



The planted matching problem: sharp threshold and infinite-order phase transition

Jian Ding¹ · Yihong Wu² · Jiaming Xu³ · Dana Yang⁴

Received: 2 May 2021 / Accepted: 28 April 2023

© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2023

Abstract

We study the problem of reconstructing a perfect matching M^* hidden in a randomly weighted $n \times n$ bipartite graph. The edge set includes every node pair in M^* and each of the $n(n-1)$ node pairs not in M^* independently with probability d/n . The weight of each edge e is independently drawn from the distribution \mathcal{P} if $e \in M^*$ and from \mathcal{Q} if $e \notin M^*$. We show that if $\sqrt{d}B(\mathcal{P}, \mathcal{Q}) \leq 1$, where $B(\mathcal{P}, \mathcal{Q})$ stands for the Bhattacharyya coefficient, the reconstruction error (average fraction of misclassified edges) of the maximum likelihood estimator of M^* converges to 0 as $n \rightarrow \infty$. Conversely, if $\sqrt{d}B(\mathcal{P}, \mathcal{Q}) \geq 1 + \epsilon$ for an arbitrarily small constant $\epsilon > 0$, the reconstruction error for any estimator is shown to be bounded away from 0 for both the sparse (fixed d) and dense (growing d) regimes, resolving the conjecture in Moharrami et al. (Ann Appl Probab 31(6):2663–2720, 2021. <https://doi.org/10.1214/20-AAP1660>) and Semerjian et al. (Phys Rev E 102:022304, 2020. <https://doi.org/10.1103/PhysRevE.102.022304>). Furthermore, in the special case of complete exponentially weighted graph with $d = n$, $\mathcal{P} = \exp(\lambda)$, and $\mathcal{Q} = \exp(1/n)$, for which the sharp threshold simplifies to $\lambda = 4$,

Jian Ding is supported by the NSF Grants DMS-1757479 and DMS-1953848. Yihong Wu is supported in part by the NSF Grant CCF-1900507, an NSF CAREER award CCF-1651588, and an Alfred Sloan fellowship. Jiaming Xu is supported by NSF grants IIS-1838124, CCF-1850743, CCF-1856424 and an NSF CAREER award CCF-2144593.

✉ Dana Yang
dana.yang@cornell.edu

Jian Ding
dingjian@math.pku.edu.cn

Yihong Wu
yihong.wu@yale.edu

Jiaming Xu
jx77@duke.edu

¹ School of Mathematical Sciences, Peking University, Beijing, China

² Department of Statistics and Data Science, Yale University, New Haven, CT, USA

³ The Fuqua School of Business, Duke University, Durham, NC, USA

⁴ Department of Statistics and Data Science, Cornell University, Ithaca, NY, USA

we prove that when $\lambda = 4 - \epsilon$, the optimal reconstruction error is $\exp(-\Theta(1/\sqrt{\epsilon}))$, confirming the conjectured infinite-order phase transition in Semerjian et al. (2020).

Keywords Planted matching recovery · Information-theoretic threshold · Phase transition · Linear assignment · Bhattacharyya coefficient

Mathematics Subject Classification 60C05 · 94A15

1 Introduction

Let \mathcal{M} denote the set of perfect matchings on the complete bipartite graph $K_{n,n}$, with left-hand vertices labeled as $[n] = \{1, \dots, n\}$ and right-hand vertices labeled as $[n]' = \{1', \dots, n'\}$. We denote the weights on the edge $e = (i, j')$ by $W_{i,j'} = W_e$. Each perfect matching M is understood as a set of edges in $K_{n,n}$.

Definition 1 (*Planted matching recovery*) Consider a weighted bipartite graph G randomly generated as follows. First sample M^* uniformly at random from the set \mathcal{M} of all perfect matchings on $K_{n,n}$ and connect all pairs in M^* . Then for every node pair (i, j') not in M^* , connect them independently with probability $\frac{d}{n}$. Finally, for edges in M^* , the edge weights are drawn independently from \mathcal{P} . The remaining edge weights are drawn independently from \mathcal{Q} . The goal is to reconstruct the hidden matching M^* based on G .

This problem is first proposed by [5] and motivated from tracking moving objects in a video, such as flocks of birds, motile cells, or particles in a fluid. A slight variation of Definition 1 is studied in [25] for unipartite graphs, where M^* is chosen uniformly at random from the set of all perfect matchings on the complete unipartite graph K_n (with even n) and the edge set of G includes all $n/2$ node pairs in M^* and each of the $\binom{n}{2} - n/2$ node pairs not in M^* independently with probability $\frac{d}{n}$; the edge weights are still independently distributed according to \mathcal{P} for edges in M^* and \mathcal{Q} otherwise. In this paper, we present our results and analysis for bipartite graphs; nevertheless, the proof techniques can be straightforwardly extended to the unipartite version and all conclusions hold verbatim.

Of particular interest are the following two regimes, which are the focus of [25] and the present paper:

- **Sparse model:** The average degree d is a constant. In this case, the null (i.e. unplanted) distribution \mathcal{Q} and planted distribution \mathcal{P} can be arbitrary distributions independent of n .
- **Dense model:** The average degree $d \equiv d(n) \rightarrow \infty$ as n grows. In particular, when $d = n$, we observe a complete bipartite graph with weights correlated with the hidden matching. In this case, we focus on the following special case of weight distributions, where the planted distribution \mathcal{P} has a fixed density p (with respect to the Lebesgue measure) and the null distribution \mathcal{Q} has a density q of the following form:

$$q(x) = \frac{1}{d} \rho\left(\frac{x}{d}\right) \quad (1)$$

where ρ is some fixed density on \mathbb{R} with $\rho(0) > 0$.¹ This scaling is natural and meaningful, because the magnitude of each unplanted edge weight is $O(d)$ on average, and each vertex is incident to an average of d unplanted edges. Under this scaling, the minimum magnitude of the unplanted edge weights incident to a given vertex is on the same order of $O(1)$ as the planted edge weight. Of special interest is the complete graph with exponential weights [21, 25], which we refer to as the *exponential model*, where $d = n$, $\mathcal{P} = \exp(\lambda)$ and $\mathcal{Q} = \exp(\frac{1}{n})$.

Note that its unplanted version is the celebrated random assignment model studied in [1, 14, 18, 19, 23, 27, 29]. Another special case was studied by in [5] where $\mathcal{P} = |\mathcal{N}(0, \kappa)|$ is a folded Gaussian and \mathcal{Q} is uniform over $[0, n]$.

Let $\widehat{M} \equiv \widehat{M}(G)$ denote an estimator of M^* that is a set of edges on $K_{n,n}$. The reconstruction error, namely, the fraction of misclassified edges is

$$\ell(\widehat{M}, M^*) = \frac{1}{n} |M^* \Delta \widehat{M}|, \quad (2)$$

where Δ denotes the symmetric set difference. We say that \widehat{M} achieves *almost perfect recovery* if $\mathbb{E}[\ell(M^*, \widehat{M})] = o(1)$. It can be shown (see e.g. [13, Appendix A]) that achieving a vanishing reconstruction error in expectation is equivalent to that with high probability. Note that for any estimator $\widehat{M} = \widehat{M}(G)$, we have

$$\mathbb{E}[\ell(M^*, \widehat{M})] = \frac{1}{n} \sum_{e \in E(K_{n,n})} \mathbb{P}\{e \in M^* \Delta \widehat{M}\}. \quad (3)$$

Thus the average reconstruction error is minimized by the *marginal maximum a posteriori (MAP)* estimator, where $e \in \widehat{M}$ if and only if $\mathbb{P}\{e \in M^*|G\} \geq \mathbb{P}\{e \notin M^*|G\}$, since it minimizes each summand in (3). Note that in general the marginal MAP need not be a matching or even be of size n . Nevertheless, it is easy to see that one can project any estimator to the set of perfect matchings (in distance metric ℓ) while increasing the average reconstruction error by a factor of two.

1.1 Main results

The information-theoretic threshold of the planted matching model is determined by the following key quantity, known has the Bhattacharyya coefficient (or Hellinger affinity) [4]:

$$B(\mathcal{P}, \mathcal{Q}) \triangleq \int \sqrt{fg} \, d\mu, \quad (4)$$

where f, g denote the relative density of \mathcal{P}, \mathcal{Q} with respect to some common dominating measure μ , respectively. It is well-known that $0 \leq B(\mathcal{P}, \mathcal{Q}) \leq 1$, with $B(\mathcal{P}, \mathcal{Q})$ equals 0 (resp. 1) when \mathcal{P} and \mathcal{Q} are mutually singular (resp. identical). For simplicity,

¹ If $\rho(0) = 0$, then for each vertex, among its incident edges, the planted edge weight has the smallest magnitude with high probability, in which case almost perfect recovery is trivially achievable.

we assume throughout the main body of the paper that $\mathcal{P} \ll \mathcal{Q}$, so that their relative density, denoted henceforth by $\frac{\mathcal{P}}{\mathcal{Q}}$, is well-defined. Nevertheless, even when $\mathcal{P} \not\ll \mathcal{Q}$ (such as the weight distributions considered in [5]), all results continue to hold (see Appendices E.2 and E.3 for justification).

We first give a sufficient condition for the maximum likelihood estimator (MLE) to achieve almost perfect recovery. The MLE reduces to the max-weighted matching on $K_{n,n}$ where each edge e is weighted by the corresponding log likelihood ratio, namely,

$$\widehat{M}_{\text{ML}} \in \arg \max_{M \in \mathcal{M}} \sum_{e \in M} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e), \quad (5)$$

which can be computed in polynomial time (as linear assignment). Note that the MLE need not be unique. For example, in the unweighted model (see Remark 2), the MLE is any maximum matching in the observed bipartite graph.

Theorem 1 *Assume that*

$$\sqrt{d} B(\mathcal{P}, \mathcal{Q}) \leq 1 + \epsilon \quad (6)$$

for some $\epsilon \geq 0$. Then there exists universal constant $C > 0$, such that for large enough n ,

$$\mathbb{E} [\ell(M^*, \widehat{M}_{\text{ML}})] \leq C \max \left\{ \log(1 + \epsilon), \sqrt{\frac{\log n}{n}} \right\}.$$

Note that Theorem 1 allows ϵ to take arbitrary values, including 0. In particular, if $\sqrt{d} B(\mathcal{P}, \mathcal{Q}) \leq 1$, then $\mathbb{E}[\ell(M^*, \widehat{M}_{\text{ML}})] \leq C \sqrt{\log n / n}$; if (6) holds with $\epsilon \rightarrow 0$ as $n \rightarrow \infty$, then \widehat{M}_{ML} achieves almost perfect recovery.

Next we proceed to negative results, which are the main focus of this paper. The following theorem shows the tightness of the condition (6):

Theorem 2 *Assume that*

$$\sqrt{d} B(\mathcal{P}, \mathcal{Q}) \geq 1 + \epsilon. \quad (7)$$

for some arbitrary constant $\epsilon > 0$. Under the dense model assume that ρ is continuous at 0 and $\rho(0) < \infty$. Then in both the sparse and the dense model, for any estimator \widehat{M} and large n ,

$$\mathbb{E}[\ell(M^*, \widehat{M})] \geq c, \quad (8)$$

where $c > 0$ is a constant independent of n . In the sparse model, c only depends on $\epsilon, \mathcal{P}, \mathcal{Q}$; in the dense model, c only depends on $\epsilon, \mathcal{P}, \rho$.

Theorems 1 and 2 together establish $\sqrt{d} B(\mathcal{P}, \mathcal{Q}) = 1$ as the sharp threshold for almost perfect recovery in both the sparse and dense model, proving the conjecture in

[25].² In the dense model with the scaling (1), the condition (7) simplifies to

$$\int_0^\infty \sqrt{p(x)} dx \geq \frac{1+\epsilon}{\sqrt{\rho(0)}}. \quad (9)$$

Note that this condition depends on the density function ρ in the null case only through its value at zero. In the special case of exponential weights of $p(x) = \lambda e^{-\lambda x}$ and $\rho(x) = e^{-x}$, the condition (9) further simplifies to

$$\lambda \leq 4 - \epsilon. \quad (10)$$

In view of the positive result in [21], this establishes $\lambda = 4$ as the sharp threshold of almost perfect recovery, resolving a conjecture in [21]. The next result, specialized to the exponential model, shows that the optimal reconstruction error is in fact $e^{-\Theta(\frac{1}{\sqrt{\epsilon}})}$, resolving a conjecture in [25]. Interestingly, this shows that the phase transition in the average reconstruction error is of infinite order,³ unlike other well-known planted problems such as the stochastic block model with two groups (second order) or with four or more groups (first order) [22].

Theorem 3 (Optimal reconstruction error for exponential model) *There exist an absolute constant C_0 such that the following holds. Suppose that $\lambda = 4 - \epsilon$ for some arbitrary constant $\epsilon > 0$. Then there exists $n_0 = n_0(\epsilon)$, such that for all $n \geq n_0$ and for any estimator $\widehat{M} = \widehat{M}(W)$,*

$$\mathbb{E}[\ell(M^*, \widehat{M})] \geq e^{-\frac{C_0}{\sqrt{\epsilon}}}. \quad (11)$$

Furthermore, let \widehat{M}_{ML} be given in (5) (which in this case coincides with the min-weight bipartite matching). Then

$$\mathbb{E}[\ell(M^*, \widehat{M}_{\text{ML}})] \leq \frac{C_0}{\epsilon^3} e^{-\frac{2\pi}{\sqrt{\epsilon}}}. \quad (12)$$

We end this section with two remarks on the universality of the sharp threshold and infinite-order phase transitions.

Remark 1 The sharp threshold $\sqrt{d} B(\mathcal{P}, \mathcal{Q}) = 1$ does not hold universally for all $d, \mathcal{P}, \mathcal{Q}$. Here is a simple example where this condition is not tight. Consider a complete graph with Gaussian edge weights drawn from either $\mathcal{P} = N(\mu, 1)$ or $\mathcal{Q} = N(0, 1)$. This model does not follow the scaling in (1) for the dense regime. In the Gaussian model, $\sqrt{n} B(\mathcal{P}, \mathcal{Q}) = 1$ simplifies to $\mu^2 = 4 \log n$. However, the sharp threshold in

² To be precise, the conjecture given in [25, eq. (45) and (40)] is stated under the aforementioned unipartite version of the planted matching problem. Nevertheless, our proof techniques as well as the sharp thresholds in Theorems 1–3 continue to hold for the unipartite version.

³ In statistical physics parlance, a phase transition is called *first-order* if the order parameter (in this case, the average reconstruction error) is discontinuous at the threshold, and of *pth order* for $p \geq 2$ if its $(p-2)$ th derivative is continuous but its $(p-1)$ -th derivative is discontinuous [25].

fact occurs at $\mu^2 = 2 \log n$. Indeed, when $\mu^2 \geq (2 + \epsilon) \log n$, a simple thresholding algorithm finds a matching that differs from the planted matching by $o(n)$ edges with high probability; conversely, using the mutual information argument in [10], it is easy to show that almost perfect recovery is impossible if $\mu^2 \leq (2 - \epsilon) \log n$ for any constant $\epsilon > 0$.

Remark 2 (Finite-order phase transition for unweighted graphs) Although it is possible that the infinite-order phase transition established in Theorem 3 holds beyond exponential density, we cannot expect it to be a completely universal phenomenon. In fact, for sparse unweighted graphs the phase transition is of a *finite order*. In this case, the observed graph is a bipartite Erdős-Rényi graph $G(n, n, \frac{d}{n})$ with a planted perfect matching. Applying Theorems 1 and 2 with $\mathcal{P} = \mathcal{Q}$, we conclude that the almost perfect recovery is possible if and only if $d \leq 1$. When $d = 1 + \epsilon$ for small ϵ , on the one hand, Theorem 1 shows that the average reconstruction error of MLE is at most $O(\epsilon)$. On the other hand, by slightly modifying the proof of Theorem 2, in Appendix F we show that the average reconstruction error is at least $\Omega(\epsilon^8)$. Determining the exact order of the phase transition for unweighted Erdős-Rényi model is an open problem.

1.2 Proof techniques

The proof of Theorem 1 is a simple application of large-deviation analysis and the union bound, similar to that of [21, Theorem 1]. The bulk of the paper is devoted to proving the negative results of Theorems 2 and 3, which are much more challenging.

Our starting point is the simple observation that to prove the impossibility of almost perfect recovery, it suffices to consider the random matching sampled from the posterior distribution, which reduces the problem to studying the typical behavior of this Gibbs distribution. We aim to show that with high probability, there is more posterior mass over the *bad* matchings (those far away from the hidden one) than that over the *good* matchings (those near the hidden one) in the posterior distribution. Via a first-moment calculation with proper truncation, it is not hard to bound from above the total posterior mass of good matchings. To bound from below the posterior mass of bad matchings, a key observation is that for a perfect matching M , the symmetric difference $M^* \Delta M$ consists of a disjoint union of even cycles which alternate between planted and unplanted edges. Therefore, a natural idea is to show the existence of many alternating cycles of length $\Theta(n)$ that are *augmenting*, that is, cycles for which the unplanted edges have a total log-likelihood that exceeds that of the planted edges. Unfortunately, a straightforward second-moment calculation fundamentally fails, due to the excessive correlations among long augmenting cycles. To construct the desired long augmenting cycles, we instead proceed in two steps: First, we construct many disjoint alternating paths of constant lengths; then we connect them to form exponentially many distinct augmenting cycles using the remaining edges via sprinkling. The first step is achieved by greedily exploring the local neighborhoods in analogy to a super-critical branching process, and the second step can be attained by reducing it to a problem of finding long cycles in a super-critical Erdős-Rényi bipartite graph with a planted perfect matching. This two-stage cycle finding scheme suffices to prove the sharp threshold in Theorem 2.

Although the above construction suffices for determining the sharp threshold, the local neighborhood exploration is too wasteful to extract sufficiently long alternating paths and falls short of proving the optimal reconstruction error bound in Theorem 3 for the exponential model. To resolve this inefficiency, in the first stage, following the program in [8], we use the truncated first and second moment methods to show the existence of many alternating paths that are sufficiently long, then applying Turán’s theorem to extract a large disjoint subcollection. Notably, we further impose extra *uniformity* constraints introduced in [7, 9] on the weights of the alternating paths to reduce the correlations in the second moment calculation, while at the same time keeping the first moment large.

In passing, we remark that our proof strategy of the impossibility results significantly deviates from most existing approaches in the literature. In many planted problems such as community detection in stochastic block models [6] or sparse PCA [17], the optimal overlap (one minus reconstruction error) exhibits a sharp transition from zero to strictly positive. Such correlated recovery threshold can be established via either mutual information arguments or reduction to detection (hypothesis testing) – see [30] for a survey; however, these techniques are either too loose or inapplicable for our model, where the optimal overlap undergoes a phase transition from strictly less than one to one.

1.3 Discussion and open problems

A recent line of work [16, 28] studies the planted matching problem in a geometric setting, where vertices correspond to random points (e.g. Gaussian) in \mathbb{R}^d , and edge weights are functions (e.g. Euclidean distance, inner product) of the coordinates of the end points. As such, the edge weights are correlated, in contrast to the independent edge weights assumed in this paper. The proof techniques for the positive results and the behavior of the reconstruction thresholds also differ significantly. On the other hand, the information-theoretic lower bound for bounded dimension d in [28, Theorem 4] is established using a similar but simpler argument to that of the present paper. It turns out that for geometric models it suffices to consider augmenting 4-cycles as opposed to long cycles. The correlation between these augmenting 4-cycles is mild enough so that a straightforward second-moment calculation followed by an application of Turán’s theorem suffices to extract $\Omega(n)$ vertex-disjoint augmenting 4-cycles, giving rise to exponentially many matchings that differ from the true one by $\Omega(n)$ vertices.

In this paper we show that the phase transition under the exponential model (resp. unweighted model) is of infinite order (resp. finite order), by characterizing the optimal reconstruction error near the sharp threshold. As mentioned in Remark 2, it is an open problem to determine whether such infinite-order phase transitions hold for weight densities other than exponential. Furthermore, for the exponential model, determining the value of the optimal limiting overlap as a function of $\lambda < 4$ is an interesting and challenging question, which likely requires going beyond the MLE (max-weight matching).

Finally, we briefly discuss the planted k -factor model recently studied in [26]. The special case of $k = 1$ is a variant of the planted matching model studied in this

paper. The conjectured threshold for almost perfect recovery is at $\sqrt{kd} B(\mathcal{P}, \mathcal{Q}) = 1$. The positive direction of this conjecture can be established by extending the proof of Theorem 1. Extending the impossibility results of Theorem 2 to the planted k -factor model is an interesting future direction.

1.4 Organization

The rest of the paper is organized as follows. In Sect. 2 we prove the positive result in Theorem 1 by analyzing the MLE. In Sect. 3, we outline the proof of the negative results in Theorems 2 and 3, which are the main results of this paper. In this section, we focus on the sparse model and the exponential model (with general dense model deferred till Appendix E.1). The negative results are proved by analyzing the posterior distribution, which relies on a cycle-finding scheme involving a path construction stage (detailed in Sects. 4 and 7 for sparse and exponential models, respectively) and a sprinkling stage (specified in Sect. 5). Combining results from Sects. 4 and 5, in Sect. 6 we prove the key Lemma 2 previously stated in Sect. 3, thereby finishing the proof of Theorem 2 under the sparse model. In Sect. 7, we give the details for the exponential model and complete the proof of the negative part of Theorem 3.

The appendix contains auxiliary technical results and postponed proofs. Appendix A presents a large deviation result used throughout the paper. Next, two auxiliary results crucial for proving the negative result under the exponential model (Theorem 3) are given: Appendix B contains Chernoff bounds for the Erlang distribution, and Appendix C recalls a technical lemma from [8] for controlling the deviation of Exp-minus-one random bridges. The positive part of Theorem 3 is proved in Appendix D. Appendix E contains all reduction-type arguments used in our proofs: In Appendix E.1, we prove Theorem 2 under the dense model by reducing it to the sparse model; in Appendix E.2 and Appendix E.3, we drop the absolute continuity conditions on distributions \mathcal{P}, \mathcal{Q} , which are assumed in the proof given in the main part of the paper. Appendix F proves the finite-order phase transition for unweighted graphs previously announced in Remark 2.

2 Positive results via maximum likelihood

To prove Theorem 1, we first consider a general case of complete graph ($d = n$) with arbitrary weight distributions, and then deduce the result for the general planted matching model by refining the weight distribution to incorporate edges not in G .

Theorem 4 *For the dense model with parameters $(n, \mathcal{P}, \mathcal{Q})$, suppose that*

$$\sqrt{n} B(\mathcal{P}, \mathcal{Q}) \leq 1 + \epsilon. \quad (13)$$

for some $\epsilon \geq 0$. Then there exists a universal constant $C > 0$, such that for large enough n ,

$$\mathbb{E}[\ell(M^*, \hat{M}_{\text{ML}})] \leq C \max \left\{ \log(1 + \epsilon), \sqrt{\frac{\log n}{n}} \right\}.$$

Proof of Theorem 4 Recall our standing assumption that $\mathcal{P} \ll \mathcal{Q}$. We give in Appendix E.2 a reduction-based argument that handles the case where \mathcal{P} is not absolutely continuous with respect to \mathcal{Q} . Under this assumption, the likelihood-ratio \mathcal{P}/\mathcal{Q} is well-defined and we can apply the large deviation result (83). We also assume WLOG that $\epsilon \leq 1$. For a fixed $M \in \mathcal{M}$ for which $M^* \Delta M$ contains $2t$ edges, we have

$$\mathbb{P} \left\{ \sum_{e \in M} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e) \geq \sum_{e \in M^*} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e) \right\} = \mathbb{P} \left\{ \sum_{i=1}^t (Y_i - X_i) \geq 0 \right\},$$

where X_i 's and Y_i 's be two independent sequences of random variables such that X_i 's are i.i.d. copies of $\log(\mathcal{P}/\mathcal{Q})$ under distribution \mathcal{P} and Y_i 's are i.i.d. copies of $\log(\mathcal{P}/\mathcal{Q})$ under distribution \mathcal{Q} . Using standard large-deviation estimates (see (83) in Appendix A), the RHS of the inequality above is upper bounded by $e^{-t\alpha}$, where

$$\alpha \triangleq -2 \log B(\mathcal{P}, \mathcal{Q}) \quad (14)$$

is the Rényi divergence of order $\frac{1}{2}$ between distributions \mathcal{P} and \mathcal{Q} . Since there are at most $\binom{n}{t} t! \leq n^t e^{-t(t-1)/2n}$ perfect matchings differing from the true matching M^* by $2t$ edges, it follows from a union bound and the assumption (13) that

$$\begin{aligned} \mathbb{P} \{ |M^* \Delta \hat{M}_{\text{ML}}| \geq \beta n \} &\leq \sum_{t \geq 2\beta n} \binom{n}{t} t! e^{-t\alpha} \leq e^{1/2} \sum_{t \geq \beta n} e^{-t\alpha + t \log n - t^2/(2n)} \\ &\stackrel{(a)}{\leq} e^{1/2} \sum_{t \geq \beta n} \left((1 + \epsilon)^2 e^{-\beta/2} \right)^t \\ &\stackrel{(b)}{\leq} e^{1/2} \frac{e^{-\beta^2 n/4}}{1 - e^{-\beta/4}}, \end{aligned} \quad (15)$$

where (a) follows from the assumption $\sqrt{n}B(\mathcal{P}, \mathcal{Q}) \leq 1 + \epsilon$; (b) holds for all strictly positive $\beta \geq 8 \log(1 + \epsilon)$ so that $(1 + \epsilon)^2 e^{-\beta/2} \leq e^{-\beta/4} < 1$. Thus,

$$\begin{aligned} \mathbb{E}[|M^* \Delta \hat{M}_{\text{ML}}|] &= \mathbb{E} \left[|M^* \Delta \hat{M}_{\text{ML}}| \mathbf{1}_{\{|M^* \Delta \hat{M}_{\text{ML}}| \leq 2\beta n\}} \right] + \mathbb{E} \left[|M^* \Delta \hat{M}_{\text{ML}}| \mathbf{1}_{\{|M^* \Delta \hat{M}_{\text{ML}}| > 2\beta n\}} \right] \\ &\leq 2\beta n + 2n \mathbb{P} \{ |M^* \Delta \hat{M}_{\text{ML}}| \geq \beta n \} \\ &\leq \beta 2n + 2e^{1/2} \frac{e^{-\beta^2 n/4}}{1 - e^{-\beta/4}} n. \end{aligned}$$

Next, by choosing

$$\beta = \max \left\{ 8 \log(1 + \epsilon), 2 \sqrt{\frac{\log n}{n}} \right\},$$

we have $e^{-\beta^2 n/4} \leq \beta^2/4$ and $1 - e^{-\beta/4} \geq \beta/8$. Therefore,

$$\begin{aligned}\mathbb{E}[\ell(M^*, \widehat{M}_{\text{ML}})] &= \frac{1}{n} \mathbb{E}[|M^* \Delta \widehat{M}_{\text{ML}}|] \leq 2\beta + 2e^{1/2} \frac{\beta^2/4}{\beta/8} \leq 9\beta \\ &\leq 18 \max \left\{ 4 \log(1 + \epsilon), \sqrt{\frac{\log n}{n}} \right\}.\end{aligned}$$

□

Theorem 1 follows from Theorem 4 as a corollary.

Proof of Theorem 1 To apply Theorem 4, let us first reformulate the planted matching model in Definition 1 in a more convenient form. Recall that the observed weight on the edge e is drawn independently from \mathcal{P} or \mathcal{Q} depending on whether e belongs to the planted matching M^* or not, where \mathcal{P} and \mathcal{Q} are arbitrary probability measures on some space \mathcal{X} . Let us use a special symbol $\star \notin \mathcal{X}$ to signify an edge e that is not in G and write $W_e = \star$. Therefore, for $e \in M^*$, W_e is drawn from \mathcal{P} independently; for each $e \in E(K_{n,n}) \setminus M^*$, with probability $1 - \frac{d}{n}$, $W_e = \star$; with probability $\frac{d}{n}$, W_e is drawn from \mathcal{Q} independently. In other words, $W_e \stackrel{\text{i.i.d.}}{\sim} \mathcal{Q}' \triangleq (1 - \frac{d}{n})\delta_\star + \frac{d}{n}\mathcal{Q}$ for $e \notin M^*$. The model is thus reformulated into a dense model with parameters $(n, \mathcal{P}, \mathcal{Q}')$.

By Theorem 4, \widehat{M}_{ML} achieves $\mathbb{E}[\ell(M^*, \widehat{M}_{\text{ML}})] \leq C \max \left\{ \log(1 + \epsilon), \sqrt{\frac{\log n}{n}} \right\}$, provided that $\sqrt{n}B(\mathcal{P}, \mathcal{Q}') \leq 1 + \epsilon$. Since $\sqrt{n}B(\mathcal{P}, \mathcal{Q}') = \sqrt{d}B(\mathcal{P}, \mathcal{Q})$, Theorem 1 readily follows. □

Remark 3 We note that when $\alpha - \log n \rightarrow +\infty$, from (15) we have $\widehat{M}_{\text{ML}} = M^*$ with high probability, i.e., \widehat{M}_{ML} achieves exact recovery. This coincides with the exact recovery threshold for the hidden Hamiltonian cycle problem [3], in which a (unipartite) weighted graph is observed such that the edge weights on the planted Hamiltonian cycle are drawn from \mathcal{P} and other weights are drawn from \mathcal{Q} .

3 Negative results via analyzing posterior distribution

In this section, we prove the impossibility results in Theorems 2 and 3 by directly analyzing the posterior distribution. For Theorem 2, we focus on the sparse model where the average degree d is a constant. The impossibility result under the dense model follows from a reduction argument (see Appendix E.1).

3.1 Proof outline of negative results

In this subsection we outline the proof of the impossibility results in Theorem 2 under the sparse model, and Theorem 3.

