



HAL
open science

Partial Recovery in the Graph Alignment Problem

Georgina Hall, Laurent Massoulié

► **To cite this version:**

Georgina Hall, Laurent Massoulié. Partial Recovery in the Graph Alignment Problem. *Operations Research*, 2022, 10.1287/opre.2022.2355 . hal-03938708

HAL Id: hal-03938708

<https://hal.science/hal-03938708v1>

Submitted on 13 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Partial Recovery in the Graph Alignment Problem

Georgina Hall* and Laurent Massoulié†

Abstract

In this paper, we consider the graph alignment problem, which is the problem of recovering, given two graphs, a one-to-one mapping between nodes that maximizes edge overlap. This problem can be viewed as a noisy version of the well-known graph isomorphism problem and appears in many applications, including social network deanonymization and cellular biology. Our focus here is on *partial recovery*, i.e., we look for a one-to-one mapping which is correct on a fraction of the nodes of the graph rather than on all of them, and we assume that the two input graphs to the problem are correlated Erdős-Rényi graphs of parameters (n, q, s) . Our main contribution is then to give necessary and sufficient conditions on (n, q, s) under which partial recovery is possible with high probability as the number of nodes n goes to infinity. In particular, we show that it is possible to achieve partial recovery in the $nqs = \Theta(1)$ regime under certain additional assumptions.

Keywords: Graph Alignment, Correlated Erdős-Rényi graphs, Partial Recovery

1 Introduction

The goal of *graph alignment* is to infer, given two graphs, a one-to-one mapping between the nodes of the two graphs such that the number of edge overlaps is maximized. Graph alignment problems appear in a variety of different applications, including computer vision [8], natural language processing [31, 29], biology [17], and social network deanonymization [35]. In biology, graph alignment is used for example to match proteins across species. This is generally difficult to do via a direct comparison of the amino acid sequences that make up the proteins as mutations of the species' genotype can lead to significant differences in the sequences. However, while their sequences may vary greatly, proteins generally retain similar functions in each species' metabolism. By drawing up a Protein-Protein Interaction graph for each species and aligning these graphs across species, biologists are able to match proteins to better understand, e.g., the evolution of protein complexes. We refer the reader to [17] for more information on the topic. In the context of social networks, graph alignment has been successfully used to deanonymize private social networks. Arguably, the best-known example is that of the deanonymization of the Netflix Prize database using data from IMDb [34] as it led to a court case against Netflix and the end of the Netflix Prize. Other instances of social network deanonymization via graph alignment have been investigated as well, and comprise e.g. the deanonymization of Twitter using Flickr [35]. In both cases, the ideas are similar: if a user has an account on two social networks that serve similar purposes (e.g., both Netflix and IMDb enable movie reviewing), then the content shared on the two platforms by the user would be quite similar. In other words, if one was to create a graphical representation of the content shared

*Georgina Hall is with the department of Decision Sciences at INSEAD. Email: georgina.hall@insead.edu

†Laurent Massoulié is with the DYOGENE Group at INRIA and the MSR-Inria Joint Centre. Email: laurent.massoulie@inria.fr

on the platform by users, then these graphical representations would be very correlated across social networks. (Consider in the Netflix and IMDb case, a bipartite graph that links a user to a movie he or she has reviewed.) If one of the social platforms is public (as is IMDb) and the other, though being private, releases partial or anonymized information about its user content (as Netflix did during the Netflix Prize), then one can simply construct a similar graphical representation of the information, align the graphs obtained, and recover information on the private network from the public network, which leads to a breach of privacy of the private network’s user’s content.

The set-up we consider here is a stylized mathematical model of problems of this type: we assume that the input graphs are generated via a probabilistic model. More specifically, these input graphs are obtained by constructing two correlated Erdős-Rényi graphs with parameters (n, q, s) (see Section 2 for a mathematical description of correlated Erdős-Rényi graphs, roughly speaking, n is the number of nodes of the graphs, q is a measure of their sparsity, and s describes how correlated they are), selecting one of the two graphs at random and randomly permuting its labels, keeping the other graph as is. The question we wish to answer is then: given two graphs generated thus, under what conditions on (n, q, s) are we able to recover the underlying permutation? Or, to be more precise, under what conditions on (n, q, s) are we able to *partially* recover the underlying permutation, i.e., recover a permutation that overlaps with the permutation that was used on just a constant fraction α of its vertices? The notion of partial recovery is in contrast with the notions of exact (resp. almost exact) recovery which would aim to recover the correct permutation for all nodes (resp. all but a vanishing fraction of nodes) of the graph. While information-theoretical results such as ours exist in the literature for these two types of recovery, the notion of partial recovery has, to the best of our knowledge, not yet been considered (see Section 2 for a literature review). This is in spite of its relevance in applications: for example, it would be enough in the case of social network deanonymization to provide a partially recovered permutation to prove that a privacy breach has occurred. This paper aims to fill this gap in our collective knowledge.

In it, we provide both necessary and sufficient conditions for partial recovery of a graph alignment between two correlated Erdős-Rényi graphs. Our necessary condition states that if the difference between two probability distributions measured via the Kullback-Leibler divergence is too small—in fact, small in front of 1—then one cannot hope to recover a partial alignment. The two probability distributions that are contrasted in the result are the joint distribution of two independent edges (one in each graph) and the joint distribution of two correlated edges (again, one in each graph). This result is consequently very easy to interpret: if two correlated edges too closely mimic in probabilistic terms two independent edges, then there is no hope to recover any meaningful alignment. For our sufficient condition, we show that if $nqs = \Theta(1)$ and some other conditions relating to sparsity of the graphs hold, then partial recovery is possible. Though this result is perhaps less intuitively interpreted, it has the advantage of closely paralleling existing information-theoretic results for other types of recovery. Roughly speaking, it is known that the regime $nqs = \log(n) + \omega(1)$ leads to exact recovery and that the regime $nqs = \omega(1)$ leads to almost exact recovery: our result goes a step further—we now know that the regime $nqs = \Theta(1)$ is a regime in which partial recovery is possible. Though our necessary and sufficient conditions are not always easy to compare, it should be noted that in the case where $s = s_0$ and $q = n^{-\beta}$ for some positive β , they nearly read as counterparts to one another (see Section 2.5 for a more detailed analysis).

Before moving onto the outline of the paper, we briefly comment on the proofs of our results to give the reader a flavor of the rest of the paper. The proof for necessity is based on a generalization of Fano’s inequality. The proof for sufficiency involves designing and analyzing an algorithm which is guaranteed to return a permutation that overlaps with the true permutation on a fraction α of the nodes if our sufficient conditions hold. The algorithm we propose has some commonalities with that proposed in [12] in the sense that it also involves the study of the degrees of the nodes

of a particular intersection graph. However, our analysis is necessarily more involved in order to capture results in a regime ($nqs = \Theta(1)$) which [12] does not cover.

The rest of the paper is structured thus: in Section 2, we describe the mathematical formulation of our model, introduce our notation, present our main results, and discuss practical applications and future results. We also provide a literature review of existing work on the graph alignment problem and contrast these results to ours. In Section 3, we give the proof of validity of our necessary conditions and in Section 4, that of our sufficient conditions.

2 Mathematical description of our problem

The goal of this section is to set the stage for the rest of the paper. In Section 2.1, we present the mathematical formulation of the graph alignment problem and briefly review the literature on the topic. In Section 2.2, we focus more specifically on the case where the two input graphs are assumed to be correlated Erdős-Rényi graphs and review the literature that is specific to this case. In Section 2.3, we provide a short reminder on Landau notation and, in Section 2.4, we give our main results. We end with Section 2.5 which discusses the practical implications of our work, how the bounds compare, and future directions.

2.1 The graph alignment problem

Let $G_A = (V_A, E_A)$ and $G_B = (V_B, E_B)$ be two undirected unweighted graphs such that $|V_A| = |V_B| = n$. Recall that the adjacency matrix of a graph G on n vertices is an $n \times n$ binary matrix with entry (i, j) equal to 1 if there is an edge between node i and node j in the graph, and 0 otherwise. We denote by A and B the adjacency matrices of G_A and G_B respectively. The goal of graph alignment is to recover a permutation π of the labels of G_B such that the overlap between G_A and G_B is maximized. If P_n is the set of permutations over the set $\{1, \dots, n\}$, this can be written as the following optimization problem:

$$\max_{\pi \in P_n} \sum_{1 \leq i, j \leq n} A_{ij} B_{\pi(i)\pi(j)} + \sum_{1 \leq i, j \leq n} (1 - A_{ij})(1 - B_{\pi(i)\pi(j)}),$$

i.e., we maximize the number of overlapping edges and the number of non-overlapping edges between G_A and the graph obtained by relabeling G_B . This optimization problem is equivalent to $\max_{\pi \in P_n} \sum_{1 \leq i, j \leq n} A_{ij} B_{\pi(i)\pi(j)}$, which corresponds to maximizing the number of overlapping edges only. Note that this equivalence is not true anymore when we consider a generalization of the previous model to a case where the number of nodes in G_A and the number of nodes in G_B are different, see [20]. One can rewrite the previous optimization problem in yet another way by using permutation matrices rather than permutations. Recall that if π is a permutation over $\{1, \dots, n\}$, then the corresponding permutation matrix Π is an $n \times n$ binary matrix with entry (i, j) equal to 1 if and only if $\pi(i) = j$. We use the convention throughout the paper that lower-case π refers to the permutation, and upper-case Π refers to the corresponding permutation matrix. Furthermore, we denote by P_n^M the set of $n \times n$ permutation matrices. With this notation in mind, the graph alignment problem can be written

$$\max_{\Pi \in P_n^M} \text{Tr}(A\Pi B\Pi^T). \tag{1}$$

This problem is related to a couple of other well-known combinatorial problems [8]. The graph isomorphism problem, e.g., is a particular case of the graph alignment problem, where there exists

a permutation matrix Π such that $\text{Tr}(A\Pi B\Pi^T) = \text{Tr}(A^2)$, i.e., the graphs G_A and G_B overlap exactly up to a relabeling of the nodes of G_B . This problem has generated a lot of interest as it sits in a particular computational complexity class: it is not known if deciding whether two graphs are isomorphic is polynomial-time solvable or NP-hard. We refer the reader to [38] for a survey of the topic. At the other end of the spectrum, the quadratic assignment problem (QAP) is a generalization of the graph alignment problem, where the matrices A and B in (1) are not limited to adjacency matrices of graphs [39]. The quadratic assignment problem is known to be NP-hard to solve [40]. However, it does not follow from this result that the subcase of graph alignment problems is NP-hard to solve as we are imposing additional structure on the matrices A and B . In fact, the instances that are shown to be hard to solve in [40] explicitly preclude the proof from transferring to the graph alignment setting as it relies on the fact that A and B are non-binary matrices. A different reduction based on the maximum common subgraph problem can be used however to show NP-hardness of the graph alignment problem [5].

As a consequence of its intractability, heuristics for solving the graph alignment problem (i.e., problem (1)) abound in the literature. These heuristics can be divided roughly into four categories: methods based on network topological similarity, message-passing methods, spectral methods, and convex relaxation-based methods. Some of these methods do have guarantees under the assumption that the input graphs are randomly generated following the correlated Erdős-Renyi model, a model which we formally describe in Section 2.2. We consequently revisit some of the literature we describe here in Section 2.2. The methods based on network topological similarity are probably the most numerous and are especially prevalent in biology applications. They typically involve identifying and matching in both graphs either subgraphs of specific shapes (also known as *graphlets*) or nodes with specific neighborhood structures. Algorithms of this type appear in [28, 32, 54, 4, 15]. A more exhaustive list can be found in [8]. Message-passing methods include belief propagation approaches such as the one given in [5] as well as percolation-based methods which assume an initial seed set of matched nodes [52, 24, 33, 14]. For some of the latter approaches, the requirement that a seed set be given can be somewhat relaxed by considering a phase 1/phase 2-type algorithm. In the first phase, the algorithm infers from the graphs a reasonable seed set by using high-degree nodes; this seed set is then used to proceed with the algorithm in phase 2. Methods which use spectral properties of the graphs are also popular; see, e.g. [46] which relies on a PageRank-style algorithm, or [47, 20, 18, 19] which focus on aligning the eigenvectors of the adjacency matrices of the graphs. Finally, some methods involve approximating the optimization problem given in (1) by a convex optimization problem. Indeed, (1) is equivalent to $\min_{\Pi \in P_n^M} \|A - \Pi B \Pi^T\|_F^2$, where $\|\cdot\|_F$ is the Frobenius norm. By invariance of the Frobenius norm with respect to orthogonal matrices, this can be written as $\min_{\Pi \in P_n^M} \|A\Pi - \Pi B\|_F^2$. A natural convex relaxation can then be obtained by replacing the set of permutation matrices by its convex hull, the set of doubly stochastic matrices, i.e., the set of matrices

$$DS_n := \{D \in \mathbb{R}^{n \times n} \mid D\mathbf{1} = \mathbf{1}, D^T\mathbf{1} = \mathbf{1}, D \geq 0\},$$

where $\mathbf{1}$ is the all-ones vector. The problem then becomes $\min_{D \in DS_n} \|AD - DB\|_F^2$, which is a quadratic program to solve. Unfortunately, the matrix obtained is rarely a permutation matrix—in fact, it can be shown [30] that, under the correlated Erdős-Renyi model that we describe in Section 2.2, this is very unlikely to occur. A second step consequently involves the projection of the doubly stochastic matrix obtained to the set of permutation matrices, a projection that can be done in polynomial time, via, e.g., the Hungarian algorithm. Quadratic programming relaxations of this type are studied in [22, 48, 26, 16, 20, 21, 53]. These relaxations sometimes also involve a quadratic regularization term added to the objective, which corresponds to the Frobenius norm of the doubly stochastic matrix. An alternative convex relaxation for (1) which gives rise to a linear program

rather than a quadratic program can be obtained by replacing the Frobenius norm squared by the 1-norm applied to the vectorized matrix, i.e.,

$$\|A\Pi - \Pi B\|_1 = \sum_{i,j} |(A\Pi - \Pi B)_{ij}|.$$

This approach is much less prevalent in the literature and to the best of our knowledge only appears in [2]. There also exist convex relaxations of (1) based on semidefinite programming. These are relaxations that have more often than not been developed for the QAP and are then leveraged on this subcategory of problems [44, 55]. Though they can be quite powerful, their main caveat remains their prohibitively high running-time.

