) \in \Theta(G) := \left\{ \boldsymbol{\theta} \in [q]^V : \hat{v}_{\boldsymbol{\theta}}^G = \bar{v}_e \right\}.$$

(This set is non-empty since $\boldsymbol{\theta}_0 \in \Theta(G)$ due to the constraints (9.1).) We have an analogue of Theorem C:

Theorem 1 *Let $\theta_1 \sim \pi$ and $\theta_2 \mid \theta_1 \sim M(\theta_1, \cdot)$. If*

$$\frac{k}{2} I(\theta_1; \theta_2) \geq H(\theta_1) + \delta \tag{9.2}$$

for some $\delta > 0$ then there exists $\delta' > 0$ and a constant $c_0 > 0$ such that

$$\mathbb{P}\left(d_{\mathrm{TV}}(\hat{\omega}_{\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0}, \bar{v} \times \bar{v}) \geq \delta' \right) \geq 1 - c_0^{-1} e^{-c_0 n}.$$

We note that if the approximation details described in the previous remark are implemented, the above probability bound will deteriorate to $1 - c_0^{-1} \exp\{-c_0 (\log n)^2\}$, which is the probability of $(G_n, \boldsymbol{\theta}_0)$ being a good instance, just as is the case for Theorem C.

We can now look at a few special cases:

- In the balanced two-color case $q = 2$, $\pi = (\frac{1}{2}, \frac{1}{2})$ and $M = \begin{pmatrix} 1 - \theta & \theta \\ \theta & 1 - \theta \end{pmatrix}$, the simple computation shows that the condition (9.2) reduces to $k(\log 2 - h(\theta)) > 2 \log 2$ where $h(\theta) = -\theta \log \theta - (1 - \theta) \log(1 - \theta)$. For θ close to a half, this reduces to $k(1 - 2\theta)^2 + O(k(1 - 2\theta)^3) > 8 \log 2 \simeq 2.4$. In contrast the Kesten-Stigum bound is $k(1 - 2\theta)^2 > 1$.

- More generally, in the balanced symmetric case $\pi = (\frac{1}{q}, \cdots, \frac{1}{q})$, and

$$M = \lambda I + \frac{(1 - \lambda)}{q} \mathbf{1} \mathbf{1}^\top, \quad \lambda \in \left[-\frac{1}{q - 1}, 1 \right],$$

the condition (9.2) reduces to $k f_q(\lambda) > 2q \log q$ with

$$f_q(\lambda) = (1 + (q - 1)\lambda) \log(1 + (q - 1)\lambda) + (q - 1)(1 - \lambda) \log(1 - \lambda).$$

This is the condition stated in [12, Theorem 3] in the context of hypothesis testing for the usual SBM.

Proof (sketch) The proof of this theorem follows exactly the same scheme used in Sect. 7, so we will only highlight the main differences. For $\mathcal{S} \subseteq \mathscr{P}([q] \times [q])$, let

$$Z(G_n, \theta_0, \mathcal{S}) := \left| \left\{ \boldsymbol{\theta} \in \Theta(G_n) : \hat{\omega}_{\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0} \in \mathcal{S} \right\} \right|.$$

we have

$$\mathbb{E}\big[Z(G_n, \boldsymbol{\theta}_0, \mathcal{S})\big] = \frac{\sum_{G, \boldsymbol{\theta}_0, \boldsymbol{\theta}} \mathbf{1}\{\hat{v}_{\boldsymbol{\theta}_0}^{G_n} = \bar{v}_{\mathrm{e}},\ \hat{v}_{\boldsymbol{\theta}}^{G_n} = \bar{v}_{\mathrm{e}},\ \hat{\omega}_{\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_0} \in \mathcal{S}\}}{\binom{n}{n\pi(\cdot)} N(\pi, M)}, \tag{9.3}$$

where $N(\pi, M)$ is the number of graphs from the configuration model such that Eq. (9.1) holds, i.e, such that $\hat{v}_{\boldsymbol{\theta}_0}^{G_n} = \bar{v}_{\mathrm{e}}$. A counting argument similar to Lemma 3 yields

$$N(\pi, M) = \prod_{x \in [q]} \frac{[nk\pi(x)]!}{\prod_{x' \in [q]}[m\bar{v}_{\mathrm{e}}(x, x')]! \cdot [m\bar{v}_{\mathrm{e}}(x', x)]!} \prod_{x, x' \in [q]} [m\bar{v}_{\mathrm{e}}(x, x')]!$$

$$= \exp\Big(\frac{nk}{2} H(\bar{v}_{\mathrm{e}}) - nkH(\pi) + \frac{nk}{2} \log\Big(\frac{2nk}{e}\Big) + \mathcal{O}(\log n)\Big). \tag{9.4}$$

