

THE PLANTED MATCHING PROBLEM: PHASE TRANSITIONS AND EXACT RESULTS

BY MEHRDAD MOHARRAMI¹, CRISTOPHER MOORE² AND JIAMING XU³

¹*Coordinated Science Lab, University of Illinois at Urbana-Champaign, moharrami@illinois.edu*

²*Santa Fe Institute, moore@santafe.edu*

³*The Fuqua School of Business, Duke University, jx77@duke.edu*

We study the problem of recovering a planted matching in randomly weighted complete bipartite graphs $K_{n,n}$. For some unknown perfect matching M^* , the weight of an edge is drawn from one distribution P if $e \in M^*$ and another distribution Q if $e \notin M^*$. Our goal is to infer M^* , exactly or approximately, from the edge weights. In this paper we take $P = \exp(\lambda)$ and $Q = \exp(1/n)$, in which case the maximum-likelihood estimator of M^* is the minimum-weight matching $M_{\min}$. We obtain precise results on the overlap between M^* and $M_{\min}$, that is, the fraction of edges they have in common. For $\lambda \geq 4$ we have almost perfect recovery, with overlap $1 - o(1)$ with high probability. For $\lambda < 4$ the expected overlap is an explicit function $\alpha(\lambda) < 1$: we compute it by generalizing Aldous' celebrated proof of the $\zeta(2)$ conjecture for the unplanted model, using local weak convergence to relate $K_{n,n}$ to a type of weighted infinite tree, and then deriving a system of differential equations from a message-passing algorithm on this tree.

CONTENTS

1. Introduction	2664
2. Almost perfect recovery for $\lambda \geq 4$	2666
3. Exact results for the expected overlap when $\lambda < 4$	2667
4. Open questions	2672
5. Analysis of the system of ODEs	2674
5.1. Basic properties of the solution	2675
5.2. Monotonicity to the initial condition	2676
5.3. Limiting behavior of (U, V, W)	2678
5.4. Basins of attraction	2679
5.5. Proof of Theorem 5.1	2680
6. Planted networks and local weak convergence	2681
7. The PWIT and planted PWIT	2686
8. The optimal involution invariant matching on the planted PWIT	2687
8.1. The message-passing algorithm	2688
8.2. A rigorous construction of $\mathcal{M}_{\infty, \text{opt}}$	2689
8.3. Optimality of $\mathcal{M}_{\infty, \text{opt}}$	2692
8.4. Uniqueness of the solution of RDEs	2696
9. Convergence of the minimum matching on $(K_{n,n}, \ell_n)$ to $\mathcal{M}_{\infty, \text{opt}}$	2697
9.1. The easy half: A simple compactness argument	2699
9.2. The harder half	2699
9.2.1. The bi-infinite planted PWIT	2700
9.2.2. The unfolding map	2703
9.2.3. Assigning values to the edges of $(K_{n,n}, \ell_n)$	2704
9.2.4. Construction of the matching	2706
Appendix A: Proof of Lemma 2.2	2709

Received April 2020; revised October 2020.

MSC2020 subject classifications. Primary 90C27, 68Q87; secondary 05C80, 62F15, 05C70, 82B26.

Key words and phrases. Random graphs, combinatorial optimization, phase transitions, planted problems, local weak convergence, message-passing algorithms.

Appendix B: Proof of Theorem 7.1 2709

Appendix C: Proof of Proposition 8.4 2712

Appendix D: Proof of Lemma 8.8 2714

 D.1. A lemma on T_u and T_p 2716

 D.2. Calculating with the bi-infinite tree 2717

Acknowledgments 2719

Funding 2719

References 2719

1. Introduction. Consider a weighted complete bipartite graph $K_{n,n}$ with an unknown perfect matching M^* , where for each edge e the weight w_e is independently distributed according to P when $e \in M^*$ and Q when $e \notin M^*$. The goal is to recover the “hidden” or “planted” matching M^* from the edge weights.

This problem is inspired by the long history of planted problems in computer science, where an instance of an optimization or constraint satisfaction problem is built around a planted solution in some random way. As we vary the parameters used to generate these instances, such as the size of a hidden clique or the density of communities in the stochastic block model of social networks, we encounter phase transitions in our ability to find this planted solution, exactly or approximately. In an inference problem, the instance corresponds to some noisy observation, such as a data set produced by a generative model, and the planted solution corresponds to the ground truth—the underlying structure we are trying to discover.

More concretely, we are motivated by the problem of tracking moving objects in a video, such as flocks of birds, motile cells, or particles in a fluid. Figure 1, taken from [12], shows two frames of such a video, where each particle has moved from its original position by some amount. Our goal is then to find the most-likely matching between the two frames, assuming some probability distribution of these displacements.

For many planted problems such as Hidden Clique (e.g., [10]) or community detection in the stochastic block model (e.g., [1, 26]), there are two types of thresholds: information-theoretic and computational. When these are distinct, the region in between them has the interesting property that finding the planted solution, or at least approximating it better than chance, is information-theoretically possible but (conjecturally) computationally hard. These regions are also known as statistical-computational gaps.

In the planted matching problem, one obvious estimator to try is the minimum weight matching (a.k.a. the linear assignment problem) which can be found in polynomial time. The natural question is then, as a function of the distributions P and Q on the planted and




FIG. 1. On the left, the positions of particles in two frames of a video, with one frame red and the other blue. On the right, an inferred matching, hypothesizing how each particle has moved from one frame to the next. Taken from [12].

unplanted edges, how much the minimum matching $M_{\min}$ has in common with the planted matching M^* . In general, we define the *overlap* of an estimator M' with M^* as (assuming that $|M'| = n$)

$$(1.1) \quad \text{overlap}(M^*, M') = 1 - \frac{1}{2n} |M^* \Delta M'| = \frac{1}{n} |M^* \cap M'|.$$

We say that M' achieves *almost perfect recovery* if $\mathbb{E}[\text{overlap}(M^*, M')] = 1 - o(1)$, or equivalently if $\text{overlap}(M^*, M') = 1 - o(1)$ with high probability. We say that M achieves *partial recovery* if $\mathbb{E}[\text{overlap}(M^*, M')] > 0$ as $n \rightarrow \infty$.

Chertkov et al. [12] studied the case where $P = |\mathcal{N}(0, \kappa)|$ is a folded Gaussian and Q is the uniform distribution over $[0, n]$. When $\kappa = O(1)$, the planted edges are competitive with the lightest unplanted edges at each vertex, which have expected weight 1. This suggests a phase transition in this regime, and indeed they predicted a transition from almost perfect recovery to partial recovery at $\kappa \approx 0.17$ using the cavity method of statistical physics.

We focus on exponential weight distributions, $P = \exp(\lambda)$ and $Q = \exp(1/n)$, so that the planted and unplanted weights have expectation $1/\lambda$ and n respectively. For this family of distributions we obtain exact results, proving a transition from almost perfect recovery to partial recovery at $\lambda = 4$, and determining the expected overlap between M^* and $M_{\min}$ for $\lambda < 4$.