The boundaries of these four classes of approaches for the graph alignment problem are of course quite porous with many links between them. For example, the topological similarity method described in [15], which involves matching nodes via the number of neighbors of their neighbors, can be viewed as a second-order approximation of the doubly stochastic relaxation. Indeed, it is known that two graphs G_A and G_B have the same *ultimate degree sequence* if and only if there exists a doubly stochastic matrix D such that $AD = DB$ [43]. As the ultimate degree sequence of a graph is the ultimate degree sequence of its nodes, and the ultimate degree sequence of a node is an infinite list comprising in first position the degree of the node, in second position the degree of its neighbors, in third position the degree of the neighbors of its neighbors, and so on, it is not difficult to see why the algorithm in [15] is a second-order approximation of the doubly stochastic relaxation. Likewise, the spectral method presented in [19] can be viewed as a less-constrained version of the doubly stochastic relaxation; see [19] for details. Further note that we have focused above on methods that return graph alignments under the assumption that the graphs G_A and G_B are unweighted graphs on the same number of nodes. Extensions to the graph alignment problem which cover cases where the number of nodes in each graph are different [25] or where the graphs are weighted [45] have also been developed.

2.2 The correlated Erdős-Rényi model

The goal of this paper is not to study the model in its full generality, nor from a computational perspective. Rather, we focus on a particular random generative model for the input graphs, the correlated Erdős-Rényi model [36], which is very common in the literature, and wish to understand the conditions under which one can or cannot recover an alignment in this model from an information-theoretic perspective.

We start this section with a description of the Erdős-Rényi model. Let n be an integer, and q, s scalars in $[0, 1]$ with $q < s$. Let G be an Erdős-Rényi graph with parameters n and $q/s \leq 1$. In other words, G is a graph on n nodes, two nodes of which are independently connected with probability q/s . Two copies of this graph are created and each copy is subsampled with probability s , which means that each edge present in the copies is maintained with probability s and deleted with probability $1 - s$ independently and at random. We denote by G_A and $G_{B'}$ the two graphs obtained and by A and B' their adjacency matrices. They are both Erdős-Rényi graphs with parameter q , but entries (i, j) of A and B' are correlated Bernoulli random variables with correlation $\frac{s-q}{1-q}$. This is what gives this generation process the name of correlated Erdős-Rényi model. We then let π^* be a permutation over $\{1, \dots, n\}$ with associated permutation matrix Π^* . We obtain G_B by permuting the labels of $G_{B'}$ with π^* , i.e., $B_{\pi^*(i)\pi^*(j)} = B'_{ij}$ or equivalently, $\Pi^* B' \Pi^{*T} = B'$. The input to the graph alignment problem is then G_A and G_B and our goal is to recover π^* (or equivalently Π^*) from these observed graphs. In this framework, there is a notion of ground truth, given by π^* , which we refer to as *the* graph alignment in the rest of the paper. It is related to the definition given

in Section 2.1 (which is the one used in the absence of a generative model and a ground truth) in the sense that, in the correlated Erdős-Rényi model, any matrix $\hat{\Pi}$ solution to (1) is a maximum a posteriori estimator of Π^* provided that the correlation between edges is positive, i.e., $s > q$.

As mentioned previously, the parameter s quantifies the amount of noise added to the model. If $s = 1$, G_A and G_B are isomorphic. In this setting, one can recover exactly the permutation both information-theoretically and computationally if $\log(n) + \omega(1) \leq nq \leq n - \log(n) - \omega(1)$ [6, 50, 13]. The proofs of this result rely on showing that the automorphism group of G_A is trivial under these conditions.

In our case, we assume $s < 1$ and wish to recover the permutation π^* *partially*. This means that we wish to find a permutation $\hat{\pi}$ such that the overlap between $\hat{\pi}$ and the permutation π^* ,

$$\beta(\hat{\pi}, \pi^*) := \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\pi(i)=\pi^*(i)}$$

is greater than some constant $\alpha > 0$ with high probability as n goes to infinity. We refer to such a permutation $\hat{\pi}$ as a partial alignment of the graphs G_A and G_B , or alternatively we say that we have partially recovered an alignment of G_A and G_B . This concept is defined as an analog to the concept of partial recovery in the stochastic block model for community detection [1]. As far as we know, this paper is the first to cover this notion of recovery for graph alignment problems in the correlated Erdős-Rényi model from an information theoretic perspective (some computational results for sparse graphs can be found in [23]).

A vast majority of prior papers have instead focused on *exact recovery*, i.e., the case where one wishes to find $\hat{\pi}$ such that $\mathbb{P}(\beta(\hat{\pi}, \pi^*) = 1) = 1 - o(1)$. It is known that exact recovery can be achieved if $nqs \geq \log(n) + \omega(1)$ together with some conditions on the sparsity of the graphs and the strength of correlation of the edges. A converse result which states that one cannot hope to exactly recover π^* if $nqs \leq \log(n) - \omega(1)$ and the correlation between the edges is strictly less than one is also given [10, 11]. The algorithm that shows that exact recovery can be achieved under the aforementioned conditions is exponential time. Many papers have consequently focused on providing polynomial-time (or quasi-polynomial-time) algorithms that exactly recover π^* under certain conditions on the parameters [4, 33, 14, 15]. [4] propose a quasi-polynomial-time algorithm that exactly recovers π^* under the assumption that $nq/s \in [n^{o(1)}, n^{1/153}] \cup [n^{2/3}, n^{1-\delta}]$ for small $\delta > 0$. Mossel and Xu provide instead a seeded algorithm: under the assumption that the seed set is large enough (i.e., the fraction of seeded nodes is larger than $n^{-1/2+3\epsilon}$ for fixed $\epsilon > 0$), then exact recovery is possible all the way down to the information-theoretic threshold (under the condition that $nq/s \leq n^{1/2+\epsilon}$). [14] extend an algorithm called *canonical labeling*, which proved to be successful for graph isomorphism, to the graph alignment setting. The algorithm relies on a first step, which generates a surrogate of a seed set, to move on to the second set where these seeds are used to align the rest of the graph. Under the assumption that one is able to generate a sufficiently large set in the first step, then exact recovery can be achieved under the condition that $nqs \geq \omega(n^{4/5} \log(n)^{7/5})$ and $q(1-s) \leq o((qs)^5 / \log(n)^6)$. Finally, [15] analyze a polynomial-time algorithm based on neighbors of neighbors of nodes in the graph and are able to show its success when $nq = \Omega(\log(n)^2)$ and $1-s = O(1/\log(n)^2)$. [18, 19] propose an algorithm called GRAMPA for graph alignment: in particular, they show that if two graphs are correlated Erdős-Rényi graphs with edge correlation coefficient $1 - \sigma^2$ and average degree at least $\text{polylog}(n)$, then GRAMPA exactly recovers the alignment whp when $\sigma \lesssim 1/\text{polylog}(n)$.

Another type of recovery that has been investigated in the literature is that of *almost exact recovery*, i.e., the case where one wishes to recover $\hat{\pi}$ in such a way that $P(\beta(\hat{\pi}, \pi^*) = 1 - o(1)) = 1 - o(1)$. This is a stronger notion of recovery than the one we are interested in though it also

appears in the literature under the name of partial recovery [12]. For the moment, only information-theoretic results have been shown for this case: using the notion of k -core alignments, [12] show that almost exact recovery is possible when $nqs \geq \omega(1)$, $\frac{q^2(1-s)^2}{qs(1-2q+qs)} + 2q(1-s) \leq n^{-\Omega(1)}$ and $qs \leq \frac{1}{8e^3}$. A converse result in the same paper states that almost exact recovery is impossible when $nqs \leq O(1)$ (and correlation between the edges in the graphs is strictly less than one.) As this paper is the closest in the literature to ours, we further contrast our results to the ones appearing in [12] in Section 2.4, where we state our main results.

2.3 A reminder regarding Landau notation

As Landau notation is prevalent in the rest of the paper, we briefly remind the reader of the conventions used in this notation. Let f and g be two functions of n . We say that $f(n) \sim g(n)$ if $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 1$. We further say that $f(n) = o(g(n))$ if $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 0$ and that $f(n) = \omega(g(n))$ if $\lim_{n \rightarrow \infty} \left| \frac{f(n)}{g(n)} \right| = \infty$. We finally say that $f(n) = O(g(n))$ (resp. $f(n) = \Omega(g(n))$) if there exists $k > 0$ and $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$, $|f(n)| \leq k \cdot g(n)$ (resp. $f(n) \geq k \cdot g(n)$), and that $f(n) = \Theta(g(n))$ if $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$.

2.4 Main results

The focus of this paper is on the case of *partial recovery*, i.e., for fixed $\alpha \in (0, 1)$, we wish to recover $\hat{\pi}$ such that

$$\mathbb{P}(\beta(\hat{\pi}, \pi^*) > \alpha) = 1 - o(1).$$

A running assumption throughout the paper is that $s > q$. This implies that the covariance of (A_{ij}, B'_{ij}) , which we denote by σ^2 , is positive, or equivalently that the correlation between A_{ij} and B'_{ij} , which we denote by ρ , is positive. Indeed,

$$\sigma^2 = q(s - q) \text{ and } \rho = \frac{q(s - q)}{q(1 - q)}. \quad (2)$$

2.4.1 Impossibility result.

The statement of the theorem requires us to first define two probability distributions over $\{0, 1\} \times \{0, 1\}$, which will appear frequently in the rest of the paper. These are the probability distributions of two correlated edges (A_{ij}, B'_{ij}) in G_A and $G_{B'}$, and of two independent edges (A_{ij}, B'_{kl}) , $(i, j) \neq (k, l)$, in G_A and $G_{B'}$. Let P be the distribution of (A_{ij}, B'_{ij}) . We use the shorthand p_{xy} to denote the probability that $A_{ij} = x$ and $B'_{ij} = y$. By definition of the correlated Erdős-Renyi model, we have that

$$p_{00} = 1 - 2q + qs, \quad p_{01} = q(1 - s), \quad p_{10} = q(1 - s), \quad p_{11} = qs. \quad (3)$$

Likewise, let Q be the distribution of (A_{ij}, B'_{kl}) . Using similar notation, we have that:

$$q_{00} = 1 - 2q + q^2, \quad q_{01} = q(1 - q), \quad q_{10} = q(1 - q), \quad q_{11} = q^2. \quad (4)$$

Intuitively, our success in finding a graph alignment should hinge on how similar these two distributions are: if they are too similar for example, then we cannot hope to recover the underlying graph alignment. In this direction we introduce a measure of disparity between probability distributions,

known as the *Kullback-Leibler* divergence. Recall that the Kullback-Leibler divergence between two discrete probability distributions P and Q over \mathcal{X} is given by

$$D(P||Q) = \sum_{x \in \mathcal{X}} P(x) \log \left(\frac{P(x)}{Q(x)} \right).$$

It is always nonnegative and equals zero when $P = Q$ almost everywhere. We are now ready to state our theorem.

Theorem 2.1. *Let G_A and G_B be two graphs generated via the correlated Erdős-Renyi model (n, q, s) with π^* selected uniformly at random. Let $\alpha \in (0, 1)$, possibly dependent on n . If*

$$\frac{D(P||Q)}{\alpha} = o\left(\frac{\log(n)}{n}\right)$$

then, for any algorithm which infers $\hat{\pi}$ from G_A and G_B , we have $\mathbb{P}[\beta(\hat{\pi}, \pi^) < \alpha] \rightarrow 1$ when $n \rightarrow \infty$. In other words, no learning algorithm will be able to infer, from G_A and G_B , a permutation $\hat{\pi}$ that overlaps with π^* on more than a fraction α of its vertices.*

The proof of this result is given in Section 3 and relies on a generalization of Fano's inequality. Theorem 2.1 is easy to interpret: when P and Q are close, then $D(P||Q)$ is smaller, which makes this condition easier to meet. In other words, as P and Q get more and more similar, it becomes harder to recover the graph alignment. This is also true as α gets closer to 1, i.e., when we are trying to recover a more significant fraction of the nodes. Theorem 2.1 implies the following corollary, which enables comparisons with [12] and Theorem 2.3.

Corollary 2.2. *Let G_A and G_B be two graphs generated via the correlated Erdős-Renyi model (n, q, s) with π^* selected uniformly at random. Let $\alpha \in (0, 1)$, possibly dependent on n . If*

$$q = \Omega(n^{-k}) \text{ for some } k > 0, \quad q = o(n^{-1/2}), \quad \text{and } nqs = o(1), \quad (5)$$

then, for any algorithm which infers $\hat{\pi}$ from G_A and G_B , we have $\mathbb{P}[\beta(\hat{\pi}, \pi^) < \alpha] \rightarrow 1$ when $n \rightarrow \infty$.*

The proof of this result is also given in Section 3. In this form, it is much simpler to compare the impossibility result to the one that features in [12] for almost exact recovery. As a reminder, this is the closest result to ours in the literature, and it states that almost exact recovery is impossible when $nqs = O(1)$. Thus, when the conditions on q are met, our result is more restrictive than in [12], but the conclusion we get is correspondingly stronger. In [12], one cannot hope to obtain a permutation that agrees with π^* on all but a vanishing fraction of vertices; in our case, we also rule out the possibility of obtaining a permutation that agrees with π^* on more than a fraction α of vertices.