The negative results are proved by studying the posterior distribution of the hidden matching. Recall in the proof of Theorem 1 the reformulation of the planted matching

model in Definition 1 in the complete graph model with weight distributions $(\mathcal{P}, \mathcal{Q}')$. In particular, $W_e \stackrel{\text{i.i.d.}}{\sim} \mathcal{Q}' \triangleq (1 - \frac{d}{n})\delta_{\star} + \frac{d}{n}\mathcal{Q}$ for $e \notin M^*$. Here for simplicity we again assume $\mathcal{P} \ll \mathcal{Q}$ (the general case is handled in Appendix E.3). Under this assumption, $\mathcal{P} \ll \mathcal{Q}'$ and the density \mathcal{P}/\mathcal{Q}' is well-defined. The likelihood function of $W = (W_e : e \in E(K_{n,n}))$ given $M^* = m$ is given by

$$\mathbb{P}\{W \mid M^* = m\} = \prod_{e \in m} \mathcal{P}(W_e) \prod_{e \notin m} \mathcal{Q}'(W_e) \propto \prod_{e \in m} \frac{\mathcal{P}}{\mathcal{Q}'}(W_e) = \exp \left(\underbrace{\sum_{e \in m} \log \frac{\mathcal{P}}{\mathcal{Q}'}(W_e)}_{\triangleq L(m)} \right), \quad (16)$$

where \propto denotes equality up to multiplicative factors that do not depend on m , $\log \frac{\mathcal{P}}{\mathcal{Q}'}$ takes extended real-values in $\mathbb{R} \cup \{-\infty\}$, and $L(m) = \sum_{e \in m} \log \frac{\mathcal{P}}{\mathcal{Q}'}(W_e)$ is the total log-likelihood ratio on a set m of edges. Thus, conditioned on W , the posterior distribution of M^* is a Gibbs distribution, given by

$$\mu_W(m) = \frac{1}{Z(W)} \exp(L(m)), \quad m \in \mathcal{M}, \quad (17)$$

where $Z(W) = \sum_{m \in \mathcal{M}} \exp(L(m))$ is the normalization factor.

In order to reduce the impossibility proof to a statement on the posterior distribution, the first observation is that it suffices to consider the estimator \tilde{M} which is sampled from the posterior distribution μ_W . Indeed, given any estimator $\hat{M} = \hat{M}(W)$, we have $(\hat{M}, M^*) \stackrel{L}{=} (\tilde{M}, \tilde{M})$ (in distribution) and hence

$$\mathbb{E}|\tilde{M} \Delta M^*| \leq \mathbb{E}|\tilde{M} \Delta \hat{M}| + \mathbb{E}|M^* \Delta \hat{M}| = 2\mathbb{E}|M^* \Delta \hat{M}|,$$

which shows that \tilde{M} is optimal within a factor of two. Thus it suffices to bound $\mathbb{E}|\tilde{M} \Delta M^*|$ from below. To this end, fix some δ to be specified later and define the sets of good and bad solutions respectively as

$$\begin{aligned} \mathcal{M}_{\text{good}} &= \{M \in \mathcal{M} : \ell(M, M^*) < 2\delta\} \\ \mathcal{M}_{\text{bad}} &= \{M \in \mathcal{M} : \ell(M, M^*) \geq 2\delta\}. \end{aligned}$$

By the definition of \tilde{M} , we have

$$\mathbb{E}[\ell(\tilde{M}, M^*)] \geq 2\delta \cdot \mathbb{E}[\mu_W(\mathcal{M}_{\text{bad}})].$$

Next we show

Lemma 1 *Assume that (7) holds for some arbitrary constant $\epsilon > 0$. There exist constants $n_0 = n_0(\epsilon)$ and $c = c(\delta, \epsilon)$, such that for all $n \geq n_0$, in both the sparse and the dense model, with probability at least $1 - e^{-cn/\log n}$,*

$$\frac{\mu_W(\mathcal{M}_{\text{good}})}{\mu_W(M^*)} \leq 2e^{c_1 n}, \quad (18)$$

where $c_1 = 7\epsilon\delta$.

Lemma 2 Consider the sparse model such that (7) holds for some arbitrary constant $\epsilon > 0$. There exist constants c_0, c, c_2, n_0 that only depend on $\epsilon, \mathcal{P}, \mathcal{Q}$, such that for all $n \geq n_0$, and $\delta \leq c_2$, with probability at least $1 - \frac{c}{n}$,

$$\frac{\mu_W(\mathcal{M}_{\text{bad}})}{\mu_W(M^*)} \geq e^{c_0 n}. \quad (19)$$

Furthermore, in the exponential model with $d = n$ and $\mathcal{P} = \exp(\lambda)$ and $\mathcal{Q} = \exp(1/n)$ where $\lambda \leq 4 - \epsilon$ for some arbitrary constant $\epsilon > 0$, (19) continues to hold with probability at least $\frac{1}{2} - \frac{c}{n}$, where all constants c_0, c, c_2, n_0 depend only on ϵ and $c_0, c_2 = e^{-O(1/\sqrt{\epsilon})}$.

Given the above two lemmas, Theorems 2 and 3 readily follow. Indeed, combining Lemma 1 and Lemma 2 and choosing $\delta = \min\{c_2, c_0/(14\epsilon)\}$ yields that $c_0 \geq 2c_1$ and hence $\mu_W(\mathcal{M}_{\text{bad}}) \geq \frac{e^{c_1 n}}{2 + e^{c_1 n}}$ with probability at least $\frac{1}{2} - o(1)$, which shows that $\mathbb{E}[\ell(\tilde{M}, M^*)] \gtrsim \delta$ as desired in both the sparse and the exponential model on complete graph.

We prove Lemma 1 in Sect. 3.2, and outline the proof of Lemma 2 in Sect. 3.3. For the rest of the proof, we will assume WLOG⁴ that (7) holds with equality. That is,

$$\sqrt{d} B(\mathcal{P}, \mathcal{Q}) = 1 + \epsilon, \quad (20)$$

for some arbitrary constant $\epsilon > 0$.

3.2 Upper bounding the posterior mass of good matchings

In this section, we prove Lemma 1. Before proceeding with the proof, let us first introduce some notation that will be used throughout the remainder of the paper. Fix the true matching M^* . Without loss of generality, we assume that M^* corresponds to the identity matching, so that each left vertex i is uniquely matched to i' on the right. We shall represent the planted edges (those in M^*) and the unplanted edges (those outside M^*) as red and blue edges, respectively. A cycle in $K_{n \times n}$ is called *alternating* if it is an even cycle and alternates between red and blue edges. An important observation is that the symmetric difference between the truth M^* and another perfect matching M is always a disjoint union of alternating cycles. See Fig. 1 for an example.

For any set T of edges in $K_{n,n}$, let $\mathbf{r}(T) \triangleq T \cap M^*$ and $\mathbf{b}(T) \triangleq T \setminus M^*$ denote the set of red and blue edges in T , respectively. Define the *excess weight* of T as the total log-likelihood of blue edges minus that of the red edges:

⁴ Indeed, suppose that $\sqrt{d} B(\mathcal{P}, \mathcal{Q}) = c > 1 + \epsilon$. Consider the model parametrized with $(d', \mathcal{P}, \mathcal{Q})$ where $d' = d(1 + \epsilon)^2/c^2 < d$ so that $\sqrt{d'} B(\mathcal{P}, \mathcal{Q}) = 1 + \epsilon$. From an observed graph G generated from the $(d', \mathcal{P}, \mathcal{Q})$ model, one can “densify” G by adding edges independently with edge weight drawn from \mathcal{Q} to arrive at an instance of the $(d, \mathcal{P}, \mathcal{Q})$ model. Therefore, the lower bound on the average reconstruction error carries over to the $(d, \mathcal{P}, \mathcal{Q})$ model.

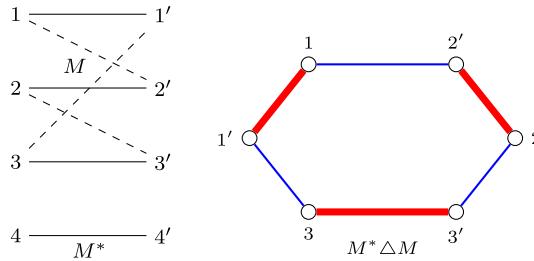


Fig. 1 Example of matching and alternating cycle ($n = 4$, $\ell = 3$). The matching M^* and M matches 1234 to $1'2'3'4'$ and $2'3'1'4'$, respectively. The difference of their corresponding permutation is a cyclic shift of length three, i.e., $\pi^{*-1}\pi' = (123)$ in the notation of cycle decomposition, and the symmetric difference graph $M^* \Delta M$ is an alternating 6-cycle that contains 3 planted (red, thick) edges and 3 unplanted (blue, thin) edges (color figure online)

$$\Delta(T) = \sum_{e \in b(T)} \log \frac{\mathcal{P}}{\mathcal{Q}'}(W_e) - \sum_{e \in r(T)} \log \frac{\mathcal{P}}{\mathcal{Q}'}(W_e). \quad (21)$$

For a perfect matching $M \in \mathcal{M}$, we will denote $\Delta(M \Delta M^*)$ simply as $\Delta(M)$. Recall the posterior distribution of M^* given W in (17). Therefore, if $M^* = m^*$, we have

$$\Delta(m) = \sum_{e \in m \setminus m^*} \log \frac{\mathcal{P}}{\mathcal{Q}'}(W_e) - \sum_{e \in m^* \setminus m} \log \frac{\mathcal{P}}{\mathcal{Q}'}(W_e) = L(m) - L(m^*) = \log \frac{\mu_W(m)}{\mu_W(m^*)}.$$

Note that in the display above, the likelihood ratio \mathcal{P}/\mathcal{Q}' can be replaced by \mathcal{P}/\mathcal{Q} if m only contains edges in G . To see this, note the fact that $|m \setminus m^*| = |m^* \setminus m| = \ell$ and $\log \frac{\mathcal{P}}{\mathcal{Q}'}(w) = \log \frac{\mathcal{P}}{\mathcal{Q}}(w) + \log \frac{n}{d}$ whenever $w \neq \star$. Moreover, if $W_e = \star$ for some $e \in m$ then $\mu_W(m) = 0$ and $\Delta(m) = -\infty$. Thus, we have

$$\Delta(m) = \begin{cases} \sum_{e \in m \setminus m^*} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e) - \sum_{e \in m^* \setminus m} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e) & W_e \neq \star, \forall e \in m \setminus m^* \\ -\infty & \text{else.} \end{cases} \quad (22)$$

We have the following lemma:

Lemma 3 *Let $m \in \mathcal{M}$ be such that $|m \Delta m^*| = 2\ell$. For each $x \geq 0$,*

$$\mathbb{P} \{ \Delta(m) \geq x\ell \mid M^* = m^* \} \leq \left(\frac{d}{n} e^{-(\alpha + x/2)} \right)^\ell.$$

Proof From (22), we have

$$\mathbb{P} \{ \Delta(m) \geq x\ell \mid M^* = m^* \}$$

$$= \mathbb{P} \left\{ \sum_{e \in m \setminus m^*} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e) - \sum_{e \in m^* \setminus m} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e) \geq x\ell \mid M^* = m^*, W_e \neq \star, \forall e \in m \right\} \left(\frac{d}{n} \right)^\ell$$

$$\begin{aligned}
&\stackrel{(a)}{=} \mathbb{P} \left\{ \sum_{e \in m \setminus m^*} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e) - \sum_{e \in m^* \setminus m} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e) \geq x\ell \mid M^* = m^*, W_e \neq \star, \forall e \in m \right\} \left(\frac{d}{n} \right)^\ell \\
&\stackrel{(b)}{=} \mathbb{P} \left\{ \sum_{i=1}^{\ell} (Y_i - X_i) \geq x\ell \right\} \left(\frac{d}{n} \right)^\ell \stackrel{(c)}{\leq} e^{-(\alpha+x/2)\ell} \left(\frac{d}{n} \right)^\ell,
\end{aligned}$$

where (a) follows from (22); in (b) we write X_i and Y_i are iid copies of $\log \frac{\mathcal{P}}{\mathcal{Q}}$ under \mathcal{P} and \mathcal{Q} respectively; (c) follows from the large deviation bound (83) in Appendix A. \square

We now proceed to the proof of (18). Recall that our standing assumption from (20) is that

$$\sqrt{d} B(\mathcal{P}, \mathcal{Q}) = 1 + \epsilon,$$

for some small ϵ . Throughout the proof we condition on $M^* = m^*$.

Note that

$$\frac{\mu_W(\mathcal{M}_{\text{good}})}{\mu_W(M^*)} = \sum_{m: |m \Delta m^*| < 2\delta n} e^{\Delta(m)} = R_1 + R_2,$$

where

$$\begin{aligned}
R_1 &\triangleq \sum_{m: |m \Delta m^*| < 2\beta n / \log n} e^{\Delta(m)} \\
R_2 &\triangleq \sum_{m: \frac{2\beta n}{\log n} \leq |m \Delta m^*| < 2\delta n} e^{\Delta(m)}
\end{aligned}$$

for some β to be specified. Next we bound R_1 and R_2 separately.

First, we note that given a perfect matching m , the number of matchings that differ from m by 2ℓ edges is

$$|\{m \in \mathcal{M} : |m \Delta m^*| = 2\ell\}| = !\ell \cdot \binom{n}{\ell}, \quad (23)$$

where $!l$ denotes the number of derangements of ℓ elements, given by

$$!l = l! \sum_{i=0}^{\ell} \frac{(-1)^i}{i!} = \left[\frac{\ell!}{e} \right],$$

and $[\cdot]$ denotes rounding to the nearest integer. Thus

$$\frac{1}{2e} n(n-1) \cdots (n-\ell+1) \leq |\{m \in \mathcal{M} : |m \Delta m^*| = 2\ell\}| \leq \frac{2}{e} n(n-1) \cdots (n-\ell+1). \quad (24)$$

Furthermore, for any m ,

$$\mathbb{E}[\exp(\Delta(m))] = 1. \quad (25)$$

This follows from the definition of log likelihood ratio as

$$\mathbb{E}[\exp(\Delta(m))] = \underbrace{\prod_{e \in m \setminus m^*} \mathbb{E}_{W_e \sim \mathcal{Q}'} \left[\frac{\mathcal{P}}{\mathcal{Q}'}(W_e) \right]}_{=1} \underbrace{\prod_{e \in m^* \setminus m} \mathbb{E}_{W_e \sim \mathcal{P}} \left[\frac{\mathcal{Q}'}{\mathcal{P}}(W_e) \right]}_{=1} = 1.$$

To bound R_1 , using (24) and (25) we have

$$\mathbb{E}[R_1] = \sum_{|m \Delta m^*| < \frac{2\beta n}{\log n}} \mathbb{E}[e^{\Delta(m)}] \leq \sum_{\ell < \frac{\beta n}{\log n}} \frac{2}{e} n^\ell \leq \frac{2n}{e} \exp(\beta n).$$

By Markov's inequality,

$$\mathbb{P}\{R_1 \geq e^{2\beta n}\} \leq \frac{2n}{e} \exp(-\beta n). \quad (26)$$

To bound R_2 , the calculation above shows that directly applying the Markov inequality is too crude since $\mathbb{E}[R_2] = e^{\Theta(n \log n)}$. Note that although $\Delta(m)$ is negatively biased, when $\Delta(m)$ is atypically large it results in an excessive contribution to the exponential moments. Thus we truncate on the following event:

$$\mathcal{E}_r \triangleq \bigcap_{m: 2\beta n / \log n \leq |m \Delta m^*| < 2\delta n} \{\Delta(m) \leq r|m \Delta m^*|/2\}$$

for some constant r to be chosen. Then

$$\begin{aligned} & \mathbb{P}\{R_2 \geq e^{c'n}\} \\ & \leq \mathbb{P}\{\mathcal{E}_r^c\} + \mathbb{P}\{R_2 \geq e^{c'n} \cap \mathcal{E}_r\} \\ & \leq \mathbb{P}\{\mathcal{E}_r^c\} + \mathbb{P}\left\{\sum_{\frac{2\beta n}{\log n} \leq |m \Delta m^*| < 2\delta n} e^{\Delta(m)} \mathbf{1}_{\{\Delta(m) \leq r|m \Delta m^*|/2\}} \geq e^{c'n}\right\} \\ & \leq \mathbb{P}\{\mathcal{E}_r^c\} + e^{-c'n} \sum_{\frac{2\beta n}{\log n} \leq |m \Delta m^*| < 2\delta n} \mathbb{E}[e^{\Delta(m)} \mathbf{1}_{\{\Delta(m) \leq r|m \Delta m^*|/2\}}]. \end{aligned} \quad (27)$$

To bound the first term, note the fact that if $|m \Delta m^*| = 2\ell$, by Lemma 3, we have

$$\mathbb{P}\{\Delta(m) \geq x\ell\} \leq \left(\frac{d}{n} e^{-(\alpha+x/2)}\right)^\ell,$$

Therefore, it follows from a union bound and the fact that $\alpha = -2 \log \int \sqrt{\mathcal{P}\mathcal{Q}}$ that

$$\begin{aligned}\mathbb{P}\{\mathcal{E}_r^c\} &\leq \sum_{\frac{2\beta n}{\log n} \leq |m\Delta m^*| < 2\delta n} \mathbb{P}\{\Delta(m) \geq r|m\Delta m^*|/2\} \\ &\leq \frac{2}{e} \sum_{\frac{\beta n}{\log n} \leq \ell < \delta n} n^\ell \left(\frac{d}{n} \left(\int \sqrt{\mathcal{P}\mathcal{Q}}\right)^2\right)^\ell e^{-r\ell/2} \\ &= \frac{2}{e} \sum_{\frac{\beta n}{\log n} \leq \ell < \delta n} \left((1+\epsilon)^2 e^{-r/2}\right)^\ell\end{aligned}$$

Choose $r = 8\epsilon$. We have

$$\mathbb{P}\{\mathcal{E}_r^c\} \leq e^{-c\beta n/\log n} \quad (28)$$

for some $c = c(\epsilon)$ and all sufficiently large n .

For the second term in (27), we bound the truncated MGF as follows:

$$\begin{aligned}\mathbb{E}\left[e^{\Delta(m)} \mathbf{1}_{\{\Delta(m) \leq r\ell\}}\right] &\leq \mathbb{E}\left[\exp\left(\frac{1}{2}(\Delta(m) + r\ell)\right)\right] \\ &= \mathbb{E}\left[e^{\Delta(m)/2} \mid W_e \neq \star, \forall e \in m \setminus m^*\right] \left(\frac{d}{n}\right)^\ell e^{r\ell/2} \\ &\stackrel{(a)}{=} \prod_{e \in m \setminus m^*} \mathbb{E}_{W_e \sim \mathcal{Q}} \left[\sqrt{\frac{\mathcal{P}}{\mathcal{Q}}(W_e)} \right] \prod_{e \in m^* \setminus m} \mathbb{E}_{W_e \sim \mathcal{P}} \left[\sqrt{\frac{\mathcal{Q}}{\mathcal{P}}(W_e)} \right] \left(\frac{d}{n}\right)^\ell e^{r\ell/2} \\ &= \left(\frac{d}{n} B(\mathcal{P}, \mathcal{Q})^2\right)^\ell e^{r\ell/2} \leq \left(\frac{e^{6\epsilon}}{n}\right)^\ell\end{aligned}$$

where (a) follows from (22). Combining the above with (24), we have

$$\sum_{\frac{2\beta n}{\log n} \leq |m\Delta m^*| < 2\delta n} \mathbb{E}\left[e^{\Delta(m)} \mathbf{1}_{\{\Delta(m) \leq r|m\Delta m^*|/2\}}\right] \leq \frac{2}{e} \delta n e^{6\epsilon \delta n}$$

Choosing $c' = 7\epsilon\delta$, we get that

$$e^{-c'n} \sum_{\frac{2\beta n}{\log n} \leq |m\Delta m^*| < 2\delta n} \mathbb{E}\left[e^{\Delta(m)} \mathbf{1}_{\{\Delta(m) \leq r|m\Delta m^*|/2\}}\right] \leq e^{-c\delta n}. \quad (29)$$

for some $c = c(\epsilon)$ and all large n . Substituting (28) and (29) into (27), we get

$$\mathbb{P}\left\{R_2 \geq e^{7\epsilon\delta n}\right\} \leq e^{-\Omega(n/\log n)}$$

Combining this with (26) and upon choosing $\beta = \epsilon\delta$, we have $\mathbb{P}\{R_1 + R_2 \geq 2e^{7\epsilon\delta n}\} \leq e^{-\Omega(n/\log n)}$, concluding the proof.

3.3 Lower bounding the posterior mass of bad matchings

In this section, we outline the proof of Lemma 2. Recall that

$$\frac{\mu_W(\mathcal{M}_{\text{bad}})}{\mu_W(M^*)} = \sum_{m: |m \Delta m^*| \geq 2\delta n} e^{\Delta(m)}.$$

For any perfect matching m such that $|m \Delta m^*| = 2\ell$, the set difference $m \Delta m^*$ can be represented by a disjoint union of *alternating* cycles, denoted by C , where the edges in $m^* \setminus m$ and $m \setminus m^*$ are colored red and blue respectively, so that in total there are ℓ red edges and ℓ blue edges.

Recall from (21) that the excess weight $\Delta(C)$ of a cycle C denotes the difference between the total blue and red edge (log-likelihood) weights. In order to lower bound the posterior mass of bad matchings, we show that with probability at least $1 - c_1/n$ under sparse model ($1/2 - c_1/n$ under exponential model), there exist at least e^{nc_2} distinct alternating cycles C of length at least nc_3 , so that $\Delta(C) \geq nc_4$, for some constants c_1, c_2, c_3, c_4 that are independent of n and only depend on $\epsilon, \mathcal{P}, \mathcal{Q}$. Note that $c_1, c_2, c_3 > 0$, while c_4 is non-negative and is 0 when $\mathcal{P} = \mathcal{Q}$. Since each alternating cycle C corresponds to a perfect matching in \mathcal{M}_{bad} , we have

$$\frac{\mu_W(\mathcal{M}_{\text{bad}})}{\mu_W(M^*)} \geq e^{n(c_2+c_4)}. \quad (30)$$

Let us point out a simple yet useful observation: if a perfect matching m contains any edges not in G (e with $W_e = \star$), then it has zero posterior mass. Thus, all alternating cycles that give rise to a perfect matching in \mathcal{M}_{bad} with positive posterior mass must consist only of edges in G . Therefore, to show (30), it suffices to consider the alternating cycles in G . Let us also remark that by the same reasoning as (22), for any set T of edges in G such that $|\mathbf{b}(T)| = |\mathbf{r}(T)|$,

$$\Delta(T) = \sum_{e \in \mathbf{b}(T)} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e) - \sum_{e \in \mathbf{r}(T)} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e).$$

Since for all m on G , $m \Delta m^*$ contains the same number of red and blue edges, for the remainder of the proof, we will work directly with \mathcal{P}/\mathcal{Q} instead of \mathcal{P}/\mathcal{Q}' .

To show the existence of many alternating cycles in G , one natural idea is to define S as the set of alternating cycles C satisfying the aforementioned length and weight requirements; and then bound the cardinality of S from below using the first and second-moment methods. This boils down to proving that $\text{Var}(|S|) \lesssim (\mathbb{E}[|S|])^2$. Unfortunately, this idea fails because the variance of $|S|$ turns out to be exponentially larger than $(\mathbb{E}[|S|])^2$, due to the excessive correlations among these long alternating

cycles. A similar phenomenon is also observed in counting the number of long cycles in Erdős-Rényi graphs [20, Section 3].

Alternatively, recall that the set difference $m \Delta m^*$ is allowed to be a *disjoint union* of alternating cycles rather than a *single* alternating cycle. Thus, one can resort to showing the existence of disjoint unions of many but short alternating cycles satisfying the total length and weight requirements. However, using the first and second moment methods, at best we can show that there exists a disjoint union of $\Theta(n/\log^2 n)$ desired alternating cycles of length $\Theta(\log n)$; thus the total length is only $\Theta(n/\log n)$, falling short of meeting the total length requirement of $\Omega(n)$.

To construct the desired alternating cycles, we instead proceed in two steps. We first reserve a set of vertices and construct many but short alternating paths with the desired total weight in the subgraph induced by the non-reserved vertices. Then we use the edges incident to the reserved vertices to connect these paths to form the desired alternating cycles. At a high level, our two-stage cycle finding scheme is inspired by the previous work [8] in a different context: The goal therein is to find a long path whose average weight is below a certain threshold in a complete graph with i.i.d. exponentially weighted edges. More broadly, our second step is similar in spirit to the sprinkling idea commonly used in random graph theory (see, e.g. [2, Section 11.9] and [15, Section 2.1.4]).

The specific construction is described in Algorithm 1. Therein, V denotes the set of reserved left vertices from $[n]$, where $|V| = \gamma n$ and $\gamma = \gamma_0 \epsilon$ for some small constant $\gamma_0 > 0$. Denote by $V' \subset [n]'$ its counterpart on the right side defined by the red edges. We write $V^c = [n] \setminus V$ and therefore $(V^c)' = [n]' \setminus V'$. Recall that we have assumed WLOG that M^* is the identity matching that matches each i to i' . As such, the cycle-finding algorithm also assumes knowledge of the true matching M^* .

Algorithm 1 Two-stage cycle finding algorithm

Input: Weighted bipartite graph G on $[n] \times [n]'$ with weight vector w , thresholds τ_{red} , τ_{blue} , and parameters $\gamma, s, \Delta_0, c_5, c_6, c_7$.

- Step 1: Path construction.** Reserve an arbitrary set of vertices V of size γn , e.g. by taking $V = [\gamma n]$. Let G_1 denote the weighted subgraph of G induced by the edges in $V^c \times (V^c)'$. Construct a family of disjoint sets $L_k \subset V^c$ of left vertices and $R_k \subset (V^c)'$ of right vertices for $k \in \mathcal{K}_1 \subset [K]$ such that: (1) $|L_k| = |R_k| = s$ and $K_1 \triangleq |\mathcal{K}_1| \geq nc_5$; (2) each pair of vertices $u \in L_k$ and $v' \in R_k$ are connected via an alternating path P that starts and ends in red edges, with $\Delta(P) \geq \Delta_0$, where $\Delta(P)$ is defined in (21).
- Step 2: Sprinkling.** Let $V^* = \{i \in V : \log(\mathcal{P}/\mathcal{Q})(W_{i,i'}) \leq \tau_{\text{red}}\}$. Let G_2 be the subgraph of G that contains every red edge in $V^* \times (V^*)'$ and every blue edge e in $V^* \times (V^*)', V^c \times (V^*)'$, or $V^* \times (V^c)'$, if and only if $\log(\mathcal{P}/\mathcal{Q})(W_e) \geq \tau_{\text{blue}}$. Let $\{U_k : k \in \mathcal{K}_1\}$ (resp. $\{V'_k : k \in \mathcal{K}_1\}$) be a collection of disjoint subsets of left (resp. right) vertices, such that every vertex in U'_k is connected to L_k by at least one blue edge in G_2 , and every vertex in V_k is connected to R_k by at least one blue edge in G_2 . Next, on a subset $\mathcal{K}_2 \subset \mathcal{K}_1$ of size $K_2 \geq K_1/16$, define a bipartite “super graph” G_{super} with vertex sets \mathcal{K}_2 and \mathcal{K}'_2 . In G_{super} , there is a red edge between k and k' for every $k \in \mathcal{K}_2$, and a blue edge between i and j' if and only if there is at least a blue edge in G_2 connecting U_i and V'_j . Construct $e^{c_6 K_1}$ distinct alternating cycles in G_{super} , each of length at least $2c_7 K_1$ for some universal constants $c_6, c_7 > 0$.

Output: Expand each alternating cycle on G_{super} into an alternating cycle on G , by replacing each red edge (k, k') in G_{super} by an alternating path between U_k and V'_k that starts and ends in red edges and connects L_k and R_k . See Fig. 2 for an illustration. Output all the resulting alternating cycles on G .

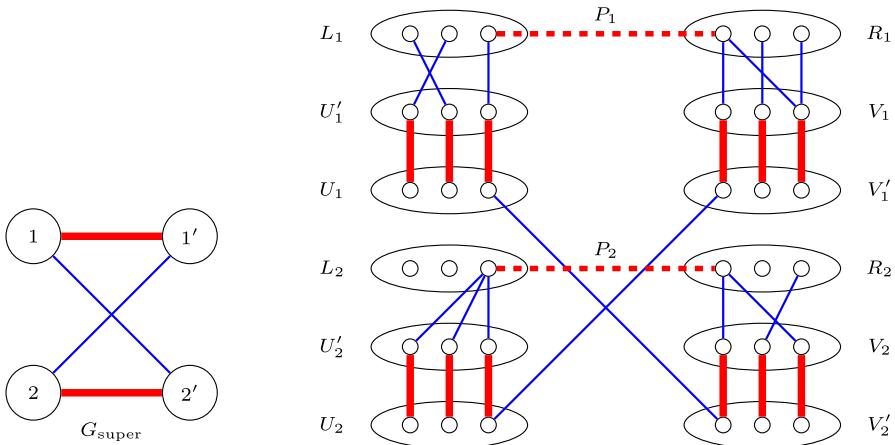


Fig. 2 Example of the sprinkling stage of Algorithm 1. To the left is an alternating cycle $C_{\text{super}} = (1, 1', 2, 2')$ on the super graph G_{super} . To the right is its expansion to an alternating cycle on G . The two blue edges in G_{super} correspond to the two long crossed blue edges in G : the blue edge $(1, 2')$ in G_{super} means that there exist $u_1 \in U_1, v'_2 \in V'_2$ such that (u_1, v'_2) is a blue edge in G_2 . Similarly, (u_2, v'_1) is a blue edge for some $u_2 \in U_2$ and $v'_1 \in V'_1$. The dashed red edges represent alternating paths P_1 and P_2 that start and end in red edges. The existence of P_1 and P_2 follows from the first stage of Algorithm 1 (color figure online)

We now explain the main idea of Algorithm 1. After disjoint alternating paths P_1, \dots, P_K are found in step 1 (path construction), step 2 (sprinkling) connects a subset of these paths into an alternating cycle. As illustrated in Fig. 2, in the sprinkling step we connect one endpoint of P_i (in L_i) and another endpoint of P_j (in R_j) via an alternating path that goes through the sets $L_i, U'_i, U_i, V'_j, V_j, R_j$ in this order (see Fig. 2). The definition of U'_i (resp. V_j) ensures the existence of blue edges from L_i to every vertex in U'_i (resp. from every vertex in V_j to R_j). To study the connectivity from U_i to V'_j , we introduce the notion of the super graph G_{super} , which is the key step in Algorithm 1. Whenever there exists a blue edge from i to j' in G_{super} , there exists a blue edge from U_i to V'_j in G , which completes the alternating path between P_i and P_j . In this way, each alternating cycle on G_{super} corresponds to a distinct alternating cycle on G .