Many of our results apply more generally for any distribution of unplanted edge weights with density $Q'(0) = 1/n$, such as when Q is uniform in the interval $[0, n]$. However, our assumption that the planted weights P are exponentially distributed is important for two reasons. First, it makes possible to exactly analyze a message-passing algorithm, and obtain precise results for the expected overlap. Second, it has the pleasing consequence of making $M_{\min}$ the maximum-likelihood estimator for M^* . To see this, note that all $n!$ matchings are equally likely a priori. Let G denote the observed complete bipartite graph with edge weights W . The posterior probability for a given matching M' , that is, $\mathbb{P}[M^* = M' \mid G]$, is proportional to the density

$$(1.2) \quad \begin{aligned} \mathbb{P}[G \mid M'] &= \prod_{e \in M'} P(w_e) \prod_{e \notin M'} Q(w_e) \propto \prod_{e \in M'} \exp(-(\lambda - 1/n)w_e) \\ &= \exp\left(-(\lambda - 1/n) \sum_{e \in M'} w_e\right). \end{aligned}$$

Thus maximizing the likelihood is equivalent to minimizing the total weight of M' .

Our main results are as follows.

- In Theorem 2.1, we show that the minimum matching $M_{\min}$ achieves almost perfect recovery with high probability whenever $\lambda \geq 4$. This proof is a simple first-moment argument using the expected number of augmenting cycles of each length.
- In Theorem 3.1, we compute the expected overlap between M^* and $M_{\min}$ for $\lambda < 4$, showing that it is an explicit function $\alpha(\lambda)$ given by a system of differential equations.

The proof of Theorem 3.1 takes up most of the paper. Our proof is inspired by Aldous' analysis of the minimum matching in the unplanted case where all edges have the same weight distribution with $Q'(0) = 1/n$. Using the machinery of local weak convergence [2, 5, 6] Aldous gave a rigorous justification for the cavity method of statistical physics [25], modeling $K_{n,n}$ as a Poisson-weighted infinite tree (PWIT). The cost of matching a vertex with one of its children then follows a probability distribution which is the fixed point of a recursive distributional equation (RDE) which can then be transformed into an ordinary differential equation (ODE). Solving this ODE proves the conjecture of Mézard and Parisi [25] that the expected cost per vertex is $\zeta(2) = \pi^2/6$.

Generalizing Aldous’ analysis to the planted case presents several challenges. We now have an infinite weighted tree we call the *planted PWIT* with two types of edges and two types of vertices, since the partner of a vertex in M^* can be its parent or one of its children. The cost of matching a vertex with a child follows a pair of probability distributions fixed by a system of RDEs, which (when P is exponential) we can transform into a system of four coupled ODEs. We use techniques from dynamical systems to show that this system has a unique solution consistent with its boundary conditions, and express the expected overlap $\alpha(\lambda)$ as an integral involving this solution.

While we focus on the case where P is exponential, we claim that a qualitatively similar picture to Theorems 2.1 and Theorem 3.1 holds for other distributions of planted weights. Indeed, much of our proof applies to any distribution P , including the general framework of a message-passing algorithm on the planted PWIT, and the resulting system of RDEs. Thus while the location of the threshold and the overlap would change, in any one-parameter family of distributions P we expect there to be a phase transition from almost-perfect to partial recovery when P ’s expectation crosses some critical value.

2. Almost perfect recovery for $\lambda \geq 4$. We start by proving that the minimum matching achieves almost perfect recovery whenever $\lambda \geq 4$.

THEOREM 2.1. *For any $\lambda \geq 4$, we have $\mathbb{E}[\text{overlap}(M^*, M_{\min})] = 1 - o(1)$. In particular, $\mathbb{E}[|M^* \triangle M_{\min}|]$ is $O(1)$ for $\lambda > 4$ and $O(\sqrt{n})$ for $\lambda = 4$.*

To prove Theorem 2.1, we use the following Chernoff-like bound on the probability that one Erlang random variable exceeds another. The proof is elementary and appears in Appendix A.

LEMMA 2.2. *Suppose X_1 is the sum of t independent exponential random variables with rate λ_1 , and X_2 is the sum of t independent exponential random variables with rate λ_2 (and independent of X_1) where $\lambda_1 > \lambda_2$. Then*

$$\mathbb{P}[X_1 > X_2] \leq \left(\frac{4\lambda_1\lambda_2}{(\lambda_1 + \lambda_2)^2}\right)^t \leq \left(\frac{4\lambda_2}{\lambda_1}\right)^t.$$

PROOF OF THEOREM 2.1. An *alternating cycle* is a cycle in $K_{n,n}$ that alternates between planted and unplanted edges, and an *augmenting cycle* is an alternating cycle C where the total weight of its planted edges $C \cap M^*$ exceeds that of its unplanted edges $C \setminus M^*$.

Now recall that the symmetric difference $M^* \triangle M_{\min}$ is a disjoint union of augmenting cycles. The number of cyclic permutations of t things is $(t - 1)!$. Thus the number of alternating cycles of length $2t$, that is, containing t planted edges and t unplanted edges, is at most

$$(2.1) \qquad \binom{n}{t} (t - 1)! = \frac{1}{t} n^t \left(1 - \frac{1}{n}\right) \left(1 - \frac{2}{n}\right) \cdots \left(1 - \frac{t - 1}{n}\right) \leq \frac{1}{t} n^t e^{-t(t-1)/(2n)}.$$

Applying Lemma 2.2 with $\lambda_1 = \lambda$ and $\lambda_2 = 1/n$, the probability that a given alternating cycle of length $2t$ is augmenting is at most $(4/(\lambda n))^t$.

Now the size of the symmetric difference $|M^* \triangle M_{\min}|$ is at most the total length of all augmenting cycles. By the linearity of expectation, its expectation is bounded by

$$\mathbb{E}[|M^* \triangle M_{\min}|] \leq \sum_{t=1}^n 2t \left(\frac{4}{\lambda n}\right)^t \frac{1}{t} n^t e^{-t(t-1)/(2n)} \leq 2e^{1/2} \sum_{t=1}^\infty \left(\frac{4}{\lambda}\right)^t e^{-t^2/(2n)}.$$

When $\lambda > 4$ the geometric sum $\sum_{t=1}^{\infty} (4/\lambda)^t$ converges, giving $\mathbb{E}[|M^* \Delta M_{\min}|] = O(1)$. When $\lambda = 4$, we have $\sum_{t=1}^{\infty} e^{-t^2/(2n)} \leq \int_0^{\infty} e^{-t^2/(2n)} dt = \sqrt{\pi n/2}$, so $\mathbb{E}[|M^* \Delta M_{\min}|] = O(\sqrt{n})$.

To complete the proof, let $\omega(1)$ be any function of n that tends to infinity. By Markov's inequality, with high probability $|M^* \Delta M_{\min}|$ is less than $\omega(1)$ times its expectation, and (1.1) gives w.h.p. $\text{overlap}(M^*, M_{\min}) = 1 - o(1)$. $\square$

We note that when $\lambda > 4$ is sufficiently large we have $\mathbb{E}[|M^* \Delta M_{\min}|] < 1$, implying that $M_{\min}$ achieves perfect recovery, that is, $M_{\min} = M^*$, with positive probability. We also note that a similar argument shows that, for $\lambda < 4$, the overlap is w.h.p. at least $1 - 2 \log \frac{4}{\lambda}$. But this bound is far from tight, and below we give much more precise results.

3. Exact results for the expected overlap when $\lambda < 4$. In this section we provide a characterization of the asymptotic overlap of $M_{\min}$, showing exactly how well $M_{\min}$ achieves partial recovery when $\lambda < 4$.

THEOREM 3.1. *Suppose $0 < \lambda < 4$ is a fixed constant. Then the expected overlap between the minimum matching and the planted one is*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[|M_{\min} \cap M^*|] = \alpha(\lambda),$$

where

$$(3.1) \quad \alpha(\lambda) = 1 - 2 \int_0^{\infty} (1 - F(x))(1 - G(x))V(x)W(x) dx < 1,$$

and where (F, G, V, W) is the unique solution to the coupled system of ordinary differential equations (3.8)–(3.11) given below with boundary conditions (3.12)–(3.13).