2.4.2 Possibility result

The possibility result is given next.

Theorem 2.3. *Let $0 < \alpha < 1$. If, there exist $\beta > 0$, $\gamma > 0$ such that for large enough n ,*

$$\begin{aligned} nqs &\geq \max \left\{ 20, 84 \cdot \log \left(\frac{2}{1-\alpha} \right), \frac{16}{\min(\gamma, \beta)(1-\alpha)} \right\} \\ s &> \frac{8}{1-\alpha} \cdot q \\ \frac{2q(1-s^2)}{s} &\leq n^{-\beta} \\ qs &\leq n^{-2\gamma} \end{aligned} \quad (6)$$

then there exists an (exponential-time) algorithm which returns a permutation $\hat{\pi}$ such that

$$\mathbb{P}[\beta(\hat{\pi}, \pi^*) \geq \alpha] \rightarrow 1 \text{ when } n \rightarrow \infty.$$

Crucially, our assumptions cover the case where $nqs = \Theta(1)$ as the case where $nqs = \omega(1)$ is already covered by [12]. It is shown there that when $nqs = \omega(1)$, almost exact recovery is guaranteed to happen.

The complete proof of this result is given in Section 4. The algorithm we propose stems from the idea that if we were able to select π^* correctly among all permutations, then the intersection graph we would observe between G_A and G_B permuted by π^* would be an Erdős-Rényi graph of parameters (n, qs) . Our choice of a permutation π is then such that the intersection graph obtained with π has some commonalities with an $ER(n, qs)$ graph. In particular, in our regime of interest ($nqs = \Theta(1)$), it is known that if nqs is above some constant threshold, a positive fraction of the nodes in $ER(n, qs)$ (with this fraction increasing in nqs) has degree greater than or equal to $nqs/2$. As we are only interested in returning a permutation π that overlaps with π^* on a fraction α of the nodes, our algorithm will output a permutation which gives rise to an intersection graph where only a little more than a fraction α of the nodes have degree greater than or equal to $nqs/2$.

In spirit, our algorithm has many commonalities with the algorithm given in [12]—it was in fact inspired by the authors’ results. The main difference is of course the regime under consideration. In their case, they consider $nqs = \omega(1)$ and show that almost exact recovery can be achieved. In our case, we show that in the regime $nqs = \Theta(1)$, we can achieve partial recovery. Combined with [11] which discusses exact recovery, this leads to a nearly complete picture of which type of recovery is possible in which regime. In terms of proof techniques, as our algorithm shares many commonalities with the algorithm given in [12], the key ideas behind our proof do overlap considerably with those in [12]. There are two main differences. First, the change of regime requires a much finer analysis to show that the algorithm that we propose terminates. Second, we propose an alternative approach to analyzing the moment generating function of the sum of correlated random variables $\sum_{i \neq j} A_{ij} B_{\pi(i)\pi(j)}$, which eschews the many sets, graph constructions, and functions introduced in [12]. Rather, it relies on elementary algebra and recurrent sequences, which we would argue are easier to apprehend and use. In this spirit, we are interested to note that some of the computational techniques developed in this paper have come in useful in more recent publications. For example, Lemma B.1 and an analog to Lemma B.2 from this paper are leveraged by [51] in their proof of the positive part of their Theorem 4 establishing the sharp threshold for exact recovery.

The conditions in Theorem 2.3 merit a brief discussion. The first condition is the key condition here as it confirms that partial recovery is feasible in the $nqs = \Theta(1)$ regime. It should be noted though that, as our goal was to establish this latter point, we did not make a particular effort to drive the constant that appears in this condition down to its smallest value. This could be the subject of future work. The second condition amounts to requiring a stronger correlation between pairs (A_{ij}, B'_{ij}) than being simply positive: it is automatically satisfied if $q = o(s)$. The third and fourth conditions are sparsity conditions and it can be observed that there is a trade-off between how sparse the graph is and how large nqs should be. It may be that these last conditions are artefacts of the analysis rather than the algorithm in itself: in other words, it may be possible to remove them with a more nuanced analysis.

2.5 Discussion and future directions

Comparison of the possibility and impossibility results in the partial alignment case. We contrast the necessary and sufficient conditions obtained in this paper. To do so, we make use

of Corollary 2.2 and Theorem 2.3, which are easier to compare. We remark that in the case where $n^{-k} \leq q \ll n^{-1/2}$ for some $k > 0$, and conditions $s > \frac{8}{1-\alpha}q$ and $2q\frac{1-s^2}{s} \leq n^{-\beta}$ for some $\beta > 0$ hold, the remaining conditions (nqs larger than a constant and $nqs = o(1)$) nearly partition the space.

Relevance of the results in practice. As noted above, partial recovery is a weaker form of recovery than exact or almost exact recovery, as it returns an alignment which only partially coincides with the true alignment. The counterpart to this is that it can be achieved under weaker conditions, namely when $nqs > C$ where C is some constant (together with some other more technical constraints as given in Theorem 2.3). Two particular regimes, which are of possible interest in applications, are covered by these weaker conditions. The first one is the case where the graphs G_A and G_B are sparse (e.g., $q \sim 1/n$ and $s = \Theta(1)$). Sparse graphs are quite prevalent in real-world applications, particularly as models of social networks and recommender databases, and thus our results can be used to tackle problems on these graphs, such as e.g. the problem of anonymizing/de-anonymizing these structures. The second case corresponds to a setting where the graphs are denser (e.g., $q \sim \frac{\log(n)}{n}$) but we wish to allow for more noise (e.g. $1-s \sim 1 - \frac{1}{\log(n)}$). This can be relevant, for example, when aligning Protein-Protein Interaction graphs for species whose most recent common ancestor was more distant.

Future research directions. As discussed above, there is a gap between the upper and lower bounds given in this paper. Obtaining matching upper and lower bounds would thus be of great future interest, as would be a polynomial-time algorithm which would recover a partial alignment under similar assumptions to the exponential one in this paper. Another interesting research direction would be to study how one can leverage m correlated Erdős-Renyi graphs, rather than 2, to obtain a permutation π which overlaps with π^* on a larger fraction of nodes than what is currently guaranteed under the conditions of Theorem 2.3. Finally, the techniques developed in this paper to analyze sums of correlated random variables may be more widely applicable.

3 Impossibility result

In this section, we give the proof of Theorem 2.1. We start off with some preliminaries in Section 3.1 before moving onto the core of the proof in Section 3.2.

3.1 Preliminaries: Fano’s generalized inequality and permutation results

The proof of the impossibility result relies on a generalization of Fano’s inequality which appears in e.g. [41]. The formulation we use here comes from [3]. This generalization differs from the classical Fano’s inequality in the sense that it uses a wider definition of what a successful learning algorithm constitutes. In the classical Fano’s inequality, an algorithm is unsuccessful if there is one difference or more between the ground truth and the algorithm output. In its generalization, we allow for some amount of differences before classifying the algorithm as unsuccessful.

It should be noted that the presentation of the theorem in [3] is tailored quite specifically to a privatization setting. We rewrite the theorem here in fuller generality. With the assumptions we take, the proof of the theorem that is given in [3] goes through exactly and hence we do not repeat it here.

Theorem 3.1 (Lemma 20 in [3]). *Let \mathcal{H} be a hypothesis class of cardinality M with Z being a hypothesis drawn uniformly at random from \mathcal{H} . Let \hat{X} and $\hat{Z} \in \mathcal{H}$ be random variables defined in such a way that $Z \rightarrow \hat{X} \rightarrow \hat{Z}$ constitutes a Markov Chain. We define $d : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}_+$ to be a*

distance function, and, for a given $d > 0$, we say that a learner is successful if $d(Z, \hat{Z}) < d$. For any $h \in \mathcal{H}$, we further define $B_d(h) = \{h' \in \mathcal{H} \mid d(h, h') \leq d\}$ and we let $M_d = \max_{h \in \mathcal{H}} |B_d(h)|$. Then, we have

$$\mathbb{P}[d(\hat{Z}, Z) > d] \geq 1 - \frac{I(Z; \hat{X}) + 1}{\log(M/M_d)},$$

where $I(Z; \hat{X})$ is the mutual information between the pair Z and \hat{X} . In other words, for any learning algorithm, the average error probability will be greater than $1 - \frac{I(Z; \hat{X}) + 1}{\log(M/M_d)}$.

It is quite straightforward to transpose this to our setting, noting that $\pi^* \in \mathcal{P}_n \rightarrow (A, B) \rightarrow \hat{\pi}$ is a Markov chain and that one can define a distance metric over $P_n \times P_n$ by taking $d(\pi, \pi^*) = 1 - \beta(\pi, \pi^*)$. Given $\alpha > 0$, we then say that a learner is successful if $\beta(\pi, \pi^*) > \alpha$. We thus obtain the following corollary.

Corollary 3.2. *Let π^* be a permutation drawn uniformly from P_n and let G_A and G_B be two graphs generated using π^* as described in Section 2.1, with adjacency matrices A and B . For any learning algorithm (i.e., any algorithm which infers from A and B a permutation $\hat{\pi}$), we have that*

$$\mathbb{P}(\beta(\pi, \pi^*) < \alpha) \geq 1 - \frac{I(\pi^*; (A, B)) + 1}{\log(M/M_\alpha)},$$

where $M = |P_n|$ and $M_\alpha = \max_{\pi^* \in P_n} |\{\pi \mid \beta(\pi, \pi^*) \geq \alpha\}|$.

This corollary is the cornerstone of our proof. As can be seen from its statement, it involves the ratio M/M_α . Before computing this quantity in Proposition 3.3, we first remind the reader of a few properties of permutations, which will come in useful in the proof of Proposition 3.3.

Recall that the set of permutations over $\{1, \dots, n\}$, denoted by P_n , is a group when equipped with the composition operation \circ . Its size is $n!$. As P_n is a group, it follows that if π_1 and π_2 are in P_n , then $\pi_1 \circ \pi_2$ is again in P_n . Likewise, each permutation π has an inverse permutation π^{-1} which is such that the application of $\pi^{-1} \circ \pi$ (or $\pi \circ \pi^{-1}$) to $\{1, \dots, n\}$ leaves it unchanged. A notion that will be key in Proposition 3.3 is that of *rencontres numbers*. These are the number of permutations over n elements that have k fixed points, i.e., k integers which are mapped to themselves. We refer to the rencontres numbers as $D_{n,k}$ and we use the fact that

$$D_{n,k} \sim e^{-1} \cdot \frac{n!}{k!}.$$

We are now able to show Proposition 3.3.

Proposition 3.3. *Let M and M_α be defined as in Corollary 3.2. We have that: $\frac{M}{M_\alpha} = \Omega((n\alpha)!)$.*

Proof. Proof. As π^* and π are in P_n , then $\tilde{p} := \pi^* \circ \pi$ is again a permutation and the overlap between π^* and π can be characterized via the fixed points of \tilde{p} :

$$\beta(\pi, \pi^*) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\pi(i)=\pi^*(i)} = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\tilde{p}(i)=i}, \quad (7)$$

i.e., the overlap between π^* and π is simply the number of fixed points of \tilde{p} (divided by n). From this it follows that

$$|\{\pi \in P_n \mid \beta(\pi, \pi^*) \geq \alpha\}| = |\{\tilde{p} \in P_n \mid \sum_{i=1}^n \mathbf{1}_{\tilde{p}(i)=i} \geq \lceil n\alpha \rceil\}| = \sum_{k=\lceil n\alpha \rceil}^n D_{n,k}.$$

Note that this value does not change with π^* , and so

$$M_\alpha = \sum_{k=\lceil n\alpha \rceil}^n D_{n,k} \sim \sum_{k=\lceil n\alpha \rceil}^n e^{-1} \frac{n!}{k!}.$$

Using Taylor's theorem on e^1 with e.g. the mean-value form of the remainder, we get that

$$\sum_{k=\lceil n\alpha \rceil}^n \frac{1}{k!} = O\left(\frac{1}{\lceil n\alpha \rceil!}\right) = O\left(\frac{1}{(n\alpha)!}\right)$$

which implies that $M_\alpha = O(n!/(n\alpha)!)$ and that $\frac{M}{M_\alpha} = \Omega((n\alpha)!)$. \square

3.2 Proof of Theorem 2.1

We now give the proof of Theorem 2.1. At a high level, it involves tightly bounding each quantity that appears in Corollary 3.2.