We then argue that there are exponentially many alternating cycles on G_{super} by showing that it is very well-connected. First, we show that the blue edges in G_{super} are independent. To see this, note that in the construction of L_k, R_k, U_k, V_k , only the edges in $V^c \times (V^c)', V \times (V^c)'$ and $V^c \times V'$ are inspected. Therefore, whether a blue edge appears in the super graph is independent of all the steps prior, since it only depends on edges in $V \times V'$. Furthermore, the sets $\{U_k\}$ and $\{V_k\}$ are disjoint. Thus the blue edges in the super graph are all formed independently. Therefore, G_{super} can be viewed as (or at least contains) an Erdős-Renyi bipartite graph with a planted perfect matching. We further show that G_{super} contains exponentially many alternating cycles as long as we are in the very super-critical regime, i.e. the average degree of G_{super} is a constant much larger than one. To that end, in the path construction stage, we require that $\{L_k\}, \{R_k\}$ to contain sufficiently many vertices so that the sets $\{U_i\}$ and $\{V'_j\}$ are sufficient large, leading to the super-criticality of G_{super} .

In addition to finding exponentially many alternating cycles, we also need that they correspond to matchings with sufficiently large posterior probability. In other words, we need a lower bound on the total weight (in log-likelihood) $\Delta(C)$ for each alternating cycle C . This task is accomplished as follows: In the path construction stage, each path P_k satisfies the constraint $\Delta(P_k) \geq \Delta_0$. We then choose the thresholds τ_{red} small enough and τ_{blue} large enough so that the contributions from the sprinkling edge weights do not reduce Δ_0 by too much. Note that τ_{red} and τ_{blue} also need to be chosen carefully so as to not compromise the super-criticality of G_{super} .

In summary, we show that Algorithm 1 finds $e^{\mathcal{Q}(K_2)} = e^{\mathcal{Q}(n)}$ many alternating cycles each of length $\mathcal{Q}(n)$. Each cycle C is formed by connecting paths P_1, \dots, P_r , and we have $\Delta(C) \approx \sum_{k \leq r} \Delta(P_k) = \mathcal{Q}(n)$, resulting in the desired (30). In Sects. 4 and 5, we spell out the details of the path construction and the sprinkling respectively under the sparse model. In Sect. 6, we specify the choices of the parameters including τ_{red} , τ_{blue} , and combine the results in Sects. 4 and 5 to prove Lemma 2 under the sparse model. The analysis for the exponential model is contained in Sect. 7.

We remark that the path construction schemes actually differ for the sparse model (Theorem 2) and the exponential model (Theorem 3). Specifically, for the sparse model with bounded average degree, the desired alternating paths are found by exploring the local neighborhood in G_1 using breadth-first-search and constructing two-sided trees T_k where L_k and R_k correspond to the set of leaves of the left- and right-sided tree, respectively. This suffices to prove the sharp threshold in Theorem 2 for the sparse model (and, by a reduction argument, for the dense model as well). However, this path finding scheme via constructing two-sided trees is wasteful in the sense that the length of the path extracted is only about the depth of the tree, which is much smaller than the size of the tree; thus the obtained cycles are not long enough to yield the optimal reconstruction error bound in Theorem 3 for the exponential model. To remedy this, we take a more direct approach: Following the program in [8], we use first and second moment methods combined with Turán's theorem to show the existence of many disjoint, short alternating paths in G_1 of desired total weights. It is worth noting that this method results in fewer paths in total and a smaller super graph. Nevertheless, for the exponential model we can appropriately choose a large enough threshold for the blue edge weights in G_2 in the sprinkling stage, so that the super graph is still very supercritical (cf. (61)). In contrast, under the sparse model with bounded average degree, we no longer have the freedom to truncate the blue edge weights at a high enough threshold to boost the blue edge probability in G_2 and hence cannot recoup of the loss on the number of paths incurred by applying Turán's theorem. That is why we need to resort to neighborhood-exploration to find paths in the sparse model.

4 Path construction (under the sparse model)

In this section we construct K disjoint subgraphs on G_1 , which contain the disjoint alternating short paths we will later use to form the long cycles. In particular, each subgraph consists of two trees on G_1 whose root nodes are connected via a red edge. We will refer to the subgraphs as *two-sided trees* denoted as T_1, \dots, T_K . See Fig. 3 for an illustration. For each k , our construction is such that on either side of T_k , the

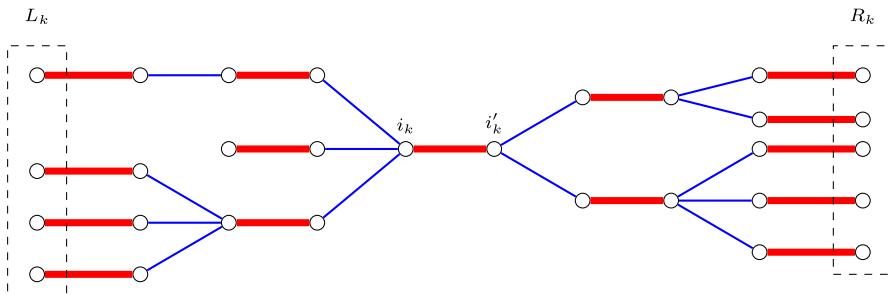


Fig. 3 Example of a two-sided tree T_k . For each $u \in L_k$ and $v \in R_k$, there is an alternating path from u to v that starts and ends in red edges, and passes through the red edge (i_k, i'_k) (color figure online)

path from any leaf node to its root node is of alternating color that starts in a red edge and ends in a blue edge. As such, the path from any leaf node on one side to any leaf node on the other side is also an alternating path. The two-sided trees are constructed via a greedy neighborhood exploration process. Before elaborating on the exploration process, the following are a few desirable features that we aim for.

1. The trees are not too small: on either side, we want the exploration process to survive long enough to yield sufficiently long paths;
2. There is a large number of leaf nodes on each side of the trees, so that there are plenty of paths to choose from when forming the long cycles via sprinkling.
3. The trees are not “overgrown”: since we need the K trees to be vertex-disjoint, none of them should have too many vertices;
4. The edge weights on all the alternating paths are well behaved. The definition of “well behaved” will become clear later. This is to ensure that the long cycles constructed from these paths occupy sufficiently large posterior probability.

In the remainder of this section, we give the precise construction of the two-sided trees, and show that sufficiently many of them fulfill all the features above.

4.1 Construction of two-sided trees

As shown in Fig. 3, each two-sided tree is centered at a red edge. We will refer to the subtree rooted at its left (resp. right) end point as the left (resp. right) subtree of this two-sided tree. We construct each left or right subtree via a neighborhood exploration process that starts from the root vertex. To ensure that the paths are alternating in color, each explored blue edge must be followed by (the unique) red edge (see Fig. 3). In other words, the vertices are always explored in *pairs*, where the pairs are identified by the red edges. Recall that G_1 is a bipartite graph with $(1 - \gamma)n$ left and right vertices. We will have γn pairs of vertices actively participate in the construction, and leave $(1 - 2\gamma)n$ pairs unexplored in the neighborhood exploration process.

The local graph neighborhoods are explored in a fashion analogous to the breadth-first search (BFS). However, as opposed to the vanilla BFS algorithm, we design a two-stage exploration-selection scheme, where the neighborhoods (left and right subtrees) are grown in epochs, where each epoch contains a few exploration steps and

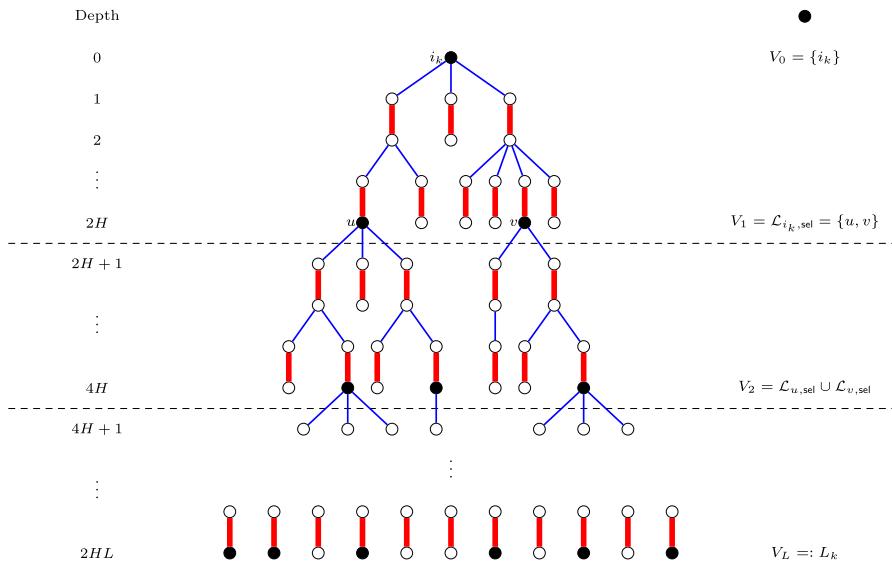


Fig. 4 Example of the left tree of T_k rooted at i_k , with $H = 2$. The shaded vertices represent the set of selected vertices in each epoch. The epochs are separated by the horizontal dashed lines. In the figure, after the first epoch, the set of selected vertices V_1 consists of two vertices u, v . Hence the exploration step in the second epoch is run with root vertices u, v . The set L_k is defined as the set of selected vertices after the L 'th epoch (color figure online)

a leaf node selection step. The selection step is necessary to ensure that the weights of the paths, and the resulting cycles, fulfill the weight requirement of bad matchings. Since we need to construct many disjoint two-sided trees, the selection step is done periodically to prevent each tree from overgrowing and using up too many vertices.

An illustration of the construction of the left tree of T_k is given in Fig. 4. In step 1, we introduce a set \mathcal{U} so that $1\{i \in \mathcal{U}\}$ tracks whether the pair (i, i') has been visited (note that the algorithm always explores vertices in pairs). Every time a vertex is added to the tree, it is marked as explored so that no edge gets inspected more than once. This ensures that the edge weights in the constructed trees are independent. In step 1, we also set a parameter m which controls the maximum number of vertices each tree explores, in order to prevent “overgrowth”. The tree is initialized at the middle edge (i_k, i'_k) (step 2). We then finish growing the left tree rooted at i_k (steps 3–4), followed by the right tree rooted at i'_k (step 5).

The specific exploration-selection tree construction is described in step 3. In step 3(a), we grow the tree iteratively using BFS, where in each iteration, each leaf node v finds the set of its unexplored neighbors O'_v . Since the pair (v, v') would have already been marked as explored, we cannot have $v \in O'_v$. Thus for each $u' \in O'_v$, (u', v) is a blue edge. We append this blue edge to the tree, along with the red edge (u, u') attached to u' . As seen in Fig. 4, the algorithm always explores the edges in alternating colors.

To prevent each tree from using up too many vertices, we introduce the selection step 3(b), where we select a subset of the leaf nodes every $2H$ layers (per epoch) and

Algorithm 2 Construction of two-sided trees

Input: n, γ , a bipartite graph G_1 that contains $(1 - \gamma)n$ pairs of vertices, threshold ζ , and parameters H, L, ϵ, α .

1: Initialize $\mathcal{U} = \{\text{all left vertices of } G_1\}$ as the set of unexplored left vertices. Define an integer⁵

$$m = (1 + \epsilon)^{2HL} \exp(3H\alpha + \epsilon H)$$

that stands for the number of left vertices each of the left and right subtrees will contain. Let $K = \frac{\gamma n}{2m}$ be an integer. For $k = 1, 2, \dots, K$, repeat the following steps 3-6 to construct T_k .

2: Let i_k be the member of \mathcal{U} with the smallest index. Update $\mathcal{U} \leftarrow \mathcal{U} \setminus \{i_k\}$ to mark i_k as explored. Initialize $T_k = \{(i_k, i'_k)\}$ to be a tree containing only one red edge.

3: Construct the left tree of T_k via L epochs of color-alternating breadth-first search (BFS) on G_1 . Define $V_0 = \{i_k\}$. In the i 'th epoch, the BFS starts from the set of vertices in V_{i-1} . Repeat the following two steps for $i = 1, \dots, L$:

(a) Exploration: for each $v \in V_{i-1}$, grow a color-alternating subtree in G_1 with v as its root. Concretely, define the offsprings of v as

$$O'_v = \{u' \in \mathcal{U} : (u', v) \in E(G_1)\}.$$

For all $u' \in O'_v$, append edges (u', v) and (u, u') to T_k . Update $\mathcal{U} \leftarrow \mathcal{U} \setminus O_v$ to mark all members of O_v as explored.

Grow the next 2 layers of the left tree similarly: sequentially (ordered by the vertex indices) for each $u \in O_v$, define its offsprings as the set of all unexplored vertices that are connected to u via a blue edge in G_1 ; append to T_k all the blue edges from u to the offsprings and their corresponding red edges; and mark all the offsprings as explored. Repeat this exploration step above until the tree is of depth $2H$, unless the exploration process becomes extinct or the following termination condition is met:

Terminate the growth of the left tree of T_k (go to step 5) (31)
when the number of left vertices in the tree exceeds m .

(b) Selection of leaf nodes: for each $v \in V_{i-1}$, let \mathcal{L}_v denote the set of all leaf nodes at depth $2H$ in the subtree rooted at v . Among those, select

$$\mathcal{L}_{v,\text{sel}} = \{u \in \mathcal{L}_v : \Delta(P_{u,v}) \geq \zeta H\}, \quad (32)$$

where

$$P_{u,v} = \text{the path from } u \text{ to } v \text{ on the subtree rooted at } v.$$

Let $V_i = \cup_{v \in V_{i-1}} \mathcal{L}_{v,\text{sel}}$ be the set of selected leaf nodes from the i 'th epoch. The next epoch of the construction will start from V_i .

4: Let \tilde{m} be the total pairs of vertices used in the construction of the left tree. Let $\tilde{\mathcal{U}}$ be the first $m - \tilde{m}$ indices in \mathcal{U} , and update \mathcal{U} to be $\mathcal{U} \setminus \tilde{\mathcal{U}}$.

5: Construct the right tree via the same scheme, starting from the root i'_k .

continue to grow the tree only from those nodes. The selected leaves are required to have the highest depth, whose paths to the previous root has a sufficiently large excess weight Δ .

We repeat this exploration-selection procedure for L epochs. Unless the exploration process becomes extinct before completing the L epochs, or no leaf node is selected

⁵ Strictly speaking m needs to be rounded to the nearest integer. To lighten the notation, from here on we will not explicitly specify the rounding step as it only affects constant factors.

$(V_i = \emptyset)$ for some $i \leq L$, or the tree becomes overgrown per condition (31), the exploration-selection process yields a tree of depth $2HL$, rooted at i_k . As we will see in Sect. 4.2, the parameters H and L need to be chosen appropriately to ensure that existence of many two-sided trees that are not overgrown, yet contain sufficiently many leaf nodes. On one hand, we cannot perform the selection step too frequently, i.e., H must be large enough. This is to guarantee that the large deviation analysis is tight, so that $\mathbb{P}\{\Delta(P) \geq \zeta H\} \approx e^{-\alpha H}$ for each alternating path of length $2H$ – see (36). On the other hand, we also need to select the leaf nodes sufficiently often, i.e. L must be large enough. This is to ensure that the growth of each tree does not use up too many vertices, so that the exploration process yields many two-sided trees.

In step 4 of Algorithm 2, if a tree uses up $\tilde{m} < m$ vertices, we mark $m - \tilde{m}$ other arbitrary vertices as explored. This ensures that each two-sided tree uses up *exactly* $2m$ pairs of vertices, so that the structure (i.e. the isomorphism class) of the K two-sided trees are independent. Moreover, from the definition

$$K = \frac{\gamma n}{2m} = \frac{\gamma}{2(1+\epsilon)^{2HL} \exp(3H\alpha + \epsilon H)} \cdot n, \quad (33)$$

we have that after all trees are constructed, the total number of unexplored vertex pairs, i.e., the size of \mathcal{U} , equals $(1 - 2\gamma)n$.

Finally, we remark that the exploration process is where we use the assumption $\sqrt{d}B(\mathcal{P}, \mathcal{Q}) \geq 1 + \epsilon$. Since all the red edges are already planted, the survival and extinction of the trees depend only on the blue edges. Roughly speaking, the average number of blue neighbors of each vertex is $(1 - \gamma)n \cdot d/n \approx d$. Thus each epoch yields approximately d^H leaves before the selection criterion in step 3(b) is applied. The probability that each leaf is selected is roughly $\mathbb{P}\{\Delta(H \text{ blue edges}) - \Delta(H \text{ red edges}) \geq \zeta H\} \approx B(\mathcal{P}, \mathcal{Q})^{2H}$ for H large and ζ small, so that the number of selected leaves is roughly $(dB(\mathcal{P}, \mathcal{Q})^2)^H$. Thus to ensure there is at least a constant proportion of trees survive through all L epochs, we need $dB(\mathcal{P}, \mathcal{Q})^2 > 1$. These heuristic calculations will be made precise in Theorem 5.

4.2 Existence of many two-sided trees

Set $s = (1 + 3\epsilon/4)^{2HL}$ as the desired number of leaf vertices. For each two-sided tree T_k constructed from Algorithm 2, define

$$\begin{aligned} \tilde{L}_k &= \text{the set of selected leaf vertices at depth } 2HL \text{ in the left subtree of } T_k, \\ \tilde{R}_k &= \text{the set of selected leaf vertices at depth } 2HL \text{ in the right subtree of } T_k, \end{aligned}$$

and let

$$L_k = \begin{cases} \text{an arbitrary size-}s \text{ subset of } \tilde{L}_k & \text{if } |\tilde{L}_k| \geq s; \\ \tilde{L}_k & \text{if } |\tilde{L}_k| < s. \end{cases}$$

Similarly define R_k . Since the left tree and the right tree are rooted at $i_k \in V^c$ and $i'_k \in (V^c)'$ respectively, we have $L_k \subset V^c$ and $R_k \subset (V^c)'$. From the definition of

the two-sided trees, each pair of vertices $i \in L_k$ and $j \in R_k$ are connected via an alternating path P of length $4HL + 1$ on T_k . In this subsection, we show that at least $K_1 = \Omega(n)$ of the K two-sided trees yield sufficiently large sets L_k , R_k , and large enough weight $\Delta(P)$.

Theorem 5 *Suppose that the condition (7) holds and Algorithm 2 is run with parameters*

$$\gamma \leq \gamma_0 = \epsilon/32, \quad \zeta = \min \{\epsilon/32, D(\mathcal{P}\|\mathcal{Q}) + D(\mathcal{Q}\|\mathcal{P})\}, \quad (34)$$

and H is a large constant depending only on ϵ , \mathcal{P} , \mathcal{Q} . Then with probability at least $1 - e^{-\Omega(n)}$, there exists $\mathcal{K}_1 \subset [K]$ with

$$K_1 = |\mathcal{K}_1| = \frac{\gamma}{16(1+\epsilon)^{2HL} \exp(7H\alpha + 3\epsilon H/2)} \cdot n,$$

such that for all $k \in \mathcal{K}_1$, T_k satisfies

1. $|L_k| = |R_k| = s = (1 + 3\epsilon/4)^{2HL}$;
2. for all $i \in L_k$ and $j \in R_k$, the path P between i and j on T_k has weight $\Delta(P) \geq 2\zeta HL - \tau_{\text{red}}$, where $\tau_{\text{red}} = \inf\{x : \mathcal{P}[\log(\mathcal{P}/\mathcal{Q}) \leq x] \geq 1/2\}$ is the median of $\log(\mathcal{P}/\mathcal{Q})$ under \mathcal{P} .

Before proving Theorem 5, we introduce the following auxiliary lemma. The lemma states that for each two-sided tree T_k , with sufficiently large probability, the termination condition (31) is not hit and there are sufficiently many leaf nodes at depth $2HL$ in both the left and the right subtrees. As a consequence, we get large sets L_k and R_k as desired in the statement of Theorem 5.

Lemma 4 *Under the assumptions of Theorem 5, we have for each $k = 1, \dots, K$, with probability at least $\frac{1}{2} \exp(-4\alpha H - \epsilon H/2)$, the following hold for both the left and the right subtree in the two-sided tree T_k .*

- (a) (not overgrown) *The construction of the tree is not terminated by hitting condition (31).*
- (b) (many leaf nodes) *The number of selected leaf nodes at depth $2HL$ is at least $s = (1 + 3\epsilon/4)^{2HL}$.*

With Lemma 4, we are ready to prove Theorem 5. The proof of Lemma 4 is deferred to the end of this subsection.

Proof of Theorem 5 For each k , each pair of $i \in L_k$ and $j \in R_k$, the path P between i and j on T_k consists of the central edge (i_k, i'_k) and $2L$ subpaths, each of length H . By the leaf node selection step (32), each subpath has weight at least ζH . Therefore,

$$\Delta(P) \geq 2\zeta HL - \log \frac{\mathcal{P}}{\mathcal{Q}} \left(W_{i_k, i'_k} \right).$$

From the definition of τ_{red} ,

$$\mathbb{P} \left\{ \log \frac{\mathcal{P}}{\mathcal{Q}} \left(W_{i_k, i'_k} \right) \leq \tau_{\text{red}} \right\} \geq 1/2.$$

Therefore, with probability at least 1/2, the weight of all paths from L_k to R_k have weight at least $2\zeta HL - \tau_{\text{red}}$.

Let $\tilde{\mathcal{K}}$ be the set of all $k \leq K$ such that $|L_k| \geq (1+3\epsilon/4)^{2HL}$, $|R_k| \geq (1+3\epsilon/4)^{2HL}$, and $\Delta(P) \geq 2\zeta HL - \tau_{\text{red}}$ for all P from L_k to R_k . Since the construction of the trees are completely independent of the weight of the central edges, we have from property (b) in Lemma 4,

$$\mathbb{P} \{k \in \tilde{\mathcal{K}}\} \geq \frac{1}{4} \exp(-4\alpha H - \epsilon H/2).$$

Moreover, recall that we deliberately ensured that the construction of each tree uses up a deterministic number of vertices. As a result, the isomorphism class of each two-sided tree T_k , in particular its total number of vertices and the number of leaf nodes, are independent of the prior T_1, \dots, T_{k-1} . That, combined with the independence of the red edge weights, yields independence of the events $\{k \in \tilde{\mathcal{K}}\}_{k \leq K}$. Thus

$$\begin{aligned} & \mathbb{P} \left\{ |\tilde{\mathcal{K}}| < \frac{1}{8} \exp(-4\alpha H - \epsilon H/2) K \right\} \\ & \leq \mathbb{P} \left\{ \text{Binom} \left(K, \frac{1}{4} \exp(-4\alpha H - \epsilon H/2) \right) < \frac{1}{8} \exp(-4\alpha H - \epsilon H/2) K \right\} = \exp(-\Omega(K)) \end{aligned}$$

by Hoeffding's inequality. Combine with (33) to conclude that with probability at least $1 - e^{-\Omega(n)}$,

$$|\tilde{\mathcal{K}}| \geq \frac{1}{8} \exp(-4\alpha H - \epsilon H/2) K = \frac{\gamma}{16(1+\epsilon)^{2HL} \exp(7H\alpha + 3\epsilon H/2)} \cdot n \triangleq c_5 n.$$

To finish the proof of Theorem 5, let \mathcal{K}_1 be an arbitrary subset of $\tilde{\mathcal{K}}$ of size exactly $c_5 n$. \square

Proof of Lemma 4 For a fixed k , let

$$\begin{aligned} \mathcal{A}_{L,k} & \triangleq \{\text{The left tree of } T_k \text{ satisfies condition (a) (not overgrown)}\}, \\ \mathcal{B}_{L,k} & \triangleq \{\text{The left tree of } T_k \text{ satisfies condition (b) (many leaf nodes)}\}. \end{aligned}$$

Similarly define events $\mathcal{A}_{R,k}$, $\mathcal{B}_{R,k}$. To prove Lemma 4, it suffices to show that for each k ,

$$\mathbb{P} (\mathcal{A}_{L,k} \cap \mathcal{B}_{L,k} \cap \mathcal{A}_{R,k} \cap \mathcal{B}_{R,k}) \geq \frac{1}{2} e^{-4\alpha H - \epsilon H/2}. \quad (35)$$

Let us first focus on the analysis of the left tree. Recall that in Algorithm 2, in each step of the exploration process from vertex v , the set of offsprings of v is defined as all members of \mathcal{U} that are incident to v via a blue edge. Thus, the offspring distribution is $\text{Binom}(|\mathcal{U}|, d/n)$. The randomness in $|\mathcal{U}|$ results in dependence between different steps of the exploration process, which is technically inconvenient. To resolve this problem, we utilize the fact that $|\mathcal{U}|$ is always between $(1 - 2\gamma)n$ and n , and couple the tree with two two-sided trees $T_k^{(L)} = T_k^{(\text{Lower})}$ and $T_k^{(U)} = T_k^{(\text{Upper})}$, whose offspring distributions follow $\text{Binom}((1 - 2\gamma)n, d/n)$ and $\text{Binom}(n, d/n)$, respectively.

More specifically, in each epoch of the growth of the left tree of $T_k^{(L)}$, the exploration process follows a branching process with offspring distribution $\text{Binom}((1 - 2\gamma)n, d/n)$, such that it is homomorphic to a subtree of T_k ; the leaf node selection is done in a similar way that preserves the stochastic ordering, such that each vertex is selected with probability

$$p_{\text{sel}} \triangleq \mathbb{P} \left\{ \sum_{j \leq H} Y_j - \sum_{j \leq H} X_j \geq \xi H \right\}, \quad (36)$$

where X_i 's are i.i.d. copies of $\log(\mathcal{P}/\mathcal{Q})$ under distribution \mathcal{P} and Y_i 's are i.i.d. copies of $\log(\mathcal{P}/\mathcal{Q})$ under distribution \mathcal{Q} . Construct the right tree of $T_k^{(L)}$, and $T_k^{(U)}$ in similar fashion. One caveat is that the tree T_k may hit the termination condition (31), while both $T_k^{(L)}$ and $T_k^{(U)}$ are assumed to be “free-growing” without any termination conditions. As a result, on the event that neither subtree of T_k is terminated by hitting the termination condition (31), we have

$$T_k^{(L)} \subset T_k \subset T_k^{(U)},$$

where \subset denotes stochastic dominance of the three trees. Let $\mathcal{A}_{L,k}^{(L)}, \mathcal{A}_{R,k}^{(L)}, \mathcal{B}_{L,k}^{(L)}, \mathcal{B}_{R,k}^{(L)}$ denote the analogue events on the tree $T_k^{(L)}$, and similarly define the events on $T_k^{(U)}$. A key observation is that

$$\mathcal{A}_{L,k}^{(U)} \cap \mathcal{B}_{L,k}^{(L)} \subset \mathcal{A}_{L,k} \cap \mathcal{B}_{L,k}.$$

To see why, note that $T_k \subset T_k^{(U)}$ always holds. Therefore on $\mathcal{A}_{L,k}^{(U)}$, we have that the left subtree of T_k never hits the termination condition (31). As a result, $T_k^{(L)} \subset T_k$ also holds, and thus $\mathcal{B}_{L,k}^{(L)}$ implies $\mathcal{B}_{L,k}$. A similar relationship holds for the right tree. It follows that

$$\begin{aligned} \mathbb{P}(\mathcal{A}_{L,k} \cap \mathcal{B}_{L,k} \cap \mathcal{A}_{R,k} \cap \mathcal{B}_{R,k}) &\geq \mathbb{P}(\mathcal{A}_{L,k}^{(U)} \cap \mathcal{B}_{L,k}^{(L)} \cap \mathcal{A}_{R,k}^{(U)} \cap \mathcal{B}_{R,k}^{(L)}) \\ &\stackrel{(a)}{=} \mathbb{P}(\mathcal{A}_{L,k}^{(U)} \cap \mathcal{B}_{L,k}^{(L)})^2 \\ &\stackrel{(b)}{\geq} \left[\mathbb{P}(\mathcal{A}_{L,k}^{(U)}) + \mathbb{P}(\mathcal{B}_{L,k}^{(L)}) - 1 \right]^2, \end{aligned}$$

where (a) is from symmetry and independence of the left and right subtrees; (b) is from the basic inequality $P(AB) \geq P(A) + P(B) - 1$ for all probability events A, B , and that $\mathbb{P}(\mathcal{A}_{L,k}^{(U)}) + \mathbb{P}(\mathcal{B}_{L,k}^{(L)}) - 1 \geq 0$ from the following two inequalities (37),(38) that we will show.