Denote the weight of the minimum matching by $w(M_{\min}) \triangleq \sum_{e \in M_{\min}} w_e$, where w_e is the weight of the edge e . We also derive the asymptotic value of $(1/n)\mathbb{E}[w(M_{\min})]$ for $0 < \lambda < 4$.

COROLLARY 3.2. *For $0 < \lambda < 4$, the weight of the minimum matching is*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[w(M_{\min})] = \beta_p(\lambda) + \beta_u(\alpha),$$

where $\beta_p(\lambda)$ and $\beta_u(\lambda)$ are the contributions of planted and unplanted edges to the weight of $M_{\min}$ respectively:

$$\begin{aligned} \beta_p(\lambda) &= \int_{-\infty}^{\infty} (1 - F(x))(1 - G(x))V(x) \left[\int_0^{\infty} \lambda t e^{-\lambda t} (1 - F(t - x)) dt \right] dx, \\ \beta_u(\lambda) &= \int_{-\infty}^{\infty} (1 - F(y))[(1 - G(y))V(y)W(y) - \lambda(G(y) - W(y))] \\ &\quad \times \left[\int_0^{\infty} t(1 - F(t - y))W(t - y) dt \right] dy. \end{aligned}$$

PROOF. We start by relating the planted model $(\ell_n, K_{n,n})$ where ℓ_n denotes the random edge weights, to a type of weighted infinite tree $(\ell_{\infty}, T_{\infty})$ as Aldous did for the unplanted model [2, 6]. This tree corresponds to the neighborhood of a uniformly random vertex, where “local” is defined in terms of shortest path length (sum of edge weights). While $K_{n,n}$ has plenty of short loops, this neighborhood is locally treelike since it is unlikely to have any short loops consisting entirely of low-weight edges.




FIG. 2. The planted Poisson weighted infinite tree (planted PWIT) (ℓ_∞, T_∞) with the labeling scheme described in the text. The bold red edges are planted edges and the solid blue edges are unplanted. The root is the empty string $\emptyset$. Appending 0 to the label of a vertex indicates its planted child, if any, while appending $i \geq 1$ indicates its unplanted child with the i th lightest edge.

Starting at a root vertex $\emptyset$, we define the tree T_∞ shown in Figure 2. The root has a planted child, that is, a child connected to it by a planted edge (bold in red), and a series of unplanted children (solid blue). We label these vertices with strings of integers as follows: the root is labeled with the empty string $\emptyset$. Appending 0 to a label indicates the planted child of that parent, if it has one—that is, if its partner in the planted matching is a child rather than its parent. We indicate the unplanted children by appending i for $i \in \{1, 2, 3, \dots\}$.

We sort the unplanted children of each vertex so that the one labeled with i is the i th lightest, that is, has the i th lightest edge. Since the distribution of unplanted weights has density $Q'(0) = 1/n$ at 0, these weights are asymptotically described by the arrivals of a Poisson process with rate 1, while the weight of the planted edges are distributed as $\exp(\lambda)$. We call the resulting structure the planted Poisson weighted infinite tree, or planted PWIT, and use ℓ_∞ to denote its edge weights. We define all this formally in Section 6 and Section 7, and prove that the finite planted model $(\ell_n, K_{n,n})$ weakly converges to (ℓ_∞, T_∞) .

Following Aldous [6], in Section 8 we then construct a matching $\mathcal{M}_{\infty,\text{opt}}$ on the planted PWIT. Crucially, it has a symmetry property called *involution invariance*, which roughly speaking means that it treats the root just like any other vertex in the tree. We prove that it is the unique involution invariant matching that minimizes the expected cost at the root.

We define $\mathcal{M}_{\infty,\text{opt}}$ in terms of the fixed point of a message-passing algorithm that computes, for each vertex v , the cost of matching v with its best possible child. This cost is the minimum over v ’s children w of the weight of the edge between them, minus the analogous cost for w :

$$X_v = \min_{\text{children } w \text{ of } v} (\ell_\infty(v, w) - X_w).$$

Now suppose that the X_w ’s are independent, and our goal is to compute the distribution of X_v . Unlike the unplanted model, the two types of children will have their X_w drawn from two different distributions. In the first case, w is v ’s planted child, and w ’s children are all unplanted. In the second case, w is an unplanted child of v , and has a planted child of its own. Let X and Y denote the distributions of X_w in these two cases. Then assuming that X_v obeys the appropriate distribution gives the following system of recursive distributional equations (RDEs):

(3.2)
$$X \stackrel{d}{=} \min\{\zeta_i - Y_i\}_{i=1}^\infty,$$

(3.3)
$$Y \stackrel{d}{=} \min(\eta - X, \{\zeta_i - Y_i\}_{i=1}^\infty) \stackrel{d}{=} \min(\eta - X, X'),$$

where the Y_i 's are i.i.d., X and X' are i.i.d., $\eta \sim \exp(\lambda)$, and the ζ_i for $i = 1, 2, \dots$ are jointly distributed as the arrivals of a Poisson process of rate 1.

In general, analyzing recursive distributional equations (RDEs) is very challenging, since they act on the infinite-dimensional space of probability distributions over the reals. However, it is sometimes possible to “collapse” them into a finite-dimensional system of ordinary differential equations. For the unplanted case of the random matching problem, Aldous [6] derived a single differential equation whose solution is the logistic distribution. For the planted case, we use a similar approach, but arrive at a more complicated system of four coupled ODEs.

LEMMA 3.3. *Let $f_X, f_Y, F_X(x) = \mathbb{P}[X < x]$ and $F_Y(y) = \mathbb{P}[Y < y]$ denote the probability density functions and cumulative distribution functions (CDFs) of X and Y , and let $\bar{F}_X = 1 - F_X$ and $\bar{F}_Y = 1 - F_Y$. If (3.2)–(3.3) have a solution, then*

$$(3.4) \quad \frac{dF_X(x)}{dx} = \bar{F}_X(x)\bar{F}_X(-x)\mathbb{E}[F_X(\eta + x)].$$

PROOF. First note that (3.3) gives

$$(3.5) \quad \bar{F}_Y(y) = \bar{F}_X(y)\mathbb{E}[F_X(\eta - y)].$$

Now the pairs $\{(\zeta_i, Y_i)\}$ in (3.2) form a two-dimensional Poisson point process $\{(z, y)\}$ on $\mathbb{R}_+ \times \mathbb{R}$ with density $f_Y(y) dz dy$. We have $X > x$ if and only if none of these points have $z - y < x$, so

$$(3.6) \quad \bar{F}_X(x) = \exp\left(-\iint_{z-y < x} f_Y(y) dz dy\right) = \exp\left(-\int_{z=-x}^{\infty} \bar{F}_Y(z) dz\right).$$

Taking derivatives of both sides of this equation with respect to x and using (3.5) gives

$$f_X(x) = \frac{dF_X(x)}{dx} = -\frac{d\bar{F}_X(x)}{dx} = \bar{F}_X(x)\bar{F}_Y(-x) = \bar{F}_X(x)\bar{F}_X(-x)\mathbb{E}[F_X(\eta + x)]. \quad \square$$

For the sake of simplicity, we omit the subscript X in $F_X(\cdot)$ in the sequel. Define

$$(3.7) \quad G(x) = F(-x), \quad V(x) = \mathbb{E}[F(\eta + x)], \quad W(x) = V(-x).$$