Proof. Proof of Theorem 2.1. We use Corollary 3.2 to find conditions on n, q , and s under which $\mathbb{P}(\beta(\pi, \pi^*) < \alpha) = 1 - o(1)$. From Proposition 3.3, we have that $\frac{M}{M_\alpha} = \Omega((n\alpha)!)$ and so, using Stirling's formula, we deduce that $\log(M/M_\alpha) = \Omega(n\alpha \log(n\alpha))$. It now remains to compute an upper bound on $I(\pi^*, (A, B))$ to conclude. Recall that π^* is selected uniformly at random and that G_A and G_B are two graphs generated via the correlated Erdős-Rényi model (n, q, s) . Let $\tilde{\pi} \in P_n$ and \tilde{A}, \tilde{B} be the adjacency matrices of two undirected graphs on n nodes. Following standard information inequalities (see, e.g., [9]), we have:

$$I(\pi^*; (A, B)) = D\left(\mathbb{P}(\pi^* = \tilde{\pi}, (A, B) = (\tilde{A}, \tilde{B})) \parallel \mathbb{P}(\pi^* = \tilde{\pi}) \cdot \mathbb{P}\left((A, B) = (\tilde{A}, \tilde{B})\right)\right).$$

As $\mathbb{P}\left(\pi^* = \tilde{\pi}, (A, B) = (\tilde{A}, \tilde{B})\right) = \mathbb{P}\left((A, B) = (\tilde{A}, \tilde{B}) \mid \pi^* = \tilde{\pi}\right) \cdot \mathbb{P}(\pi^* = \tilde{\pi})$, it follows that:

$$I(\pi^*; (A, B)) = \sum_{\tilde{\pi}} \mathbb{P}(\pi^* = \tilde{\pi}) \cdot \sum_{(\tilde{A}, \tilde{B})} \mathbb{P}\left((A, B) = (\tilde{A}, \tilde{B}) \mid \pi^* = \tilde{\pi}\right) \log\left(\frac{\mathbb{P}\left((A, B) = (\tilde{A}, \tilde{B}) \mid \pi^* = \tilde{\pi}\right)}{\mathbb{P}\left((A, B) = (\tilde{A}, \tilde{B})\right)}\right)$$

Letting \mathcal{P} be the joint distribution of (A, B) under the correlated Erdős-Rényi model (n, q, s) and \mathcal{Q} be the joint distribution of adjacency matrices (A, B) of two independent Erdős-Rényi graphs, G_A and G_B , of parameters (n, q) , we write:

$$\begin{aligned} I(\pi^*; (A, B)) &= \mathbb{E}_{\pi^*} \left[\sum_{(\tilde{A}, \tilde{B})} \mathbb{P}\left((A, B) = (\tilde{A}, \tilde{B}) \mid \pi^* = \tilde{\pi}\right) \cdot \log\left(\frac{\mathbb{P}\left((A, B) = (\tilde{A}, \tilde{B}) \mid \pi^* = \tilde{\pi}\right)}{\mathcal{Q}(\tilde{A}, \tilde{B})} \cdot \frac{\mathcal{Q}(\tilde{A}, \tilde{B})}{\mathcal{P}(\tilde{A}, \tilde{B})}\right) \right] \\ &= \mathbb{E}_{\pi^*} \left[\sum_{(\tilde{A}, \tilde{B})} \mathbb{P}\left((A, B) = (\tilde{A}, \tilde{B}) \mid \pi^* = \tilde{\pi}\right) \cdot \log\left(\frac{\mathbb{P}\left((A, B) = (\tilde{A}, \tilde{B}) \mid \pi^* = \tilde{\pi}\right)}{\mathcal{Q}(\tilde{A}, \tilde{B})}\right) \right] \\ &\quad - \sum_{(\tilde{A}, \tilde{B})} \mathcal{P}(\tilde{A}, \tilde{B}) \cdot \log\left(\frac{\mathcal{P}(\tilde{A}, \tilde{B})}{\mathcal{Q}(\tilde{A}, \tilde{B})}\right) \\ &\leq \mathbb{E}_{\pi^*} \left[D\left(\mathbb{P}\left((A, B) = (\tilde{A}, \tilde{B}) \mid \pi^* = \tilde{\pi}\right) \parallel \mathcal{Q}(\tilde{A}, \tilde{B})\right) \right], \end{aligned} \tag{8}$$

where we have used the total law of probability in the second equality and the fact that

$$\sum_{(\tilde{A}, \tilde{B})} \mathcal{P}(\tilde{A}, \tilde{B}) \cdot \log \left(\frac{\mathcal{P}(\tilde{A}, \tilde{B})}{\mathcal{Q}(\tilde{A}, \tilde{B})} \right) = D(\mathcal{P}||\mathcal{Q}) \geq 0$$

in the inequality. Now, note that

$$\begin{aligned} \mathbb{P} \left((A, B) = (\tilde{A}, \tilde{B}) \mid \pi^* = \tilde{\pi} \right) &= \mathbb{P} \left((A_{12}, B_{\tilde{\pi}(1)\tilde{\pi}(2)}) = (\tilde{A}_{12}, \tilde{B}_{\tilde{\pi}(1)\tilde{\pi}(2)}), \dots \mid \pi^* = \tilde{\pi} \right) \\ &= P \left(\tilde{A}_{12}, \tilde{B}_{\tilde{\pi}(1)\tilde{\pi}(2)} \right) \dots P \left(\tilde{A}_{\tilde{\pi}(n-1)\tilde{\pi}(n)}, \tilde{B}'_{\tilde{\pi}(n-1)\tilde{\pi}(n)} \right) \end{aligned}$$

and $\mathcal{Q}(\tilde{A}, \tilde{B}) = Q(\tilde{A}_{12}, \tilde{B}_{12}) \dots Q(\tilde{A}_{(n-1)n}, \tilde{B}_{(n-1)n})$. Using these results in (8), we obtain:

$$I(\pi^*; (A, B)) \leq \binom{n}{2} D(P||Q),$$

from whence, $\mathbb{P}(\beta(\pi, \pi^*) < \alpha) \geq 1 - \Omega \left(\frac{\binom{n}{2} D(P||Q) + 1}{n\alpha \log(n\alpha)} \right)$. Hence, if

$$\frac{D(P||Q)}{\alpha} = o \left(\frac{\log n}{n} \right),$$

it follows that $\mathbb{P}(\beta(\pi, \pi^*) < \alpha) = 1 - o(1)$. □

3.3 Proof of Corollary 2.2

We provide a short proof of Corollary 2.2 in this section.

Proof. Proof of Corollary 2.2. We make use of the following reverse Pinsker inequality (see, e.g., [42, Section 6.C]): for two probability measures P and Q defined on a common finite set \mathcal{A} , and Q is strictly positive on \mathcal{A} , which has more than one element, we have

$$D(P||Q) \leq \log \left(\frac{1}{Q_{min}} \right) \cdot |P - Q|,$$

where $|\cdot|$ is the total variation between P and Q , and $Q_{min} = \min_{a \in \mathcal{A}} Q(a) \leq \frac{1}{2}$. In our setting, we have that $Q_{min} = q^2 \leq \frac{1}{2}$ for large enough n , and $|P - Q| = q(s - q)$. We now show that, under the conditions given in (5),

$$\frac{D(P||Q)}{\alpha} = o \left(\frac{\log(n)}{n} \right),$$

which concludes the proof. We have

$$\frac{D(P||Q)}{\alpha} \cdot \frac{n}{\log(n)} \leq \log \left(\frac{1}{Q_{min}} \right) \cdot |P - Q| \cdot \frac{n}{\alpha \log(n)} = \log \left(\frac{1}{q^2} \right) \cdot \frac{nqs - nq^2}{\alpha \log(n)}.$$

It is easy to see using (5) that $\log \left(\frac{1}{q^2} \right) = O(\log(n))$, that $nqs = o(1)$, and that $nq^2 = o(1)$. Thus,

$$\frac{D(P||Q)}{\alpha} \cdot \frac{n}{\log(n)} \leq O(\log(n)) \cdot \frac{o(1)}{\alpha \log(n)} = o(1).$$

□

4 Possibility result

In this section, we give a proof of Theorem 2.3. To do this, we need to provide an algorithm that outputs a permutation π which overlaps with π^* over at least $n\alpha$ vertices whp. In our case, the algorithm involves enumerating each permutation and checking whether it is *good* (see Definition 2). We give the proof that this algorithm terminates whp (i.e., that there exists a permutation that is good whp) and that any good permutation must overlap with π^* on at least $n\alpha$ vertices whp in Section 4.2. Before getting to that however, we briefly discuss, in Section 4.1, some properties of permutations and k -cores of graphs which will come in useful in the proofs of Section 4.2.

4.1 Preliminaries: k -cores and more permutation results

4.1.1 Definition and size of a k -core

A k -core of a graph G is the largest induced subgraph of G with minimum degree at least k . (Note that it may be empty.) The k -core is unique and can simply be obtained by iteratively removing all nodes of degree less than k . It is easy to see from this definition that if a graph has a k -core of size N (i.e., the k -core contains N nodes) then G has at least N nodes of degree greater than or equal to k . Conditions under which a k -core in an Erdős-Rényi(n, p) graph is guaranteed to be non-empty are known: in fact, if $k \geq 3$ and the k -core is non-empty, then it contains a positive fraction of the vertices of the graph. We give the exact statement of this result here as it will be useful to us later on.

Lemma 4.1 (Theorem 2 in [37]). *Let $\mu > 0$ and let $Po(\mu)$ denote a Poisson random variable of parameter μ . We define*

$$\psi_j(\mu) = \mathbb{P}(Po(\mu) \geq j) = \sum_{i=j}^{\infty} \frac{\mu^i e^{-\mu}}{i!}, \text{ for an integer } j \geq 0.$$

We further let $k \geq 3$ and take

$$c_k := \inf_{\mu > 0} \frac{\mu}{\psi_{k-1}(\mu)}.$$

For $\lambda \geq c_k$, define $\mu_k(\lambda) > 0$ to be the largest solution to $\frac{\mu}{\psi_{k-1}(\mu)} = \lambda$. Now, consider the random graph $G(n, \lambda/n)$ where λ is fixed. If $\lambda > c_k$, then whp the k -core of G contains a fraction $\psi_k(\mu_k(\lambda))$ of the nodes.

4.1.2 Additional properties of permutations.

As done in Section 3, we consider in this section the permutation $p := \pi \circ \pi^{*-1}$ and its corresponding permutation matrix $P := \Pi \cdot \Pi^{*T}$. Using this notation, we have that $B_{\pi(i)\pi(j)} = B'_{p(i)p(j)}$. Of interest to us will be two sets related to the fixed points of p .

Definition 1. Let π^* be a permutation in P_n and let π be a permutation such that $\beta(\pi, \pi^*) = \epsilon$. For $p = \pi \circ \pi^{*-1}$, define the sets:

$$S := \{(i, j) \in \{1, \dots, n\}^2 \mid i \neq j\}, \quad S_1 := \{(i, j) \in S \mid i = p(i)\}, \quad \text{and } S_2 := S \setminus S_1 \quad (9)$$

and let $|S|$, $|S_1|$ and $|S_2|$ denote their respective cardinalities. It is immediate to see that

$$|S| = n(n-1), \quad |S_1| = n\epsilon(n-1), \quad |S_2| = n(1-\epsilon)(n-1).$$

Note that from these definitions, if $(i, j) \in S_2$, then it cannot be the case that $(i, j) = (p(i), p(j))$; however, S_2 does contain all pairs (i, j) such that $(i, j) = (p(j), p(i))$. This leads us to further subdivide the set S_2 : let

$$S_2^1 = \{(i, j) \in S_2 \mid (p(i), p(j)) = (j, i)\} \text{ and } S_2^2 = S_1 \setminus S_2^1.$$

We also introduce a set \bar{S}_2^2 , containing *unordered pairs*,

$$\bar{S}_2^2 := \{\{i, j\} \in \{1, \dots, n\}^2 \mid i \neq j, \{i, j\} \neq \{p(i), p(j)\}\},$$

and related to S_2^2 in the sense that $S_2^2 \subseteq \{(i, j) \in S \mid (i, j) \neq (p(i), p(j)), (i, j) \neq (p(j), p(i))\}$, which is the ordered counterpart of \bar{S}_2^2 . It is easy to see that

$$|S_2^1| \leq n(1 - \epsilon) \text{ and } |\bar{S}_2^2| \leq \binom{n}{2} - \binom{n\epsilon}{2} \leq \frac{n^2}{2}(1 - \epsilon^2). \quad (10)$$

We now define the concept of a *cycle*, which will play an important role in Section 4.2. A cycle C is a set of unordered pairs $\{\{i, j\}\}$ such that if $\{i, j\} \in C$, for any other element $\{k, l\} \in C$, there exists $s \in \mathbb{N}$ such that $\{k, l\} = \{p^s(i), p^s(j)\}$. It is not hard to see that the cycle relationship between $\{i, j\}$ and $\{k, l\} \in \bar{S}_2^2$ defines an equivalence relationship over \bar{S}_2^2 . Thus, \bar{S}_2^2 can be partitioned into cycles. It naturally follows that a pair $\{i, j\}$ appears in at most one cycle and only once in that cycle. Furthermore, by definition of \bar{S}_2^2 , all cycles in \bar{S}_2^2 must be of length greater or equal to 2. With this in mind, we use the notation \mathcal{C}_k to denote the set of cycles in \bar{S}_2^2 of size k and ℓ_k to denote the number of cycles of size k belonging to \bar{S}_2^2 . We have $2 \leq k \leq |\bar{S}_2^2|$,

$$\bar{S}_2^2 = \bigcup_{k=2}^{|\bar{S}_2^2|} \mathcal{C}_k \text{ and } |\bar{S}_2^2| = \sum_{k=1}^{|\bar{S}_2^2|} k \cdot \ell_k. \quad (11)$$

4.2 Proof of Theorem 2.3

We prove Theorem 2.3 in two steps. First, we define the notion of a *good permutation* and we show that under the conditions of Theorem 2.3, whp π^* is a good permutation (Proposition 4.2). The algorithm then involves enumerating all permutations until we obtain a good permutation: this will happen with high probability as one at least, π^* , exists. Secondly, we show that, under the conditions of Theorem 2.3, any good permutation returned is guaranteed to overlap with π^* on a fraction α of its nodes (Proposition 4.4). These two propositions together imply the conclusions of Theorem 2.3.