Next, we compute the numerator. We define

$$\widehat{\Omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}}^G = \frac{1}{|E|} \sum_{(u,v) \in E} \delta_{\theta_{0,u}, \theta_u, \theta_{0,v}, \theta_v},$$

and let $\mathcal{Q} \subseteq \mathscr{P}([q]^4)$ be the subset of $\Omega \in \mathscr{P}([q]^4)$ with entries in $\mathbb{Z}_+/m$ satisfying the constraints

$$\begin{cases} \frac{1}{2}(\pi_1\Omega + \pi_2\Omega) \in \mathcal{S}, \\ \Big(\sum_{\tilde{x}_1, \tilde{x}_2 \in [q]} \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2)\Big)_{x_1, x_2} = \bar{v}_{\mathrm{e}}, \\ \Big(\sum_{x_1, x_2 \in [q]} \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2)\Big)_{\tilde{x}_1, \tilde{x}_2} = \bar{v}_{\mathrm{e}}. \end{cases} \tag{9.5}$$

We recall the notation $(\pi_1\Omega)(x_1, \tilde{x}_1) = \sum_{x_2, \tilde{x}_2} \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2)$ and $(\pi_2\Omega)(x, x') = \sum_{x_2, \tilde{x}_2} \Omega(x_2, \tilde{x}_2, x_1, \tilde{x}_1)$. We also let $\omega(x, x') := \frac{1}{2}(\pi_1\Omega + \pi_2\Omega)(x, x')$. We have

$$\sum_{G, \boldsymbol{\theta}_0, \boldsymbol{\theta}} \mathbf{1}\{\hat{v}_{\boldsymbol{\theta}_0}^{G_n} = \bar{v}_{\mathrm{e}},\ \hat{v}_{\boldsymbol{\theta}}^{G_n} = \bar{v}_{\mathrm{e}}\} = \sum_{\Omega \in \mathcal{Q}} \sum_{G, \boldsymbol{\theta}_0, \boldsymbol{\theta}} \mathbf{1}\{\widehat{\Omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}}^G = \Omega\}$$

$$= \sum_{\Omega \in \mathcal{Q}} \binom{n}{n\pi(\cdot)} \prod_{x \in [q]} \binom{n\pi(x)}{n\omega(x, \cdot)} N(\Omega).$$

In the above sum, the first multinomial accounts for the number of ways of choosing $\boldsymbol{\theta}_0$, the second product account for the number of ways of choosing $\boldsymbol{\theta}$ in agreement with the overlap distribution $\hat{\omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}} = \omega$, and third term $N(\Omega)$ is the number of graphs G from the configuration model such that $\widehat{\Omega}_{\boldsymbol{\theta}_0, \boldsymbol{\theta}}^G = \Omega$. Using the identity in Eq. (9.4), we have

$$N(\Omega) = \prod_{x, x' \in [q]} \frac{[nk\omega(x, x')]!}{\prod_{\tilde{x}, \tilde{x}' \in [q]}[m\Omega(x, x', \tilde{x}, \tilde{x}')]![m\Omega(\tilde{x}, \tilde{x}', x, x')]!}$$

$$\cdot \prod_{x,x',\tilde{x},\tilde{x}'\in[q]} [m\Omega(x, x', \tilde{x}, \tilde{x}')]!$$

$$= \exp\left(-nkH(\omega) + \frac{nk}{2}H(\Omega) + \frac{nk}{2}\log\left(\frac{2nk}{e}\right) + \mathcal{O}(\log n)\right). \qquad (9.6)$$