LEMMA 3.4. *When $\eta \sim \exp(\lambda)$, F is a solution to (3.4) if and only if (F, G, V, W) is a solution to the following four-dimensional system of ordinary differential equations (ODEs):*

$$(3.8) \quad \frac{dF}{dx} = (1 - F(x))(1 - G(x))V(x),$$

$$(3.9) \quad \frac{dV}{dx} = \lambda(V(x) - F(x)),$$

$$(3.10) \quad \frac{dG}{dx} = -(1 - F(x))(1 - G(x))W(x),$$

$$(3.11) \quad \frac{dW}{dx} = \lambda(G(x) - W(x))$$

with the boundary conditions

$$(3.12) \quad \begin{aligned} F(-\infty) &= V(-\infty) = G(+\infty) = W(+\infty) = 0, \\ F(+\infty) &= V(+\infty) = G(-\infty) = W(-\infty) = 1, \end{aligned}$$

and

$$(3.13) \quad 0 \leq F, G \leq 1, \quad 0 < V, W \leq 1.$$

PROOF. For one direction, suppose F is a solution to (3.4). Then (3.8) and (3.10) directly follow from (3.4) by plugging in the definition of (F, G, V, W) ; thus they hold for any distribution of η . In contrast, (3.9) and (3.11) are derived via integration by parts under the assumption that $\eta \sim \exp(\lambda)$. The conditions (3.12) and (3.13) hold because F must be a valid CDF. Note that $V(x), W(x) > 0$ for any finite x by definition, as η is larger than any fixed threshold with a positive probability.

For the other direction, suppose F is a solution to the system of ODEs (3.8)–(3.11) with conditions (3.12)–(3.13). Clearly F satisfies (3.4). We only need to verify that F is a valid CDF, which is equivalent to checking (1) F is nondecreasing; (2) $F(+\infty) = 1$ and $F(-\infty) = 0$; and (3) F is right continuous. All these properties are satisfied automatically. $\square$

We comment that RDEs can be solved exactly for some other problems with random vertex or edge weights in the case of the exponential distribution, such as maximum weight independence sets and maximum weight matching in sparse random graphs [15–17]. In some cases this is simply because the minimum of a set of exponential random variables is itself an exponential random variable. To our knowledge our situation involving integration by parts is more unusual.

An interesting consequence of (3.8)–(3.13) is the following conservation law:

$$(3.14) \quad F(x)W(x) + G(x)V(x) - V(x)W(x) = 0.$$

Since $F(0) = G(0)$ and $V(0) = W(0)$, this also implies that

$$(3.15) \quad V(0) = 2F(0).$$

Surprisingly, we find that the system (3.8)–(3.11) exhibits a sharp phase transition at $\lambda = 4$. On the one hand, when $\lambda \geq 4$, they have no solution consistent with (3.12)–(3.13), corresponding to Theorem 2.1 that we have almost perfect recovery in that case. To see this, assume that $V(x) \neq 0$ and introduce a new function $U(x)$ as

$$U(x) = \frac{F(x)}{V(x)}.$$

Then $U(x)$ is differentiable and satisfies

$$(3.16) \quad \frac{dU}{dx} = -\lambda U(1 - U) + (1 - F)(1 - G).$$

LEMMA 3.5. *If $\lambda \geq 4$, then the system of ODEs (3.8)–(3.11) with conditions (3.12)–(3.13) has no solution.*

PROOF. We prove by contradiction. Suppose the system of ODEs (3.8)–(3.11) has a solution satisfying the conditions (3.12)–(3.13). Then $U(x) \rightarrow 1$ as $x \rightarrow +\infty$. Since (3.15) gives $U(0) = F(0)/V(0) = 1/2$, this implies that there is some $x_0 \geq 0$ such that $U(x_0) = 1/2$ and $U'(x_0) \geq 0$. Since $U(x_0) = F(x_0)/V(x_0)$ and $V(x_0) > 0$, we also have $F(x_0) = V(x_0)/2 > 0$, and we also have $G(x_0) \geq 0$. But then (3.16) gives

$$U'(x_0) = -\lambda/4 + (1 - F(x_0))(1 - G(x_0)) < -\lambda/4 + 1 \leq 0.$$

This contradicts $U'(x_0) \geq 0$, and shows that U can never exceed its initial value of $1/2$ or tend to 1 as $x \rightarrow +\infty$. $\square$

On the other hand, Theorem 5.1 in Section 5 proves that for all $\lambda < 4$, there is a unique solution to (3.8)–(3.11) consistent with the conditions (3.12)–(3.13), and hence giving the CDFs of X and Y . The idea hinges on a dynamical fact, namely that the $(U, V, W) = (1, 1, 0)$

is a saddle point, and there is a unique initial condition that approaches it as $x \rightarrow \infty$ along its unstable manifold.

Along with Lemma 3.4, this unique solution to the ODEs gives the unique solution to the RDEs (3.2) and (3.3). Moreover, Theorem 9.1 in Section 9 tells us that the expected overlap of $M_{\min}$ converges to that of $\mathcal{M}_{\infty, \text{opt}}$, which in turn is the probability that the edge weight of a planted edge is less than the cost of matching its endpoints to other vertices:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[|M_{\min} \cap M^*|] = \alpha(\lambda) = \mathbb{P}[\eta < X + X'],$$

where X and X' are i.i.d. with CDF given by F and $\eta \sim \exp(\lambda)$ is independent. Finally, we compute $\mathbb{P}[\eta < X + \hat{X}]$ as follows:

$$\begin{aligned} \mathbb{P}[\eta < X + \hat{X}] &= 1 - \mathbb{E}_{\eta} \left[\int_{-\infty}^{+\infty} f(x) F(\eta - x) dx \right] \\ &= 1 - \int_{-\infty}^{+\infty} \frac{dF(x)}{dx} \mathbb{E}_{\eta} [F(\eta - x)] dx \\ &= 1 - \int_{-\infty}^{+\infty} (1 - F(x))(1 - G(x))V(x)W(x) dx \\ &= 1 - 2 \int_0^{+\infty} (1 - F(x))(1 - G(x))V(x)W(x) dx, \end{aligned}$$

where in the last line we used the fact that the integrand is an even function of x .

Similarly, by Theorem 9.1, Lemma 8.3 and Corollary 8.2, we have

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\sum_{e \in M_{\min}} \ell_n(e) \right] &= \lim_{n \rightarrow \infty} \mathbb{E}[\ell_n(1, M_{\min}(1))] = \mathbb{E}[\ell_{\infty}(\emptyset, \mathcal{M}_{\infty, \text{opt}}(\emptyset))] \\ &= \mathbb{E}[\eta \mathbf{1}(\eta < X + X')] + \mathbb{E}[\zeta \mathbf{1}(\zeta < Y + Y')], \end{aligned}$$

where X and X' are i.i.d. with CDF given by F , Y and Y' are i.i.d. with CDF given by $1 - (1 - F)W$, $\eta \sim \exp(\lambda)$ is independent of X and X' . Finally ζ is “uniformly distributed” over $\mathbb{R}_+$, that is, with Lebesgue measure on $\mathbb{R}_+$, and is independent of Y and Y' . Note that

$$\begin{aligned} \beta_p(\lambda) &\triangleq \mathbb{E}[\eta \mathbf{1}(\eta < X + X')] = \int_0^{\infty} \lambda t e^{-\lambda t} \int_{-\infty}^{\infty} \frac{dF(x)}{dx} (1 - F(t - x)) dx dt \\ &= \int_{-\infty}^{\infty} (1 - F(x))(1 - G(x))V(x) \int_0^{\infty} \lambda t e^{-\lambda t} (1 - F(t - x)) dt dx, \end{aligned}$$

and

$$\begin{aligned} \beta_u(\lambda) &\triangleq \mathbb{E}[\zeta \mathbf{1}(\zeta < Y + Y')] \\ &= \int_0^{\infty} t \int_{-\infty}^{\infty} \frac{dF_Y(y)}{dy} (1 - F_Y(t - y)) dy dt \\ &= \int_0^{\infty} t \int_{-\infty}^{\infty} \left(\frac{dF(y)}{dy} W(y) - (1 - F(y)) \frac{dW(y)}{dy} \right) (1 - F(t - y)) W(t - y) dy dt \\ &= \int_{-\infty}^{\infty} (1 - F(y)) [(1 - G(y))V(y)W(y) - \lambda(G(y) - W(y))] \\ &\quad \times \left[\int_0^{\infty} t (1 - F(t - y)) W(t - y) dt \right] dy. \end{aligned}$$