To prove (35) it suffices to show the following two claims:

$$\mathbb{P}(\mathcal{A}_{L,k}^{(U)}) \geq 1 - \frac{1}{4}e^{-2\alpha H - \epsilon H/4}, \quad (37)$$

$$\mathbb{P}(\mathcal{B}_{L,k}^{(L)}) \geq e^{-2\alpha H - \epsilon H/4}. \quad (38)$$

We start with the proof of (37). First, we can control the probability p_{sel} defined in (36) from above via the large deviation bound (83):

$$p_{\text{sel}} \leq \mathbb{P} \left\{ \sum_{j \leq H} Y_j - \sum_{j \leq H} X_j \geq 0 \right\} \leq \exp(-\alpha H). \quad (39)$$

Let $W_i^{(U)}$ denote the number of vertices that are used in the exploration process at the i 'th epoch (before selection) in the left subtree of $T_k^{(U)}$. Let $Z_i^{(U)}$ denote the number of leaf nodes at depth iH that are selected. We have for all $i \leq L$,

$$\begin{aligned} \mathbb{E}(W_i^{(U)}) &= \mathbb{E}(Z_{i-1}^{(U)}) \cdot \sum_{j=1}^H d^j \\ &= \left[\left(n \cdot \frac{d}{n} \right)^H p_{\text{sel}} \right]^{i-1} \cdot \sum_{j=1}^H d^j \\ &\leq [de^{-\alpha}]^{H(i-1)} \cdot \sum_{j=1}^H d^j \\ &\leq [de^{-\alpha}]^{H(i-1)} d^H \frac{d}{d-1}. \end{aligned}$$

Recall that $de^{-\alpha} = (1 + \epsilon)^2$. Therefore, the expected total number of vertices in the left tree

$$\begin{aligned} \mathbb{E} \left(1 + \sum_{i=1}^L W_i^{(U)} \right) &\leq 1 + \sum_{i=1}^L (1 + \epsilon)^{2H(i-1)} d^H \frac{d}{d-1} \\ &\stackrel{(a)}{\leq} \sum_{i=0}^L (1 + \epsilon)^{2H(i-1)} d^H \frac{d}{d-1} \\ &\stackrel{(b)}{\leq} \frac{d}{d-1} d^H (1 + \epsilon)^{2H(L-1)} \frac{1}{1 - (1 + \epsilon)^{-H}} \\ &\stackrel{(c)}{\leq} \frac{2d}{d-1} (1 + \epsilon)^{2HL} e^{H\alpha}, \end{aligned} \quad (40)$$

where (a) is from $(1 + \epsilon)^{-2H} d^H d / (d - 1) = e^{\alpha H} d / (d - 1) \geq 1$ since $\alpha \geq 0$ and $d / (d - 1) \geq 1$; (b) is by bounding the finite geometric sum by an infinite sum; (c) is by choosing H to be a sufficiently large constant so that $(1 + \epsilon)^{-H} \leq 1/2$. Therefore,

$$\begin{aligned} 1 - \mathbb{P}(\mathcal{A}_{L,k}^{(U)}) &\stackrel{(a)}{\leq} \mathbb{P} \left\{ 1 + \sum_{i=1}^L W_i^{(U)} \geq (1 + \epsilon)^{2HL} \exp(3H\alpha + \epsilon H) \right\} \\ &\stackrel{(b)}{\leq} \frac{\mathbb{E} \left(1 + \sum_{i=1}^L W_i^{(U)} \right)}{(1 + \epsilon)^{2HL} \exp(3H\alpha + \epsilon H)} \\ &\stackrel{(c)}{\leq} \frac{2d}{d - 1} \exp(-2H\alpha - \epsilon H) \\ &\stackrel{(d)}{\leq} \frac{1}{4} \exp(-2H\alpha - \epsilon H/4), \end{aligned}$$

where (a) is from the definition of $\mathcal{A}_{L,k}^{(U)}$; (b) is from Markov's inequality; (c) is from (40); (d) is because by the assumption $d e^{-\alpha} = (1 + \epsilon)^2$, and that $e^{-\alpha} = B(\mathcal{P}, \mathcal{Q})^2 \leq 1$, we have $d \geq (1 + \epsilon)^2$. Thus $2d / (d - 1) \leq 2(1 + \epsilon)^2 / (2\epsilon + \epsilon^2) \leq \exp(3\epsilon H/4) / 4$ by choosing H large enough.

Next we show (38). As before, let $Z_i^{(L)}$ denote the number of leaf nodes at depth iH that are selected in the left subtree of $T_k^{(L)}$. We have

$$\mathbb{P}(\mathcal{B}_{L,k}^{(L)}) = \mathbb{P} \left\{ Z_L^{(L)} \geq (1 + 3\epsilon/4)^{2HL} \right\}. \quad (41)$$

We will bound $\mathbb{P}\{\mathcal{B}_{L,k}^{(L)}\}$ by analyzing the first and second moments of $Z_L^{(L)}$. Note that for all i ,

$$\mathbb{E}(Z_i^{(L)}) = \left[((1 - 2\gamma)d)^H p_{\text{sel}} \right]^i \triangleq \mu^i.$$

We claim the following inequalities on p , μ , and the second moment of $Z_L^{(L)}$.

$$p_{\text{sel}} \geq \exp(-\alpha H - \epsilon H/16), \quad (42)$$

$$\mu \geq 2 \left(1 + \frac{3}{4}\epsilon \right)^{2H}, \quad (43)$$

$$\mathbb{E} \left(\left[Z_L^{(L)} \right]^2 \right) \leq \mu^{2L} + \left[(1 - 2\gamma)d \right]^{2H} \frac{(1 - 2\gamma)d}{(1 - 2\gamma)d - 1} \cdot \frac{\mu^{2L-2}}{1 - \mu^{-1}}. \quad (44)$$

Assuming that (42), (43), (44) all hold, we first finish the proof of (38). By combining (41) and (43), we have

$$\mathbb{P} \left\{ \mathcal{B}_{L,k}^{(L)} \right\} \geq \mathbb{P} \left\{ Z_L^{(L)} > \frac{1}{2} \mathbb{E}(Z_L^{(L)}) \right\}. \quad (45)$$

By the Paley-Zygmund inequality,

$$\mathbb{P}\left\{Z_L^{(\text{L})} > \frac{1}{2}\mathbb{E}(Z_L^{(\text{L})})\right\} \geq \frac{1}{4} \frac{\left[\mathbb{E}(Z_L^{(\text{L})})\right]^2}{\mathbb{E}\left(Z_L^{(\text{L})}\right)^2} \geq \frac{1}{4} \frac{1}{1 + \left[(1-2\gamma)d\right]^{2H} \frac{(1-2\gamma)d}{(1-2\gamma)d-1} \frac{1}{\mu(\mu-1)}}, \quad (46)$$

where the last inequality follows from (44). From (43), we can choose H large enough such that $\mu - 1 \geq \mu/2$. Therefore, the right-hand side of (46) is bounded from below by

$$\begin{aligned} & \frac{1}{4} \frac{1}{1 + 2 \frac{(1-2\gamma)d}{(1-2\gamma)d-1} \left[(1-2\gamma)d\right]^{2H} \mu^{-2}} \\ & \stackrel{(a)}{\geq} \frac{1}{4 \left(1 + 2 \frac{(1-2\gamma)d}{(1-2\gamma)d-1}\right)} p_{\text{sel}}^2 \\ & \stackrel{(b)}{\geq} \frac{1}{4 \left(1 + 2 \frac{(1-\epsilon/16)(1+\epsilon)^2}{(1-\epsilon/16)(1+\epsilon)^2-1}\right)} \exp(-2\alpha H - \epsilon H/8) \\ & \stackrel{(c)}{\geq} \exp(-2\alpha H - \epsilon H/4), \end{aligned}$$

where (a) is from the definition of μ , and $p_{\text{sel}} \leq 1$; (b) is from (42), and $\gamma \leq \epsilon/32$, $d \geq (1+\epsilon)^2$ (which follows from (7)); (c) is by choosing H large enough so that the fractional factor is absorbed into $\exp(\epsilon H/8)$. Combine the display above with (45) and (46) to finish the proof of (38).

It remains to prove (42), (43) and (44).

Proof of (42): By (84), we have that for ζ chosen as in (34),

$$p_{\text{sel}} \geq \exp(-\alpha H - \zeta H + o(H)).$$

(42) follows by choosing H to be a large enough constant only depending on ϵ , \mathcal{P} , \mathcal{Q} .

Proof of (43): By definition of μ and the inequality (42), we have

$$\begin{aligned} \mu & \geq \left[(1-2\gamma)d\right]^H \exp(-H\alpha - \epsilon H/16) \\ & \geq \left[(1-\epsilon/16)(1+\epsilon)^2\right]^H \exp(-\epsilon H/16) \geq 2 \left(1 + \frac{3}{4}\epsilon\right)^{2H} \end{aligned}$$

by choosing H to be a large enough constant.

Proof of (44): Following the arguments in [11, Theorem 2.1.6], we can control the second moment of $Z_i^{(\text{L})}$ with

$$\mathbb{E}\left(Z_i^{(\text{L})}\right)^2 \leq \mu^{2i} + \sigma^2 \frac{\mu^{2i-2}}{1-\mu^{-1}}, \quad (47)$$

where $\sigma^2 = \text{Var}(Z_1^{(\text{L})}) \leq \mathbb{E}([Z_1^{(\text{L})}]^2)$. For completeness we include a short proof of (47) here. First, note that

$$\mathbb{E}(Z_i^{(\text{L})} \mid Z_{i-1}^{(\text{L})}) = \mu Z_{i-1}^{(\text{L})}, \quad \text{and}$$

$$\text{Var}(Z_i^{(\text{L})} \mid Z_{i-1}^{(\text{L})}) = \sum_{j \leq Z_{i-1}^{(\text{L})}} \sigma^2 = Z_{i-1}^{(\text{L})} \sigma^2.$$

Therefore

$$\mathbb{E}\left(\left[Z_i^{(\text{L})}\right]^2 \mid Z_{i-1}^{(\text{L})}\right) = \mu^2 \left(Z_{i-1}^{(\text{L})}\right)^2 + \text{Var}(Z_i^{(\text{L})} \mid Z_{i-1}^{(\text{L})}) = \mu^2 \left(Z_{i-1}^{(\text{L})}\right)^2 + Z_{i-1}^{(\text{L})} \sigma^2.$$

Take expectation on both sides. We have

$$\mathbb{E}\left(\left[Z_i^{(\text{L})}\right]^2\right) = \mu^2 \mathbb{E}\left(\left[Z_{i-1}^{(\text{L})}\right]^2\right) + \mu^{i-1} \sigma^2.$$

The induction above with $Z_0^{(\text{L})} = 1$ yields

$$\mathbb{E}\left(\left[Z_i^{(\text{L})}\right]^2\right) = \mu^{2i} + \sigma^2 \sum_{j=i-1}^{2i-2} \mu^j \leq \mu^{2i} + \sigma^2 \frac{\mu^{2i-2}}{1 - \mu^{-1}}.$$

To bound σ^2 , we upper bound $Z_1^{(\text{L})}$ by the number of leaves at depth H before selection. Following the same argument as in the derivation of (47), we have

$$\sigma^2 \leq [(1-2\gamma)d]^{2H} + [(1-2\gamma)d]^{2H-1} \frac{1}{1 - [(1-2\gamma)d]^{-1}} = [(1-2\gamma)d]^{2H} \frac{(1-2\gamma)d}{(1-2\gamma)d - 1}. \quad (48)$$

Combine (47) and (48) to yield (44). \square

5 Sprinkling stage

Recall that we have constructed a family of disjoint sets L_k of left vertices and R_k of right vertices for $k \in \mathcal{K}_1$ such that each pair of vertices $i \in L_k$ and $j \in R_k$ are connected via an alternating path P of length ℓ through red edge (i_k, i'_k) in G_1 , where $|L_k| = |R_k| = s$ and $K_1 \triangleq |\mathcal{K}_1|$. Crucially the construction of L_k and R_k does not involve the vertices in V . Recall from Algorithm 1 that G_2 is an unweighted subgraph of G that does not contain any edges in $V^c \times (V^c)'$. A blue edge e appears in G_2 if and only if it appears in G , and $\log(\mathcal{P}/\mathcal{Q})(W_e) \geq \tau_{\text{blue}}$. Therefore, the blue edges in G_2 are independently generated with probability η/n , where $\eta = d \cdot \mathcal{Q}[\log(\mathcal{P}/\mathcal{Q})(W_e) \geq \tau_{\text{blue}}]$.

Recall from Algorithm 1 that

$$V^* = \{i \in V : \log(\mathcal{P}/\mathcal{Q})(W_{i,i'}) \leq \tau_{\text{red}}\}$$

is a subset of the reserved vertices whose incident red edge weight is below threshold. In this section, we connect the alternating paths between L_k and R_k through vertices in V^* and $(V^*)'$ to form long alternating cycles in G . This scheme is referred to as “sprinkling”, which is detailed in Algorithm 3.

Algorithm 3 Sprinkling

Input: Parameters n, η, s , sets $V^* \subset V \subset [n]$, a bipartite graph G_2 whose blue edges are independently generated with probability η/n , and disjoint sets $\{L_k\}_{k \in \mathcal{K}_1}$ of left vertices and $\{R_k\}_{k \in \mathcal{K}_1}$ of right vertices with $|L_k| = |R_k| = s$ for all $k \in \mathcal{K}_1$.

1: Define

$$\begin{aligned} A'_k &= \{v' \in (V^*)' : \exists u \in L_k, \text{ s.t. } (u, v') \in E(G_2)\}, \\ B_k &= \{v \in V^* : \exists u' \in R_k, \text{ s.t. } (u', v) \in E(G_2)\}. \end{aligned}$$

2: For each $v \in V^*$, let

$$d_v = \sum_{k \in \mathcal{K}_1} \mathbf{1}_{\{v' \in A'_k\}} + \sum_{k \in \mathcal{K}_1} \mathbf{1}_{\{v \in B_k\}}.$$

Define $A_{\text{overlap}} = \{v \in V^* : d_v \geq 2\}$. For all $k \in \mathcal{K}_1$, define

$$U'_k = A'_k \setminus A'_{\text{overlap}}, \quad V_k = B_k \setminus A_{\text{overlap}}.$$

Define $\beta = |V^*|/n$, and let

$$\mathcal{K}_2 = \left\{ k \in \mathcal{K}_1 : |U'_k| \geq \frac{\beta s \eta}{4}, |V_k| \geq \frac{\beta s \eta}{4} \right\}.$$

3: Define a bipartite graph G_{super} on $\mathcal{K}_2 \times (\mathcal{K}_2)'$, where there is a red edge between i and i' for every $i \in \mathcal{K}_2$ and there is a blue edge between i and j' if and only if U_i and $V'_{j'}$ are connected via a blue edge in G_2 .

4: For each alternating cycle $C_{\text{super}} = (i_1, i'_1, i_2, i'_2, \dots, i_r, i'_r)$ in G_{super} , extend it to an alternating cycle in G as

$$C = (v'_1, v_1, P_1, u'_1, u_1, v'_2, v_2, P_2, u'_2, u_2, \dots, v'_r, v_r, P_r, u'_r, u_r),$$

where u_k and v'_{k+1} are two nodes in U_{i_k} and $V'_{i_{k+1}}$ that are connected by a blue edge; P_k is an alternating path connecting L_{i_k} and R_{i_k} ; and u'_k (resp. v_k) and P_k are connected by a blue edge according to the definition of U'_k (resp. V_k).

In step 1 of Algorithm 3, we find all vertices in $(V^*)'$ (resp. V^*) that are connected to A'_k (resp. B_k) through a blue edge. Note that these sets may contain overlapping vertices, e.g. v could belong to two different sets B_{k_1} and B_{k_2} , or we could have $v \in B_{k_1}$ and $v' \in A'_{k_2}$. To make sure the algorithm yields legitimate cycles, in step 2, we remove these overlaps. We define A_{overlap} as the set of all $v \in V^*$ such that itself combined with its pair v' have more than one connections to $\{A'_k\}, \{B_k\}$. We removed those overlapping vertices to form the sets $\{U'_k\}, \{V_k\}$. We then only keep those $k \in \mathcal{K}_1$

with large enough sizes of U'_k and V_k , which helps ensure the super graph has a high enough average degree.

See Fig. 2 for an illustration of steps 4–5 of Algorithm 3, which extends each alternating cycle on the super graph into an alternating cycle on G . The following theorem provides a sufficient condition under which with high probability, G_{super} contains exponentially many distinct alternating cycles, which correspond to distinct alternating cycles on G . Note that the theorem also states that $K_2 = |\mathcal{K}_2|$ is of order $\mathcal{O}(K_1)$. Thus, the sprinkling scheme yields $e^{\mathcal{O}(K_2)} = e^{\mathcal{O}(K_1)}$ distinct alternating cycles, each of length $3K_2/4 = \mathcal{O}(K_1)$.

Theorem 6 (Sprinkling) *Let*

$$\beta = |V^*|/n, \quad \kappa = \frac{2K_1 s \eta}{n}, \quad b = \frac{\beta s \eta}{4}, \quad d_{\text{super}} = \frac{1}{32n} K_1 b^2 \eta.$$

Suppose $b \geq 4$, $K_1 \geq 8400$, $\kappa \leq 1/16^2$, and $d_{\text{super}} \geq 256 \log(32e)$. Then conditional on a fixed V^* , we have with probability at least $\left(1 - e^{-K_1/32} - \frac{2}{\beta n \kappa^3}\right) \left(1 - e^{-d_{\text{super}} K_1/2^{18}}\right)$, G_{super} obtained by Algorithm 3 contains at least $\exp(K_2/20)$ distinct alternating cycles of length at least $\frac{3}{4}K_2$, and $K_2 \geq K_1/16$.

The proof of Theorem 6 consists of two major steps. First, we show that with high probability, $|U'_k| \geq b$ and $|V_k| \geq b$ for all $k \in \mathcal{K}_2$ for some \mathcal{K}_2 of size $K_2 \geq K_1/16$. Then conditional on $\{U_k, V'_k\}_{k \in \mathcal{K}_2}$, we show that G_{super} contains exponentially many long alternating cycles with high probability. In the sequel, we present the detailed proof.

5.1 Existence of large \mathcal{K}_2

To show that there exist many indices $k \in \mathcal{K}_1$ with $|U'_k| \geq b$ and $|V_k| \geq b$, we first show in the following Lemma 5 that there exist many indices with large $|A'_k|$ and $|B_k|$. We then show in Lemma 6 that there are not many indices who lose a large fraction of vertices when the overlapping vertices are removed.

Lemma 5 *If $\beta s \eta \geq 4$ and $s \eta \leq n/4$, then*

$$\mathbb{P} \left\{ \sum_{k \in \mathcal{K}_1} \mathbf{1}_{\{|A'_k| \geq \frac{\beta s \eta}{2}, |B_k| \geq \frac{\beta s \eta}{2}\}} \leq \frac{K_1}{8} \right\} \leq \exp(-K_1/32)$$

Lemma 6 *If $\kappa \triangleq 2K_1 s \eta / n \leq \frac{1}{16^2}$. Then*

$$\mathbb{P} \left\{ \sum_{k \in \mathcal{K}_1} \mathbf{1}_{\{|A'_k| - |U'_k| \geq \frac{\beta s \eta}{4}, \text{ or } |B_k| - |V_k| \geq \frac{\beta s \eta}{4}\}} \leq \frac{K_1}{16} \right\} \geq 1 - \frac{2}{\beta n \kappa^3}.$$

Combining Lemma 5 and 6, we conclude that there exists some \mathcal{K}_2 of size $K_2 \geq K_1/16$, such that for all $k \in \mathcal{K}_2$, $|U'_k| \geq \frac{\beta s \eta}{2} - \frac{\beta s \eta}{4} = \frac{\beta s \eta}{4} = b$, and $|V_k| \geq b$.

Proof of Lemma 5 For each $k \in \mathcal{K}_1$ and each $v' \in (V^*)'$,

$$\mathbb{P}\{\exists u \in L_k, \text{ s.t. } (u, v') \in E(G_2)\} = 1 - \left(1 - \frac{\eta}{n}\right)^{|L_k|} \geq 1 - \left(1 - \frac{\eta}{n}\right)^s.$$

Therefore, $|A'_k|$ is stochastically dominant over a random variable distributed $\text{Binom}(\beta n, 1 - (1 - \frac{\eta}{n})^s)$. Since the mean and median of a binomial distribution differ by at most 1, the median of $|A'_k|$ is lower bounded by

$$\beta n \left[1 - \left(1 - \frac{\eta}{n}\right)^s\right] - 1 \stackrel{(a)}{\geq} \frac{3\beta s \eta}{4} - 1 \geq \frac{\beta s \eta}{2},$$

where (a) holds due to $(1 - x)^s \leq e^{-sx} \leq 1 - 3sx/4$ when $0 \leq sx \leq 1/4$ and the assumption that $s\eta \leq n/4$; the last inequality follows from the assumption that $\beta s \eta \geq 4$. Thus $\mathbb{P}\{|A'_k| \geq \frac{\beta s \eta}{2}\} \geq 1/2$. Similarly argue that $\mathbb{P}\{|B_k| \geq \frac{\beta s \eta}{2}\} \geq 1/2$. By independence of $|A'_k|$ and $|B_k|$, we have

$$\mathbb{P}\left\{|A'_k| \geq \frac{\beta s \eta}{2}, |B_k| \geq \frac{\beta s \eta}{2}\right\} \geq \frac{1}{4}.$$

Combined with independence across all $k \in \mathcal{K}_1$, we have

$$\begin{aligned} & \mathbb{P}\left\{\sum_{k \in \mathcal{K}_1} \mathbf{1}_{\{|A'_k| \geq \frac{\beta s \eta}{2}, |B_k| \geq \frac{\beta s \eta}{2}\}} \leq \frac{K_1}{8}\right\} \\ & \leq \mathbb{P}\{\text{Binom}(K_1, 1/4) \leq K_1/8\} \leq \exp(-K_1/32) \end{aligned}$$

by Hoeffding's inequality. \square

Proof of Lemma 6 Recall that for each $v \in V^*$,

$$d_v = \sum_{k \in \mathcal{K}_1} \mathbf{1}_{\{v' \in A'_k\}} + \sum_{k \in \mathcal{K}_1} \mathbf{1}_{\{v \in B_k\}}.$$

Since $v \in A_{\text{overlap}}$ if and only if $d_v \geq 2$, we have

$$\begin{aligned} \sum_{v \in V^*} d_v \mathbf{1}_{\{d_v \geq 2\}} &= \sum_{v \in A_{\text{overlap}}} d_v \\ &= \sum_{v \in A_{\text{overlap}}} \sum_{k \in \mathcal{K}_1} (\mathbf{1}_{\{v' \in A'_k\}} + \mathbf{1}_{\{v \in B_k\}}) \\ &= \sum_{k \in \mathcal{K}_1} \left(\sum_{v \in A_{\text{overlap}}} \mathbf{1}_{\{v' \in A'_k\}} + \sum_{v \in A_{\text{overlap}}} \mathbf{1}_{\{v \in B_k\}} \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{k \in \mathcal{K}_1} (|A'_k| - |U'_k| + |B_k| - |V_k|) \\
&\geq \frac{\beta s \eta}{4} \sum_{k \in \mathcal{K}_1} \mathbf{1}_{\{|A'_k| - |U'_k| + |B_k| - |V_k| \geq \frac{\beta s \eta}{4}\}} \\
&\geq \frac{\beta s \eta}{4} \sum_{k \in \mathcal{K}_1} \mathbf{1}_{\{|A'_k| - |U'_k| \geq \frac{\beta s \eta}{4}, \text{ or } |B_k| - |V_k| \geq \frac{\beta s \eta}{4}\}}.
\end{aligned}$$

In other words, $\sum_{v \in V^*} d_v \mathbf{1}_{\{d_v \geq 2\}}$ controls the number of $k \in \mathcal{K}_1$ for which A'_k or B_k loses over $\frac{\beta s \eta}{4}$ vertices when A_{overlap} is removed. It remains to prove that with high probability

$$\sum_{v \in V^*} d_v \mathbf{1}_{\{d_v \geq 2\}} \leq \frac{\beta s \eta}{4} \times \frac{K_1}{16}.$$

First, note that the random variables $\{d_v\}_{v \in V^*}$ are independent since the blue edges in G_2 are independent. Moreover, for each $v \in V^*$,

$$\begin{aligned}
d_v &= \sum_{k \in \mathcal{K}_1} \mathbf{1}_{\{\exists u \in L_k, s.t. (u, v') \in E(G_2)\}} + \sum_{k \in \mathcal{K}_1} \mathbf{1}_{\{\exists u' \in R_k, s.t. (v', u) \in E(G_2)\}} \\
&\sim \text{Binom}\left(2K_1, 1 - \left(1 - \frac{\eta}{n}\right)^s\right).
\end{aligned}$$

Using $(1 - x)^s \geq 1 - sx$, we get that $1 - \left(1 - \frac{\eta}{n}\right)^s \leq s\eta/n$. Hence, there exist random variables $\tilde{d}_v \stackrel{i.i.d.}{\sim} \text{Binom}(2K_1, s\eta/n)$ such that

$$\sum_{v \in V^*} d_v \mathbf{1}_{\{d_v \geq 2\}} \leq \sum_{v \in V^*} \tilde{d}_v \mathbf{1}_{\{\tilde{d}_v \geq 2\}}. \quad (49)$$

Note that

$$\begin{aligned}
\mathbb{E}[\tilde{d}_v \mathbf{1}_{\{\tilde{d}_v \geq 2\}}] &= \mathbb{E}[\tilde{d}_v] - \mathbb{P}\{\tilde{d}_v = 1\} \\
&= 2K_1 \frac{s\eta}{n} \left[1 - \left(1 - \frac{s\eta}{n}\right)^{2K_1-1}\right] \leq (2K_1 s\eta/n)^2 \triangleq \kappa^2.
\end{aligned}$$

where the last inequality follows from using $(1 - x)^a \geq 1 - ax$ for $0 \leq x \leq 1$ and $a \geq 1$.

Moreover,

$$\text{Var}[\tilde{d}_v \mathbf{1}_{\{\tilde{d}_v \geq 2\}}] \leq \mathbb{E}[\tilde{d}_v^2] \leq \kappa + \kappa^2 \leq 2\kappa,$$

where the last inequality follows from $\kappa \leq 1$. Thus by Chebyshev's inequality and (49), we have that with probability at least $1 - 2/(\beta n \kappa^3)$,

$$\sum_{v \in V^*} d_v \mathbf{1}_{\{d_v \geq 2\}} \leq \sum_{v \in V^*} \tilde{d}_v \mathbf{1}_{\{\tilde{d}_v \geq 2\}} \leq 2\beta n \kappa^2 = 4K_1 \beta s \eta \kappa \leq \frac{K_1 \beta s \eta}{4 \times 16},$$

where the last inequality holds by the assumption that $\kappa \leq \frac{1}{16^2}$. \square

5.2 Construction of exponentially many alternating cycles

In this subsection we show the existence of many long alternating cycles in G_{super} . Recall that in G_{super} there is a planted (red) edge between i and i' for every $i \in \mathcal{K}_2$. Moreover, there is an unplanted (blue) edge between i and j' for $i \neq j$ if and only if U_i and V'_j are connected via an edge in G_2 . In the previous steps, we have not inspected any of the edges between U_i and V'_j . Therefore

$$\mathbb{P}\{(i, j') \in E(G_{\text{super}})\} = 1 - \left(1 - \frac{\eta}{n}\right)^{|U_i| \cdot |V'_j|} \geq \frac{1}{2} \frac{\eta}{n} b^2 \geq \frac{d_{\text{super}}}{K_2}, \quad (50)$$

where the first inequality holds by $(1 - x)^a \leq 1 - ax/2$ for $0 \leq x \leq 1/a$, and the assumptions that $|U_i|, |V'_j| \geq b$ and $b^2 \eta \leq n$; the last inequality holds by the definition of $d_{\text{super}} = \frac{K_1 \eta b^2}{32n}$, and $K_2 \geq K_1/16$ as shown in Sect. 5.1. Moreover, from the independence of the blue edges in G_1 , all edges (i, j') in G_{super} appear independently. Thus, G_{super} is a bipartite graph on $\mathcal{K}_2 \times (\mathcal{K}_2)'$ with planted red edges between i' and i and unplanted blue edges between i and j' appearing independently with probability at least $\frac{d_{\text{super}}}{K_2}$. Lemma 7 shows that we have G_{super} contains exponentially many long alternating cycles for d_{super} large enough.

Lemma 7 *Let G be a bi-colored bipartite graph on $[n] \times [n]'$ whose n red edges are defined by a perfect matching, and blue edges are generated from a bipartite Erdős-Rényi graph with edge probability D/n . If $n \geq 525$ and $D \geq 256 \log(32e)$, then with probability at least $1 - \exp(-Dn/2^{14})$, G contains $\exp(n/20)$ distinct alternating cycles of length at least $3n/4$.*

Proof In order to show the existence of exponentially many long cycles in G , we construct exponentially many subsets of $[n]$ with relatively small overlaps. For ease of presentation, we assume n is even; otherwise we replace $n/2$ by $\lfloor n/2 \rfloor$ in the following proof.

First, it is well-known that there exists a collection \mathbb{V} of subsets of $[n]$ of cardinality $n/2$,

such that for any distinct $S, T \in \mathbb{V}$, $|S \Delta T| \geq n/3$ and

$$|\mathbb{V}| \geq \frac{\binom{n}{n/2}}{\sum_{i=0}^{n/3} \binom{n}{i}} \geq \binom{n}{n/2} e^{-n(\log 3 - \frac{2}{3} \log 2)} \geq e^{n/20}, \quad (51)$$

where the first inequality is by a volume argument (Gilbert-Varshamov bound)⁵; the second inequality follows from the Chernoff bound $2^{-n} \sum_{i=0}^{n/3} \binom{n}{i} = \mathbb{P}\{\text{Binom}(n, 1/2) \leq n/3\} \leq e^{-nD(\text{Bern}(1/3)\|\text{Bern}(1/2))}$; the last inequality is from the bound on the binomial coefficient $\binom{n}{n/2} \geq \frac{1}{\sqrt{2n}} 2^n$ by Stirling's approximation [24], and the fact that $\frac{1}{\sqrt{2n}} \exp(-n(\log 3 - \frac{5}{3} \log 2)) \geq e^{n/20}$ for all $n \geq 525$.

Next we show that for each size- $(n/2)$ subset V of $[n]$, $G[V \times V']$ contains a long alternating cycle with high probability. By the same argument as in [12, Theorem 6.8], we claim that the graph $G[V \times V']$ contains an alternating cycle of length at least $3n/4$, if for all subsets $S_1 \times S_2 \subset V \times V'$ such that $|S_1|, |S_2| \geq \frac{n}{32} - 1 \geq \frac{n}{64}$, there is at least one pair of $u \in S_1, v' \in S_2$ that are connected by a blue edge. For completeness we include a proof of this claim below.

The claim is shown by constructing a long alternating path on $G[V \times V']$ with the depth-first search (DFS) algorithm. Trace the DFS algorithm with the variables

1. U = The set of unexplored left vertices;
2. D = The set of dead (fully explored) left vertices;
3. $P = (v'_1, v_1, v'_2, v_2, \dots, v'_r, v_r) =$ the current path.