Plugging this in Eq. (9.3), we end up with

$$\mathbb{E}\big[Z(G_n, \boldsymbol{\theta}_0, \mathcal{S})\big] = \sum_{\Omega\in\mathcal{Q}} \exp\left(-n(k-1)H(\omega) + \frac{nk}{2}H(\Omega) - \frac{nk}{2}H(\overline{v}_e)\right.$$

$$\left. + n(k-1)H(\pi) + \mathcal{O}(\log n)\right)$$

$$\leq n^C \exp\left(n\max_{\omega\in\mathcal{S}} S_*(\omega) + \mathcal{O}(\log n)\right),$$

where, recalling that $\overline{v}_e(x, x') = \pi(x)M(x, x')$,

$$S_*(\omega) \equiv \max_{\Omega\in\mathscr{P}([q]^4)} S(\Omega),$$

$$\text{subj. to} \quad (\pi_1\Omega + \pi_2\Omega)/2 = \omega,$$

$$\sum_{\tilde{x}_1,\tilde{x}_2\in\mathcal{X}} \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2) = \overline{v}_e(x_1, x_2),$$

$$\sum_{x_1,x_2\in\mathcal{X}} \Omega(x_1, \tilde{x}_1, x_2, \tilde{x}_2) = \overline{v}_e(\tilde{x}_1, \tilde{x}_2), \qquad (9.7)$$

and

$$S(\Omega) \equiv \frac{k}{2}H(\Omega) - (k-1)H\left(\frac{1}{2}(\pi_1\Omega + \pi_2\Omega)\right) - \frac{k}{2}H(\overline{v}_e) + (k-1)H(\pi).$$

For $\omega = \pi \times \pi$, let Ω_* the maximizer in (9.7), and let $(X_1, \tilde{X}_1, X_2, \tilde{X}_2) \sim \Omega_*$, then $(X_1, \tilde{X}_1) \sim \overline{v}_e$ and $(X_2, \tilde{X}_2) \sim \overline{v}_e$. Therefore, by subadditivity of the entropy,

$$S_*(\pi \times \pi) = \frac{k}{2}H(X_1, \tilde{X}_1, X_2, \tilde{X}_2) - (k-1)H(\pi \times \pi) - \frac{k}{2}H(X_1, X_2) + (k-1)H(\pi)$$

$$\leq \frac{k}{2}H(X_1, \tilde{X}_1) + \frac{k}{2}H(X_2, \tilde{X}_2) - (k-1)H(X_1) - \frac{k}{2}H(X_1, \tilde{X}_1)$$

$$= \frac{k}{2}H(\tilde{X}_1|X_1) + \frac{k}{2}H(X_1) - (k-1)H(X_1)$$

$$= -\frac{k}{2}I(\tilde{X}_1; X_1) + H(X_1).$$

The rest of the proof follows verbatim from Sect. 7. $\square$

10 Amenable graphs: some omitted proofs

10.1 Proof of Proposition 1

The proof is based on a decoupling principle under ε-perturbation of a general observation channel. This principle is given in Lemma 3.1 in [44], which once specialized to our setting, takes the following form:

Lemma 10 (Lemma 3.1 [44]) *For all $\varepsilon > 0$, it holds that*

$$\frac{1}{n} \int_0^\varepsilon \sum_{u,v \in V_n} I\left(\theta_u; \theta_v \big| Y_{G_n}^{(\varepsilon')}\right) d\varepsilon' \leq 2 \log |\mathcal{X}|.$$

This is very similar to our Lemma 2. In fact the latter follows the same line of proof.

Recall the definition of the decoupled estimator

$$\begin{aligned}
\widehat{X}_{uv}^{(\mathrm{dec})} &:= \mathbb{E}\left[f(\hat{\theta}_u) \big| Y_{G_n}^{(\varepsilon)}\right] \cdot \mathbb{E}\left[f(\hat{\theta}_v) \big| Y_{G_n}^{(\varepsilon)}\right] \\
&= \left(\sum_{x \in \mathcal{X}} \mu_{u,G_n}(x) f(x)\right) \cdot \left(\sum_{x \in \mathcal{X}} \mu_{v,G_n}(x) f(x)\right), \quad u, v \in V_n.
\end{aligned}$$

For a pair of vertices $u, v \in V_n$ we let $\mu_{u,v,G_n}(x, x') := \mathbb{P}\left(\theta_u = x, \theta_v = x' \big| Y_{G_n}^{(\varepsilon)}\right)$, for $x, x' \in \mathcal{X}$. Expanding the squares and cancelling equal terms we have

$$\mathcal{R}_n(\widehat{X}^{(\mathrm{dec})}; f) - \mathcal{R}_n^{\mathrm{Bayes}}(f) = \frac{1}{n^2}\left(\mathbb{E}\|X^{\mathrm{Bayes}}\|_F^2 - \mathbb{E}\|\widehat{X}^{(\mathrm{dec})}\|_F^2\right).$$