This completes the proof of Theorem 3.1 and Corollary 3.2. $\square$




FIG. 3. The solid red line is the theoretical value computed by numerically solving the system of ODEs (3.8)–(3.11). The blue dots, is the empirical mean of the corresponding quantity on bipartite graphs generated by the planted model with $n = 1000$. Each dot is the average of 100 independent trials. The error bars show the 95% confidence interval of the distribution, that is, the range into which 95 out of 100 trials fell, suggesting that both quantities are concentrated around their expectations.

To illustrate our results, we plot the function $\alpha(\lambda)$ and $\beta_p(\lambda) + \beta_u(\lambda)$ in Figure 3, and compare with experimental results from finite graphs with $n = 1000$. As the plot shows, when $\lambda \geq 4$ and the overlap is $1 - o(1)$ w.h.p., $\beta_p(\lambda) + \beta_u(\lambda) = 1/\lambda$, that is, the expected weight of a planted edge. Conversely when $\lambda \rightarrow 0$ and the overlap is $o(1)$ w.h.p., $\beta_p(\lambda) + \beta_u(\lambda) \rightarrow \zeta(2)$ as in the unplanted model, since $M_{\min}$ consists almost entirely of unplanted edges.

We comment that the connection between the finite planted model and the planted PWIT is an integral part of the above argument. We explore this connection in detail in Sections 6–9. The results presented in these sections are true for any distribution of unplanted edge weights with density $Q'(0) = 1/n$, and any distribution of planted edge weights P .

4. Open questions. We conclude the discussion of our main results with some open questions.

1. We have computed the expected overlap and cost (per vertex) of the minimum matching. However, in Figure 3 both quantities appear to concentrate on their expected values, and we conjecture that they both have variance $O(1/n)$. We give numerical evidence for this in Figure 4, where we plot empirical values of these variances for fixed $\lambda = 1$ and varying n . For the unplanted model with exponentially distributed edge weights, Wästlund [34, 36] derived the precise asymptotic formula $(4\zeta(2) - 4\zeta(3))/n + O(1/n^2)$ for $\text{Var}[\frac{1}{n}w(M_{\min})]$; it is not known if this holds, for example, for the uniform distribution. In the planted case, we can adapt the arguments in [32], Section 10, from the unplanted model to show that

$$\text{Var}\left[\frac{1}{n}w(M_{\min})\right] = O\left(\frac{(\log n)^4}{N(\log \log n)^2}\right)$$

via Talagrand’s inequality.¹ We do not know how to tighten this to $O(1/n)$, or to prove a similar bound for the overlap. We believe both would follow from correlation decay of messages in the planted PWIT, in which case distant pairs of edges in $K_{n,n}$ would be asymptotically independent.

¹The only minor adaptation is to show [32], Proposition 10.3, continues to hold for the planted model, that is, the bipartite graph $K_{n,n}$ after truncating edges with weights above $2u \log n$ is $(u \log n)$ -expanding with high probability. This can be proved by first coupling the planted model to the unplanted one in a way that every edge weight in the planted model is no larger and then directly invoking Propsition.10.3.




FIG. 4. The blue dots are computed by finding the variance of the corresponding quantity on bipartite graphs generated by the planted model with $\lambda = 1$. Each dot is the average of 1000 independent trials. For both the overlap and the cost the variance appears to decrease as $O(1/n)$ as shown by the dashed lines with slope -1 on this log-log plot.

2. We have focused here on the maximum-likelihood estimator, which for the exponential distribution is also the minimum-weight matching. In physical terms, we have considered this problem at zero temperature. In contrast, the posterior distribution $\mathbb{P}[M^* = M' \mid G]$ given in (1.2) is a Gibbs distribution at nonzero temperature. The estimator with the largest expected overlap would then be the maximum marginal estimator, that is, the set of edges e for which $\mathbb{P}[e \in M^* \mid G] \geq \mathbb{P}[e \notin M^* \mid G]$. This estimator is not generally a matching or even of size n ; nevertheless, one can restrict to estimators which are perfect matchings while increasing the expected misclassification rate $|M' \triangle M^*|/(2n)$ at most by a factor of two. We leave for future work the problem of computing the expected overlap of this estimator. It is possible that it achieves almost-perfect recovery for some $\lambda < 4$, that is, that the information-theoretic threshold for almost-perfect recovery is different from the threshold we have computed here, but we conjecture this is not the case.

3. In physics, a phase transition is called *continuous* if the order parameter (in this case, the overlap) is continuous at the threshold, and as *p th order* if its $(p - 2)$ th derivative is continuous. Although we have not proved this, $\alpha(\lambda)$ in Figure 3 appears to have zero derivative at $\lambda = 4$. This suggests that the transition in the optimal overlap is of third or higher order, unlike other well-known problems in random graphs such as the emergence of the giant component (second order) [11], the stochastic block model with two groups (second order) or with four or more groups (first order) [26], or the appearance of the k -core for $k \geq 3$ (first order) [28]. Very recently, a nonrigorous argument that neglects small terms in the RDE [31] was given that suggests the transition is infinite order, that is, with all derivatives continuous at $\lambda = 4$. Proving this is an attractive open problem.

4. A related question is how the minimum matching changes when the graph undergoes a small perturbation. Aldous and Percus [4] introduced this problem formally and classified combinatorial optimization problems based on how the cost of the optimal solution scales with the size of the perturbation. Using a cavity-based analysis and Monte Carlo simulation, they suggested that the minimum cost among all perfect matchings that differ from the minimum matching by at least δn edges is larger than the cost of the minimum matching by $\Theta(\delta^3)$. This framework has been studied rigorously in [7] and [8] for different combinatorial optimization problems. It would be interesting to explore this same kind of perturbation in the planted model, where we study the minimum cost among all matchings that differ from the planted matching by at least δn edges.

5. Can Lemma 3.5 be turned into a proof of almost perfect recovery? More generally, when the RDEs (3.2) and (3.3) lack a solution supported on $\mathbb{R}$ (i.e., excluding weights in $\{\pm\infty\}$) does this imply almost perfect recovery?

6. We have given two proofs that the RDEs have a unique solution if $\lambda < 4$. Theorem 5.1 uses the dynamics of the ODEs, while Theorem 8.11 uses the uniqueness of $\mathcal{M}_{\infty,\text{opt}}$ on the planted PWIT. These two types of reasoning seem completely orthogonal, but they must be connected. When do the properties of the optimum involution invariant object on an appropriate type of infinite tree imply the dynamical fact that a system of RDEs has a unique fixed point?