Definition 2. Let $0 < \alpha < 1$. A permutation π is said to be *good* if

$$\sum_{i=1}^n \mathbf{1}_{\{\sum_{j \neq i} A_{ij} B_{\pi(i)\pi(j)} \geq \frac{n\alpha}{2}\}} \geq n \cdot \frac{1 + \alpha}{2},$$

where $\mathbf{1}$ denotes the indicator function.

In other words, if we look at the intersection graph of G_A and G_B to which π has been applied, there are at least $n \cdot \frac{1 + \alpha}{2}$ nodes that have degree greater than or equal to $\frac{n\alpha}{2}$.

Proposition 4.2. *For any (A, B) generated as described in Section 2 and under the conditions given in Theorem 2.3, π^* is whp a good permutation.*

The proof of this proposition is given in the Electronic Companion to this paper. The key idea is that the intersection graph of G_A and G'_B is Erdős-Rényi of parameters (n, qs) . Hence, one can use the k -core results given in Lemma 4.1. In particular, the proof in the Electronic Companion involves showing that for any $0 < \alpha < 1$, and under the conditions of Theorem 2.3, we have that $nqs > c_{nqs/2}$ and that $\psi_{nqs/2}(\mu_{nqs/2}(nqs)) \geq \frac{1+\alpha}{2}$. From Lemma 4.1, this implies that there exists with high probability in this intersection graph an $nqs/2$ -core of size greater than or equal to $n \cdot \frac{1+\alpha}{2}$, which in turn implies that π^* is good. Proposition 4.2 serves to show that the algorithm will terminate and output a good permutation with high probability.

We now move onto Proposition 4.4. Note that the proofs of some technical lemmas used to show Proposition 4.4 have been deferred to the Electronic Companion for clarity. We also make use of the following result that appears in [12], the proof of which we also include in the Electronic Companion (Section C) for completeness.

Lemma 4.3 (Lemma 3 [12]). *Let $\tau \geq 0$, $q_1 \geq 0$, and $q_2 \geq 0$. Then*

$$\min_{z \geq 0} z^{-\tau} e^{q_2(z^2-1)+q_1(z-1)} \leq \zeta^\tau \quad (12)$$

where $\zeta := \max\{\sqrt{2}e\frac{q_1}{\tau}, 4e\left(\frac{q_2}{\tau}\right)^{1/2}\}$. Furthermore, if z^* is the minimizer of this expression,

$$z^{*2} = \frac{\tau - q_1 z^*}{2q_2} \leq \frac{\tau}{2q_2}. \quad (13)$$

Proposition 4.4. *Let $0 < \alpha < 1$, $\beta > 0$, $\gamma > 0$. If the conditions of Theorem 2.3 hold, then a good permutation π agrees whp with the ground truth π^* on at least a fraction α of the nodes in the graph. In other words,*

$$\mathbb{P}(\forall \pi, [\pi \text{ good} \Rightarrow \beta(\pi, \pi^*) > \alpha]) \rightarrow 1 \text{ when } n \rightarrow \infty. \quad (14)$$

Proof. Proof. We show that when conditions (6) hold, we have $\lim_{n \rightarrow \infty} \mathbb{P}(\exists \pi : \pi \text{ good} \cap \beta(\pi, \pi^*) \leq \alpha) \rightarrow 0$, which implies (14). Note that

$$\mathbb{P}(\exists \pi : \pi \text{ good} \cap \beta(\pi, \pi^*) \leq \alpha) = \mathbb{P}(\cup_{\{\pi: \beta(\pi, \pi^*) \leq \alpha\}} \{\pi : \pi \text{ good}\}) \leq \sum_{\{\pi: \beta(\pi, \pi^*) \leq \alpha\}} \mathbb{P}(\pi \text{ good}). \quad (15)$$

Let π be such that $\beta(\pi, \pi^*) = \epsilon$ where $\epsilon \leq \alpha$. Our goal is to upper bound $\mathbb{P}(\pi \text{ good})$. To do this, we make use of the following fact: if π is good, then among all $n(1 - \epsilon)$ nodes that π and π^* do not overlap on, at least $\frac{n(1+\alpha)}{2} - n\epsilon = \frac{n(1+\alpha-2\epsilon)}{2} \geq \frac{n(1-\alpha)}{2}$ nodes must have degree greater than or equal to $\frac{nqs}{2}$, i.e., $\sum_{j \neq i} A_{ij} B_{\pi(i)\pi(j)} \geq \frac{nqs}{2}$ for a set of indices i in $\{i \mid \pi(i) \neq \pi^*(i)\}$ of cardinality greater than or equal to $n \cdot \frac{1-\alpha}{2}$. This further implies that

$$\sum_{\{i \mid \pi(i) \neq \pi^*(i)\}} \sum_{j \neq i} A_{ij} B_{\pi(i)\pi(j)} = \sum_{(i,j) \in S_2} A_{ij} B'_{p(i)p(j)} \geq \frac{nqs \cdot n(1-\alpha)/2}{2},$$

where $p = \pi \circ \pi^{*-1}$ and S_2 is the set given in Section 4.1.2. This implication enables us to upper bound $\mathbb{P}(\pi \text{ good})$ in the following fashion:

$$\mathbb{P}(\pi \text{ good}) \leq \mathbb{P}\left(\sum_{(i,j) \in S_2} A_{ij} B'_{p(i)p(j)} \geq \frac{nqs \cdot n(1-\alpha)/2}{2}\right).$$

Recalling that $S_2 = S_2^1 \cup S_2^2$, we have that

$$\mathbb{P}(\pi \text{ good}) \leq \mathbb{P} \left(\sum_{(i,j) \in S_2^1} A_{ij} B'_{p(i)p(j)} + \sum_{(i,j) \in S_2^2} A_{ij} B'_{p(i)p(j)} \geq \frac{nqs \cdot n(1-\alpha)}{4} \right).$$

As it turns out, the variables $\{A_{ij} B'_{p(i)p(j)}\}_{(i,j) \in S_2^2}$ are independent from $\{A_{ij} B'_{p(i)p(j)}\}_{(i,j) \in S_2^1}$. This gives rise to the following classical Chernoff inequality:

$$\mathbb{P}(\pi \text{ good}) \leq \inf_{t>0} e^{-tu} \mathbb{E} \left[e^{t \sum_{(i,j) \in S_2^1} A_{ij} B'_{p(i)p(j)}} \right] \cdot \mathbb{E} \left[e^{t \sum_{(i,j) \in S_2^2} A_{ij} B'_{p(i)p(j)}} \right], \quad (16)$$

where $u := \frac{nqs \cdot n(1-\alpha)}{4}$. It is easy to show that for large enough n , $u - \mathbb{E}[\sum_{(i,j) \in S_2} A_{ij} B'_{p(i)p(j)}] > 0$. Indeed, using (10) and the assumptions $s > \frac{8}{1-\alpha} \cdot q$ and $n(1-\epsilon)qs = o(nqs \cdot n(1-\alpha)/4)$ from Theorem 2.3, we have

$$\begin{aligned} u - \mathbb{E} \left[\sum_{(i,j) \in S_2} A_{ij} B'_{p(i)p(j)} \right] &= \frac{nqs \cdot n(1-\alpha)}{4} - \sum_{(i,j) \in S_2^1} \mathbb{E}[A_{ij} B'_{p(i)p(j)}] - \sum_{(i,j) \in S_2^2} \mathbb{E}[A_{ij} B'_{p(i)p(j)}] \\ &\geq \frac{nqs \cdot n(1-\alpha)}{4} - n(1-\epsilon)qs - n(1-\epsilon)(n-1)q^2 > 0 \end{aligned}$$

for large enough n . It remains to compute the two moment generating functions that appear in (16). It is quite straightforward to see that, for $t > 0$,

$$\begin{aligned} \mathbb{E} \left[e^{t \sum_{(i,j) \in S_2^1} A_{ij} B'_{p(i)p(j)}} \right] &= \mathbb{E} \left[e^{t \sum_{(i,j) \in S_2^1} A_{ij} B'_{ji}} \right] \leq \mathbb{E} \left[e^{t(A_{ij} B'_{ji} + A_{ji} B'_{ij})} \right]^{\frac{|S_2^1|}{2}} \leq \mathbb{E} \left[e^{2t(A_{ij} B'_{ji} + A_{ji} B'_{ij})} \right]^{\frac{|S_2^1|}{4}} \\ &= (p_{11} e^{4t} + 1 - p_{11})^{\frac{|S_2^1|}{4}} \\ &= e^{\frac{|S_2^1|}{4} \cdot \log(1+p_{11}(e^{4t}-1))}. \end{aligned}$$

We now focus on $\mathbb{E} \left[e^{t \sum_{(i,j) \in S_2^2} A_{ij} B'_{p(i)p(j)}} \right]$: this expression is not as straightforward to analyze as the variables involved all depend on one another in a more complicated way. A first step in simplifying this analysis is to use \bar{S}_2^2 rather than S_2^2 as it enables us to use the concept of *cycles* introduced in Section 4.1.2. As the matrices A and B' are symmetric and nonnegative, and $t > 0$, and in light of the definition of \bar{S}_2^2 , we have:

$$t \sum_{(i,j) \in S_2^2} A_{ij} B'_{p(i)p(j)} \leq 2t \sum_{\{i,j\} \in \bar{S}_2^2} A_{ij} B'_{p(i)p(j)}.$$

Together with (11), this implies

$$\mathbb{E} \left[e^{t \sum_{(i,j) \in S_2^2} A_{ij} B'_{p(i)p(j)}} \right] \leq \mathbb{E} \left[e^{2t \sum_{\{i,j\} \in \bar{S}_2^2} A_{ij} B'_{p(i)p(j)}} \right] \leq \mathbb{E} \left[e^{2t \sum_{k=2}^{|\bar{S}_2^2|} \sum_{C \in \mathcal{C}_k} \sum_{\{i,j\} \in C} A_{ij} B'_{p(i)p(j)}} \right].$$

Given $\{i,j\} \in \bar{S}_2^2$, $A_{\{i,j\}} = A_{ij} = A_{ji}$ and $B'_{\{i,j\}} = B'_{ij} = B'_{ji}$ will appear exactly once and their appearance will be in the *same* cycle. As $A_{\{i,j\}}$ and $B'_{\{i,j\}}$ are independent of all other random variables, it follows that the cycles are all independent of one another. Thus, we obtain:

$$\mathbb{E} \left[e^{t \sum_{(i,j) \in S_2^2} A_{ij} B'_{p(i)p(j)}} \right] \leq \prod_{k=2}^{|\bar{S}_2^2|} \prod_{C \in \mathcal{C}_k} \mathbb{E} \left[e^{2t \sum_{\{i,j\} \in C} A_{ij} B'_{p(i)p(j)}} \right]. \quad (17)$$

The difficulty then becomes computing the moment generating function

$$\mathbb{E} \left[e^{2t \sum_{\{i,j\} \in C} A_{ij} B'_{p(i)p(j)}} \right].$$

As it turns out, this can be done by noting the special structure of the sum: the pair $(A_{\{i,j\}}, B'_{\{i,j\}})$ only appears once, at the extremity of the sum, similarly for $(A_{\{p(i),p(j)\}}, B'_{\{p(i),p(j)\}})$, though they are once removed from the extremity of the sum, and so on and so forth. By conditioning recursively on these pairs, one can obtain the moment generating function of interest as a function of four sequences defined via recursive relationships. Solving each of the four recursions then gives us the moment generating function, which only depends on the length of the cycle C . We omit the details of this proof in the main text: it can be found in Section B of the Electronic Companion. We give here only the end result, namely that for a cycle C of length k ,

$$\mathbb{E} \left[e^{2t \sum_{\{i,j\} \in C} A_{ij} B'_{p(i)p(j)}} \right] = \left(\frac{T + \sqrt{T^2 - 4D}}{2} \right)^k + \left(\frac{T - \sqrt{T^2 - 4D}}{2} \right)^k$$

where $T = p_{11}(e^{2t} - 1) + 1$ and $D = \sigma^2(e^{2t} - 1)$. Using this result in (17), we obtain:

$$\mathbb{E} \left[e^{t \sum_{(i,j) \in S_2^2} A_{ij} B'_{p(i)p(j)}} \right] \leq \prod_{k=2}^{|\bar{S}_2^2|} \left(\left(\frac{T + \sqrt{T^2 - 4D}}{2} \right)^k + \left(\frac{T - \sqrt{T^2 - 4D}}{2} \right)^k \right)^{\ell_k}.$$

We then leverage Lemma B.1 from the Electric Companion to further upper bound this quantity:

$$\begin{aligned} \mathbb{E} \left[e^{t \sum_{(i,j) \in S_2^2} A_{ij} B'_{p(i)p(j)}} \right] &\leq \prod_{k=2}^{|\bar{S}_2^2|} \left(\left(\frac{T + \sqrt{T^2 - 4D}}{2} \right)^2 + \left(\frac{T - \sqrt{T^2 - 4D}}{2} \right)^2 \right)^{k \cdot \ell_k / 2} \\ &= \left(\left(\frac{T + \sqrt{T^2 - 4D}}{2} \right)^2 + \left(\frac{T - \sqrt{T^2 - 4D}}{2} \right)^2 \right)^{\sum_{k=2}^{|\bar{S}_2^2|} k \cdot \ell_k / 2}. \end{aligned}$$