Moreover,

$$\mathbb{E}\|\widehat{X}^{(\mathrm{dec})}\|_F^2 = \sum_{u,v \in V_n} \mathbb{E}\left[\mathbb{E}\left[f(\theta_u) \big| Y_{G_n}^{(\varepsilon)}\right]^2 \cdot \mathbb{E}\left[f(\theta_v) \big| Y_{G_n}^{(\varepsilon)}\right]^2\right],$$

and

$$\mathbb{E}\|\widehat{X}^{\mathrm{Bayes}}\|_F^2 = \sum_{u,v \in V_n} \mathbb{E}\left[\mathbb{E}\left[f(\theta_u) f(\theta_v) \big| Y_{G_n}^{(\varepsilon)}\right]^2\right].$$

Therefore,

$$\begin{aligned}
&\frac{1}{n^2}\left(\mathbb{E}\|\widehat{X}^{\mathrm{Bayes}}\|_F^2 - \mathbb{E}\|\widehat{X}^{(\mathrm{dec})}\|_F^2\right) \\
&\leq \frac{2\|f\|_\infty^2}{n^2} \sum_{u,v \in V_n} \mathbb{E}\left[\left|\mathbb{E}\left[f(\theta_u) f(\theta_v) \big| Y_{G_n}^{(\varepsilon)}\right] - \mathbb{E}\left[f(\theta_u) \big| Y_{G_n}^{(\varepsilon)}\right]\mathbb{E}\left[f(\theta_v) \big| Y_{G_n}^{(\varepsilon)}\right]\right|\right] \\
&\leq \frac{2\|f\|_\infty^2}{n^2} \sum_{u,v \in V_n} \sum_{x,x' \in \mathcal{X}} |f(x) f(y)| \mathbb{E}\left[\left|\mu_{u,v,G_n}(x, x') - \mu_{u,G_n}(x)\mu_{v,G_n}(x')\right|\right]
\end{aligned}$$

$$\leq \frac{4\|f\|_\infty^4}{n^2} \sum_{u,v \in V_n} \mathbb{E}\big[d_{\mathrm{TV}}(\mu_{u,v,G_n}, \mu_{u,G_n} \times \mu_{v,G_n})\big]$$

$$\leq 4\|f\|_\infty^4 \sqrt{\frac{1}{2n^2} \sum_{u,v \in V_n} I\Big(\theta_u; \theta_v \big| Y_{G_n}^{(\varepsilon)}\Big)}.$$

We used Pinsker's inequality and Jensen's inequality in the last line. We apply Lemma 10 and Jensen's inequality and obtain for all $\varepsilon > 0$,

$$\int_0^\varepsilon \Big\{ \mathcal{R}_n(\widehat{\boldsymbol{X}}^{(\mathrm{dec})}; f) - \mathcal{R}_n^{\mathrm{Bayes}}(f) \Big\} d\varepsilon' \leq 4\|f\|_\infty^4 \sqrt{\frac{\varepsilon \log |\mathcal{X}|}{n}} \longrightarrow 0.$$

Since the integrand is non-negative, it too converges to zero almost everywhere.

10.2 Proof of Corollary 1

We lower bound the overlap as by taking the identity permutation:

$$\mathbb{E}\big[\mathtt{overlap}(\hat{\boldsymbol{\theta}}^{(l)}, \boldsymbol{\theta})\big] \geq \frac{1}{|V_n|} \sum_{u \in V_n} \mathbb{P}(\hat{\theta}_u = \theta_u)$$

$$= \frac{1}{|V_n|} \sum_{u \in V_n} \sum_{x \in \mathcal{X}} \mathbb{E}\big[\mathbf{1}_{\hat{\theta}_u = x} \mathbf{1}_{\theta_u = x}\big]$$

$$= \frac{1}{|V_n|} \sum_{u \in V_n} \sum_{x \in \mathcal{X}} \mathbb{E}\Big[\mathbb{E}\big[\mathbf{1}_{\hat{\theta}_u = x} \mathbf{1}_{\theta_u = x} | Y_{B_{G_n}(u,l)}^{(\varepsilon)}\big]\Big]$$

$$\overset{(a)}{=} \frac{1}{|V_n|} \sum_{u \in V_n} \sum_{x \in \mathcal{X}} \mathbb{E}\big[\hat{\mu}_{u,l,G_n}^2(x)\big]$$

$$\overset{(b)}{\longrightarrow} \sum_{x \in \mathcal{X}} \mathbb{E}\big[\mu_{G,o}(x)^2\big].$$

We in (a) used the fact that $\hat{\theta}_u$ and θ_u are independent and identically distributed conditional on $Y_{B_{G_n}(u,l)}^{(\varepsilon)}$. In (b) we take $n \to \infty$ followed by $l \to \infty$. The statement follows from local weak convergence of (G, o_n) towards (G, o), o_n being a uniformly random root of G_n.