7. What can we say about distributions $P(w)$ of planted weights other than exponential? For what distributions is it possible to collapse the RDEs into a finite-dimensional system of ODEs? As stated above, Chertkov et al. [12] studied the folded Gaussian distribution $P = |\mathcal{N}(0, \kappa)|$, but we have been unable to reduce the RDEs to ODEs in this case. Nevertheless, in the spirit of universality classes in physics, we expect any reasonable family of distributions P to undergo a phase transition similar to what we have shown here for the exponential distribution, namely from almost-perfect to partial recovery at some critical value of P 's expectation (where this critical value may depend on the shape of the distribution P). Moreover, with respect to question #3 above, we expect the order of this phase transition, and other scaling properties in its vicinity, to be robust as long as $P'(0) > 0$.

8. Finally, what about planted models with spatial structure, as in the original problem of particle tracking from [12]?

5. Analysis of the system of ODEs. In this section, we state and prove Theorem 5.1.

THEOREM 5.1. *When $\lambda < 4$, the system of ODEs (3.8)–(3.11) has a unique solution (F, G, V, W) satisfying conditions (3.12)–(3.13).*

When $V(x) \neq 0$, recall that

$$U(x) = \frac{F(x)}{V(x)}.$$

Thus, when $V(x) \neq 0$ and $W(x) \neq 0$, the conservation law $FW + GV - VW = 0$ is equivalent to

$$\frac{G(x)}{W(x)} = 1 - U(x).$$

Also, recall that from the conservation law we have $V(0) = W(0) = 2F(0) = 2G(0)$. Hence, the previous 4-dimensional system of ODEs (3.8)–(3.11) with conditions (3.12)–(3.13) reduces to the following 3-dimensional system of ODEs:

(5.1)

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

with initial condition

(5.2)

$$U(0) = \frac{1}{2}, \quad V(0) = W(0) = \epsilon, \quad \epsilon \in [0, 1].$$

Note that the partial derivatives of the right-hand side of (5.1) with respect to (U, V, W) are continuous. Therefore, by the standard existence and uniqueness theorem for solutions of systems of ODEs (see, e.g., [29], Theorem 2) it follows that the system (5.1) with the initial condition (5.2) has a unique solution in a neighborhood of 0 for a fixed $\epsilon \in [0, 1]$. We write this unique solution as $U(x, \epsilon)$, $V(x, \epsilon)$, and $W(x, \epsilon)$, which we abbreviate as (U, V, W) whenever the context is clear. We extend the neighborhood to the maximum interval I_ϵ in which the solution $(U(x, \epsilon), V(x, \epsilon), W(x, \epsilon))$ is finite, so that the existence and uniqueness theorem applies to I_ϵ . In particular, if the solution $(U(x, \epsilon), V(x, \epsilon), W(x, \epsilon))$ is finite for all $x \in [0, \infty)$, then $I_\epsilon = [0, \infty)$; otherwise, at least one of U , V , W converges to ∞ as x converges to some finite x_0 and then $I_\epsilon = [0, x_0)$. In the latter case, we say the solution is $\pm\infty$ for $x \geq x_0$.

Therefore, to prove Theorem 5.1, it suffices to show that the system of ODEs (5.1) with the initial condition (5.2) has a solution $(U(x, \epsilon_0), V(x, \epsilon_0), W(x, \epsilon_0))$ for $x \in [0, +\infty)$ satisfying the boundary condition $U(+\infty, \epsilon_0) = V(+\infty, \epsilon_0) = 1$ and $W(+\infty, \epsilon_0) = 0$ for a unique $\epsilon_0 \in [0, 1]$. Geometrically speaking, this is due to the fact that $(U = 1, V = 1, W = 0)$ is a *saddle point*, and there is a unique choice of ϵ_0 so that the trajectory $(U(x, \epsilon_0), V(x, \epsilon_0), W(x, \epsilon_0))$ falls into the stable manifold, that is, set of initial conditions $(U(0, \epsilon_0), V(0, \epsilon_0), W(0, \epsilon_0))$ such that $(U(x, \epsilon_0), V(x, \epsilon_0), W(x, \epsilon_0)) \rightarrow (1, 1, 0)$ as $x \rightarrow +\infty$. For any other choice of $\epsilon \neq \epsilon_0$, the trajectory $(U(x, \epsilon), V(x, \epsilon), W(x, \epsilon))$ veers away from $(1, 1, 0)$ to infinity.

The outline of the proof is as follows. We first prove some basic properties satisfied by the solution (U, V, W) in Section 5.1. Then based on these properties, in Section 5.2 we prove that the solution satisfies some monotonicity properties with respect to ϵ by studying the sensitivity of the solution to the initial condition. Next, in Section 5.3 we characterize the limiting behavior of the solution depending on whether it hits 1 or not. The monotonicity properties and the limiting behavior enable us to completely characterize the basins of attraction in Section 5.4. In particular, we show that the basin of attraction for $(U = 1, V = 1, W = 0)$ is a singleton, that is, there is a unique choice of $\epsilon_0 \in [0, 1]$ such that $(U(x, \epsilon_0), V(x, \epsilon_0), W(x, \epsilon_0)) \rightarrow (1, 1, 0)$ as $x \rightarrow +\infty$. Finally, we connect the 3-dimensional system of ODEs (5.1) back to the 4-dimensional system of ODEs (3.8)–(3.11) and finish the proof of Theorem 5.1 in Section 5.5.

5.1. Basic properties of the solution. In the following lemma, we prove some basic properties of the solution.

LEMMA 5.2. *Fix any $\epsilon \in (0, 1]$. Then for any $x \in [0, +\infty)$ such that the unique solution (U, V, W) is well-defined (not equal to $\pm\infty$), it holds that*

$$V(x) = W(x)e^{\lambda x}, \quad UV < 1, \quad (1 - U)W < 1, \quad U, V, W > 0.$$

PROOF. It follows from (5.1) that

$$(5.3) \quad V(x) = \epsilon \exp\left(\lambda \int_0^x (1 - U(y)) dy\right) > 0,$$

$$(5.4) \quad W(x) = \epsilon \exp\left(-\lambda \int_0^x U(y) dy\right) > 0.$$

Hence $V(x) = W(x)e^{\lambda x}$. Thus the conservation law $FW + GV - VW = 0$ implies that

$$V = F + Ge^{\lambda x}, \quad W = Fe^{-\lambda x} + G.$$

Recall that $F = UV$ and $G = (1 - U)W$. Then

$$(5.5) \quad \begin{aligned} \frac{dF}{dx} &= (1 - F)(1 - G)V = (1 - F)(1 - G)(F + Ge^{\lambda x}), \\ \frac{dG}{dx} &= -(1 - F)(1 - G)W = -(1 - F)(1 - G)(Fe^{-\lambda x} + G). \end{aligned}$$

For the sake of contradiction, suppose $\max\{F(x), G(x)\} \geq 1$ for some finite $x > 0$. Since $F(x)$ and $G(x)$ are continuous in x and $F(0) = G(0) < 1$, there is an $x_0 > 0$ such that $\max\{F(x_0), G(x_0)\} = 1$. Define $\tilde{F}(x) \equiv F(x_0)$ and $\tilde{G}(x) \equiv G(x_0)$. Then $(\tilde{F}(x), \tilde{G}(x))$ is a solution to ODE (5.5) in $x \in [0, x_0]$ running backward with its initial value at $x = x_0$ given by $(F(x_0), G(x_0))$. Note that $(F(x), G(x))$ is also a solution to ODE (5.5) in the backward time $x \in [0, x_0]$ with its initial value at $x = x_0$ given by $(F(x_0), G(x_0))$. Also note that the right-hand side of ODE (5.5) is continuous in x and the partial derivatives with respect to F and G are continuous. By existence and uniqueness ([29], Theorem 2), it follows that $F(x) \equiv F(x_0)$ and $G(x) \equiv G(x_0)$ for $x \in [0, x_0]$. Hence, $\max\{F(0), G(0)\} = 1$, which contradicts that $F(0) = G(0) = \epsilon/2 < 1$. Thus, $F = UV < 1$ and $G = (1 - U)W < 1$.