Recalling (11) and as

$$\left(\frac{T + \sqrt{T^2 - 4D}}{2} \right)^2 + \left(\frac{T - \sqrt{T^2 - 4D}}{2} \right)^2 = T^2 - 2D,$$

it follows that

$$\begin{aligned} \mathbb{E} \left[e^{t \sum_{(i,j) \in S_2^2} A_{ij} B'_{p(i)p(j)}} \right] &\leq (T^2 - 2D)^{|\bar{S}_2^2|/2} \\ &= (p_{11}^2(e^{2t} - 1)^2 + 2q_{11}(e^{2t} - 1) + 1)^{|\bar{S}_2^2|/2} \\ &= (p_{11}^2(e^{4t} - 1) + 2(q_{11} - p_{11}^2)(e^{2t} - 1) + 1)^{|\bar{S}_2^2|/2} \\ &= e^{\frac{|\bar{S}_2^2|^2}{2} \cdot \log(1 + p_{11}^2(e^{4t} - 1) + 2(q_{11} - p_{11}^2)(e^{2t} - 1))}. \end{aligned}$$

Putting everything together in (16), we get:

$$\begin{aligned}
\mathbb{P}(\pi \text{ good}) &\leq \inf_{t>0} e^{-t \frac{np_{11} \cdot n(1-\alpha)}{4}} \cdot e^{\frac{|S_2^1|}{4} \cdot \log(1+p_{11}(e^{4t}-1))} \cdot e^{\frac{|\bar{S}_2^2|^2}{2} \cdot \log(1+p_{11}^2(e^{4t}-1)+2(q_{11}-p_{11}^2)(e^{2t}-1))} \\
&\leq \inf_{t>0} e^{-t \frac{np_{11} \cdot n(1-\alpha)}{4}} + \frac{n}{4} \cdot \log(1+p_{11}(e^{4t}-1)) + \frac{n^2}{4} \cdot \log(1+p_{11}^2(e^{4t}-1)+2(q_{11}-p_{11}^2)(e^{2t}-1)) \\
&\leq \inf_{t>0} e^{-t \frac{np_{11} \cdot n(1-\alpha)}{4}} e^{\frac{n}{4} p_{11}(e^{4t}-1)} \cdot e^{\frac{n^2}{4} \cdot (p_{11}^2(e^{4t}-1)+2(q_{11}-p_{11}^2)(e^{2t}-1))} \\
&\leq \left(\inf_{t>0} e^{-t(1-\alpha)} e^{\frac{e^{4t}-1}{n}} \cdot e^{p_{11}(e^{4t}-1)+2\frac{(q_{11}-p_{11}^2)}{p_{11}}(e^{2t}-1)} \right)^{\frac{n \cdot np_{11}}{4}} \\
&=_{z=e^{2t}} \left(\inf_{z>0} z^{-(1-\alpha)/2} e^{\frac{z^2-1}{n}} e^{p_{11}(z^2-1)+2\frac{q_{11}-p_{11}^2}{p_{11}}(z-1)} \right)^{\frac{n \cdot np_{11}}{4}},
\end{aligned}$$

where we have used (10) in the second inequality to obtain $|S_1^1| \leq n(1-\epsilon) \leq n$ and $|\bar{S}_2^2| \leq \frac{n^2}{2}$, and the fact that $\log(1+x) \leq x$ for $x \geq 0$ for the third inequality. We now leverage Lemma 4.3 to conclude. In our case,

$$\tau = \frac{1-\alpha}{2}, \quad q_1 = 2\frac{q_{11}-p_{11}^2}{p_{11}}, \quad q_2 = p_{11}. \quad (18)$$

Let z^* be as defined in Lemma 4.3: from (13), we have $z^{*2} \leq \frac{\tau - q_1 z^*}{2q_2} \leq \frac{\tau}{2q_2}$ and so

$$\frac{z^{*2}-1}{n} \leq \frac{z^{*2}}{n} \leq \frac{1-\alpha}{4np_{11}}.$$

Leveraging this to upper bound $e^{\frac{z^2-1}{n}}$ independently of z and then using (12) with the definitions in (18) in mind, it follows that

$$\mathbb{P}(\pi \text{ good}) \leq e^{\frac{n(1-\alpha)}{16}} \left(\frac{1}{\zeta} \right)^{-\frac{n(1-\alpha)np_{11}}{8}},$$

where ζ is defined as in Lemma 4.3, i.e.,

$$\frac{1}{\zeta} = \min \left\{ \frac{\tau}{\sqrt{2}eq_1}, \frac{\tau^{1/2}}{4eq_2^{1/2}} \right\}.$$

We show that $\log\left(\frac{1}{\zeta}\right) \geq \min\{\gamma, \beta\} \log n$ by lower bounding each term under the assumptions of Theorem 2.3. First,

$$\log\left(\frac{\tau}{\sqrt{2}eq_1}\right) = \log\left(\frac{p_{11} \cdot (1-\alpha)}{4\sqrt{2} \cdot e(q_{11}-p_{11}^2)}\right) \geq \beta \log(n) + O(1).$$

Secondly, noting that for the assumptions in (6) to make sense, it must be the case that $\gamma < \frac{1}{2}$ and so $n^{1/2} \geq n^\gamma$:

$$\log\left(\frac{\tau^{1/2}}{4eq_2^{1/2}}\right) = \log\left(\frac{\sqrt{1-\alpha}}{4e\sqrt{2}\sqrt{qs}}\right) \geq \gamma \log(n) + O(1).$$

As $|\{\pi \mid \beta(\pi, \pi^*) \leq \alpha\}| \leq |\{\pi \in P_n\}| = n! \leq n^{\log n}$, plugging everything back into (15) gives us

$$\mathbb{P}(\exists \pi : \pi \text{ good} \cap \beta(\pi, \pi^*) < \alpha) \leq e^{n \log n + \frac{n(1-\alpha)}{16} - \frac{n(1-\alpha)n p_{11}(\min\{\gamma, \beta\} \log n + O(1))}{8}} \leq e^{-n \log n + O(n)},$$

where the final inequality is implied by the assumptions in Theorem 2.3. As the right hand side of this inequality converges to 0 when $n \rightarrow \infty$, we get our result. \square

A Proof of Proposition 4.2

The proof of Proposition 4.2 makes use of the following lemmas.

Lemma A.1 (Berry-Esseen Theorem; see, e.g., [27]). *Let X_1, X_2, \dots be independent and identically distributed random variables such that $\mathbb{E}[X_1] = 0$, $\mathbb{E}[X_1^2] = \sigma^2 > 0$ and $\mathbb{E}[|X_1|^3] < \beta^3 < \infty$. For $n \geq 1$, define*

$$S_n = \frac{X_1 + \dots + X_n}{\sigma\sqrt{n}}$$

and let $F_n(x) = \mathbb{P}[S_n \leq x]$ be its cdf. Let $\phi(x)$ be the cdf of a standard normal distribution. For any x and any $n \geq 1$, it holds

$$|F_n(x) - \phi(x)| \leq \frac{0.34445\beta^3 + 0.16844}{\sqrt{n}}.$$

Lemma A.2. *Let $\mu > 0$ and recall the notation $\psi_j(\mu) = \mathbb{P}(Po(\mu) \geq j)$, where $Po(\mu)$ is a Poisson random variable of parameter μ and j is a nonnegative integer. Denote by ϕ the cdf of a standard normal. We have*

$$\psi_j(\mu) \geq 1 - \phi\left(\frac{j - \mu}{\sqrt{\mu}}\right) - \frac{0.55}{\sqrt{\lceil\mu\rceil}}.$$

In particular, for $\lambda \geq 20$,

$$\psi_{\lambda/2-1}(2\lambda/3) \geq 0.7$$

Proof. Proof. We can write $Po(\mu) = \sum_{i=1}^{\lceil\mu\rceil} P_i$ where $P_i, i = 1, \dots, \mu$ are independent Poisson random variables with parameter $\mu/\lceil\mu\rceil \leq 1$. Let $X_i = P_i - \frac{\mu}{\lceil\mu\rceil}$: we have that $\mathbb{E}[X_i] = 0$, $\mathbb{E}[X_i^2] = \mu/\lceil\mu\rceil$, and $\mathbb{E}[X_i^3] = \mu/\lceil\mu\rceil$. From this, we have that

$$\begin{aligned} \psi_j(\mu) &= \mathbb{P}(Po(\mu) \geq j) = \mathbb{P}\left(\sum_{i=1}^{\lceil\mu\rceil} (P_i - \mathbb{E}[P_i]) \geq j - \mu\right) \\ &= \mathbb{P}\left(\frac{\sum_{i=1}^{\lceil\mu\rceil} X_i}{\sqrt{\mu/\lceil\mu\rceil}\sqrt{\lceil\mu\rceil}} \geq \frac{j - \mu}{\sqrt{\mu/\lceil\mu\rceil}\sqrt{\lceil\mu\rceil}}\right) \\ &= \mathbb{P}\left(\frac{\sum_{i=1}^{\lceil\mu\rceil} X_i}{\sqrt{\mu}} \geq \frac{j - \mu}{\sqrt{\mu}}\right) \\ &= 1 - F_{\lceil\mu\rceil}\left(\frac{j - \mu}{\sqrt{\mu}}\right). \end{aligned}$$

Using Lemma A.1 and the fact that $\mathbb{E}[|X_i|^3] = \frac{\mu}{\lceil\mu\rceil} \leq 1$, we have that for any x ,

$$|1 - F_{\lceil\mu\rceil}(x)| - |1 - \phi(x)| \leq |F_{\lceil\mu\rceil}(x) - \phi(x)| \leq \frac{0.55}{\sqrt{\lceil\mu\rceil}},$$

which implies $1 - F_{\lceil\mu\rceil}(x) \geq 1 - \phi(x) - \frac{0.55}{\sqrt{\lceil\mu\rceil}}$ and so,

$$\psi_j(\mu) \geq 1 - \phi\left(\frac{j - \mu}{\sqrt{\mu}}\right) - \frac{0.55}{\sqrt{\lceil\mu\rceil}}.$$

Now, assume that $\lambda \geq 20$. We have that

$$\psi_{\lambda/2-1}(2\lambda/3) \geq 1 - \phi\left(\frac{\lambda/2 - 1 - 2\lambda/3}{\sqrt{2\lambda/3}}\right) - \frac{0.55}{\sqrt{\lceil 2\lambda/3 \rceil}} = 1 - \phi\left(\frac{-\lambda/6 - 1}{\sqrt{2\lambda/3}}\right) - \frac{0.55}{\lceil \sqrt{2\lambda/3} \rceil}.$$

When $\lambda \geq 20$, we have that $\lambda \mapsto -\frac{\lambda/6+1}{\sqrt{2\lambda/3}}$ is decreasing and that $\lambda \mapsto -\frac{0.55}{\lceil \sqrt{2\lambda/3} \rceil}$ is increasing. From this, we get

$$\psi_{\lambda/2-1}(2\lambda/3) \geq 1 - \phi\left(\frac{-20/6 - 1}{\sqrt{2 \cdot 20/3}}\right) - \frac{0.55}{\sqrt{\lceil 2 \cdot 20/3 \rceil}} \geq 0.7.$$

□

Proof. Proof of Proposition 4.2. The permutation π^* is good if the intersection graph of G_A and $G_{B'}$ has at least $n \cdot \frac{1+\alpha}{2}$ nodes that have degree greater than or equal to $\frac{nqs}{2}$. By definition of a k -core, if this intersection graph has an $nqs/2$ -core of size greater than or equal to $n(1+\alpha)/2$ whp, then π^* is good whp.

The intersection graph of G_A and $G_{B'}$ is a graph over n nodes such that there is an edge between nodes i and j if and only if there is an edge in G_A and $G_{B'}$ between i and j . As $\mathbb{P}((A_{ij}, B'_{ij}) = (1, 1)) = qs$ and (A_{ij}, B'_{ij}) is independent of (A_{kl}, B'_{kl}) for any $(k, l) \neq (i, j)$, it follows that the intersection graph of G_A and $G_{B'}$ is an Erdős-Rényi graph of parameters (n, qs) . If we are able to show that under the conditions of Theorem 2.3,

$$(i) \quad nqs > c_{nqs/2}$$

$$(ii) \quad \psi_{nqs/2}(\mu_{nqs/2}(nqs)) \geq n \cdot \frac{1+\alpha}{2}$$

then from Lemma 4.1, the $nqs/2$ -core of the intersection graph of G_A and $G_{B'}$ contains a fraction greater than or equal to $\frac{1+\alpha}{2}$ of the nodes in the graph whp and Proposition 4.2 follows.

We start by proving (i). Recall that $c_{nqs/2} = \inf_{\mu>0} \frac{\mu}{\psi_{nqs/2-1}(\mu)}$. Hence, for any $\mu > 0$, in particular for $\mu = 2nqs/3 > 0$,

$$c_{nqs/2} \leq \frac{2nqs/3}{\psi_{nqs/2-1}(2nqs/3)} \leq \frac{2nqs/3}{0.7} < nqs,$$

where we have used Lemma A.2 for the second inequality as the conditions of Theorem 2.3 imply that $\lambda \geq 20$.

We now prove (ii) by showing that under the conditions of Theorem 2.3, $\frac{2nqs}{3} \leq \mu_{nqs/2}(nqs) \leq nqs$ where $\mu_{nqs/2}(nqs)$ is the largest solution to $\frac{\mu}{\psi_{nqs/2-1}(\mu)} = nqs$. This implies that

$$\begin{aligned} \psi_{nqs/2}(\mu_{nqs/2}(nqs)) &= \mathbb{P}(Po(\mu_{nqs/2}(nqs)) \geq nqs/2) \\ &= 1 - \mathbb{P}(Po(\mu_{nqs/2}(nqs)) < nqs/2) \\ &\leq 1 - \mathbb{P}(Po(\mu_{nqs/2}(nqs)) \leq \mu_{nqs/2}(nqs) - nqs/2) \\ &\geq 1 - \mathbb{P}(Po(\mu_{nqs/2}(nqs)) \leq \mu_{nqs/2}(nqs) - nqs/6) \\ &\geq 1 - e^{-(nqs/6)^2/2(\mu_{nqs/2}(nqs)+nqs/6)} \\ &\geq 1 - e^{-(nqs)^2/(36 \cdot 7/3 \cdot nqs)} = 1 - e^{-nqs/84}, \end{aligned}$$

where we have used a standard Chernoff bound for Poisson random variables for the second inequality, see, e.g. [7]. As the conditions of Theorem 2.3 require that $nqs \geq 84 \cdot \log\left(\frac{2}{1-\alpha}\right)$, we get that $\psi_{nqs/2}(\mu_{nqs/2}(nqs)) \geq \frac{1+\alpha}{2}$.