References

1. Abbe, E.: Community detection and stochastic block models: recent developments. J. Mach. Learn. Res. **18**(1), 6446–6531 (2017)

2. Abbe, E., Boix, E.: An information-percolation bound for spin synchronization on general graphs. arXiv:1806.03227 (2018)
3. Abbe, E., Boix, E., Ralli, P., Sandon, C.: Graph powering and spectral robustness. arXiv:1809.04818 (2018)
4. Abbe, E., Massoulie, L., Montanari, A., Sly, A., Srivastava, N.: Group synchronization on grids. arXiv:1706.08561 (2017)
5. Abbe, E., Sandon, C.: Proof of the achievability conjectures for the general stochastic block model. Commun. Pure Appl. Math. **71**(7), 1334–1406 (2018)
6. Achlioptas, D., Naor, A.: The two possible values of the chromatic number of a random graph. In: Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of computing, pp. 587–593. ACM (2004)
7. Aldous, D., Lyons, R.: Processes on unimodular random networks. Electron. J. Probab. **12**(54), 1454–1508 (2007)
8. Alon, N., Krivelevich, M., Vu, V.H.: On the concentration of eigenvalues of random symmetric matrices. Israel J. Math. **131**(1), 259–267 (2002)
9. Amini, A.A., Wainwright, M.J.: High-dimensional analysis of semidefinite relaxations for sparse principal components. Ann. Stat. **37**(5B), 2877–2921 (2009)
10. Bandeira, A.S., Banks, J., Kunisky, D., Moore, C., Wein, A.S.: Spectral planting and the hardness of refuting cuts, colorability, and communities in random graphs. arXiv preprint arXiv:2008.12237 (2020)
11. Banks, J., Kleinberg, R., Moore, C.: The lov\'asz theta function for random regular graphs and community detection in the hard regime. arXiv preprint arXiv:1705.01194 (2017)
12. Banks, J., Moore, C., Neeman, J., Netrapalli, P.: Information-theoretic thresholds for community detection in sparse networks. In: Conference on Learning Theory, pp. 383–416. PMLR (2016)
13. Barak, B., Hopkins, S.B., Kelner, J., Kothari, P., Moitra, A., Potechin, A.: A nearly tight sum-of-squares lower bound for the planted clique problem. In: 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS), pp. 428–437. IEEE (2016)
14. Barbier, J., Krzakala, F., Macris, N., Miolane, L., Zdeborová, L.: Optimal errors and phase transitions in high-dimensional generalized linear models. Proc. Natl. Acad. Sci. **116**(12), 5451–5460 (2019)
15. Benjamini, I., Schramm, O.: Recurrence of distributional limits of finite planar graphs. Electron. J. Probab. **6**, 1–13 (2001)
16. Berthet, Q., Rigollet, P.: Complexity theoretic lower bounds for sparse principal component detection. In: Conference on Learning Theory, pp. 1046–1066 (2013)
17. Bollobás, B.: A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. Eur. J. Comb. **1**(4), 311–316 (1980)
18. Celentano, M., Montanari, A.: Fundamental barriers to high-dimensional regression with convex penalties. arXiv:1903.10603 (2019)
19. Coja-Oghlan, A., Krzakala, F., Perkins, W., Zdeborová, L.: Information-theoretic thresholds from the cavity method. Adv. Math. **333**, 694–795 (2018)
20. Csiszár, I., Shields, P.C.: Information theory and statistics: A tutorial (2004)
21. Decelle, A., Krzakala, F., Moore, C., Zdeborová, L.: Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. Phys. Rev. E **84**(6), 066106 (2011)
22. Dembo, A., Montanari, A., Sun, N.: Factor models on locally tree-like graphs. Ann. Probab. **41**(6), 4162–4213 (2013)
23. Deshpande, Y., Abbe, E., Montanari, A.: Asymptotic mutual information for the balanced binary stochastic block model. Inform. Inference: J. IMA **6**(2), 125–170 (2017)
24. Deshpande, Y., Montanari, A.: Finding hidden cliques of size $sqrt{N/e}$ in nearly linear time. Foundations of Computational Mathematics , 1–60 (2013)
25. Evans, W., Kenyon, C., Peres, Y., Schulman, L.J.: Broadcasting on trees and the ising model. Ann. Appl. Probab. **10**(2), 410–433 (2000)
26. Fan, Z., Montanari, A.: How well do local algorithms solve semidefinite programs? In: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, pp. 604–614. ACM (2017)
27. Feige, U., Krauthgamer, R.: Finding and certifying a large hidden clique in a semirandom graph. Random Struct. Algor. **16**(2), 195–208 (2000)
28. Gallager, R.G.: Information Theory and Reliable Communication, vol. 2. Springer, New York (1968)