Initialize at $U = V \setminus \{v_1\}$, $D = \emptyset$ and $P = (v'_1, v_1)$ where v_1 is an arbitrary member of V , e.g., the one with the smallest index. At each step of the DFS algorithm, we proceed according to the following two cases:

1. If there is some $u' \in U'$ such that (v_r, u') is a blue edge, we update

$$v_{r+1} \leftarrow u'; \quad P \leftarrow (P, u', u); \quad U \leftarrow U \setminus \{u\}; \quad r \leftarrow r + 1.$$

2. If no vertex in U' is incident to v_r in G , we update

$$D \leftarrow D \cup \{v_r\}; \quad P \leftarrow P \setminus \{v_r, v'_r\}; \quad r \leftarrow r - 1.$$

By definition of the DFS algorithm, a vertex v is only added to D if v is not incident to any vertex in U' . Since the algorithm never adds any new vertices to U , the set D and U' are always disconnected in G at any stage of the algorithm. Furthermore the size of U is non-increasing, and the size of D is non-decreasing, such that $|V| = r + |D| + |U|$ always holds. Therefore at some time point in the DFS algorithm, we have $|D| = |U| = |U'|$. Assuming that all pairs of $S_1 \times S_2 \subset V \times V'$ of size $S_1 \geq n/32 - 1, S_2 \geq n/32 - 1$ are connected by at least one blue edge, we must have $|D| = |U| \leq n/32 - 1$, hence $r = n/2 - |D| - |U| \geq 7n/16$.

Apply the assumption again with $S_1 = \{v_{r-n/32+1}, \dots, v_r\}$, $S_2 = \{v'_1, \dots, v'_{n/32}\}$. There exists $v_i \in S_1, v'_j \in S_2$ such that (v_i, v'_j) is a blue edge in G . We have constructed a cycle

$$(v'_j, v_j, v'_{j+1}, v_{j+1}, \dots, v'_i, v_i, v'_j)$$

⁵ Indeed, construct \mathbb{V} greedily until no more subsets can be added. Then for any T with $|T| = n/2$, there exists some $S \in \mathbb{V}$ such that $|S \Delta T| < n/3$. Since for each fixed S , the number of T such that $|S \Delta T| < n/3$ is at most $\sum_{i=0}^{n/3} \binom{n}{i}$, the lower bound on $|\mathbb{V}|$ follows.

of length at least $2|i - j| \geq 2(r - n/16) \geq 3n/4$. It follows that

$$\begin{aligned}
& \mathbb{P}\{G[V, V'] \text{ contains an alternating cycle of length } 3n/4\} \\
& \geq \mathbb{P}\{\forall S_1 \times S_2 \subset V \times V' \text{ such that} \\
& \quad |S_1|, |S_2| \geq \frac{n}{64}, \exists u \in S_1, v' \in S_2, \text{ such that } (u, v') \in E(G)\} \\
& \geq 1 - \left(\frac{n/2}{n/64}\right)^2 \left(1 - \frac{D}{n}\right)^{(n/64)^2} \\
& \geq 1 - \left[(32e)^2 \left(1 - \frac{D}{n}\right)^{n/64}\right]^{n/64} \\
& \geq 1 - \left[\exp\left(2 + 2 \log 32 - \frac{D}{64}\right)\right]^{n/64} \\
& \geq 1 - \exp\left(-Dn/2^{13}\right). \tag{52}
\end{aligned}$$

where the last inequality holds by the assumption that $D/128 \geq 2 + 2 \log 32$.

Finally, since $D/2^{14} \geq 1/20$, combining (52) with (51) and applying a union bound, we get that with probability at least $1 - e^{-Dn/2^{14}}$, $G[V \times V']$ contains an alternating cycle of length $3n/4$ for every subset $V \in \mathbb{V}$. Since the number of left vertices visited by each alternating cycle is at least $3n/8$, which exceeds the maximum overlap (at most $n/3$ by construction) between distinct subsets in \mathbb{V} , these alternating cycles must be distinct. This complete the proof. \square

We now finish the proof of Theorem 6 by combining the previous results. First, combining Lemma 5 and Lemma 6 yields that with probability at least $1 - e^{-K_1/32} - \frac{2}{\beta n k^3}$, there exists a subset $\mathcal{K}_2 \subset \mathcal{K}_1$ such that $|\mathcal{K}_2| = K_2 \geq K_1/16$, and $|U_k| \geq b$ and $|V'_k| \geq b$ for all $k \in \mathcal{K}_2$. Thus (50) holds. Conditioning on $\{U_k, V'_k\}_{k \in \mathcal{K}_2}$ and applying Lemma 7, we get that with probability at least $1 - \exp(-d_{\text{super}} K_2/2^{14})$, G_{super} contains $\exp(K_2/20)$ distinct alternating cycles of length at least $3K_2/4$. Thus the conclusion of Theorem 6 readily follows.

6 Proof of Lemma 2 under the sparse model

In this section, we prove Lemma 2, which, as mentioned in Sect. 3.1, completes the proof of Theorem 2 for the sparse model and, in turn, also for the dense model in view of the reduction in Appendix E.1.

To prove Lemma 2, we apply Algorithm 1 on G with $w_e = \log(\mathcal{P}/\mathcal{Q})(W_e)$ for e in G . Define thresholds

$$\tau_{\text{red}} \triangleq \inf \left\{ x : \mathcal{P} \left(\log \frac{\mathcal{P}}{\mathcal{Q}} \leq x \right) \geq 1/2 \right\} \tag{53}$$

$$\tau_{\text{blue}} \triangleq \sup \left\{ x : \mathcal{Q} \left(\log \frac{\mathcal{P}}{\mathcal{Q}} \geq x \right) \geq 1/2 \right\}. \tag{54}$$

Note that τ_{red} and τ_{blue} are well defined under the assumption that $\mathcal{P} \ll \mathcal{Q}$ and $\mathcal{Q} \ll \mathcal{P}$, which can be assumed WLOG in view of the reduction argument in Appendix E.3.

Assume that Algorithm 1 succeeds. It returns at least $e^{c_6 K_1}$ distinct alternating cycles C of the form

$$C = (v'_1, v_1, P_1, u'_1, u_1, v'_2, v_2, P_2, u'_2, u_2, \dots, v'_r, v_r, P_r, u'_r, u_r),$$

where $r \geq c_7 K_1$ for universal constants c_6, c_7 . Here each P_k is an alternating path of length $\ell \triangleq 4HL + 1$ in G_1 , with $\Delta(P_k) \geq \zeta(\ell - 1)/2 - \tau_{\text{red}}$. All the other edges in C are contained in G_2 . Since $K_1 = c_5 n$, it follows that $|C| \geq r\ell \geq nc_2$ for some constant c_2 . Let m be the perfect matching such that $m\Delta m^* = C$. By choosing $\delta \leq c_2$, we get that $m \in \mathcal{M}_{\text{bad}}$. Moreover, by construction

$$\begin{aligned} \Delta(C) &\geq 3r\tau_{\text{blue}} - 2r\tau_{\text{red}} + \sum_{k=1}^r \Delta(P_k) \geq 3r\tau_{\text{blue}} \\ &\quad - 3r\tau_{\text{red}} + \zeta r(\ell - 1)/2 \geq \zeta r(\ell - 1)/4 \geq nc_4, \end{aligned}$$

for some constant $c_4 > 0$, where the last inequality holds by choosing ℓ large enough such that $\zeta(\ell - 1) \geq 12(\tau_{\text{red}} - \tau_{\text{blue}})$. Since there are at least $e^{c_6 K_1} = e^{c_5 c_6 n}$ distinct such alternating cycles C , the desired (30) follows. To complete the proof, it suffices to show that Algorithm 1 succeeds with probability at least $1 - O(1/n)$.

Path construction Let the family of disjoint sets $L_k \subset V^c$ and $R_k \subset (V^c)'$ be defined in Sect. 4.2. By Theorem 5, with probability $1 - e^{\mathcal{O}(n)}$, there exists $\mathcal{K}_1 \subset [K]$, such that $K_1 = |\mathcal{K}_1| = c_5 n$ with constant

$$c_5 = \frac{\gamma}{16(1 + \epsilon)^{2HL} \exp(7H\alpha + 3\epsilon H/2)}.$$

For all $k \in \mathcal{K}_1$, we have $|L_k| = |R_k| = s$ where $s = (1 + 3\epsilon/4)^{2HL}$. Moreover, each pair of vertices $u \in L_k$ and $v \in R_k$ are connected via an alternating path P of length $\ell = 4HL + 1$ through red edge (i_k, i'_k) and $\Delta(P) \geq 2\zeta HL - \tau_{\text{red}}$.

Sprinkling We need to check that the sprinkling step yields $e^{c_6 K_1}$ distinct alternating cycles of the form

$$C = (v'_1, v_1, P_1, u'_1, u_1, v'_2, v_2, P_2, u'_2, u_2, \dots, v'_r, v_r, P_r, u'_r, u_r)$$

for $r \geq c_7 K_1$. We show this using Theorem 6. We start by specifying the parameters β and η that appear in the statement of Theorem 6.

Recall from Algorithm 1 that $V^* = \{i \in V : \log(\mathcal{P}/\mathcal{Q})(W_{i,i'}) \leq \tau_{\text{red}}\}$, and $W_{i,i'} \sim \mathcal{P}$ for all i . By definition of τ_{red} given in (53) and the right-continuity of the cumulative distribution function, we have $\mathbb{P}\{i \in V^*\} \geq 1/2$ for all i . By the independence of the edge weights, Hoeffding's inequality yields

$$\mathbb{P}\left\{|V^*| < \frac{\gamma n}{4}\right\} \leq \mathbb{P}\left\{\text{Binom}(\gamma n, 1/2) \leq \frac{\gamma n}{4}\right\} \leq e^{-\gamma n/8}. \quad (55)$$

Therefore with probability at least $1 - e^{-\gamma n/8}$, we have $\beta = |V^*|/n \geq \frac{\gamma}{4}$.

To bound the edge probability η/n of the blue edges in G_2 , note that e appears in G_2 as a blue edge if and only if it is a blue edge in G , and $\log(\mathcal{P}/\mathcal{Q})(W_e) \geq \tau_{\text{blue}}$. Therefore $\eta/n \leq d/n$, and

$$\frac{\eta}{n} = \frac{d}{n} \cdot \mathcal{Q} \left(\log \frac{\mathcal{P}}{\mathcal{Q}} \geq \tau_{\text{blue}} \right) \geq \frac{d}{2n}, \quad (56)$$

where the inequality is from the definition of τ_{blue} given in (54) and right-continuity of the cumulative distribution function. In summary, we have parameters

$$\beta \geq \frac{\gamma}{4}, \quad \frac{d}{2} \leq \eta \leq d, \quad s = (1 + 3\epsilon/4)^{2HL}, \quad K_1 = c_5 n.$$

Next, we check that these parameters lead to b , K_1 , κ and d_{super} that satisfy the assumptions of Theorem 6. Indeed, since $d \geq 1$,

$$b = \frac{\beta s \eta}{4} \geq \frac{1}{32} \gamma d (1 + 3\epsilon/4)^{2HL} \geq \frac{1}{32} \gamma (1 + 3\epsilon/4)^{2HL} \geq 4$$

by choosing H large enough; $K_1 = c_5 n \geq 8400$ for large enough n ;

$$\begin{aligned} \kappa &= \frac{2K_1 s \eta}{n} \leq \frac{(1 + 3\epsilon/4)^{2HL} d \gamma}{8(1 + \epsilon)^{2HL} \exp(7H\alpha + 3\epsilon H/2)} \\ &\stackrel{(a)}{\leq} \frac{(1 + 3\epsilon/4)^{2HL} (1 + \epsilon)^2 \gamma}{8(1 + \epsilon)^{2HL} \exp(6H\alpha + 3\epsilon H/2)} \stackrel{(b)}{\leq} \frac{1}{16^2}, \end{aligned}$$

where (a) is from $de^{-\alpha} = (1 + \epsilon)^2$, (b) holds by choosing H large enough;

$$\begin{aligned} d_{\text{super}} &= \frac{1}{32n} K_1 b^2 \eta = \frac{2K_1 \beta^2 s^2 \eta^3}{32^2 n} \\ &\geq \frac{\gamma^3 (1 + 3\epsilon/4)^{4HL} d^3}{32^4 (1 + \epsilon)^{2HL} \exp(7H\alpha + 3\epsilon H/2)} \\ &= \left[\frac{(1 + 3\epsilon/4)^2}{1 + \epsilon} \right]^{2HL} \frac{\gamma^3 (1 + \epsilon)^6}{32^4 \exp(4H\alpha + 3\epsilon H/2)} \geq 256 \log(32e), \end{aligned}$$

where the last inequality is by choosing L to be a large enough constant, since $(1 + 3\epsilon/4)^2/(1 + \epsilon) > 1$. We have checked that all the assumptions of Theorem 6 are satisfied. Thus for all V^* with $|V^*| \geq \gamma n/4$, Theorem 6 gives that conditional on V^* , G_{super} contains at least $e^{K_2/20} \geq e^{c_6 K_1}$ distinct alternating cycles of length at least $3K_2/4 \geq c_7 K_1$ for universal constants c_6, c_7 , with (conditional) probability at least

$$\begin{aligned} \left(1 - e^{-K_1/32} - \frac{2}{\beta n \kappa^3}\right) \left(1 - d_{\text{super}} K_1/2^{18}\right) &= \left(1 - e^{-\mathcal{Q}(n)} - O(1/n)\right) \left(1 - e^{-\mathcal{Q}(n)}\right) \\ &= 1 - O(1/n). \end{aligned}$$

Combined with $\mathbb{P}\{|V^*| < \gamma n/4\} = e^{-\Omega(n)}$ from (55), we have shown that the sprinkling step in Algorithm 1 goes through for constants c_6, c_7 with probability $1 - O(1/n)$.

7 Exponential model

In this section, we focus on the special case of complete graph with exponential weights, where $d = n$, $\mathcal{P} = \exp(\lambda)$, and $\mathcal{Q} = \exp(1/n)$, and prove the lower bound to the optimal reconstruction error given in (11) in Theorem 3.

As a convention, we call an alternating path a $(2\ell - 1)$ -alternating path if it consists of ℓ red edges and $\ell - 1$ blue edges. Recall that $\Delta(P) = \sum_{e \in b(P)} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e) - \sum_{e \in r(P)} \log \frac{\mathcal{P}}{\mathcal{Q}}(W_e)$ for path P in G . The following result shows that with high probability there exist many disjoint $(2\ell - 1)$ -alternating P with large $\Delta(P)$.

Theorem 7 *Suppose that $\lambda = 4 - \epsilon$. There exists absolute constants $\epsilon_0, c_1, c_2, c_3 > 0$, and $n_0 = n_0(\epsilon)$, such that for all $\epsilon \leq \epsilon_0$, $c_1/\epsilon \leq \ell \leq e^{-c_1/\sqrt{\epsilon}}\sqrt{n}$, and $n \geq n_0$, with probability at least $\frac{1}{2} - \frac{\ell^2 e^{c_2/\sqrt{\epsilon}}}{n}$, there is a set S^* of disjoint $(2\ell - 1)$ -alternating paths P with*

$$|S^*| \geq \frac{n}{\ell^2 e^{c_3/\sqrt{\epsilon}}}, \quad (57)$$

such that for every $\ell/3 \leq \ell' \leq \ell$ and every $(2\ell' - 1)$ -alternating subpath Q of P , it holds that

$$\Delta(Q) \geq (\lambda - 1/n)\xi_0\epsilon\ell'. \quad (58)$$

with $\xi_0 = \frac{1}{96}$.

Theorem 7 provides the needed ingredient for proving the negative part of Theorem 3. The proof of the positive part is deferred till Appendix D.

Proof of Theorem 3: negative part As mentioned in Sect. 3, Theorem 3 is a direct consequence of Lemma 1 and Lemma 2 with $c_0, c_2 = e^{-O(1/\sqrt{\epsilon})}$. Therefore we only need to prove Lemma 2 with the desired c_0, c_2 . We choose

$$\gamma = \frac{\epsilon}{8}, \quad \xi_0 = \frac{1}{192}, \quad \tau = e^{\tau_0/\sqrt{\epsilon}}, \quad \ell = \frac{18\tau}{\xi_0\epsilon} \quad (59)$$

for some constant τ_0 to be specified later. Without loss of generality, we assume that ℓ is a multiple of 3. By the same argument that we used in the proof of Theorem 2 to reduce condition (7) to (20), we also assume here that $\lambda = 4 - \epsilon$.

In Step 1 of the two-stage cycle finding scheme, we first apply Theorem 7 to find a set S^* of disjoint $(2\ell - 1)$ -alternating paths in $G_1 = G[V^c \times (V^c)']$ with $|V^c| = (1 - \gamma)n$. Specifically, by shrinking every edge weight in G_1 by a multiplicative factor $1 - \gamma$, we arrive at an instance of the exponential model with n' left (right) vertices, planted weight distribution $\exp(\lambda')$ and null weight distribution $\exp(\frac{1}{n'})$, where $n' = (1 - \gamma)n$

and $\lambda' = \frac{\lambda}{1-\gamma}$. Since $\lambda = 4 - \epsilon$ and $\gamma = \epsilon/8$, it follows that $\lambda' \leq 4 - \epsilon/2 \equiv 4 - \epsilon'$. Replacing (n, λ, ϵ) by $(n', \lambda', \epsilon')$, the same conclusion of Theorem 7 holds for G_1 (without weight shrinkage) with $\zeta_0 = \frac{1}{192}$.

For any alternating path P in S^* , it is centered at a red edge (k, k') with a $(\ell - 1)$ -alternating subpath on each side. Let L_k denote the set of left vertices in the first $(2\ell/3)$ -segment of P and R_k denote the right vertices in the last $(2\ell/3)$ -segment of P . Then we have $|L_k| = |R_k| = s$, where $s = \ell/3$. Moreover, each pair of vertices $u \in L_k$ and $v \in R_k$ is connected via a $(2\ell' - 1)$ -subpath Q of P through the red edge (k, k') consisting of ℓ' red edges and $\ell' - 1$ blue edges, where $\ell' \geq \ell/3$ and $\Delta(Q) \geq (\lambda - 1/n)\zeta_0\epsilon\ell'$. Let \mathcal{K}_1 denote the collection of such indices k , where $K_1 \triangleq |\mathcal{K}_1| = |S^*|$. It follows from Theorem 7 that with probability at least $1 - O(1/n)$, $K_1 \geq c_5 n$ with constant

$$c_5 \triangleq \frac{1}{\ell^2 e^{c_3/\sqrt{\epsilon}}}.$$

Following Step 2 of Algorithm 1, we connect $\{L_k, R_k\}_{k \in \mathcal{K}_1}$ to form alternating cycles in G via sprinkling. Choose

$$\tau_{\text{red}} = \log(n\lambda), \quad \text{and} \quad \tau_{\text{blue}} = \log(n\lambda) - (\lambda - 1/n)\tau.$$

Note that $\log \frac{P}{Q}(W_e) = \log(n\lambda) - (\lambda - 1/n)W_e$. Thus $V^* = V$ and equivalently subgraph G_2 is the subgraph of G that contains every red edge in $V \times V'$, and every blue edge $e \in [n] \times [n] \setminus (V^c \times V^{c'})$ if $W_e \leq \tau$. Then we apply Theorem 6 with $V^* = V$ to show there exist exponentially many distinct alternating cycles via sprinkling. We start by specifying the parameters β and η that appear in the statement of Theorem 6. Note that $\beta = \frac{|V^*|}{n} = \gamma$ and the average blue degree is

$$\eta = n\mathbb{P}\{\exp(1/n) \leq \tau\} = n(1 - e^{-\tau/n}). \quad (60)$$

Using $e^{-x} \geq 1 - x$ and $e^{-x} \leq 1 - x/2$ for $x \in [0, 1]$, and $\tau/n \leq 1$ for all sufficiently large n , we have $\tau/2 \leq \eta \leq \tau$. Next, we check the parameters above lead to b , K_1 , κ , and d_{super} that satisfy the assumptions of Theorem 6. In particular,

$$b = \frac{\beta s \eta}{4} \geq \frac{\gamma \ell \tau}{24} = \frac{\epsilon \ell \tau}{192} \geq 4,$$

by choosing the constant τ_0 in (59) sufficiently large; $K_1 = c_5 n \geq 8400$ for all large enough n ;

$$\kappa = \frac{2K_1 s \eta}{n} \leq \frac{2\ell \tau}{3\ell^2 e^{c_3/\sqrt{\epsilon}}} = \frac{\zeta_0 \epsilon}{27 e^{c_3/\sqrt{\epsilon}}} \leq \frac{1}{16^2}$$

for all sufficiently small ϵ ; and

$$d_{\text{super}} = \frac{1}{32n} K_1 b^2 \eta \geq \frac{1}{32\ell^2 e^{c_3/\sqrt{\epsilon}}} \left(\frac{\epsilon \ell \tau}{192} \right)^2 \frac{\tau}{2} = \frac{\epsilon^2 e^{3\tau_0/\sqrt{\epsilon}}}{64 \cdot (192)^2 e^{c_3/\sqrt{\epsilon}}} \geq 256 \log(32e), \quad (61)$$

by choosing τ_0 to be a sufficiently large constant. Having verified all assumptions of Theorem 6, we conclude that G_{super} contains at least $e^{K_2/20} \geq e^{c_6 K_1}$ distinct alternating cycles of length at least $3K_2/4 \geq c_7 K_1$ for universal constants c_6, c_7 , with probability at least

$$\begin{aligned} & \left(1 - e^{-K_1/32} - \frac{2}{\beta n \kappa^3} \right) \left(1 - d_{\text{super}} K_1 / 2^{18} \right) \\ &= \left(1 - e^{-\Omega(n)} - O(1/n) \right) \left(1 - e^{-\Omega(n)} \right) \\ &= 1 - O(1/n). \end{aligned}$$

In conclusion, the sprinkling step yields $e^{c_6 K_1}$ distinct alternating cycles (in fact here one such cycle suffices) $C \in \mathcal{C}$ of the form

$$C = (v'_1, v_1, P_1, u'_1, u_1, v'_2, v_2, P_2, u'_2, u_2, \dots, v'_r, v_r, P_r, u'_r, u_r)$$

for $r \geq c_7 K_1$. It follows that $|C| \geq 2r\ell/3 = c_2 n$, where

$$c_2 \triangleq 2c_7 c_5 \ell / 3 = \frac{2c_7}{3\ell e^{c_3/\sqrt{\epsilon}}} = \frac{c_7 \zeta_0 \epsilon}{27 e^{(c_3+\tau_0)/\sqrt{\epsilon}}}.$$

Moreover, by construction of G_2 ,

$$\begin{aligned} \Delta(C) &\geq 3r\tau_{\text{blue}} - 2r\tau_{\text{red}} + \sum_{k=1}^r \Delta(P_k) \geq -3r(\lambda - 1/n)\tau + (\lambda - 1/n)\zeta_0 \epsilon r \ell / 3 \\ &\geq (\lambda - 1/n)r\zeta_0 \epsilon \ell / 6, \end{aligned}$$

where the last inequality holds by the choice of $\zeta_0 \epsilon \ell = 18\tau$ in (59). It follows that

$$\frac{\mu_W(\mathcal{M}_{\text{bad}})}{\mu_W(M^*)} \geq |\mathcal{C}| e^{\Delta(C)} \geq e^{c_6 K_1} e^{(\lambda - 1/n)c_7 K_1 \zeta_0 \epsilon \ell / 6} \geq \exp(c_0 n),$$

where

$$c_0 \geq c_7 K_1 \zeta_0 \epsilon \ell / 6 \geq \frac{c_7 \zeta_0 \epsilon \ell}{6\ell^2 e^{c_3/\sqrt{\epsilon}}} = \frac{c_7 \zeta_0^2 \epsilon^2}{6 \cdot 18 e^{(c_3+\tau_0)/\sqrt{\epsilon}}}.$$

Theorem 3 then readily follows by combining Lemmas 1 and 2. \square

In the sequel, we proceed to prove Theorem 7. A more direct approach is to define S as the set of $(2\ell - 1)$ -alternating paths P with large $\Delta(P)$ and show that $\mathbb{E}[|S|]$ is large while $\text{Var}(|S|) \ll (\mathbb{E}[|S|])^2$ so that $|S|$ concentrates on its mean. Unfortunately,

this idea fails as the second moment of $|S|$ blows up for $\ell = \Theta(n)$. This is because conditioning on finding an alternating path P with large $\Delta(P)$, it is very likely to have a large number of paths P' with large $\Delta(P')$ overlapping with P . These clusters of overlapping paths induce an excessive contribution to the second moment. To address this issue, we adopt the notion of *uniformity* introduced in [9] for studying minimum mean-weight cycles.

Recall that for a set T of edges, $\mathbf{r}(T)$ and $\mathbf{b}(T)$ denote the set of red and blue edges in T , respectively. Furthermore, define

$$\mathbf{wt}_r(T) = \sum_{e \in \mathbf{r}(T)} W_e,$$

$$\mathbf{wt}_b(T) = \sum_{e \in \mathbf{b}(T)} W_e.$$

Let \mathcal{P}_ℓ denote the set of $(2\ell - 1)$ alternating paths with ℓ red edges and $\ell - 1$ blue edges. We agree upon that each $P \in \mathcal{P}_\ell$ is oriented so that it starts in the left vertex set and ends in the right vertex set. Let $\phi_1, \phi_2, \dots, \phi_{|\mathbf{r}(P)|}$ (resp. $\psi_1, \psi_2, \dots, \psi_{|\mathbf{b}(P)|}$) denote the sequence of red (resp. blue) edge weights in this order. Define

$$\begin{aligned} \mathbf{devr}(P) &= \sup_{1 \leq k \leq |\mathbf{r}(P)|} \left| \sum_{j=1}^k \frac{\phi_j |\mathbf{r}(P)|}{\mathbf{wt}_r(P)} - k \right|, \\ \mathbf{devb}(P) &= \sup_{1 \leq k \leq |\mathbf{b}(P)|} \left| \sum_{j=1}^k \frac{\psi_j |\mathbf{b}(P)|}{\mathbf{wt}_b(P)} - k \right| \end{aligned}$$

which characterize the maximum fluctuation of edge weights on path P .

Definition 2 (*Lightness and uniformity*) We say an alternating path P is (a, b, η) -light, if

$$\begin{aligned} |\mathbf{wt}_r(P) - a \cdot |\mathbf{r}(P)|| &\leq \eta/2 \\ |\mathbf{wt}_b(P) - b \cdot |\mathbf{b}(P)|| &\leq \eta/2; \end{aligned}$$

and A -uniform, if $\mathbf{devr}(P) \leq A$ and $\mathbf{devb}(P) \leq A$.

Let us proceed to the definition of S . Define

$$S \triangleq \{P \in \mathcal{P}_\ell : P \text{ is } (a, b, \eta)\text{-light and } A\text{-uniform}\}.$$

with

$$a = \frac{2}{\lambda}, \quad b = \frac{2 - \zeta}{\lambda}. \quad (62)$$

(We will choose $\eta = 1$, $\zeta = \frac{\epsilon}{4}$, and $A = \Theta(\frac{1}{\sqrt{\epsilon}})$ later.) Then for any $P \in \mathcal{P}_\ell \cap S$, we have

$$\begin{aligned}\Delta(P) &= -(\lambda - 1/n) [\text{wt}_b(P) - \text{wt}_r(P)] \geq -(\lambda - 1/n) [b(\ell - 1) - a\ell - \eta] \\ &\geq (\lambda - 1/n) \frac{\epsilon\ell}{96}.\end{aligned}$$

Furthermore, using the A -uniformity of P , it is not hard to verify that for every $(2\ell' - 1)$ -alternating subpath Q of P with $\ell/3 \leq \ell' \leq \ell$, $\Delta(Q) \geq (\lambda - 1/n) \frac{\epsilon\ell'}{96}$ (see forthcoming (82)). Thus, to prove Theorem 7, the key remaining challenge is to show S contains a large vertex-disjoint subcollection S^* with $|S^*| = \Omega(n/\ell^2)$.

The intuition that we require the mean weight of red (blue) edges in P to be around $2/\lambda$ is as follows. Given $P \in \mathcal{P}_\ell$, $\text{wt}_r(P)$ is distributed as a sum of ℓ i.i.d. $\exp(\lambda)$ random variables, while $\text{wt}_b(P)$ is distributed as a sum of $\ell - 1$ i.i.d. $n \cdot \exp(1)$. Conditional on $\text{wt}_r(P)$ being close to $\text{wt}_b(P)$, we expect⁶ that the mean weight for both $\text{wt}_r(P)$ and $\text{wt}_b(P)$ are close to $2/\lambda$. Moreover, we require the mean weight of blue edges to be slightly below $2/\lambda$ so that $\Delta(P)$ is positive.

Note that we further restrict the alternating path P in S to be uniform, in the sense that the mean weight of red (resp. blue) edges in every subpath Q of P concentrates around a (resp. b). As we will see in the next section, the uniformity can be interpreted as requiring an exp-minus-one random walk conditioned on returning to the origin to have a restricted range, which is shown to hold with a sufficiently large probability. This implies that after restricting to uniform alternating paths, the first moment $\mathbb{E}[|S|]$ is still large. Moreover, the number of uniform alternating paths has small enough variance for the second moment method to go through.

7.1 Exponential random walks

In order to study the pathwise fluctuation of the edge weights on a given path, let us consider the following problem. Let $X_1, X_2, \dots, X_\ell \stackrel{\text{i.i.d.}}{\sim} \exp(\mu)$ and let $X = \sum_{i=1}^\ell X_i$. Define a process

$$R_j = \sum_{i=1}^j \left(\frac{X_i}{X} \ell - 1 \right), \quad 0 \leq j \leq \ell.$$

Thanks to the property of the exponential distribution, conditional on any realization of X , $\{\frac{X_i}{X} : 1 \leq i \leq \ell\}$ are uniformly distributed on the $(\ell - 1)$ -dimensional probability simplex, regardless of the value of μ . In particular, $\{\frac{X_i}{X} : 1 \leq i \leq \ell\}$ and X are independent.