To show that $\frac{2nqs}{3} \leq \mu_{nqs/2}(nqs) \leq nqs$, let $f_{nqs}(\mu) = \mu - nqs \cdot \psi_{nqs/2-1}(\mu)$, which is a continuous function of μ for fixed nqs . Here, $\mu_{nqs/2}(nqs)$ is the largest solution to $f_{nqs}(\mu) = 0$. It is straightforward to see that for $\mu > nqs$, $f_{nqs}(\mu) > 0$ as $\psi_{nqs/2-1} \leq 1$ for any value that nqs takes. Furthermore, from Lemma A.2,

$$f_{nqs}(2nqs/3) = \frac{2nqs}{3} - nqs \cdot \psi_{nqs/2-1}(2nqs/3) \leq \frac{2nqs}{3} - nqs \cdot 0.7 < 0.$$

Combining these results, we get that $\mu_{nqs/2}(nqs)$ must be less than or equal to nqs and that there is a solution $f_{nqs}(\mu) = 0$ in the interval $[2nqs/3, nqs]$. Our result immediately follows. \square

B Proofs of auxiliary results for Proposition 4.4

Lemma B.1. *Let $n \in \mathbb{N}$ and let $2 \leq k \leq n$. Furthermore, let α, β be positive real numbers. We have*

$$(\alpha^k + \beta^k)^{n/k} \leq (\alpha^2 + \beta^2)^{n/2}.$$

Proof. Proof. We have

$$\frac{(\alpha^2 + \beta^2)^{n/2}}{(\alpha^k + \beta^k)^{n/k}} = \frac{(\alpha^2 + \beta^2)^{k/2 \cdot n/k}}{(\alpha^k + \beta^k)^{n/k}} = \left(\frac{(\alpha^2 + \beta^2)^{k/2}}{(\alpha^k + \beta^k)} \right)^{n/k}.$$

As $n/k > 1$, we simply need to check whether $(\alpha^2 + \beta^2)^{k/2} \geq (\alpha^k + \beta^k)$. If k is divisible by 2, this is straightforward from the Binomial theorem as $\alpha > 0$ and $\beta > 0$. If k is not divisible by 2, then $k - 1$ is divisible by 2 and $(\alpha^2 + \beta^2)^k = (\alpha^2 + \beta^2)^{(k-1)/2} \cdot \sqrt{\alpha^2 + \beta^2}$. By applying the Binomial theorem to the first part of the sum and noting that $\sqrt{\alpha^2 + \beta^2} \geq \alpha$ and β , the result follows. \square

Lemma B.2. *Let $(A_0, B_0), (A_1, B_1), \dots, (A_n, B_n)$ be $n + 1$ pairs of random variables taking values in $\{0, 1\}^2$. We suppose that (A_i, B_i) are distributed following the probability distribution P given in Section 2.4.1. In particular, the covariance between A_i and B_i is given by $\sigma^2 > 0$. We further assume that all pairs $\{(A_i, B_i)\}$ are independent of one another. Define*

$$W := A_0B_1 + A_1B_2 + \dots + A_{n-1}B_n + A_nB_0.$$

For any $t > 0$, we have that

$$\mathbb{E}[e^{tW}] = \left(\frac{T + \sqrt{T^2 - 4D}}{2} \right)^{n+1} + \left(\frac{T - \sqrt{T^2 - 4D}}{2} \right)^{n+1}$$

where $T = p_{11}(e^t - 1) + 1$ and $D = \sigma^2(e^t - 1)$.

Remark B.1. Note that if $\sigma^2 = 0$, i.e., A_i and B_i are uncorrelated, this moment generating function simplifies to that of a Binomial random variable of parameters (n, qs) , as it should.

Our proof will make use of the following result from [49].

Lemma B.3. *Let M be a 2×2 matrix with eigenvalues α and β . Then, if $\alpha \neq \beta$, for all $n \geq 1$,*

$$M^n = \alpha^n \left(\frac{M - \beta I}{\alpha - \beta} \right) + \beta^n \left(\frac{M - \alpha I}{\beta - \alpha} \right).$$

Proof. Proof of Theorem B.2. Let

$$\begin{aligned} u_n^{00} &= \mathbb{E}[e^{t(A_1 B_2 + A_2 B_3 + \dots + A_{n-2} B_{n-1} + A_{n-1} B_n)}], \\ u_n^{01} &= \mathbb{E}[e^{t(A_1 B_2 + A_2 B_3 + \dots + A_{n-2} B_{n-1} + A_{n-1} B_n + A_n)}], \\ u_n^{10} &= \mathbb{E}[e^{t(B_1 + A_1 B_2 + A_2 B_3 + \dots + A_{n-2} B_{n-1} + \dots + A_{n-1} B_n)}], \\ u_n^{11} &= \mathbb{E}[e^{t(B_1 + A_1 B_2 + A_2 B_3 + \dots + A_{n-2} B_{n-1} + \dots + A_{n-1} B_n + A_n)}]. \end{aligned}$$

Note that by conditioning on (A_0, B_0) , we have that $\mathbb{E}[e^{tW}] = p_{00}u_n^{00} + p_{01}u_n^{01} + p_{10}u_n^{10} + p_{11}u_n^{11}$. We thus study these four sequences. By conditioning on (A_1, B_1) , it is easy to see that the following formulas hold true:

$$\begin{aligned} u_n^{00} &= (p_{00} + p_{01})u_{n-1}^{00} + (p_{10} + p_{11})u_{n-1}^{10} \\ u_n^{01} &= (p_{00} + p_{01})u_{n-1}^{01} + (p_{10} + p_{11})u_{n-1}^{11} \\ u_n^{10} &= (p_{00} + p_{01}e^t)u_{n-1}^{00} + (p_{10} + p_{11}e^t)u_{n-1}^{10} \\ u_n^{11} &= (p_{00} + p_{01}e^t)u_{n-1}^{01} + (p_{10} + p_{11}e^t)u_{n-1}^{11}. \end{aligned}$$

Successive conditioning on (A_k, B_k) for $k \geq 2$ leads to these formulas holding for all $n \geq 2$. Hence, if we let

$$M := \begin{bmatrix} p_{00} + p_{01} & p_{10} + p_{11} \\ p_{00} + p_{01}e^t & p_{10} + p_{11}e^t \end{bmatrix}$$

then

$$\begin{bmatrix} u_n^{00} \\ u_n^{10} \end{bmatrix} = M \cdot \begin{bmatrix} u_{n-1}^{00} \\ u_{n-1}^{10} \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} u_n^{01} \\ u_n^{11} \end{bmatrix} = M \cdot \begin{bmatrix} u_{n-1}^{01} \\ u_{n-1}^{11} \end{bmatrix}.$$

We further have the following initial conditions

$$\begin{aligned} u_1^{00} &= \mathbb{E}[e^{t(0 \cdot B_n + A_n \cdot 0)}] = 1 \\ u_1^{01} &= \mathbb{E}[e^{t(0 \cdot B_n + A_n)}] = 1 + (p_{01} + p_{11})(e^t - 1) \\ u_1^{10} &= \mathbb{E}[e^{t(B_n + A_n \cdot 0)}] = 1 + (p_{01} + p_{11})(e^t - 1) \\ u_1^{11} &= \mathbb{E}[e^{-t(B_n + A_n)}] = 1 + 2(p_{11} + p_{01})(e^t - 1) + p_{11}(e^t - 1)^2 \end{aligned}$$

Using these conditions, we get that

$$\begin{bmatrix} u_n^{00} \\ u_n^{10} \end{bmatrix} = M^{n-1} \cdot \begin{bmatrix} u_1^{00} \\ u_1^{10} \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} u_n^{01} \\ u_n^{11} \end{bmatrix} = M^{n-1} \cdot \begin{bmatrix} u_1^{01} \\ u_1^{11} \end{bmatrix}.$$

The difficulty now is in computing M^{n-1} , which we will do via Lemma B.3. To this effect, we denote by T the trace of M and by D its determinant. One can easily compute their values:

$$T = \text{tr}(M) = p_{11}(e^t - 1) + 1 \quad \text{and} \quad D = \det(M) = \sigma^2(e^t - 1).$$

Recall that the eigenvalues of a 2×2 matrix satisfy the equation $x^2 - Tx + D = 0$. The discriminant here is $T^2 - 4D$, which is positive:

$$\Delta = T^2 - 4D = p_{11}^2(e^t - 1)^2 + 2(p_{11} - 2\sigma^2)(e^t - 1) + 1 = \left(p_{11}(e^t - 1) + \frac{p_{11} - 2\sigma^2}{p_{11}} \right)^2 + \frac{4\sigma^2(p_{11} - \sigma)}{p_{11}^2} > 0$$

Hence solutions to this equation exist. We write:

$$\alpha = \frac{T + \sqrt{T^2 - 4D}}{2} \quad \text{and} \quad \beta = \frac{T - \sqrt{T^2 - 4D}}{2}.$$

In accordance with Lemma B.3, we now work on computing $M - \beta I$ and $M - \alpha I$. We have:

$$M - \beta I = R + \frac{\sqrt{\Delta}}{2}I,$$

where

$$R := \begin{bmatrix} \frac{p_{00}-p_{11}}{2} - \frac{p_{11}}{2}(e^t - 1) & p_{10} + p_{11} \\ p_{00} + p_{01}e^t & -(\frac{p_{00}-p_{11}}{2}) + \frac{p_{11}}{2}(e^t - 1) \end{bmatrix}.$$

We denote by r_{ij} its entries. Note that $r_{11} = -r_{22}$. Likewise, we have

$$M - \alpha I = R - \frac{\sqrt{\Delta}}{2}I.$$

It follows that

$$M^n = \frac{\alpha^n}{\sqrt{\Delta}}(R + \frac{\sqrt{\Delta}}{2}I) - \frac{\beta^n}{\sqrt{\Delta}}(R - \frac{\sqrt{\Delta}}{2}I) = \left(\frac{\alpha^n - \beta^n}{\sqrt{\Delta}}\right)R + \frac{\alpha^n + \beta^n}{2} \cdot I$$

and from this, for any $n \geq 1$,

$$\begin{aligned} u_n^{00} &= \left(\frac{\alpha^{n-1} - \beta^{n-1}}{\sqrt{\Delta}}\right) (r_{11}u_1^{00} + r_{12}u_1^{10}) + \frac{\alpha^{n-1} + \beta^{n-1}}{2}u_1^{00} \\ u_n^{10} &= \left(\frac{\alpha^{n-1} - \beta^{n-1}}{\sqrt{\Delta}}\right) (r_{21}u_1^{00} + r_{22}u_1^{10}) + \frac{\alpha^{n-1} + \beta^{n-1}}{2}u_1^{10} \\ u_n^{01} &= \left(\frac{\alpha^{n-1} - \beta^{n-1}}{\sqrt{\Delta}}\right) (r_{11}u_1^{01} + r_{12}u_1^{11}) + \frac{\alpha^{n-1} + \beta^{n-1}}{2}u_1^{01} \\ u_n^{11} &= \left(\frac{\alpha^{n-1} - \beta^{n-1}}{\sqrt{\Delta}}\right) (r_{21}u_1^{01} + r_{22}u_1^{11}) + \frac{\alpha^{n-1} + \beta^{n-1}}{2}u_1^{11}. \end{aligned}$$

This can now be used to compute $\mathbb{E}[e^{tW}]$. We have

$$\begin{aligned} \mathbb{E}[e^{tW}] &= p_{00}u_n^{00} + p_{01}u_n^{01} + p_{10}u_n^{10} + p_{11}u_n^{11} \\ &= (r_{11}(p_{00}u_1^{00} + p_{01}u_1^{01} - p_{10}u_1^{10} - p_{11}u_1^{11}) + r_{12}(p_{00}u_1^{10} + p_{01}u_1^{11}) + r_{21}(p_{10}u_1^{00} + p_{11}u_1^{01})) \\ &\quad \cdot \left(\frac{\alpha^{n-1} - \beta^{n-1}}{\sqrt{\Delta}}\right) + (p_{00}u_1^{00} + p_{10}u_1^{10} + p_{01}u_1^{01} + p_{11}u_1^{11}) \cdot \left(\frac{\alpha^{n-1} + \beta^{n-1}}{2}\right). \end{aligned}$$

We compute each part separately. We have:

$$\begin{aligned} p_{00}u_1^{00} + p_{10}u_1^{10} + p_{01}u_1^{01} + p_{11}u_1^{11} &= T^2 - 2D \\ r_{11}(p_{00}u_1^{00} + p_{01}u_1^{01} - p_{10}u_1^{10} - p_{11}u_1^{11}) &= \frac{1}{2} \left(\frac{p_{00} - p_{11}}{2} - \frac{p_{11}}{2}(e^t - 1)\right)^2 \cdot T \\ r_{12}(p_{00}u_1^{10} + p_{01}u_1^{11}) + r_{21}(p_{10}u_1^{00} + p_{11}u_1^{01}) &= 2T(p_{01} + p_{11})(1 - p_{11} - p_{01} + p_{01}(e^t - 1)) \end{aligned}$$

We combine the results above to obtain:

$$r_{11}(p_{00}u_1^{00} + p_{01}u_1^{01} - p_{10}u_1^{10} - p_{11}u_1^{11}) + r_{12}(p_{00}u_1^{10} + p_{01}u_1^{11}) + r_{21}(p_{10}u_1^{00} + p_{11}u_1^{01}) = \frac{T^3}{2} - 2DT$$

which leads us to finally having

$$\mathbb{E}[e^{tW}] = \left(\frac{\alpha^{n-1} - \beta^{n-1}}{\sqrt{\Delta}}\right) \cdot \left(\frac{T^3}{2} - 2DT\right) + \left(\frac{\alpha^{n-1} + \beta^{n-1}}{2}\right) \cdot (T^2 - 2D).$$

As $\frac{T^3}{2} - 2DT = \frac{T}{2} \cdot (T^2 - 4D) = \frac{T}{2} \cdot \Delta$, we can further simplify

$$\mathbb{E}[e^{tW}] = (\alpha^{n-1} - \beta^{n-1}) \cdot \frac{T}{2} \cdot \sqrt{\Delta} + \left(\frac{\alpha^{n-1} + \beta^{n-1}}{2} \right) \cdot (T^2 - 2D) = \alpha^{n+1} + \beta^{n+1}$$

□

C Proof of Lemma 4.3

As mentioned previously, the proof of Lemma 4.3 is contained in [12]. We repeat it here for completeness.