Next, we argue that $U > 0$. Suppose not, since $U(0) = 1/2$, by the differentiability of U in x , there exists a finite $x_0 > 0$ such that $U(x_0) = 0$ and $U'(x_0) \leq 0$. However,

$$\left. \frac{dU}{dx} \right|_{x=x_0} = [1 - U(x_0)V(x_0)][1 - (1 - U(x_0))W(x_0)] = [1 - F(x_0)][1 - G(x_0)] > 0,$$

which leads to a contradiction. $\square$

5.2. Monotonicity to the initial condition. The key to our proof is to study how the solution of the system of ODEs (5.1) changes with respect to the initial condition (5.2).

Standard ODE theory (see [29], Theorem 15) shows that $U(x, \epsilon)$ is differentiable in ϵ and the mixed partial derivatives satisfy

$$\frac{\partial^2 U(x, \epsilon)}{\partial x \partial \epsilon} = \frac{\partial^2 U(x, \epsilon)}{\partial \epsilon \partial x};$$

similarly for V and W . Moreover, the partial derivatives $(\partial U/\partial \epsilon, \partial V/\partial \epsilon, \partial W/\partial \epsilon)$ satisfy the following system of equations:

$$(5.6) \quad \begin{aligned} \frac{\partial}{\partial x} \frac{\partial U}{\partial \epsilon} &= [-\lambda(1 - 2U) - V(1 - (1 - U)W) + (1 - UV)W] \frac{\partial U}{\partial \epsilon} \\ &\quad - U(1 - (1 - U)W) \frac{\partial V}{\partial \epsilon} - (1 - UV)(1 - U) \frac{\partial W}{\partial \epsilon}, \\ \frac{\partial}{\partial x} \frac{\partial V}{\partial \epsilon} &= -\lambda V \frac{\partial U}{\partial \epsilon} + \lambda(1 - U) \frac{\partial V}{\partial \epsilon}, \\ \frac{\partial}{\partial x} \frac{\partial W}{\partial \epsilon} &= -\lambda W \frac{\partial U}{\partial \epsilon} - \lambda U \frac{\partial W}{\partial \epsilon}, \end{aligned}$$

with initial condition

$$(5.7) \quad \frac{\partial U(0, \epsilon)}{\partial \epsilon} = 0, \quad \frac{\partial V(0, \epsilon)}{\partial \epsilon} = 1, \quad \frac{\partial W(0, \epsilon)}{\partial \epsilon} = 1.$$

The system of equations (5.6) is known as the *system of variational equations* and can be derived by differentiating (5.1) with respect to ϵ and interchange ∂x and $\partial \epsilon$. The initial condition (5.7) can be derived by differentiating (5.2) with respect to ϵ .

The following key lemma shows that whenever $U(x, \epsilon) \leq 1$, $U(x, \epsilon)$ is decreasing in ϵ , while $V(x, \epsilon)$ and $W(x, \epsilon)$ are increasing in ϵ .

LEMMA 5.3. Fix $\epsilon \in (0, 1)$. Suppose $U(x, \epsilon) \leq 1$ for $x \in (0, c]$ for a finite constant $c > 0$. Then for all $x \in (0, c]$,

$$(5.8) \quad \frac{\partial U(x, \epsilon)}{\partial \epsilon} < 0 \quad \text{and} \quad \frac{\partial W(x, \epsilon)}{\partial \epsilon} > 0.$$

Moreover, it follows that all $x \in (0, c]$,

$$(5.9) \quad \frac{\partial V(x, \epsilon)}{\partial \epsilon} \geq \exp\left(\lambda \int_0^x (1 - U(t, \epsilon)) dt\right) = \frac{V(\epsilon)}{\epsilon} \geq 1.$$

PROOF. We first show that (5.9) holds whenever $\partial U(x, \epsilon)/\partial \epsilon < 0$ for $x \in (0, c]$. Recall that in Lemma 5.2 we have shown that $V > 0$. It follows from (5.6) that for all $x \in (0, c]$

$$\frac{\partial}{\partial x} \frac{\partial V}{\partial \epsilon} \geq \lambda(1 - U) \frac{\partial V}{\partial \epsilon}.$$

Thus for all $x \in (0, c]$

$$\frac{\partial V(x, \epsilon)}{\partial \epsilon} \geq \exp\left(\lambda \int_0^x (1 - U(s)) ds\right) = \frac{V(x, \epsilon)}{\epsilon} \geq 1,$$

where the equality holds due to (5.3).

Next we show (5.8). For the sake of contradiction, suppose not, that is, there exists a $x_0 \in (0, c]$ such that either $\frac{\partial U(x_0, \epsilon)}{\partial \epsilon} \geq 0$ or $\frac{\partial W(x_0, \epsilon)}{\partial \epsilon} \leq 0$.

Define

$$a = \inf\left\{x \in (0, c] : \frac{\partial U(x, \epsilon)}{\partial \epsilon} \geq 0\right\}$$

and

$$b = \inf\left\{x \in (0, c] : \frac{\partial W(x, \epsilon)}{\partial \epsilon} \leq 0\right\},$$

with the convention that the infimum of an empty set is $+\infty$. Then $\min\{a, b\} \leq x_0 \leq c$.

Case 1: Suppose $a \leq b$. Due to the initial condition (5.7) and the initial condition (5.2), we have that

$$\frac{\partial U(0, \epsilon)}{\partial \epsilon} = 0, \quad \frac{\partial}{\partial x} \frac{\partial U(0, \epsilon)}{\partial \epsilon} = -\left(1 - \frac{\epsilon}{2}\right) < 0.$$

Then we have $a > 0$. Moreover, by the differentiability of $\frac{\partial U(x, \epsilon)}{\partial \epsilon}$ in x and the definition of a , we have

$$\frac{\partial U(x, \epsilon)}{\partial \epsilon} < 0 \quad \forall x \in (0, a), \quad \frac{\partial U(a, \epsilon)}{\partial \epsilon} = 0 \quad \text{and} \quad \frac{\partial}{\partial x} \frac{\partial U(a, \epsilon)}{\partial \epsilon} \geq 0.$$

It follows from our previous argument for proving (5.9) that $\partial V(x, \epsilon)/\partial \epsilon \geq 1$ for all $x \in (0, a]$. Since $a \leq b$, we also have that

$$\frac{\partial W(a, \epsilon)}{\partial \epsilon} \geq 0.$$

Recall that in Lemma 5.2 we have shown that $UV < 1$, $(1 - U)W < 1$ and $U > 0$. Moreover, by assumption we have $U \leq 1$. Thus we get from ODE (5.6) that

$$\frac{\partial}{\partial x} \frac{\partial U(a, \epsilon)}{\partial \epsilon} = -U(1 - (1 - U)W) \frac{\partial V(a, \epsilon)}{\partial \epsilon} - (1 - UV)(1 - U) \frac{\partial W(a, \epsilon)}{\partial \epsilon} < 0,$$

which contradicts $\frac{\partial}{\partial x} \frac{\partial U(a, \epsilon)}{\partial \epsilon} \geq 0$.