Lemma 8 *The process $(R_j : 0 \leq j \leq \ell)$ is independent of X . Furthermore, $(R_j : 0 \leq j \leq \ell)$ is distributed as an exp-minus-one random walk started from the origin and conditioned to return to the origin at time ℓ (known as the exp-minus-one ℓ -bridge).*

⁶ Indeed, using the density of sum of exponentials (see (87) in Appendix B), the probability that $\text{wt}_r(P)/\ell$ and $\text{wt}_r(P)/(\ell - 1)$ are both close to a given value x is proportional to $x^{2\ell-3} e^{-\ell x(\lambda + \frac{\ell-1}{n\ell})}$, which, for large n and ℓ , is approximately maximized at $x = \frac{2}{\lambda}$.

Proof Since $X = \sum_{i=1}^{\ell} X_i$ follows the Erlang(ℓ, μ) distribution as defined in Appendix B, the joint conditional density of X_1, \dots, X_{ℓ} given X takes the form

$$f(x_1, \dots, x_{\ell} | x) = \frac{\prod_{i=1}^{\ell} \mu \exp(-\mu x_i) \mathbf{1}_{\{x=\sum_{i=1}^{\ell} x_i\}}}{\mu^{\ell} x^{\ell-1} \exp(-\mu x)/(\ell-1)!} = \frac{(\ell-1)!}{x^{\ell-1}} \mathbf{1}_{\{x=\sum_{i=1}^{\ell} x_i\}}.$$

Therefore, the distribution of $(X_j/X : 1 \leq j \leq \ell)$ conditional on X is uniform on the $(\ell-1)$ -dimensional simplex and does not depend on X . Since $(R_j : 0 \leq j \leq \ell)$ is a function of $(X_j/X : 1 \leq j \leq \ell)$, it is also independent of X . We have

$$\begin{aligned} (R_j : 0 \leq j \leq \ell) &\stackrel{d}{=} (R_j : 0 \leq j \leq \ell) \text{ conditional on } X = \ell \\ &\stackrel{d}{=} \left(\sum_{i=1}^j (X_i - 1) : 0 \leq j \leq \ell \right) \text{ conditional on } \sum_{i=1}^{\ell} (X_i - 1) = 0 \end{aligned}$$

is an exp-minus-one ℓ -bridge. \square

Fix an alternating path $P \in \mathcal{P}_{\ell}$. Applying Lemma 8, we conclude that, crucially, $\text{wt}_r(P)$ and $\{W_e/\text{wt}_r(P) : e \in r(P)\}$ are independent. Furthermore, $\{W_e/\text{wt}_r(P) : e \in r(P)\}$ has the same distribution as $\{\frac{X_i}{X} : 1 \leq i \leq \ell\}$. Therefore

$$\text{devr}(P) \stackrel{d}{=} \max_{0 \leq j \leq \ell} |R_j|. \quad (63)$$

The similar conclusion applies to blue edge weights $\text{wt}_b(C)$ and $\{W_e/\text{wt}_b(C) : e \in b(C)\}$.

Adapted from [7, Lemma 2.3], the following lemma bounds the probability that the range of an exp-minus-one ℓ -bridge is at most A ; this result is crucial for lower bounding the first moment of A -uniform alternating paths. More precise version of the results can be also found in [9, Lemma 3.9 and Equation (8)].

Lemma 9 *There exist universal constants $c_0, c'_0 > 0$ such that for all $A \geq 1$ and $\ell \geq A^2$*

$$\exp\left(-\frac{c_0 \ell}{A^2}\right) \leq \mathbb{P}\left\{\max_{0 \leq j \leq \ell} |R_j| \leq A\right\} \leq \exp\left(-\frac{c'_0 \ell}{A^2}\right). \quad (64)$$

Applying (64) separately to both red and blue edges, we conclude that for any $P \in \mathcal{P}_{\ell}$,

$$p_{\ell} \triangleq \mathbb{P}\{P \text{ is } A\text{-uniform}\} \geq \exp\left(-\frac{2c_0 \ell}{A^2}\right). \quad (65)$$

Furthermore, the event that P is A -uniform is independent of $\{\text{wt}_r(P), \text{wt}_b(P)\}$ and hence also the event that P is (a, b) -light.

7.2 First moment estimates

In the remainder of this section, all the expectations are conditioned on $M^* = m$ for some fixed matching $m \in \mathcal{M}$, e.g., the identity.

Lemma 10 *Suppose $\lambda \leq 4$ and $\eta \leq 1$. There exist a universal constant $c_0 > 0$, such that for all $A \geq 1$ and $\ell \geq A^2$,*

$$\mathbb{E}[|S|] \geq n \frac{\eta^2 \lambda}{8e^3 b \ell} \left(2be^{-\frac{b}{n}}\right)^{\ell-1} e^{-\frac{\ell^2}{n}} p_\ell \quad (66)$$

$$\geq n \frac{\eta^2 \lambda}{16e^3 b^2 \ell} \left(2be^{-\frac{c_0}{A^2} - \frac{\ell+b}{n}}\right)^\ell. \quad (67)$$

Proof Recall that the planted edge weights are i.i.d. $\exp(\lambda)$ and unplanted edge weights are i.i.d. $\exp(1/n)$. Recall that for $\ell \geq 1$, \mathcal{P}_ℓ denotes the set of alternating paths of length $2\ell - 1$ with ℓ red edges and $(\ell - 1)$ blue edges. We have that

$$|\mathcal{P}_\ell| = n(n-1) \cdots (n-\ell+1).$$

To see this, note that there are $n(n-1) \cdots (n-\ell+1)$ different choices for ℓ left vertices on P . The right vertices are automatically fixed according to the underlying true matching M^* . Write $|\mathcal{P}_\ell| = n^\ell \exp\left\{\sum_{k=0}^{\ell-1} \log\left(1 - \frac{k}{n}\right)\right\}$. Note that by monotonicity,

$$-\frac{1}{\ell} \sum_{k=0}^{\ell-1} \log\left(1 - \frac{k}{n}\right) \leq \frac{n}{\ell} \int_0^{\ell/n} -\log(1-x) dx = F\left(\frac{\ell}{n}\right)$$

where

$$F(\delta) \triangleq \frac{1}{\delta}(\delta + (1-\delta)\log(1-\delta)) \quad (68)$$

is increasing in $\delta \in [0, 1]$ and satisfies $F(\delta) \leq \delta$ for all $\delta \in [0, 1]$. Then

$$\exp\left\{\sum_{k=0}^{\ell-1} \log\left(1 - \frac{k}{n}\right)\right\} \geq \exp\left\{-\ell^2/n\right\}.$$

In conclusion, we get that

$$|\mathcal{P}_\ell| \geq \left(ne^{-\ell/n}\right)^\ell. \quad (69)$$

Fix an alternating path $P \in \mathcal{P}_\ell$. Next we bound the probability that P is (a, b, η) -light. Recall that $r(P)$ denote the set of red (planted) edges and $\text{wt}_r(P)$ denote the

total of their weights. Then $|\mathbf{r}(P)| = \ell$ and $\text{wt}_r(P) \sim \text{Erlang}(\ell, \lambda)$. Using the Erlang density function in (87), we have

$$\begin{aligned}\mathbb{P}\left\{a\ell - \frac{\eta}{2} \leq \text{wt}_r(P) \leq a\ell\right\} &= \int_{a\lambda\ell - \frac{\lambda\eta}{2}}^{a\lambda\ell} \frac{x^{\ell-1} e^{-x}}{(\ell-1)!} dx \\ &\geq \frac{\eta\lambda}{2(\ell-1)!} x^{\ell-1} e^{-x} \Big|_{x=a\lambda\ell} \\ &= \frac{\eta\lambda}{2(\ell-1)!} (a\lambda\ell)^{\ell-1} e^{-a\lambda\ell} = \frac{\eta\lambda}{2(\ell-1)!} (2\ell)^{\ell-1} e^{-2\ell},\end{aligned}$$

where the inequality holds because $x^{\ell-1} e^{-x}$ is decreasing for $x \geq \ell-1$ and by the assumptions $a\lambda = 2$, $\lambda \leq 4$ and $\eta \leq 1$ so that $a\lambda\ell - \frac{\eta\lambda}{2} \geq \ell-1$.

Similarly, recall that $\mathbf{b}(P)$ denote the set of blue (unplanted) edges and $\text{wt}_b(P)$ denote the total of their weights. Then $|\mathbf{b}(P)| = \ell-1$ and $\text{wt}_b(P) \sim \text{Erlang}(\ell-1, \frac{1}{n})$. Thus

$$\begin{aligned}\mathbb{P}\left\{b(\ell-1) \leq \text{wt}_b(P) \leq b(\ell-1) + \frac{\eta}{2}\right\} &= \int_{\frac{b(\ell-1)}{n}}^{\frac{b(\ell-1)}{n} + \frac{\eta}{2n}} \frac{x^{\ell-2} e^{-x}}{(\ell-2)!} dx \\ &\geq \frac{1}{(\ell-2)!} \left(\frac{b(\ell-1)}{n}\right)^{\ell-2} \int_{\frac{b(\ell-1)}{n}}^{\frac{b(\ell-1)}{n} + \frac{\eta}{2n}} e^{-x} dx \\ &= \frac{1}{(\ell-2)!} \left(\frac{b(\ell-1)}{n}\right)^{\ell-2} e^{-\frac{b(\ell-1)}{n}} \left(1 - e^{-\frac{\eta}{2n}}\right) \\ &\geq \frac{1}{(\ell-2)!} \left(\frac{b(\ell-1)}{n}\right)^{\ell-2} e^{-\frac{b(\ell-1)}{n}} \frac{\eta}{4n},\end{aligned}$$

where the last inequality holds because $1 - e^{-x} \geq x/2$ for $0 \leq x \leq 1$.

Since $\text{wt}_r(P)$ and $\text{wt}_b(P)$ are independent, it follows from the last two displayed equations that

$$\begin{aligned}\mathbb{P}\{P \text{ is } (a, b, \eta) - \text{light}\} &\geq \mathbb{P}\left\{a\ell - \frac{\eta}{2} \leq \text{wt}_r(P) \leq a\ell, b(\ell-1) \leq \text{wt}_b(P) \leq b(\ell-1) + \frac{\eta}{2}\right\} \\ &\geq \frac{\eta^2 \lambda \ell}{8b\ell!(\ell-1)!} \left(\frac{2b\ell(\ell-1)}{n}\right)^{\ell-1} e^{-2\ell - \frac{b(\ell-1)}{n}} \\ &\stackrel{(a)}{\geq} \frac{\eta^2 \lambda}{8eb\sqrt{\ell(\ell-1)}} \left(\frac{2be^2}{n}\right)^{\ell-1} e^{-2\ell - \frac{b(\ell-1)}{n}} \\ &\geq \frac{\eta^2 \lambda}{8e^3 b \ell} \left(\frac{2b}{n} e^{-\frac{b}{n}}\right)^{\ell-1}\end{aligned}\tag{70}$$

where (a) holds due to $n! \leq e n^{n+1/2} e^{-n}$.

Finally, combining (69), (65), (70) and using the independence of the events $\{P \text{ is } (a, b, \eta)\text{-light}\}$ and $\{P \text{ is } A\text{-uniform}\}$, we get

$$\begin{aligned}\mathbb{E}[|S|] &= |\mathcal{P}_\ell| \cdot \mathbb{P}\{P \text{ is } (a, b, \eta)\text{-light}\} \cdot \mathbb{P}\{P \text{ is } A\text{-uniform}\} \\ &\geq \left(ne^{-\ell/n}\right)^\ell \cdot \frac{\eta^2 \lambda}{8e^3 b \ell} \left(\frac{2b}{n} e^{-\frac{b}{n}}\right)^{\ell-1} \cdot p_\ell \\ &\geq n \frac{\eta^2 \lambda}{16e^3 b^2 \ell} \left(2be^{-\frac{c_0}{A^2} - \frac{\ell+b}{n}}\right)^\ell.\end{aligned}$$

□

7.3 Second-moment estimates

Lemma 11 *There exist absolute constants C_1, C_2 , and ϵ_0 such that the following holds. Let $\lambda = 4 - \epsilon$ for some $0 < \epsilon \leq \epsilon_0$. Then there exists $n_0 = n_0(\epsilon)$ such that upon choosing $A = \left\lceil \frac{C_1}{\sqrt{\epsilon}} \right\rceil$, $\zeta = \epsilon/4$, and $\eta = 1$, for any ℓ with $A^2 \leq \ell \leq e^{-C_2 A} \sqrt{n}$ and $n \geq n_0$, it holds that*

$$\text{Var}(|S|) \leq \mathbb{E}[|S|] \times \left(1 + \frac{e^{C_2/\sqrt{\epsilon}} \ell^2}{n} \mathbb{E}[|S|]\right).$$

Remark 4 In the above estimate it is crucial to get $\frac{\text{Var}(|S|) - \mathbb{E}[|S|]^2}{\mathbb{E}[|S|]^2} = O\left(\frac{\ell^2}{n}\right)$, so that we can extract from S a vertex-disjoint subcollection S^* of $|S^*| = \Omega(n/\ell^2)$ in Sect. 7.4 by applying Turán's theorem. This turns out to be instrumental to ensure that the super graph in the sprinkling stage is supercritical as shown in (61). As a result, we need to be careful with terms that are polynomial in ℓ in the second moment computation.

Proof Note that

$$\begin{aligned}\text{Var}(|S|) &= \mathbb{E}\left[|S|^2\right] - (\mathbb{E}[|S|])^2 \\ &= \sum_{P, P' \in \mathcal{P}_\ell} (\mathbb{P}\{P \in S, P' \in S\} - \mathbb{P}\{P \in S\} \mathbb{P}\{P' \in S\}) \\ &\stackrel{(a)}{=} \sum_{P \in \mathcal{P}_\ell} \sum_{P' \in \mathcal{P}_\ell: |P \cap P'| \geq 1} (\mathbb{P}\{P \in S, P' \in S\} - \mathbb{P}\{P \in S\} \mathbb{P}\{P' \in S\}) \\ &\leq \sum_{P \in \mathcal{P}_\ell} \sum_{P' \in \mathcal{P}_\ell: |P \cap P'| \geq 1} \mathbb{P}\{P \in S, P' \in S\} \\ &= \sum_{P \in \mathcal{P}_\ell} \mathbb{P}\{P \in S\} \left(1 + \sum_{P' \in \mathcal{P}_\ell: |P \cap P'| \geq 1, P' \neq P} \mathbb{P}\{P' \in S \mid P \in S\}\right), \quad (71)\end{aligned}$$

where (a) holds because the weights in P and those in P' are mutually independent if $P \cap P' = \emptyset$.

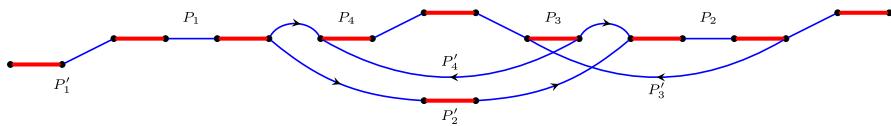


Fig. 5 An example of alternating paths P and P' , with $\ell = 8$, $k = 4$, $m = 8$. P' is divided into $2k + 1 = 9$ segments (subpaths) $\{P'_1, P'_2, P'_3, P'_4, P'_5\}$, where in this example P'_5 is empty. The subpaths P_1, P_2, P_3, P_4 in the intersection $P \cap P'$ lie on the horizontal line; below this horizontal line are the segments P'_1, P'_2, P'_3, P'_4 in $P' \setminus P$; above the horizontal line are the edges in $P \setminus P'$. From the agreement on the orientation, all the red edges are traversed from left to right in both P and P' (color figure online)

Fix any $P \in \mathcal{P}_\ell$. Recall that P is alternating paths with ℓ red edges and $\ell - 1$ blue edges which starts from a left vertex and ends with a right vertex, and whose first and last edge are both red. We group the summands of the inner sum according to how P' overlaps with P (see Fig. 5). For each $P' \neq P$ with $|P \cap P'| \geq 1$:

- $P' \cap P$ consists of k disjoint alternating paths $\mathcal{P} = (P_1, \dots, P_k)$ for some $k \geq 1$. Let $|P_i| = m_i$ and $m = \sum_{i=1}^k m_i < 2\ell$. Note that the first edge and the last edge in P_i must be red. Then $m_i \geq 1$ must be odd and P_i has $(m_i + 1)/2$ red edges and $(m_i - 1)/2$ blue edges.
- $P' \setminus P$ consists of $k + 1$ disjoint alternating paths $\mathcal{P}' \triangleq (P'_1, \dots, P'_{k+1})$, where we allow P'_1 and P'_{k+1} to be possibly empty. Let $|P'_i| = m'_i$ and $m' = \sum_{i=1}^{k+1} m'_i = 2\ell - 1 - m$. Note that for $i \in \{1, k + 1\}$, m'_i must be even and P'_i has $m'_i/2$ red edges and $m'_i/2$ blue edges. For $2 \leq i \leq k$, the first edge and the last edge in each P'_i must be blue. Hence, $m'_i \geq 1$ must be odd and P'_i has $(m'_i - 1)/2$ red edges and $(m'_i + 1)/2$ blue edges.

For the sake of enumeration, let us agree on the orientation so that all alternating paths with $2\ell + 1$ edges start from a left vertex. This way, all the red edges are traversed from left to right in both P and P' . Moreover, each P_i starts from a left vertex, and each P'_i starts from a right vertex for $2 \leq i \leq k + 1$; P_1 starts from a left vertex.

Since P is fixed, to specify P' , it suffices to specify the alternating paths \mathcal{P}' that constitute $P' \setminus P$, which further reduces to specifying for each P'_i its start and end points, as well as the internal vertices.

- We first specify the start and end points of P'_i for $2 \leq i \leq k$, which are chosen from the vertices on P . It suffices to specify P_i for $1 \leq i \leq k$. For a given sequence of lengths (m_1, m_2, \dots, m_k) , to specify P_i for $1 \leq i \leq k$, it suffices to specify the starting point of each P_i . In total, there are at most ℓ^k choices for the starting points of P_i , $i \in [k]$. Next, since $\sum_{i=1}^k m_i = m$ and $m_i \geq 1$, there are at most $\binom{m-1}{k-1}$ choices of (m_1, \dots, m_k) . Hence, in total there are at most $\binom{m-1}{k-1} \ell^k$ choices for the start and end points of P'_i for $1 \leq i \leq k$.
- Next, we specify the length configuration $(m'_1, m'_2, \dots, m'_{k+1})$ of \mathcal{P}' . Since $m' = \sum_{i=1}^{k+1} m'_i$, $m'_1, m'_{k+1} \geq 0$, and $m'_i \geq 1$ for $2 \leq i \leq k$, in total there are at most $\binom{m'+1}{k}$ choices for $(m'_1, m'_2, \dots, m'_{k+1})$.
- Next we specify the vertices of P'_1 if $m'_1 \neq 0$. Note that the end point of P'_1 is on P and has already been chosen. Thus it remains to choose the starting points of red edges in P'_1 which must be in the left vertex set per the agreed-upon orientation.

Since there are $m'_1/2$ red edges in P'_1 , there are at most $n^{m'_1/2}$ different choices for the vertices of P'_1 that are not on P . Analogously we can specify the vertices of P'_{k+1} if $m'_{k+1} \neq 0$ and there are at most $n^{m'_{k+1}/2}$ different choices for the vertices of P'_{k+1} that are not on P .

- Finally, we specify the internal vertices of P'_i for each $2 \leq i \leq k$. Note that there are $\sum_{i=2}^k (m'_i - 1)/2$ red edges in $\{P'_i : 2 \leq i \leq k\}$. For each red edge, it suffices to specify its starting point, which must be in the left vertex set per the agreed-upon orientation.

Thus in total there are at most $\prod_{i=2}^k n^{(m'_i-1)/2}$ different choices for the interval vertices of P'_i for $2 \leq i \leq k$.

Putting the above points together, we conclude that there are at most

$$\ell^k \binom{m-1}{k-1} \binom{m'+1}{k} n^{\frac{m'-(k-1)}{2}} \quad (72)$$

different choices of alternating paths P'_1, \dots, P'_{k+1} of total length m' .

Next we bound each conditional probability in (71) from above. Note that $P' \in S$ means it is (a, b, η) -light and A -uniform (see Definition 2). Since P and P' are overlapping, conditioned on the edge weights on P , the event that P' is (a, b, η) -light and the event that P' is A -uniform are no longer independent. Nevertheless, these two events imply that the i -th alternating path P'_i satisfies:

$$\begin{aligned} (|r(P'_i)| - 2A) \left(a - \frac{\eta}{2\ell} \right) &\leq \text{wt}_r(P'_i) \leq (|r(P'_i)| + 2A) \left(a + \frac{\eta}{2\ell} \right) \\ (|b(P'_i)| - 2A) \left(b - \frac{\eta}{2(\ell-1)} \right) &\leq \text{wt}_b(P'_i) \leq (|b(P'_i)| + 2A) \left(b + \frac{\eta}{2(\ell-1)} \right). \end{aligned}$$

Note that $\sum_{i=1}^{k+1} |r(P'_i)| = \frac{m'-(k-1)}{2} \triangleq L$ and $\sum_{i=1}^{k+1} |b(P'_i)| = \frac{m'+(k-1)}{2} \triangleq M$. Thus, summing over all $1 \leq i \leq k+1$ yields that

$$\begin{aligned} (L - 2(k+1)A) \left(a - \frac{\eta}{2\ell} \right) &\leq \text{wt}_r(P' \setminus P) \leq (L + 2(k+1)A) \left(a + \frac{\eta}{2\ell} \right), \\ (M - 2(k+1)A) \left(b - \frac{\eta}{2(\ell-1)} \right) &\leq \text{wt}_b(P' \setminus P) \leq (M + 2(k+1)A) \left(b + \frac{\eta}{2(\ell-1)} \right). \end{aligned}$$

Let \mathcal{E} denote the event such that the last displayed equation holds. Then

$$\begin{aligned} \mathbb{P}\{P' \in S \mid P \in S\} &= \mathbb{P}\{P' \text{ is } (a, b, \eta)\text{-light and } A\text{-uniform} \mid P \in S\} \\ &\leq \mathbb{P}\{\mathcal{E} \cap \{P' \text{ is } (a, b, \eta)\text{-light and } A\text{-uniform}\} \mid P \in S\} \\ &= \mathbb{P}\{\mathcal{E} \cap \{P' \text{ is } (a, b, \eta)\text{-light}\} \mid P \in S \times \mathbb{P}\{P' \text{ is } A\text{-uniform} \mid \mathcal{E}, P' \text{ is } (a, b, \eta)\text{-light}, P \in S\} \\ &\leq \mathbb{P}\{\mathcal{E} \mid P \in S\} \times \mathbb{P}\{P' \text{ is } A\text{-uniform} \mid \mathcal{E}, P' \text{ is } (a, b, \eta)\text{-light}, P \in S\}. \end{aligned} \quad (73)$$

We proceed to separately bound the two probability terms in the RHS of the last displayed equation. First, we bound $\mathbb{P}\{\mathcal{E} \mid P \in S\}$. Note that the edge weights in $P' \setminus P$ are independent from the edge weights in P . Thus, $\mathbb{P}\{\mathcal{E} \mid P \in S\} = \mathbb{P}\{\mathcal{E}\}$. Let

$$g_\ell(x) = e^{-x} x^{\ell-1} / (\ell - 1)! \mathbf{1}_{\{x \geq 0\}}$$

denote the probability density function of Erlang($\ell, 1$) (cf. Appendix B). Then we have that for $L \geq 1$,

$$\begin{aligned} & \mathbb{P}\left\{(L - 2(k+1)A) \left(a - \frac{\eta}{2\ell}\right) \leq \text{wt}_r(P' \setminus P) \leq (L + 2(k+1)A) \left(a + \frac{\eta}{2\ell}\right)\right\} \\ &= \int_{\lambda(L-2(k+1)A)(a-\frac{\eta}{2\ell})}^{\lambda(L+2(k+1)A)(a+\frac{\eta}{2\ell})} g_L(x) dx. \\ &\leq \frac{(\lambda(L + 2(k+1)A)(a + \frac{\eta}{2\ell}))^{L-1}}{(L-1)!} \exp\left(-\lambda(L - 2(k+1)A)(a - \frac{\eta}{2\ell})\right) \\ &\leq \frac{(\lambda a L)^{L-1}}{(L-1)!} e^{-\lambda La} e^{2(k+1)A(1+\lambda a) + \frac{\eta L}{2\ell a}(1+\lambda a)} = \frac{(2L)^{L-1}}{(L-1)!} e^{-2L} e^{6(k+1)A + \frac{3\eta\lambda L}{4\ell}}, \quad (74) \end{aligned}$$

where the inequality holds due to

$$\int_u^v g_L(x) dx \leq \frac{v^{L-1}}{(L-1)!} \int_u^v e^{-x} dx \leq \frac{v^{L-1}}{(L-1)!} e^{-u}.$$

Similarly,

$$\begin{aligned} & \mathbb{P}\left\{(M - 2(k+1)A) \left(b - \frac{\eta}{(2\ell-1)}\right) \leq \text{wt}_b(P' \setminus P) \leq (M + 2(k+1)A) \left(b + \frac{\eta}{2\ell-1}\right)\right\} \\ &\leq \int_0^{(M+2(k+1)A)(b+\frac{\eta}{2\ell})/n} g_M(x) dx \\ &\leq \frac{1}{n^M M!} \left((M + 2(k+1)A)(b + \frac{\eta}{2\ell})\right)^M \\ &\leq \frac{(Mb)^M}{n^M M!} \exp\left(2(k+1)A + \frac{\eta M}{2b\ell}\right), \quad (75) \end{aligned}$$

where the second inequality holds because

$$\int_0^u g_\ell(x) dx \leq \frac{1}{(\ell-1)!} \int_0^u x^{\ell-1} dx = \frac{u^\ell}{\ell!}, \quad \forall u \geq 0.$$

Note that $\text{wt}_r(P' \setminus P)$ and $\text{wt}_b(P' \setminus P)$ are independent. Combining (74) and (75) gives that⁷

$$\begin{aligned} \mathbb{P}\{\mathcal{E}\} &\leq \left(\frac{(2L)^{L-1}}{(L-1)!} \mathbf{1}_{\{L \geq 1\}} + \mathbf{1}_{\{L=0\}} \right) e^{-2L} e^{6(k+1)A + \frac{3\eta\lambda L}{4\ell}} \times \frac{(Mb)^M}{n^M M!} \exp\left(2(k+1)A + \frac{\eta M}{2b\ell}\right) \\ &\leq \left(\frac{\mathbf{1}_{\{L \geq 1\}}}{2\sqrt{L}} + \mathbf{1}_{\{L=0\}} \right) \frac{1}{\sqrt{M}} \left(\frac{2}{e} \right)^L \left(\frac{eb}{n} \right)^M \exp(8(k+1)A + 6\eta) \\ &= \left(\frac{\mathbf{1}_{\{k \leq m'\}}}{\sqrt{(m')^2 - (k-1)^2}} + \frac{\mathbf{1}_{\{k=m'+1\}}}{\sqrt{m'}} \right) n^{-\frac{m'+(k-1)}{2}} (2b)^{m'/2} \left(\frac{be^2}{2} \right)^{\frac{k-1}{2}} \\ &\quad \exp(8(k+1)A + 6\eta), \end{aligned} \quad (76)$$

where in the first inequality we bound the LHS of (74) by $1 \leq e^{6(k+1)A}$ when $L = 0$; the second inequality follows due to $n! \geq \sqrt{n}(n/e)^n$, $a\lambda = 2$, $b\lambda \geq 1$, and $\lambda \leq 4$.

Next, we bound $\mathbb{P}\{P' \text{ is } A\text{-uniform} \mid \mathcal{E}, P' \text{ is } (a, b, \eta)\text{-light}, P \in S\}$. For ease of notation, let us denote $\mathcal{E}' = \mathcal{E} \cap \{P' \text{ is } (a, b, \eta)\text{-light}\} \cap \{P \in S\}$. Then we have

$$\begin{aligned} &\mathbb{P}\{P' \text{ is } A\text{-uniform} \mid \mathcal{E}'\} \\ &\stackrel{(a)}{=} \mathbb{E}\left[\mathbb{P}\{P' \text{ is } A\text{-uniform} \mid \text{wt}_r(P'), \text{wt}_b(P'), \{W_e : e \in P\}\} \mid \mathcal{E}'\right] \\ &\stackrel{(b)}{\leq} \sup_{\alpha: |\alpha - a| \leq \frac{\eta}{2\ell}} \sup_{\beta: |\beta - b| \leq \frac{\eta}{2(\ell-1)}} \sup_{w_e: e \in P} \\ &\quad \mathbb{P}\{P' \text{ is } A\text{-uniform} \mid \text{wt}_r(P') = \alpha\ell, \text{wt}_b(P') = \beta\ell, \{W_e = w_e : e \in P\}\}, \end{aligned}$$

where (a) is because \mathcal{E}' is measurable with respect to $\{\text{wt}_r(P'), \text{wt}_b(P'), \{W_e : e \in P\}\}$, and (b) is because on \mathcal{E}' , P' is (a, b, η) -light so that $|\text{wt}_r(P') - a\ell| \leq \eta/2$ and $|\text{wt}_b(P') - b\ell| \leq \eta/2$.

Fix any α such that $|\alpha - a| \leq \frac{\eta}{2\ell}$ and any β such that $|\beta - b| \leq \frac{\eta}{2(\ell-1)}$. Note that $\eta = 1$, $a = 2/\lambda$ and $b = (2 - \zeta)/\lambda$ with $\zeta = \epsilon/4$ and $\lambda = 4 - \epsilon$. By choosing ϵ_0 to be sufficiently small and n_0 to be sufficiently large, we have $1/4 \leq \alpha, \beta \leq 1$ for all $\epsilon \leq \epsilon_0$ and $n \geq n_0$. We proceed to bound $\mathbb{P}\{P' \text{ is } A\text{-uniform} \mid \text{wt}_r(P') = \alpha\ell, \text{wt}_b(P') = \beta\ell, \{W_e : e \in P\}\}$. Applying [8, Lemma 3.3] (restated as Lemma 14 in Appendix C) with $\rho = \alpha$ or β , we get that

$$\begin{aligned} \mathbb{P}\{\text{dev}_r(P') \leq A \mid \text{wt}_r(P') = \alpha\ell, \{W_e : e \in P\}\} &\leq \sqrt{m} e^{c(Ak + m/A^2)} \mathbb{P}\{\text{dev}_r(P') \leq A\} \\ \mathbb{P}\{\text{dev}_b(P') \leq A \mid \text{wt}_b(P') = \beta\ell, \{W_e : e \in P\}\} &\leq \sqrt{m} e^{c(Ak + m/A^2)} \mathbb{P}\{\text{dev}_b(P') \leq A\}, \end{aligned}$$

where c is a universal constant. Hence, we get that

⁷ We emphasize that it is crucial to keep the polynomials terms in (76) so that in (79) we can get the upper bound $4(2\ell)^{k-1}$, which in turn yields the desired $\frac{\ell^2}{n}$ factor in (81).

$$\begin{aligned} & \mathbb{P} \{ P' \text{ is } A\text{-uniform} \mid \text{wt}_r(P') = \alpha\ell, \text{wt}_b(P') = \beta\ell, \{W_e : e \in P\} \} \\ & \leq m e^{2c(Ak+m/A^2)} p_\ell, \end{aligned} \quad (77)$$

where $p_\ell = \mathbb{P} \{ P' \text{ is } A\text{-uniform} \} = \mathbb{P} \{ \text{dev}_r(P') \leq A \} \times \mathbb{P} \{ \text{dev}_b(P') \leq A \}$.