29. Hopkins, S.B., Kothari, P.K., Potechin, A., Raghavendra, P., Schramm, T., Steurer, D.: The power of sum-of-squares for detecting hidden structures. In: 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), pp. 720–731. IEEE (2017)

30. Hopkins, S.B., Schramm, T., Shi, J., Steurer, D.: Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors. In: Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing, pp. 178–191. ACM (2016)

31. Hopkins, S.B., Shi, J., Steurer, D.: Tensor principal component analysis via sum-of-square proofs. In: Conference on Learning Theory, pp. 956–1006 (2015)

32. Hopkins, S.B., Steurer, D.: Efficient bayesian estimation from few samples: community detection and related problems. In: 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), pp. 379–390. IEEE (2017)

33. Janson, S., Mossel, E.: Robust reconstruction on trees is determined by the second eigenvalue. Ann. Probab. **32**(3B), 2630–2649 (2004)

34. Jerrum, M.: Large cliques elude the metropolis process. Random Struct. Algor. **3**(4), 347–359 (1992)

35. Johnstone, I.: High Dimensional Statistical Inference and Random Matrices. In: Proceedings of International Congress of Mathematicians, Madrid (2006)

36. Johnstone, I.M., Lu, A.Y.: On consistency and sparsity for principal components analysis in high dimensions. J. Am. Stat. Assoc. **104**(486), (2009)

37. Kunisky, D., Wein, A.S., Bandeira, A.S.: Notes on computational hardness of hypothesis testing: Predictions using the low-degree likelihood ratio. arXiv preprint arXiv:1907.11636 (2019)

38. Lelarge, M., Miolane, L.: Fundamental limits of symmetric low-rank matrix estimation. Probab. Theory Relat. Fields, 1–71 (2017)

39. Lyons, R., Peres, Y.: Probability on Trees and Networks, vol. 42. Cambridge University Press, Cambridge (2017)

40. Ma, T., Wigderson, A.: Sum-of-squares lower bounds for sparse pca. In: Advances in Neural Information Processing Systems, pp. 1612–1620 (2015)

41. Massoulié, L.: Community detection thresholds and the weak ramanujan property. In: Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing, pp. 694–703. ACM (2014)

42. Mézard, M., Montanari, A.: Reconstruction on trees and spin glass transition. J. Stat. Phys. **124**(6), 1317–1350 (2006)

43. Mézard, M., Montanari, A.: Information, Physics and Computation. Oxford University Press, Oxford (2009)

44. Montanari, A.: Estimating random variables from random sparse observations. Eur. Trans. Telecom. **19**, 385–403 (2008)

45. Montanari, A., Richard, E.: A statistical model for tensor pca. In: Advances in Neural Information Processing Systems, pp. 2897–2905 (2014)

46. Mossel, E., Neeman, J., Sly, A.: A proof of the block model threshold conjecture. Combinatorica **38**(3), 665–708 (2018)

47. Mossel, E., Peres, Y.: Information flow on trees. Ann. Appl. Probab. **13**(3), 817–844 (2003)

48. Newman, C., Schulman, L.: Number and density of percolating clusters. J. Phys. A: Math. Gen. **14**(7), 1735 (1981)

49. Penrose, M., et al.: Random Geometric Graphs, vol. 5. Oxford University Press, Oxford (2003)

50. Polyanskiy, Y., Wu, Y.: Application of information-percolation method to reconstruction problems on graphs. arXiv:1806.04195 (2018)

51. Sankararaman, A., Baccelli, F.: Community detection on euclidean random graphs. In: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 2181–2200. SIAM (2018)

52. Sly, A.: Reconstruction for the potts model. Ann. Probab. **39**(4), 1365–1406 (2011)

53. Wormald, N.C.: Models of random regular graphs. London Mathematical Society Lecture Note Series, pp. 239–298 (1999)