Proof. Proof of Lemma 4.3. The optimal choice of z is obtained by differentiating $z \mapsto z^{-\tau} e^{q_2(z^2-1)+q_1(z-1)}$ and setting the derivative equal to 0, i.e.,

$$(2q_2z + q_1)z^{-\tau} e^{q_2(z^2-1)+q_1(z-1)} - \tau z^{-\tau-1} e^{q_2(z^2-1)+q_1(z-1)} = 0,$$

which is equivalent to the quadratic equation

$$2q_2z^2 + q_1z - \tau = 0. \tag{19}$$

This equation has one positive root and one negative root: the positive root is z^* . From this quadratic equation, we can obtain the expression $z^{*2} = \frac{\tau - q_1z^*}{2q_2}$. As $q_1 > 0$ and $z^* > 0$, it follows that

$$z^{*2} \leq \frac{\tau}{2q_2},$$

which corresponds to (13). One can also give an explicit expression of z^* by solving the quadratic equation:

$$z^* = \frac{-q_1 + \sqrt{q_1^2 + 8\tau q_2}}{4q_2} = \frac{2\tau}{q_1 + \sqrt{q_1^2 + 8\tau q_2}}.$$

Because $q_1 \leq \sqrt{q_1^2 + 8\tau q_2}$, we have the bounds

$$\frac{\tau}{\sqrt{q_1^2 + 8\tau q_2}} \leq z^* \leq \frac{\tau}{q_1}. \tag{20}$$

Starting with one of the factors from the left side of (12), we have

$$e^{q_2(z^{*2}-1)+q_1(z^*-1)} \leq e^{\frac{q_1}{2}z^* + \frac{\tau}{2} - q_2 - q_1} \leq e^{\tau - q_1 - q_2} \leq e^\tau,$$

where we used (19) to eliminate the q_2z^{*2} term, applied the upper bound from (20), and used $q_1, q_2 \geq 0$. From the lower bound in (20),

$$z^{*-2} \leq \frac{q_1^2}{\tau^2} + \frac{8q_2}{\tau} \leq \max \left\{ \frac{2q_1^2}{\tau^2}, \frac{16q_2}{\tau} \right\}$$

so $z^{*-\tau} e^{q_2(z^{*2}-1)+q_1(z^*-1)} \leq \zeta^\tau$. □

References

- [1] E. Abbe. Community detection and stochastic block models: recent developments. *The Journal of Machine Learning Research*, 18(1):6446–6531, 2017.
- [2] H. Almohamad and S. O. Duffuaa. A linear programming approach for the weighted graph matching problem. *IEEE Transactions on pattern analysis and machine intelligence*, 15(5):522–525, 1993.
- [3] S. Banerjee, N. Hegde, and L. Massoulié. The price of privacy in untrusted recommendation engines. In *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 920–927. IEEE, 2012.
- [4] B. Barak, C.-N. Chou, Z. Lei, T. Schramm, and Y. Sheng. (nearly) efficient algorithms for the graph matching problem on correlated random graphs. *arXiv preprint arXiv:1805.02349*, 2018.
- [5] M. Bayati, D. F. Gleich, A. Saberi, and Y. Wang. Message-passing algorithms for sparse network alignment. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 7(1):3, 2013.
- [6] B. Bollobás. Distinguishing vertices of random graphs. In *North-Holland Mathematics Studies*, volume 62, pages 33–49. Elsevier, 1982.
- [7] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- [8] D. Conte, P. Foggia, C. Sansone, and M. Vento. Thirty years of graph matching in pattern recognition. *International journal of pattern recognition and artificial intelligence*, 18(03):265–298, 2004.
- [9] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [10] D. Cullina and N. Kiyavash. Improved achievability and converse bounds for erdos-rényi graph matching. *ACM SIGMETRICS Performance Evaluation Review*, 44(1):63–72, 2016.
- [11] D. Cullina and N. Kiyavash. Exact alignment recovery for correlated erdos renyi graphs. *arXiv preprint arXiv:1711.06783*, 2017.
- [12] D. Cullina, N. Kiyavash, P. Mittal, and H. V. Poor. Partial recovery of erd\h {o} sr\{e} nyi graph alignment via k -core alignment. *arXiv preprint arXiv:1809.03553*, 2018.
- [13] T. Czajka and G. Pandurangan. Improved random graph isomorphism. *Journal of Discrete Algorithms*, 6(1):85–92, 2008.
- [14] O. E. Dai, D. Cullina, N. Kiyavash, and M. Grossglauser. On the performance of a canonical labeling for matching correlated erdős-rényi graphs. *arXiv preprint arXiv:1804.09758*, 2018.
- [15] J. Ding, Z. Ma, Y. Wu, and J. Xu. Efficient random graph matching via degree profiles. *arXiv preprint arXiv:1811.07821*, 2018.
- [16] N. Dym, H. Maron, and Y. Lipman. Ds++: A flexible, scalable and provably tight relaxation for matching problems. *arXiv preprint arXiv:1705.06148*, 2017.

- [17] A. Elmsallati, C. Clark, and J. Kalita. Global alignment of protein-protein interaction networks: A survey. *IEEE/ACM transactions on computational biology and bioinformatics*, 13(4):689–705, 2015.
- [18] Z. Fan, C. Mao, Y. Wu, and J. Xu. Spectral graph matching and regularized quadratic relaxations i: The gaussian model. *arXiv preprint arXiv:1907.08880*, 2019.
- [19] Z. Fan, C. Mao, Y. Wu, and J. Xu. Spectral graph matching and regularized quadratic relaxations ii: Erdős-Rényi graphs and universality. *arXiv preprint arXiv:1907.08883*, 2019.
- [20] S. Feizi, G. Quon, M. Mendoza, M. Medard, M. Kellis, and A. Jadbabaie. Spectral alignment of graphs. *IEEE Transactions on Network Science and Engineering*, 2019.
- [21] D. E. Fishkind, S. Adali, H. G. Patsolic, L. Meng, V. Lyzinski, and C. E. Priebe. Seeded graph matching. *arXiv preprint arXiv:1209.0367*, 2012.
- [22] F. Fogel, R. Jenatton, F. Bach, and A. d’Aspremont. Convex relaxations for permutation problems. In *Advances in Neural Information Processing Systems*, pages 1016–1024, 2013.
- [23] L. Ganassali and L. Massoulié. From tree matching to sparse graph alignment. *arXiv preprint arXiv:2002.01258*, 2020.
- [24] E. Kazemi, S. H. Hassani, and M. Grossglauser. Growing a graph matching from a handful of seeds. *Proceedings of the VLDB Endowment*, 8(10):1010–1021, 2015.
- [25] E. Kazemi, L. Yartseva, and M. Grossglauser. When can two unlabeled networks be aligned under partial overlap? In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 33–42. IEEE, 2015.
- [26] I. Kezurer, S. Z. Kovalsky, R. Basri, and Y. Lipman. Tight relaxation of quadratic matching. In *Computer Graphics Forum*, volume 34, pages 115–128. Wiley Online Library, 2015.
- [27] V. Y. Korolev and I. G. Shevtsova. On the upper bound for the absolute constant in the berry–esseen inequality. *Theory of Probability & Its Applications*, 54(4):638–658, 2010.
- [28] O. Kuchaiev, T. Milenković, V. Memišević, W. Hayes, and N. Pržulj. Topological network alignment uncovers biological function and phylogeny. *Journal of the Royal Society Interface*, 7(50):1341–1354, 2010.
- [29] S. Lacoste-Julien, B. Taskar, D. Klein, and M. I. Jordan. Word alignment via quadratic assignment. In *Proceedings of the main conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, pages 112–119. Association for Computational Linguistics, 2006.
- [30] V. Lyzinski, D. Fishkind, M. Fiori, J. Vogelstein, C. Priebe, and G. Sapiro. Graph matching: Relax at your own risk. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (1):1–1, 2016.
- [31] B. MacCartney, M. Galley, and C. D. Manning. A phrase-based alignment model for natural language inference. In *Proceedings of the conference on empirical methods in natural language processing*, pages 802–811. Association for Computational Linguistics, 2008.

- [32] N. Malod-Dognin and N. Pržulj. L-graal: Lagrangian graphlet-based network aligner. *Bioinformatics*, 31(13):2182–2189, 2015.
- [33] E. Mossel and J. Xu. Seeded graph matching via large neighborhood statistics. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1005–1014. SIAM, 2019.
- [34] A. Narayanan and V. Shmatikov. How to break anonymity of the netflix prize dataset. *arXiv preprint cs/0610105*, 2006.
- [35] A. Narayanan and V. Shmatikov. De-anonymizing social networks. In *2009 30th IEEE symposium on security and privacy*, pages 173–187. IEEE, 2009.
- [36] P. Pedarsani and M. Grossglauser. On the privacy of anonymized networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1235–1243, 2011.
- [37] B. Pittel, J. Spencer, and N. Wormald. Sudden emergence of a giantk-core in a random graph. *Journal of Combinatorial Theory, Series B*, 67(1):111–151, 1996.
- [38] R. C. Read and D. G. Corneil. The graph isomorphism disease. *Journal of graph theory*, 1(4):339–363, 1977.
- [39] F. Rendl, P. Pardalos, and H. Wolkowicz. The quadratic assignment problem: A survey and recent developments. In *Proceedings of the DIMACS workshop on quadratic assignment problems*, volume 16, pages 1–42, 1994.
- [40] S. Sahni and T. Gonzalez. P-complete approximation problems. *Journal of the ACM (JACM)*, 23(3):555–565, 1976.
- [41] N. P. Santhanam and M. J. Wainwright. Information-theoretic limits of selecting binary graphical models in high dimensions. *IEEE Transactions on Information Theory*, 58(7):4117–4134, 2012.
- [42] I. Sason and S. Verdú. f -divergence inequalities. *IEEE Transactions on Information Theory*, 62(11):5973–6006, 2016.
- [43] E. R. Scheinerman and D. H. Ullman. *Fractional graph theory: a rational approach to the theory of graphs*. Courier Corporation, 2011.
- [44] C. Schellewald and C. Schnörr. Probabilistic subgraph matching based on convex relaxation. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 171–186. Springer, 2005.
- [45] F. Shirani, S. Garg, and E. Erkip. Typicality matching for pairs of correlated graphs. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 221–225. IEEE, 2018.
- [46] R. Singh, J. Xu, and B. Berger. Global alignment of multiple protein interaction networks with application to functional orthology detection. *Proceedings of the National Academy of Sciences*, 105(35):12763–12768, 2008.
- [47] S. Umeyama. An eigendecomposition approach to weighted graph matching problems. *IEEE transactions on pattern analysis and machine intelligence*, 10(5):695–703, 1988.

- [48] J. T. Vogelstein, J. M. Conroy, V. Lyzinski, L. J. Podrazik, S. G. Kratzer, E. T. Harley, D. E. Fishkind, R. J. Vogelstein, and C. E. Priebe. Fast approximate quadratic programming for graph matching. *PLOS one*, 10(4), 2015.
- [49] K. S. Williams. The n th power of a 2×2 matrix. *Mathematics Magazine*, 65(5):336–336, 1992.
- [50] E. M. Wright et al. Graphs on unlabelled nodes with a given number of edges. *Acta Mathematica*, 126:1–9, 1971.
- [51] Y. Wu, J. Xu, and S. H. Yu. Settling the sharp reconstruction thresholds of random graph matching. *arXiv preprint arXiv:2102.00082*, 2021.
- [52] L. Yartseva and M. Grossglauser. On the performance of percolation graph matching. In *Proceedings of the first ACM conference on Online social networks*, pages 119–130. ACM, 2013.
- [53] M. Zaslavskiy, F. Bach, and J.-P. Vert. A path following algorithm for the graph matching problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(12):2227–2242, 2008.
- [54] S. Zhang, J. Yang, and W. Jin. Sapper: subgraph indexing and approximate matching in large graphs. *Proceedings of the VLDB Endowment*, 3(1-2):1185–1194, 2010.
- [55] Q. Zhao, S. E. Karisch, F. Rendl, and H. Wolkowicz. Semidefinite programming relaxations for the quadratic assignment problem. *Journal of Combinatorial Optimization*, 2(1):71–109, 1998.