Combining (72), (73), (76), and (77), we get

$$\begin{aligned} & \sum_{P' \in \mathcal{P}_\ell : |P \cap P'| \geq 1, P \neq P'} \mathbb{P} \{ P' \text{ is } (a, b, \eta)\text{-light and } A\text{-uniform} \mid P \in S \} \\ & \leq \sum_{m=1}^{2\ell-2} \sum_{k=1}^m \binom{m-1}{k-1} \binom{m'+1}{k} \ell^k n^{\frac{m'-(k-1)}{2}} \\ & \quad \times \left(\frac{\mathbf{1}_{\{k \leq m'\}}}{\sqrt{(m')^2 - (k-1)^2}} + \frac{\mathbf{1}_{\{k=m'+1\}}}{\sqrt{m'}} \right) n^{-\frac{m'+(k-1)}{2}} (2b)^{m'/2} \left(\frac{be^2}{2} \right)^{\frac{k-1}{2}} e^{8(k+1)A+6\eta} \\ & \quad \times m e^{2c(Ak+m/A^2)} p_\ell. \end{aligned} \quad (78)$$

To further bound the RHS, we claim that when $\ell \geq e^2$,

$$\binom{m'+1}{k} \left(\frac{\mathbf{1}_{\{k \leq m'\}}}{\sqrt{(m')^2 - (k-1)^2}} + \frac{\mathbf{1}_{\{k=m'+1\}}}{\sqrt{m'}} \right) \leq 4(2\ell)^{k-1}. \quad (79)$$

To see this, note that if $k \geq (m'+1)/2$, then $\binom{m'+1}{k} \leq \left(\frac{e(m'+1)}{k} \right)^k \leq (2e)^k$ and the claim holds as $(m')^2 - (k-1)^2 \geq 1$ when $k \leq m'$ and $m' \geq 1$. If $k \leq (m'+1)/2$, then $\sqrt{(m')^2 - (k-1)^2} \geq (m'+1)/4$. Thus,

$$\binom{m'+1}{k} \frac{1}{\sqrt{(m')^2 - (k-1)^2}} \leq 4(m'+1)^{k-1} \leq 4(2\ell)^{k-1}.$$

Therefore, plugging (79) into (78) and recalling that $m' = 2\ell - 1 - m$, we get that

$$\begin{aligned} & \sum_{P' \in \mathcal{P}_\ell : |P \cap P'| \geq 1, P \neq P'} \mathbb{P} \{ P' \text{ is } (a, b, \eta)\text{-light and } A\text{-uniform} \mid P \in S \} \\ & \leq 4p_\ell \ell e^{(16+2c)A+6\eta} \\ & \quad \times \sum_{m=1}^{2\ell-2} m (2b)^{\ell-1/2-m/2} e^{2cm/A^2} \sum_{k=1}^m \binom{m-1}{k-1} \left(\frac{2\ell^2 \sqrt{b}}{n\sqrt{2}} e^{(8+2c)A+1} \right)^{k-1} \\ & = 4p_\ell \ell e^{A(16+2c)+6\eta+2c/A^2} (2b)^{\ell-1} \sum_{m=1}^{2\ell-2} m \kappa^{m-1}, \end{aligned}$$

where

$$\kappa \triangleq \frac{e^{2c/A^2}}{\sqrt{2b}} \left(1 + \frac{\ell^2 \sqrt{2b}}{n} e^{(8+2c)A+1} \right). \quad (80)$$

Note that $b = \frac{2-\zeta}{\lambda}$ with $\zeta = \frac{\epsilon}{4}$ and $\lambda = 4 - \epsilon$. Also, $A = \left\lceil \frac{C_1}{\sqrt{\epsilon}} \right\rceil$ and $A^2 \leq \ell \leq \sqrt{n}e^{-C_2 A}$, where C_1, C_2 are absolute constants. By choosing C_1, C_2, n_0 to be sufficiently large and ϵ_0 to be sufficiently small, we get that $\kappa \leq 1 - \epsilon/100$ for all $n \geq n_0$ and $\epsilon \geq \epsilon_0$. Hence,

$$\sum_{m=1}^{\infty} m\kappa^{m-1} \leq \left(\sum_{m=1}^{\infty} \kappa^m \right)' = \left(\frac{\kappa}{1-\kappa} \right)' = \frac{1}{(1-\kappa)^2} \leq 10^4 \epsilon^{-2}.$$

Recall that $\eta = 1$. Then we conclude that for universal constants $c_2, c_3 > 0$,

$$\begin{aligned} & \sum_{P' \in \mathcal{P}_\ell : |P \cap P'| \geq 1, P \neq P'} \mathbb{P} \{ P' \text{ is } (a, b, \eta)\text{-light and } A\text{-uniform} \mid P \in S \} \\ & \leq p_\ell \ell e^{c_2/\sqrt{\epsilon}} (2b)^\ell \leq \frac{\ell^2 e^{c_3/\sqrt{\epsilon}}}{n} \mathbb{E}[|S|] \end{aligned} \quad (81)$$

Substituting the last displayed equation back to (71), we get

$$\text{Var}(|S|) \leq \mathbb{E}[|S|] \times \left(1 + \frac{\ell^2 e^{c_3/\sqrt{\epsilon}}}{n} \mathbb{E}[|S|] \right).$$

□

7.4 Extracting many vertex-disjoint alternating paths via Turán's Theorem

To extract many disjoint alternating paths from S , following [8], we first construct a graph under which an independent set corresponds to a set of disjoint alternating paths, and then use Turán's theorem to prove the existence of a large independent set.

Specifically, we define a graph H on the set S of alternating paths, where P and P' are adjacent if they share at least one common vertex.

Thus any collection of vertex-disjoint paths is an independent set in H . The following result due to Turán (see e.g. [2, Theorem 1, p. 95]), provides a lower bound to the size of the largest independent set in a general graph.

Lemma 12 (Turán's Theorem) *Let $G = (V, E)$ be a finite, simple graph. Then G contains an independent subset of size at least $|V|^2/(2|E| + |V|)$.*

Next we lower bound the number of vertices $|V(H)|$ and upper bound the number of edges $|E(H)|$ using Lemmas 10 and 11, and then apply Turán's theorem to finish the proof of Theorem 7.

Proof of Theorem 7 Throughout the proof, $A = \left\lceil \frac{C_1}{\sqrt{\epsilon}} \right\rceil$, $A^2 \leq \ell \leq e^{-C_2 A} \sqrt{n}$, $\eta = 1$, $a = 2/\lambda$, $b = \frac{2-\xi}{\lambda}$ with $\lambda = 4 - \epsilon$ and $\xi = \frac{\epsilon}{4}$.

We start by verifying that, by definition of S , each $(2\ell - 1)$ -alternating path $P \in S$ satisfies the subpath requirement in (58). Recall that P is oriented so that it starts with a left vertex. In this order, let $\phi_1, \phi_2, \dots, \phi_\ell$ denote the sequence of red edge weights and $\psi_1, \psi_2, \dots, \psi_{\ell-1}$ the blue edge weights. Fix any ℓ' such that $\ell/3 \leq \ell' \leq \ell$ and a $(2\ell' - 1)$ -alternating subpath Q of P that has ℓ' red edges and $\ell' - 1$ blue edges. Then Q consists of red edge weights $\phi_{k+1}, \dots, \phi_{k+\ell'}$ and blue edge weights $\psi_{k+1}, \dots, \psi_{k+\ell'-1}$ in order for some $0 \leq k \leq \ell - \ell'$. In particular,

$$\begin{aligned} \text{wt}_r(Q) &= \sum_{i=1}^{\ell'} \phi_{k+i} = \sum_{i=1}^{k+\ell'} \phi_i - \sum_{i=1}^k \phi_i \\ \text{wt}_b(Q) &= \sum_{i=1}^{\ell'-1} \psi_{k+i} = \sum_{i=1}^{k+\ell'-1} \psi_i - \sum_{i=1}^k \psi_i. \end{aligned}$$

Since P is (a, b, η) -light and A -uniform with $\eta = 1$, it follows that $\sum_{i=1}^{k+\ell} \phi_i \geq (k + \ell' - A) \frac{\text{wt}_r(P)}{\ell}$, $\sum_{i=1}^k \phi_i \leq (k + A) \frac{\text{wt}_r(P)}{\ell}$, and hence

$$\text{wt}_r(Q) \geq (\ell' - 2A) \frac{\text{wt}_r(P)}{\ell} \geq (\ell' - 2A) \left(a - \frac{1}{2\ell} \right).$$

Analogously, we have

$$\text{wt}_b(Q) \leq (\ell' + 2A) \left(b + \frac{1}{2(\ell - 1)} \right).$$

Since $\log \frac{\mathcal{P}}{\mathcal{Q}}(W_e) = \log(n\lambda) - (\lambda - 1/n)W_e$, it follows that

$$\begin{aligned} \Delta(Q) &= -(\lambda - 1/n) [\text{wt}_b(Q) - \text{wt}_r(Q)] \\ &\geq -(\lambda - 1/n) \ell' \left(b - a + \frac{1}{2\ell} + \frac{1}{2(\ell - 1)} \right) - 2(\lambda - 1/n)A \left(b + \frac{1}{2(\ell - 1)} + a - \frac{1}{2\ell} \right) \\ &\geq -(\lambda - 1/n) \left[-\ell' \frac{\epsilon}{4\lambda} + \frac{8A}{\lambda} + 2 \right] \geq (\lambda - 1/n) \frac{\epsilon}{96} \ell', \end{aligned} \tag{82}$$

where the last inequality holds as $\ell' \geq \ell/3$, $\lambda = 4 - \epsilon$, and $\ell \geq A^2$, and $A \geq C_1/\sqrt{\epsilon}$.

Next we apply Turán's theorem to the graph H to extract a vertex-disjoint subcollection S^* of S . In view of (67) and $A = \left\lceil \frac{C_1}{\sqrt{\epsilon}} \right\rceil$, by choosing C_1, n_0 to be sufficiently large and ϵ_0 to be sufficiently small, for all $0 < \epsilon \leq \epsilon_0$ and $n \geq n_0$,

$$\mathbb{E}[|V(H)|] = \mathbb{E}[|S|] \geq \frac{c_4 n}{\ell} e^{c_4 \epsilon \ell}$$

for a universal constant $c_4 > 0$. By Chebyshev's inequality and Lemma 11,

$$\mathbb{P} \left\{ |S| \leq \frac{1}{2} \mathbb{E}[|S|] \right\} \leq \frac{4 \text{Var}(|S|)}{(\mathbb{E}[|S|])^2} \leq \frac{\ell^2 e^{c_2/\sqrt{\epsilon}}}{n}$$

for a universal constant $c_2 > 0$. By (81),

$$\begin{aligned} \mathbb{E}[|E(H)|] &= \sum_{P \in \mathcal{P}_\ell} \mathbb{P}\{P \in S\} \sum_{P' \in \mathcal{P}_\ell: |P \cap P'| \geq 1, P' \neq P} \mathbb{P}\{P' \in S \mid P \in S\} \\ &\leq \frac{\ell^2 e^{c_2/\sqrt{\epsilon}}}{n} (\mathbb{E}[|S|])^2. \end{aligned}$$

By Markov's inequality,

$$\mathbb{P}\{|E(H)| \geq 2\mathbb{E}[|E(H)|]\} \leq 1/2.$$

Hence, with probability at least $1/2 - \frac{\ell^2 e^{c_2/\sqrt{\epsilon}}}{n}$, we have that $|S| \geq \frac{1}{2} \mathbb{E}[|S|]$ and $|E(H)| \leq 2\mathbb{E}[|E(H)|]$, so that Lemma 12 implies the existence of $S^* \subset S$ of disjoint alternating paths such that

$$|S^*| \geq \frac{\frac{1}{4} (\mathbb{E}[|S|])^2}{4\mathbb{E}[|E(H)|] + \frac{1}{2}\mathbb{E}[|S|]} \geq \frac{n}{\ell^2 e^{c_3/\sqrt{\epsilon}}}$$

for a universal constant $c_3 > 0$. \square

Acknowledgements JX would like to thank Christopher Moore for many inspiring discussions on the posterior sampling. JX is also grateful to Guilhem Semerjian, Gabriele Sicuro, and Lenka Zdeborová for sharing the early draft of [25]. This work was done in part while the authors were participating in a program at the Simons Institute for the Theory of Computing.

Data availability Data sharing not applicable to this article as no datasets were generated or analyzed during the current study.

Appendix A Large deviation estimates

Lemma 13 *Let \mathcal{P}, \mathcal{Q} be two probability distributions such that $\mathcal{P} \ll \mathcal{Q}$. Let X_i 's and Y_i 's be two independent sequences of random variables, where X_i 's are i.i.d. copies of $\log(\mathcal{P}/\mathcal{Q})$ under distribution \mathcal{P} and Y_i 's are i.i.d. copies of $\log(\mathcal{P}/\mathcal{Q})$ under distribution \mathcal{Q} . For all $x \geq 0$ and positive integer ℓ , we have*

$$\mathbb{P} \left\{ \sum_{i=1}^{\ell} (Y_i - X_i) \geq x\ell \right\} \leq \exp(-\ell(\alpha + x/2)), \quad (83)$$

where $\alpha = -2 \log B(\mathcal{P}, \mathcal{Q})$ as defined in (14). Furthermore, if we further assume that $\mathcal{Q} \ll \mathcal{P}$, then for all $0 \leq x \leq D(\mathcal{P} \parallel \mathcal{Q}) + D(\mathcal{Q} \parallel \mathcal{P})$ and positive integer ℓ , we have

$$\mathbb{P} \left\{ \sum_{i=1}^{\ell} (Y_i - X_i) \geq x\ell \right\} \geq \exp(-\ell(\alpha + x + o_{\ell}(1))). \quad (84)$$

Proof The proof of (83) follows from standard large deviation analysis (cf. [3, Appendix B]). Let F denote the Legendre transform of the log moment generating function of $Y_1 - X_1$, i.e.,

$$F(x) = \sup_{\theta \geq 0} \{ \theta x - \psi_{\mathcal{P}}(-\theta) - \psi_{\mathcal{Q}}(\theta) \},$$

where $\psi_{\mathcal{P}}(\theta) = \log \mathbb{E}[e^{\theta X_1}]$ and $\psi_{\mathcal{Q}}(\theta) = \log \mathbb{E}[e^{\theta Y_1}]$. Then from the Chernoff bound we have the following large deviation inequality:

$$\mathbb{P} \left\{ \sum_{i=1}^{\ell} (Y_i - X_i) \geq x\ell \right\} \leq \exp(-\ell F(x)). \quad (85)$$

Note the following facts:

1. $F(0) = -\psi_{\mathcal{P}}(-1/2) - \psi_{\mathcal{Q}}(1/2) = -2 \log \sqrt{\mathcal{P}\mathcal{Q}} = \alpha$;
2. $F(x) \geq F(0) + x/2$, for all $x \geq 0$.

Combining these facts with (85) yields (83).

Next we prove (84). By Cramér's theorem, we have

$$\mathbb{P} \left\{ \sum_{i=1}^{\ell} (Y_i - X_i) \geq x\ell \right\} = \exp(-F(x)\ell + o(\ell)), \quad (86)$$

where $o(\ell)/\ell$ converges to 0 as ℓ grows to infinity. Second, we have

$$F(x) \leq \alpha + x, \quad \forall 0 \leq x \leq D(\mathcal{P} \parallel \mathcal{Q}) + D(\mathcal{Q} \parallel \mathcal{P}).$$

To see this, note that $\psi_{\mathcal{P}}(\theta) = \psi_{\mathcal{Q}}(1 + \theta)$. Thus, the optimal θ is given by

$$x + \psi'_{\mathcal{Q}}(1 - \theta) - \psi'_{\mathcal{Q}}(\theta) = 0.$$

Note that $\psi'_{\mathcal{Q}}(0) = -D(\mathcal{Q} \parallel \mathcal{P})$ and $\psi'_{\mathcal{Q}}(1) = D(\mathcal{P} \parallel \mathcal{Q})$. Moreover, since $\psi_{\mathcal{Q}}(\theta)$ is convex, it follows that $\psi'_{\mathcal{Q}}(\theta)$ is non-decreasing in θ . Thus the optimal θ must lie in $[1/2, 1]$ when $0 \leq x \leq D(\mathcal{P} \parallel \mathcal{Q}) + D(\mathcal{Q} \parallel \mathcal{P})$. Hence,

$$\begin{aligned} F(x) &= \sup_{\theta \in [1/2, 1]} \{ \theta x - \psi_{\mathcal{P}}(-\theta) - \psi_{\mathcal{Q}}(\theta) \} \leq x + \sup_{\theta \in [1/2, 1]} \{ -\psi_{\mathcal{P}}(-\theta) - \psi_{\mathcal{Q}}(\theta) \} \\ &= x + F(0) = x + \alpha. \end{aligned}$$

Combine with (86) to finish the proof of (84). \square

Appendix B Erlang distribution and Chernoff bounds

The sum of ℓ i.i.d. $\exp(\lambda)$ random variables has an Erlang distribution with parameters ℓ and λ , denoted by $\text{Erlang}(\ell, \lambda)$, whose density is given by

$$f(x) = \frac{\lambda^\ell x^{\ell-1} e^{-\lambda x}}{(\ell-1)!}, \quad x \geq 0 \quad (87)$$

Theorem 8 *Let $X_i \stackrel{\text{i.i.d.}}{\sim} \exp(1)$. Then*

$$\begin{aligned} \mathbb{P} \left\{ \sum_{i=1}^n X_i \geq n\xi \right\} &\leq \exp(-n(\xi - \log(\xi) - 1)), \quad \forall \xi > 1 \\ \mathbb{P} \left\{ \sum_{i=1}^n X_i \leq n\xi \right\} &\leq \exp(-n(\xi - \log(\xi) - 1)), \quad \forall \xi < 1 \end{aligned}$$

Appendix C Exp-minus-one random bridge

Let $X_1, X_2, \dots, X_\ell \stackrel{\text{i.i.d.}}{\sim} \exp(\mu)$ and let $X = \sum_{i=1}^\ell X_i$. Recall from Lemma 8 the exp-minus-one ℓ -bridge R is defined as

$$R_j = \sum_{i=1}^j \left(\frac{X_i}{X} \ell - 1 \right), \quad 0 \leq j \leq \ell.$$

Define $\text{dev}(R) = \max_{0 \leq j \leq \ell} |R_j|$. The following result adapted from [8, Lemma 3.3] (which is a slight extension of [7, Lemma 3.2]) bounds the probability of $\text{dev}(R) \leq A$ conditional on the total weight of X_i and the value of X_i for a set of indices i in a union of intervals, in terms of the unconditional probability $p_\ell \triangleq \mathbb{P}\{\text{dev}(R) \leq A\}$. This result is crucial for the second moment computation in Sect. 7.3.

Lemma 14 *Let $1/4 \leq \rho \leq 1$ and $1 \leq A\rho \leq \sqrt{\ell}$. Consider the integer intervals $[a_1, b_1], \dots, [a_k, b_k]$ such that $1 \leq a_1 \leq b_1 \leq \dots \leq a_k \leq b_k \leq \ell$ and $m = \sum_{i=1}^k (b_i - a_i + 1)$. Write $J = \bigcup_{i=1}^k [a_i, b_i]$. Then*

$$\mathbb{P} \left\{ \text{dev}(R) \leq A \mid \sum_{i=1}^\ell X_i = \rho\ell, \{X_j = x_j, j \in J\} \right\} \leq c_1 A \sqrt{m \wedge (\ell - m)} p_\ell 10^{100kA} e^{c_0 m/A^2},$$

where $c_0, c_1 > 0$ are two universal constants.

Appendix D Minimum-weight matching for the exponential model

In this section, we prove the positive part (12) in Theorem 3. Namely, in the complete graph case with exponentially distributed weights, when $\lambda = 4 - \epsilon$, the minimum weighted matching \widehat{M}_{ML} (linear assignment), which corresponds to the maximum likelihood estimation, misclassifies at most $O\left(\frac{1}{\epsilon^3} e^{-\frac{2\pi}{\sqrt{\epsilon}}}\right)$ fraction of edges on average.

This together with the negative part (11) in Theorem 3 shows that the minimum weighted matching achieves the optimal rate $1/\sqrt{\epsilon}$ of the error exponent. Prior work [21] provides the exact characterization of the asymptotic error of \widehat{M}_{ML} in terms of a system of ordinary differential equations when $\lambda < 4$. Our proof follows by analyzing this system of ODEs when $\lambda = 4 - \epsilon$ for small ϵ , and is inspired by the heuristic arguments in [25, Section VI].

Proof of Theorem 3: positive part First, it has been shown in [21, Theorem 2] that

$$\lim_{n \rightarrow \infty} \mathbb{E}[\ell(M^*, \widehat{M}_{\text{ML}})] = 4 \int_0^\infty (1 - U(x)V(x)) (1 - (1 - U(x))W(x)) V(x)W(x) dx,$$

where (U, V, W) is the unique solution to the following system of equations

$$\begin{aligned} \frac{dU}{dx} &= -\lambda U(1 - U) + (1 - UV)(1 - (1 - U)W) \\ \frac{dV}{dx} &= \lambda V(1 - U) \\ \frac{dW}{dx} &= -\lambda WU \end{aligned} \tag{88}$$

with initial condition

$$U(0) = \frac{1}{2}, \quad V(0) = W(0) = \delta, \quad \delta \in (0, 1), \tag{89}$$

and δ is the unique value in $(0, 1)$ such that $U(x), V(x) \rightarrow 1$ as $x \rightarrow +\infty$.

Furthermore, it has been shown in [21, Section B] that $UV < 1$, $(1 - U)W < 1$, $0 < U, V, W < 1$, and

$$V(x) = \delta \exp\left(\lambda \int_0^x (1 - U(y)) dy\right), \tag{90}$$

$$W(x) = V(x) e^{-\lambda x}. \tag{91}$$

Therefore, we have that

$$\lim_{n \rightarrow \infty} \mathbb{E}[\ell(M^*, \widehat{M}_{\text{ML}})] \leq 4 \int_0^\infty V^2(x) e^{-\lambda x} dx,$$

Let $x_0 = \inf\{x \geq 0 : U(x) < \frac{1}{2}\}$ denote the first time that $U(x)$ crosses $1/2$. Therefore by (90) we have that for all $0 \leq x \leq x_0$, $V(x) \leq \delta e^{\lambda x/2}$. Hence for any $0 \leq \tau \leq x_0$,

$$\begin{aligned}
\int_0^\infty V^2(x)e^{-\lambda x} dx &= \int_0^\tau V^2(x)e^{-\lambda x} dx + \int_\tau^\infty V^2(x)e^{-\lambda x} dx \\
&\leq \delta^2 \tau + \int_\tau^\infty e^{-\lambda x} dx \\
&= \delta^2 \tau + \frac{1}{\lambda} e^{-\lambda \tau}.
\end{aligned}$$

We claim that $\lambda x_0 \geq 2 \log \frac{\epsilon}{8\delta}$. Suppose not. Then $e^{\lambda x_0/2} < \frac{\epsilon}{8\delta}$, in particular, $x_0 < +\infty$. By continuity, we have $U(x_0) = \frac{1}{2}$, $U'(x_0) \leq 0$. But we have

$$\begin{aligned}
U'(x_0) &= -\lambda U(x_0) (1 - U(x_0)) + (1 - U(x_0) V(x_0)) (1 - (1 - U(x_0)) W(x_0)) \\
&= -\lambda/4 + \left(1 - \frac{1}{2} V(x_0)\right) \left(1 - \frac{1}{2} W(x_0)\right) \\
&\geq -\lambda/4 + 1 - V(x_0) \\
&\geq \epsilon/4 - \delta e^{\lambda x_0/2} > \epsilon/8 > 0,
\end{aligned}$$

which is the needed contradiction. Therefore, by setting $\tau = \frac{2}{\lambda} \log \frac{\epsilon}{8\delta}$, we get that

$$\int_0^\infty V^2(x)e^{-\lambda x} dx \leq \frac{2\delta^2}{\lambda} \log \frac{\epsilon}{8\delta} + \frac{1}{\lambda} \left(\frac{8\delta}{\epsilon}\right)^2 \leq \frac{c}{\epsilon^3} \exp\left(-\frac{2\pi}{\sqrt{\epsilon}}\right),$$

where $c > 0$ is a universal constant and the last inequality follows from the claim that

$$\delta \leq \frac{c'}{\sqrt{\epsilon}} \exp\left(-\frac{\pi}{\sqrt{\epsilon}}\right), \quad (92)$$

where $c' > 0$ is a universal constant.

It remains to prove (92). In view of $V < 1$ and (90), we have that

$$\delta \exp\left(\lambda \int_0^\infty (1 - U(y)) dy\right) \leq 1. \quad (93)$$

To proceed, we derive an upper bound to $U(x)$. Let $\tilde{U}(x)$ denote the unique solution of the following ODE:

$$\frac{d\tilde{U}}{dx} = -\lambda \tilde{U}(1 - \tilde{U}) + 1, \quad \tilde{U}(0) = \frac{1}{2}. \quad (94)$$

We claim that $U(x) \leq \tilde{U}(x)$ for all $x \geq 0$. To show this, let $f = \tilde{U} - U$. Then we have (a) $f(0) = 0$; (b) $f'(0) > 0$; (c) $f'(x) > 0$ whenever $f(x) = 0$. Thus $f(\delta_0) > 0$ for some small δ_0 . Let $x_1 = \inf\{x \geq 0 : f(x) < 0\}$. Suppose for the sake of contradiction that $x_1 < \infty$. Then x_1 is the first time that f crosses zero. Thus $f'(x_1) \leq 0$. But by continuity we have $f(x_1) = 0$ and hence $f'(x_1) > 0$, which is a contradiction.

Solving the ODE (94), we get that

$$\tilde{U}(x) = \frac{1}{2} + \frac{1}{2} \sqrt{\frac{4-\lambda}{\lambda}} \tan\left(\frac{x}{2} \sqrt{\lambda(4-\lambda)}\right), \quad 0 \leq x < \frac{\pi}{\sqrt{\lambda(4-\lambda)}}.$$

Therefore, for some $0 \leq a < \frac{\pi}{\sqrt{\lambda(4-\lambda)}}$ to be determined,

$$\begin{aligned} \int_0^\infty (1 - U(y)) dy &\geq \int_0^a (1 - U(y)) dy \\ &\geq \int_0^a (1 - \tilde{U}(y)) dy \\ &= \frac{a}{2} + \frac{1}{\lambda} \log \cos\left(\frac{a}{2} \sqrt{\lambda(4-\lambda)}\right), \end{aligned}$$

where the last inequality holds due to $\int_0^x \tan(y) dy = -\log \cos(x)$ for $0 \leq x < \pi/2$. Combining the last displayed equation with (93) yields that

$$\delta \leq \frac{1}{\cos\left(\frac{a}{2} \sqrt{\lambda(4-\lambda)}\right)} \exp\left(-\lambda \frac{a}{2}\right).$$

Recall that $\lambda = 4 - \epsilon$. Choose $a = \frac{\pi - \sqrt{\epsilon}}{\sqrt{\lambda(4-\lambda)}}$. Then we get

$$\delta \leq \frac{1}{\sin\left(\frac{\sqrt{\epsilon}}{2}\right)} \exp\left(-\sqrt{\lambda} \frac{\pi - \sqrt{\epsilon}}{2\sqrt{\epsilon}}\right) \leq \frac{c'}{\sqrt{\epsilon}} \exp\left(-\frac{\pi}{\sqrt{\epsilon}}\right),$$

for a universal constant $c' > 0$. □

Appendix E Reduction arguments

E.1 Reduction from dense to sparse model

In this subsection, we prove Theorem 2 for the dense model by reducing it to the sparse model. Suppose Theorem 2 holds under the sparse model. Recall that in the dense model, the planted weight density is $p(x)$ and the null weight density is $q(x) = \frac{1}{d} \rho(\frac{x}{d})$, where p and ρ are fixed probability densities on \mathbb{R} and $d \rightarrow \infty$. In this case the impossibility condition (7) simplifies to:

$$\int_{-\infty}^{\infty} \sqrt{p(x)\rho(0)} dx \geq 1 + \epsilon. \quad (95)$$

Define a sparse model with parameter $(d', \mathcal{P}', \mathcal{Q}')$, where

$$d' = \Gamma \rho(0) \left(1 - \frac{\epsilon}{2}\right)^2, \quad \mathcal{P}' = \mathcal{P}, \quad \mathcal{Q}' = \text{Unif}[-\Gamma/2, \Gamma/2],$$

and for some positive constant Γ to be specified. Note that $d', \mathcal{P}', \mathcal{Q}'$ are all independent of n . Next we show that, given graph G' drawn from the $(d', \mathcal{P}', \mathcal{Q}')$ model with planted matching M^* , there exists a (randomized) mapping f such that $G \triangleq f(G')$ is distributed according to the $(d, \mathcal{P}, \mathcal{Q})$ model with the same planted matching M^* . Note that crucially such a mapping needs to be agnostic to the latent M^* .

To construct the mapping f , we first express \mathcal{Q} as a mixture of \mathcal{Q}' and some other distribution $\bar{\mathcal{Q}}$, that is $\mathcal{Q} = t\mathcal{Q}' + (1-t)\bar{\mathcal{Q}}$, where $t = \frac{d'}{d} = \frac{\Gamma}{d}\rho(0)\left(1 - \frac{\epsilon}{2}\right)^2$ and $\bar{\mathcal{Q}}$ has density $\bar{q}(x) \triangleq \frac{q(x) - tq'(x)}{1-t}$. We claim that \bar{q} is well-defined density for large d . Indeed, note that $q(x) = \frac{1}{d}\rho\left(\frac{x}{d}\right)$, $q'(x) = \mathbf{1}_{\{|x| \leq \Gamma/2\}}/\Gamma$ and $d \rightarrow \infty$. By continuity of ρ at 0, $\rho(0)(1 - \frac{\epsilon}{2})^2 \leq \rho\left(\frac{x}{d}\right)$ holds for all $|x| \leq \Gamma/2$, provided that d is sufficiently large. Thus $tq'(x) \leq q(x)$ for all x .

Next, given G' , we generate a denser graph $G = f(G')$ as follows. For each edge e in G' , we leave its edge weight unchanged. For each edge e not in G' , we connect it in G independently with probability r/n and draw its edge weight $W_e \sim \bar{\mathcal{Q}}$, where $r \triangleq \frac{d-d'}{1-d'/n}$. The choice of r is such that for each $e \in [n] \times [n] \setminus M^*$, the probability that e is an unplanted edge in G (namely, $e \in E(G) \setminus M^*$) equals $(1 - \frac{d'}{n})\frac{r}{n} + \frac{d'}{n} = \frac{d}{n}$. Furthermore, for each unplanted e in G , it is an unplanted edge in G' with probability $\frac{d'/n}{d/n} = \frac{d'}{d} = t$. As a result, conditioned on e being an unplanted edge in G , the distribution of W_e is $t\mathcal{Q}' + (1-t)\bar{\mathcal{Q}} = \mathcal{Q}$. Also, by construction, conditioned on the true matching M^* of G' , the edge weights of G are independent. In other words, $G = f(G')$ is distributed according the dense $(d, \mathcal{P}, \mathcal{Q})$ model with M^* being the planted matching. Thus, it remains to verify the impossibility for the sparse $(d', \mathcal{P}', \mathcal{Q}')$ model by verifying the condition of Theorem 2. Indeed,

$$\sqrt{d'} B(\mathcal{P}', \mathcal{Q}') = \sqrt{td} \int_{-\Gamma/2}^{\Gamma/2} \sqrt{p(x)} \frac{1}{\Gamma} dx = \left(1 - \frac{\epsilon}{2}\right) \sqrt{\rho(0)} \int_{-\Gamma/2}^{\Gamma/2} \sqrt{p(x)} dx \geq 1 + \frac{\epsilon}{4},$$

where the last inequality holds as a consequence of the condition (95) and by choosing a large enough Γ . This completes the proof of Theorem 2 for dense models.

E.2 Reduction for general weight distributions: positive result

In Sect. 2, we proved Theorem 4 under the assumption $\mathcal{P} \ll \mathcal{Q}$. Here, we prove the theorem for general \mathcal{P}, \mathcal{Q} . Note that our main positive result Theorem 1 follows directly from Theorem 4, per the argument in Sect. 2.

Recall that in (4), f and g denote the densities of \mathcal{P} and \mathcal{Q} with respect to a common dominating measure μ , respectively. Under this general model, the maximum likelihood estimator \widehat{M}_{ML} takes the form

$$\widehat{M}_{\text{ML}} \in \arg \max_{M \in \mathcal{M}} \prod_{e \in M} f(W_e) \prod_{e \notin M} g(W_e).$$

Let $p = \mathcal{P}\{g > 0\}$. If $p = 0$, then \mathcal{P} and \mathcal{Q} are mutually singular. In this case, it is easy to see that $\widehat{M}_{\text{ML}} = M^*$ with probability one. Next, suppose $p > 0$. Let \mathcal{P}' denote the distribution \mathcal{P} conditioned on the support of \mathcal{Q} , with density $f' = f \mathbf{1}_{\{g>0\}}/p$. Then we have $\mathcal{P}' \ll \mathcal{Q}$.

Let $V = \{i \in [n] : g(W_{i,i'}) > 0\}$. Therefore, $g(W_{i,i'}) = 0$ for all $i \in V^c$. We first observe that if a perfect matching $M \in \mathcal{M}$ does not contain all the red edges $\{(i, i') : i \in V^c\}$, then it has a zero likelihood. Thus, \widehat{M}_{ML} reduces to:

$$\widehat{M}_{\text{ML}} = \{(i, i') : i \in V^c\} \cup \widehat{M}_{\text{ML},V}, \quad (96)$$

where for \mathcal{M}_V defined as the set of perfect matchings on $V \times V'$,

$$\widehat{M}_{\text{ML},V} \in \arg \max_{M \in \mathcal{M}_V} \prod_{e \in M} f(W_e) \prod_{e \notin M} g(W_e).$$

In other words, $\widehat{M}_{\text{ML},V}$ is the maximum likelihood estimator over the subgraph $G' = G[V \times V']$. Moreover, conditional on the set V , all the red edges weights on G' are *i.i.d.* following distribution \mathcal{P}' , and all the blue edge weights are *i.i.d.* according to \mathcal{Q} . Denote $M_V^* = \{(i, i') : i \in V\}$ for the true matching on G' , and let $n' = |V|$. We can therefore apply Theorem 4 on G' with $B(\mathcal{P}', \mathcal{Q}) = B(\mathcal{P}, \mathcal{Q})/\sqrt{p}$. Under the condition $\sqrt{n}B(\mathcal{P}, \mathcal{Q}) \leq 1 + \epsilon$, we have

$$\sqrt{n'}B(\mathcal{P}', \mathcal{Q}) = \sqrt{n}B(\mathcal{P}, \mathcal{Q})\sqrt{\frac{n'}{np}} \leq (1 + \epsilon)\sqrt{\frac{n'}{np}}.$$

Theorem 4 yields

$$\mathbb{E}\left[|\widehat{M}_{\text{ML},V} \Delta M_V^*| \mid V\right] \leq Cn' \max\left\{\log(1 + \epsilon'), \sqrt{\frac{\log n'}{n'}}\right\},$$

where ϵ' is such that

$$2\log(1 + \epsilon') = 2\log(1 + \epsilon) + \log\frac{n'}{np}.$$

Thus,

$$\begin{aligned} \mathbb{E}\left[|\widehat{M}_{\text{ML},V} \Delta M_V^*| \mid V\right] &\leq C \max\left\{n' \log(1 + \epsilon) + \frac{n'}{2} \log \frac{n'}{np}, \sqrt{n' \log n'}\right\} \\ &\leq C \max\left\{n \log(1 + \epsilon) + \frac{n'}{2} \log \frac{n'}{np}, \sqrt{n \log n}\right\}, \end{aligned} \quad (97)$$

where the last inequality is from $n' = |V| \leq n$. Next, we average over V to obtain that

$$\begin{aligned}\mathbb{E}[|\widehat{M}_{\text{ML},V} \Delta M_V^*|] &\leq C\mathbb{E}\left[\max\left\{n \log(1+\epsilon) + \frac{n'}{2} \log \frac{n'}{np}, \sqrt{n \log n}\right\}\right] \\ &\leq C\left[n \log(1+\epsilon) + \mathbb{E}\left[\frac{n'}{2} \log \frac{n'}{np}\right] + \sqrt{n \log n}\right].\end{aligned}$$

Under our model, $W_{i,i'} \sim \mathcal{P}$ for all i . Therefore $n' = |V| \sim \text{Binom}(n, p)$. To bound $\mathbb{E}[n' \log n']$, note that for any $u > 0$, and $x \geq 0$, $\log(x/u) \leq x/u - 1$ so that $x \log x \leq x^2/u + x \log(u/e)$. Thus,

$$\mathbb{E}[n' \log n'] \leq \mathbb{E}[(n')^2]/u + \mathbb{E}[n'] \log(u/e) \leq \mathbb{E}[n'] \log \frac{\mathbb{E}[(n')^2]}{\mathbb{E}[n']} \leq np \log(np + 1),$$

where the first inequality holds by optimally choosing $u = \mathbb{E}[(n')^2]/\mathbb{E}[n']$. Hence, $\mathbb{E}\left[n' \log \frac{n'}{np}\right] \leq np \log(1 + 1/(np)) \leq 1$ and consequently,

$$\begin{aligned}\mathbb{E}[|\widehat{M}_{\text{ML},V} \Delta M_V^*|] &\leq C\left[n \log(1+\epsilon) + 1/2 + \sqrt{n \log n}\right] \\ &\leq C_1 \max\left\{n \log(1+\epsilon), \sqrt{n \log n}\right\}\end{aligned}$$

for some universal constant C_1 . This finishes the proof of Theorem 4 for general \mathcal{P}, \mathcal{Q} .

E.3 Reduction for general weight distributions: negative results

In Appendix E.1 we have already reduced the impossibility result Theorem 2 from the dense model to the sparse model. In this section, we show that for the sparse model $(d, \mathcal{P}, \mathcal{Q})$, we can assume WLOG that $\mathcal{P} \ll \mathcal{Q}$ and $\mathcal{Q} \ll \mathcal{P}$. In other words, assuming Theorem 2 holds under the sparse model with \mathcal{P} and \mathcal{Q} mutually absolutely continuous, we prove Theorem 2 for general (fixed) \mathcal{P}, \mathcal{Q} .

Recall that μ is a common dominating measure of \mathcal{P}, \mathcal{Q} , and f, g are densities of \mathcal{P}, \mathcal{Q} respectively. As in the reduction argument for the positive result, we first define distribution \mathcal{P}' with density $f\mathbf{1}_{\{g>0\}}/p$, and distribution \mathcal{Q}' with density $g\mathbf{1}_{\{f>0\}}/q$, where $p = \mathcal{P}\{g > 0\}$ and $q = \mathcal{Q}\{f > 0\}$ are the normalizing constants. Note that since $B(\mathcal{P}, \mathcal{Q}) > 0$, both p, q are strictly positive, so that \mathcal{P}' and \mathcal{Q}' are well-defined. Then we have $\mathcal{P}' \ll \mathcal{Q}'$, and $\mathcal{Q}' \ll \mathcal{P}'$.

We first reduce the original $(d, \mathcal{P}, \mathcal{Q})$ model to the $(d', \mathcal{P}, \mathcal{Q}')$ model where $d' = dq$. Note that for edges e that either do not appear in G (in which case recall that we set $W_e = \star$ for a special symbol \star to signify that e is not in G), or are such that $f(W_e) = 0$, we have $(f/g)(W_e) = 0$. Thus the posterior distribution of M^* under the $(d, \mathcal{P}, \mathcal{Q})$ model is identical to the posterior distribution under the $(d', \mathcal{P}, \mathcal{Q}')$ model.

Therefore, we only need to show that the conclusion of Theorem 2 holds for graph G with weights (W_e) that follow the $(d', \mathcal{P}, \mathcal{Q}')$ model. Let $V = \{i \in [n] : g(W_{i,i'}) > 0\}$. We can bound the optimal overlap as follows:

$$\begin{aligned}
\sup_{\widehat{M}} \mathbb{E} \left(\sum_{i \leq n} \mathbf{1}_{\{(i, i') \in \widehat{M}\}} \right) &\leq \sup_{\widehat{M}} \left[\mathbb{E}(V^c) + \mathbb{E} \left(\sum_{i \in V} \mathbf{1}_{\{(i, i') \in \widehat{M}\}} \right) \right] \\
&\leq n(1-p) + \mathbb{E} \left[\sup_{\widehat{M}} \mathbb{E} \left(\sum_{i \in V} \mathbf{1}_{\{(i, i') \in \widehat{M}\}} \mid V \right) \right]. \quad (98)
\end{aligned}$$

Conditional on V , the subgraph $\widetilde{G} = G[V \times V']$ follows the $(d'|V|/n, \mathcal{P}', \mathcal{Q}')$ model. To see that, note that the red edges in \widetilde{G} are distributed *i.i.d.* \mathcal{P}' ; the blue edges follow distribution \mathcal{Q}' , and appear independently with probability $d'/n = (d'|V|/n)/|V|$. To apply the impossibility result under the $(d'|V|/n, \mathcal{P}', \mathcal{Q}')$ model, we need to establish a lower bound on $|V|$. Let

$$\mathcal{A} = \left\{ |V| \geq \left(1 - \frac{\epsilon}{2}\right)^2 pn \right\}.$$

Recall that under the sparse model, \mathcal{P}, \mathcal{Q} are fixed distributions that do not depend on n . Therefore p does not depend on n , and since $|V| \sim \text{Binom}(n, p)$, we have $\mathbb{P}(\mathcal{A}) = 1 - o(1)$.

Furthermore, note that

$$B(\mathcal{P}', \mathcal{Q}') = \frac{\int \sqrt{fg} d\mu}{\sqrt{pq}} = \frac{B(\mathcal{P}, \mathcal{Q})}{\sqrt{pq}}.$$

Therefore on \mathcal{A} , we have

$$\sqrt{\frac{d'|V|}{n}} B(\mathcal{P}', \mathcal{Q}') = \sqrt{\frac{d|V|}{np}} B(\mathcal{P}, \mathcal{Q}) \geq (1 + \epsilon)(1 - \epsilon/2) \geq 1 + \epsilon/3$$

for $\epsilon \leq 1/3$. Therefore, we can invoke the conclusion of Theorem 2 under the $(d'|V|/n, \mathcal{P}', \mathcal{Q}')$ model to deduce that on \mathcal{A} , for some constant $c > 0$ and large enough n ,

$$\sup_{\widehat{M}} \mathbb{E} \left(\sum_{i \in V} \mathbf{1}_{\{(i, i') \in \widehat{M}\}} \mid V \right) \leq (1 - c)|V|,$$

where the supremum is over all (possibly random) mappings from $G[V \times V']$ to \mathcal{M}_V , the set of perfect matchings on $V \times V'$. Conditional on V , the edges in $G[V \times V']$ are independent of those in $G \setminus G[V \times V']$. Thus, we can replace the range of the supremum with all (random) mappings from G to \mathcal{M}_V without changing its value, allowing us to continue upper bounding (98) with

$$\begin{aligned}
n(1-p) + \mathbb{E}[(1-c)|V|\mathbf{1}_{\mathcal{A}}] + n\mathbb{P}(\mathcal{A}^c) &\leq ((1-p) + (1-c)p + o(1))n \\
&= (1 - cp + o(1))n.
\end{aligned}$$

In other words, the fraction of misclassified edges is lower bounded by $cp + o(1)$. Under the sparse model, cp is a positive constant that only depends on $d, \mathcal{P}, \mathcal{Q}$. We have proved Theorem 2 under the $(d', \mathcal{P}, \mathcal{Q}')$ model. The conclusion of Theorem 2 under the $(d, \mathcal{P}, \mathcal{Q})$ model immediately follows.

Appendix F Finite-order phase transition under the unweighted model

In this section, we prove that in the unweighted case $\mathcal{P} = \mathcal{Q}$, under the impossibility condition (7), which translates to $\sqrt{d} \geq 1 + \epsilon$, the minimal reconstruction error is at least $\Omega(\epsilon^8)$ for small ϵ . Similar to the proof of Theorem 2, we can assume WLOG that $\sqrt{d} = 1 + \epsilon$. As in the proof under general weight distributions, this result is again proven via the two-stage cycle finding scheme described in Algorithm 1. Recall from Algorithm 1 that we reserve a set V of γn left vertices from $[n]$. We will first construct many disjoint alternating paths on the subgraph $G_1 = G[V^c \times (V^c)']$, and then use the reserved vertices to connect the paths into alternating cycles.

F.1 Path construction

The sets L_k and R_k and the alternating paths will again be constructed using two-sided trees. However, unlike the model with general weight distribution, we do not need to keep track of the weights of the paths, hence the construction is much simpler. For example, the leaf node selection step is no longer necessary. Moreover, in the previous sections we needed the paths to be long enough, so that the weights on the paths dominate the weights of the sprinkling edges. In the unweighted case, that restriction is also lifted. Therefore we can define L_k (resp. R_k) to be all the left (resp. right) vertices in the left (resp. right) subtree, instead of only the selected leaf nodes. The detailed construction is given in Algorithm 4 below.

Algorithm 4 Construction of two-sided trees (under the unweighted model)

- 1: **Input:** n, γ , a bipartite graph G_1 that contains $(1 - \gamma)n$ pairs of vertices, and parameter ℓ .
- 2: Initialize $\mathcal{U} = \{\text{all left vertices of } G_1\}$ as the set of unexplored left vertices.
- 3: For $k = 1, 2, \dots$, repeat the following steps 4-7 to construct the two-sided tree T_k , until $|\mathcal{U}| = (1 - 2\gamma)n$.
- 4: Let i_k be the member of \mathcal{U} with the smallest index. Update $\mathcal{U} \leftarrow \mathcal{U} \setminus \{i_k\}$. Initialize $T_k = \{(i_k, i'_k)\}$ to be a tree containing only one red edge.
- 5: Construct the left tree of T_k via a color-alternating breadth-first search on G_1 . Concretely, define the offsprings of i_k as

$$O'_{i_k} = \{u' \in (\mathcal{U})' : (u', i_k) \in E(G_1)\}.$$

For all $u' \in O'_{i_k}$, append edges (u', v) and (u, u') to T_k . Update $\mathcal{U} \leftarrow \mathcal{U} \setminus O'_{i_k}$. To construct the next 2 layers of the left tree, sequentially (ordering defined by the vertex indices) for all $u \in O'_{i_k}$, define its offsprings as the set of unexplored vertices that are connected to u via a blue edge in G_1 ; append to T_k all the blue edges from u to the offsprings and their corresponding red edges; and mark all the offsprings as explored. Repeat this process until the branching process dies, or the left tree contains ℓ vertices in total.

- 6: Construct the right tree of T_k via the same scheme, starting from the vertex i'_k .
- 7: Let L_k be the set of all left vertices in the left tree of T_k ; let R_k be the set of all right vertices in the right tree of T_k .
- 8: Define $\mathcal{K}_1 = \{k : |L_k| = |R_k| = 2\ell\}$.

Note that Algorithm 4 explores γn pairs of vertices in total. Next, we show that with high probability, the size K_1 of the set \mathcal{K}_1 is at least $c_3 n / \ell$ for some constant $c_3 = \Omega(\epsilon^3)$. That is, a constant proportion of two-sided trees contain exactly 2ℓ pairs of vertices in both the left and the right tree.

By construction, for each k , $|L_k|$ can only be strictly smaller than 2ℓ if the breadth-first search cannot find more vertices to explore, namely, the branching process dies. Since the number $|\mathcal{U}|$ of unused vertices to explore is at least $(1 - 2\gamma)n$, we have

$$\begin{aligned} \mathbb{P}\{|L_k| < 2\ell\} &\leq \text{Probability of extinction for a branching process} \\ &\quad \text{with offspring distribution } \text{Binom}((1 - 2\gamma)n, d/n) =: 1 - c_4. \end{aligned}$$

By choosing $\gamma = \epsilon/2$, the mean of the offspring distribution is $(1 - 2\gamma)d = (1 - \epsilon)(1 + \epsilon)^2 \geq 1 + \epsilon/2$ for ϵ smaller than some universal constant ϵ_0 . Therefore, according to the standard Branching process theory (see e.g. [12, Theorem 25.1]), $1 - c_4$ is the unique solution $\rho < 1$ so that $\phi(\rho) = \rho$, where $\phi(\rho) = \mathbb{E}[\rho^X]$ with $X \sim \text{Binom}((1 - 2\gamma)n, d/n)$. In particular, the probability of survival c_4 is strictly positive, and it can be further shown that $c_4 \geq c_5 \epsilon$ for some universal constant c_5 for small enough ϵ .

Note that the argument would be simplified if the events $\{|L_k| = 2\ell\}$ and $\{|R_k| = 2\ell\}$ were independent. This, however, is not true since when constructing R_k , the number of unexplored vertices depends on $|L_k|$. To resolve this technicality, note that there are always at least $(1 - 2\gamma)n$ unused vertices, and that the construction never reuses vertices. Therefore we can couple the construction of L_k, R_k with two independent branching processes, each with offspring distribution $\text{Binom}((1 - 2\gamma)n, d/n)$, such that

$$\mathbb{P}\{|L_k| = |R_k| = 2\ell\} \geq \mathbb{P}\{\text{both branching processes survive}\} = c_4^2.$$

Since the algorithm uses γn pairs of vertices in total, and at most 2ℓ pairs are used for each k , we have $K \triangleq \text{the total number of two-sided trees} \geq \gamma n / (2\ell)$. Thus

$$\begin{aligned} \mathbb{P}\left\{\sum_{k \leq K} \mathbf{1}_{\{|L_k| = |R_k| = 2\ell\}} < \frac{c_4^2 \gamma n}{4\ell}\right\} &\leq \mathbb{P}\left\{\text{Binom}(K, c_4^2) < \frac{c_4^2 \gamma n}{4\ell}\right\} \\ &\leq \exp\left(-\frac{c_4^4 \gamma}{4\ell} n\right) \end{aligned}$$

by Hoeffding's inequality. We have shown that with probability $1 - \exp(-\Omega(n))$, $K_1 \geq c_3 n / \ell$, with constant $c_3 \triangleq c_4^2 \gamma / 4 \geq c_5^2 \epsilon^3 / 8$.

F.2 Sprinkling stage

In this subsection, we apply Algorithm 3 and Theorem 6 to show the existence of exponentially alternating cycles in G . Since there are no weights, we set the thresholds in Algorithm 3 as $\tau_{\text{red}} = \tau_{\text{blue}} = 0$. Since $\mathcal{P} = \mathcal{Q}$, these thresholds yield $V^* = V$, and

G_2 contains all edges in G except those in $V^c \times (V^c)'$. The parameters in Theorem 6 are specified as follows. Since $|V| = \gamma n$, we have $\beta = \gamma = \epsilon/2$. The blue edge probability in G_2 is d/n , so that $\eta = d$. In the path construction stage, we showed that the set \mathcal{K}_1 is of size at least $c_5^2 \epsilon^3 n / (8\ell)$. By taking its subset, we can assume WLOG that $K_1 = |\mathcal{K}_1| = c_5^2 \epsilon^3 n / (8\ell) = c_6 \epsilon^3 n / \ell$ for some universal constant c_6 .

Let $\ell = c_7/\epsilon^5$ for some universal constant c_7 that will be chosen later. Next, we check the assumptions of Theorem 6. Note that $s = 2\ell$ and

$$b = \frac{\beta s \eta}{4} = \frac{c_7(1+\epsilon)^2}{4\epsilon^4} \geq 4$$

for all $\epsilon < \epsilon_0$ for small enough ϵ_0 ;

$$K_1 = \frac{c_6 \epsilon^3 n}{\ell} \geq 8400$$

for large enough n ;

$$\kappa = \frac{2K_1 s \eta}{n} = 4c_6 \epsilon^3 (1+\epsilon)^2 \leq \frac{1}{16^2}$$

for small enough ϵ_0 ;

$$d_{\text{super}} = \frac{K_1 b^2 \eta}{32n} = \frac{c_6 c_7 (1+\epsilon)^6}{512} \geq 256 \log(32e)$$

for c_7 chosen large enough. We have checked that all the assumptions of Theorem 6 are satisfied. Therefore, with high probability, Algorithm 1 yields at least $\exp(K_2/20) = \exp(c_6 \epsilon^8 n / (320c_7))$ alternating cycles of length at least $3K_2/4 = 3c_6 \epsilon^8 n / (64c_7)$. Each alternating cycle corresponds to a perfect matching in G . By taking $\delta = 3K_2/8 = 3c_6 \epsilon^8 n / (128c_7)$, all these perfect matchings are in \mathcal{M}_{bad} . Note that under the unweighted model, all the perfect matchings that appear in G occupy the same posterior mass. Thus we have shown that with high probability, $\mu_W(\mathcal{M}_{\text{bad}})/\mu_W(M^*) \geq \exp(c_6 \epsilon^8 n / (320c_7))$. By Lemma 1, $\mu_W(\mathcal{M}_{\text{good}})/\mu_W(M^*) \leq 2e^{7\epsilon\delta n}$ with high probability. Thus we conclude that $\mathbb{E}[\ell(\tilde{M}, M^*)] \gtrsim \delta = \Omega(\epsilon^8)$.

References

1. Aldous, D.: The $\zeta(2)$ limit in the random assignment problem. *Random Struct. Algorithms* **18**(4), 381–418 (2001). <https://doi.org/10.1002/rsa.1015>
2. Alon, N., Spencer, J.H.: The Probabilistic Method. Wiley-Interscience Series in Discrete Mathematics and Optimization, 3rd edn (2008)
3. Bagaria, V., Ding, J., Tse, D., Wu, Y., Xu, J.: Hidden Hamiltonian cycle recovery via linear programming. *Oper. Res.* **68**(1), 53–70 (2020)

4. Bhattacharyya, A.: On a measure of divergence between two statistical populations defined by their probability distributions. *Bull. Calcutta Math. Soc.* **35**, 99–109 (1943)
5. Chertkov, M., Kroc, L., Krzakala, F., Vergassola, M., Zdeborová, L.: Inference in particle tracking experiments by passing messages between images. *PNAS* **107**(17), 7663–7668 (2010). <https://doi.org/10.1073/pnas.0910994107>
6. Decelle, A., Krzakala, F., Moore, C., Zdeborová, L.: Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Phys. Rev. E* **84**, 066106 (2011)
7. Ding, J.: Scaling window for mean-field percolation of averages. *Ann. Probab.* **41**(6), 4407–4427 (2013)
8. Ding, J., Goswami, S.: Percolation of averages in the stochastic mean field model: the near-supercritical regime. *Electron. J. Probab.* **20**, 1–21 (2015)
9. Ding, J., Sun, N., Wilson, D.B.: Supercritical minimum mean-weight cycles. *Trans. Am. Math. Soc.* (2015) (to appear). [arXiv:1504.00918](https://arxiv.org/abs/1504.00918)
10. Ding, J., Wu, Y., Xu, J., Yang, D.: Consistent recovery threshold of hidden nearest neighbor graphs. In: Proceedings of Conference on Learning Theory (COLT) (2020)
11. Durrett, R.: Random Graph Dynamics, vol. 200. Citeseer (2007)
12. Frieze, A., Karoński, M.: Introduction to Random Graphs. Cambridge University Press, Cambridge (2016)
13. Hajek, B., Wu, Y., Xu, J.: Information limits for recovering a hidden community. *IEEE Trans. Inf. Theory* **63**(8), 4729–4745 (2017)
14. Karp, R.M.: An upper bound on the expected cost of an optimal assignment. In: Johnson, D.S., Nishizeki, T., Nozaki, A., Wilf, H.S. (eds.) *Discrete Algorithms and Complexity*, pp. 1–4. Academic Press, Cambridge (1987). <https://doi.org/10.1016/B978-0-12-386870-1.50006-X>
15. Krivelevich, M.: Long paths and hamiltonicity in random graphs. *Random Graphs Geom. Asymptot. Struct.* **84**, 1 (2016)
16. Kunisky, D., Niles-Weed, J.: Strong recovery of geometric planted matchings. In: Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 834–876. SIAM (2022)
17. Lesieur, T., Krzakala, F., Zdeborová, L.: Mmse of probabilistic low-rank matrix estimation: universality with respect to the output channel. In: 2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 680–687 (2015). <https://doi.org/10.1109/ALLERTON.2015.7447070>
18. Linusson, S., Wästlund, J.: A proof of Parisi's conjecture on the random assignment problem. *Probab. Theory Relat. Fields* **128**(3), 419–440 (2004). <https://doi.org/10.1007/s00440-003-0308-9>
19. Mézard, Marc, Parisi, Giorgio: On the solution of the random link matching problems. *J. Phys. France* **48**(9), 1451–1459 (1987). <https://doi.org/10.1051/jphys:019870048090145100>
20. Marinari, E., Semerjian, G.: On the number of circuits in random graphs. *J. Stat. Mech. Theory Exp.* **2006**(06), P06019 (2006)
21. Moharrami, M., Moore, C., Xu, J.: The planted matching problem: phase transitions and exact results. *Ann. Appl. Probab.* **31**(6), 2663–2720 (2021). <https://doi.org/10.1214/20-AAP1660>
22. Moore, C.: The computer science and physics of community detection: landscapes, phase transitions, and hardness. (2017). [arXiv:1702.00467](https://arxiv.org/abs/1702.00467)
23. Nair, C., Prabhakar, B., Sharma, M.: Proofs of the Parisi and Coppersmith–Sorkin random assignment conjectures. *Random Struct. Algorithms* **27**(4), 413–444 (2005). <https://doi.org/10.1002/rsa.20084>
24. Robbins, H.: A remark on Stirling's formula. *Am. Math. Mon.* **62**(1), 26–29 (1955)
25. Semerjian, G., Sicuro, G., Zdeborová, L.: Recovery thresholds in the sparse planted matching problem. *Phys. Rev. E* **102**, 022304 (2020). <https://doi.org/10.1103/PhysRevE.102.022304>
26. Sicuro, G., Zdeborová, L.: The planted k -factor problem (2020). [arXiv:2010.13700](https://arxiv.org/abs/2010.13700)
27. Walkup, D.W.: On the expected value of a random assignment problem. *SIAM J. Comput.* **8**(3), 440–442 (1979). <https://doi.org/10.1137/0208036>
28. Wang, H., Wu, Y., Xu, J., Yolou, I.: Random graph matching in geometric models: the case of complete graphs. In: Conference on Learning Theory, pp. 3441–3488. PMLR (2022)
29. Wästlund, J.: An easy proof of the $\zeta(2)$ limit in the random assignment problem. *Electron. Commun. Probab.* **14**, 261–269 (2009). <https://doi.org/10.1214/ECP.v14-1475>

30. Wu, Y., Xu, J.: Statistical problems with planted structures: information-theoretical and computational limits. In: Eldar, Y., Rodrigues, M. (eds.) *Information-Theoretic Methods in Data Science*. Cambridge University Press, Cambridge (2020)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.