Case 2: Suppose $a > b$. Due to the initial condition (5.6), we have that

$$\frac{\partial W(0, \epsilon)}{\partial \epsilon} = 1.$$

Thus $b > 0$. By the differentiability of $\frac{\partial W(x, \epsilon)}{\partial \epsilon}$ in x , we have that

$$\frac{\partial W(b, \epsilon)}{\partial \epsilon} = 0 \quad \text{and} \quad \frac{\partial}{\partial x} \frac{\partial W(b, \epsilon)}{\partial \epsilon} \leq 0 \quad \text{and} \quad \frac{\partial U(b, \epsilon)}{\partial \epsilon} < 0.$$

Recall that in Lemma 5.2 we have shown that $W > 0$. It follows from ODE (5.6) that

$$\frac{\partial}{\partial x} \frac{\partial W(b, \epsilon)}{\partial \epsilon} = -\lambda W \frac{\partial U(b, \epsilon)}{\partial \epsilon} > 0,$$

which contradicts $\frac{\partial}{\partial x} \frac{\partial W(b, \epsilon)}{\partial \epsilon} \leq 0$. $\square$

Based on Lemma 5.3, we prove another “monotonicity” lemma, showing that if $U(x, \epsilon_0) < 1$ for all $x \geq 0$ and some $\epsilon_0 \in (0, 1)$, then $U(x, \epsilon) < 1$ for all $x \geq 0$ and all $\epsilon \in (\epsilon_0, 1)$.

LEMMA 5.4. *Suppose $U(x, \epsilon_0) < 1$ for all $x \geq 0$ and some $\epsilon_0 \in (0, 1)$. Then $U(x, \epsilon) < 1$ for all $\epsilon \in (\epsilon_0, 1)$ and all $x \geq 0$.*

PROOF. Fix an arbitrary but finite $x_0 > 0$. We claim that $U(x_0, \epsilon) < 1$ for all $\epsilon \in (\epsilon_0, 1)$. Suppose not. Then define

$$\epsilon_1 \triangleq \inf\{\epsilon \in (\epsilon_0, 1) : 1 \leq U(x_0, \epsilon) < +\infty\}.$$

Note that by assumption, $U(x_0, \epsilon_0) < 1$. By the definition of ϵ_1 and the differentiability of $U(x_0, \epsilon)$ in ϵ , we have

$$U(x_0, \epsilon_1) = 1, \quad \frac{\partial U(x_0, \epsilon_1)}{\partial \epsilon} \geq 0.$$

We claim that $U(x, \epsilon_1) < 1$ for all $x \in (0, x_0)$. If not, then there exists an $x_1 \in (0, x_0)$ such that $U(x_1, \epsilon_1) = 1$. Note that $\frac{dU(x, \epsilon_1)}{dx} > 0$ if $U(x) \geq 1$. Thus $U(x, \epsilon_1) > 1$ for all $x > x_1$, which contradicts the fact that $U(x_0, \epsilon_1) = 1$. Therefore, we can apply Lemma 5.3 with $c = x_0$ and get that

$$\frac{\partial U(x_0, \epsilon_1)}{\partial \epsilon} < 0,$$

which contradicts the fact that $\frac{\partial U(x_0, \epsilon_1)}{\partial \epsilon} \geq 0$. Since x_0 is arbitrarily chosen, we conclude that $U(x, \epsilon) < 1$ for all $\epsilon \in (\epsilon_0, 1)$ and all $x > 0$. $\square$

5.3. *Limiting behavior of (U, V, W) .* In this section, we characterize the limiting behavior of (U, V, W) , depending on whether U or V hit 1.

First, we state a simple lemma, showing that if both U and V do not hit 1 in finite time, then they converge to 1 as $x \rightarrow \infty$.

LEMMA 5.5. *If $U(x, \epsilon_0) < 1$ and $V(x, \epsilon_0) < 1$ for all $x \geq 0$ and some $\epsilon_0 \in (0, 1)$, then $U(x, \epsilon_0) \rightarrow 1$, $V(x, \epsilon_0) \rightarrow 1$, and $W(x, \epsilon_0) \rightarrow 0$ as $x \rightarrow \infty$.*

PROOF. By Lemma 5.2, we have $W(x) = V(x)e^{-\lambda x} \rightarrow 0$ as $x \rightarrow \infty$. Recall that according to (5.1),

$$\frac{dV}{dx} = \lambda V(1 - U) > 0.$$

Since $V(x) < 1$ for all $x \geq 0$, it follows that $\frac{dV}{dx} \rightarrow 0$ and hence $U(x) \rightarrow 1$ as $x \rightarrow +\infty$. Thus, as $x \rightarrow +\infty$,

$$\frac{dU}{dx} = -\lambda U(1 - U) + (1 - UV)(1 - (1 - U)W) \rightarrow 0,$$

which implies that $V(x) \rightarrow 1$ as $x \rightarrow +\infty$. $\square$

The next lemma shows the behavior of U and V if they hit 1 for finite x .

LEMMA 5.6. *Let $x_0 > 0$ be finite.*

- If $V(x_0) = 1$, then $V(x)$ monotonically increases to $+\infty$ and $U(x) \rightarrow 0$ for $x \geq x_0$.
- If $U(x_0) = 1$, then $U(x)$ monotonically increases to $+\infty$ and $V(x) \rightarrow 0$ for $x \geq x_0$.

PROOF. Suppose $V(x_0) = 1$. Recall that in Lemma 5.2, we have shown that $UV < 1$. According to ODE (5.1), we get that $dV/dx > 0$ if $V \geq 1$ as $UV < 1$. Thus $V(x)$ monotonically increases to $+\infty$ for $x \geq x_0$. Moreover $U(x) \rightarrow 0$ for $x \geq x_0$.

Suppose $U(x_0) = 1$. Recall that in Lemma 5.2, we have shown that $(1 - U)WV < 1$. According to ODE (5.1), we get that $dU/dx > 0$ if $U \geq 1$. Hence, $U(x)$ monotonically increases to $+\infty$ for $x \geq x_0$. As $UV < 1$, it further follows that $V(x) \rightarrow 0$ for $x \geq x_0$. $\square$

5.4. *Basins of attraction.* In view of Lemma 5.5 and Lemma 5.6, define the basin of attraction for $(U = 0, V = +\infty)$ as

$$S_1 = \{\epsilon \in [0, 1] : V(x, \epsilon) \geq 1 \text{ for some finite } x > 0\};$$

the basin of attraction for $(U = +\infty, V = 0)$ as

$$S_2 = \{\epsilon \in [0, 1] : U(x, \epsilon) \geq 1 \text{ for some finite } x > 0\};$$

and the basin of attraction for $(U = 1, V = 1)$ as

$$S_0 = \{\epsilon \in [0, 1] : U(x, \epsilon) < 1 \text{ and } V(x, \epsilon) < 1 \text{ for all finite } x > 0\}.$$

When ϵ is either 0 or 1, we have the following simple characterizations of the solution.

LEMMA 5.7. *Suppose $\lambda < 4$.*

- If $\epsilon = 0$, then $V(x) \equiv 0$, $W(x) \equiv 0$ and $U(x)$ monotonically increases to $+\infty$.
- If $\epsilon = 1$, then $V(x)$ monotonically increases to $+\infty$ and $U(x) \rightarrow 0$.

PROOF. First, consider the case $\epsilon = 0$. Then according to the system of ODEs (5.1), we immediately get that $V(x) \equiv 0$, $W(x) \equiv 0$. Thus

$$\frac{dU}{dx} = -\lambda U(1 - U) + 1 > 0,$$

where the last inequality holds due to $\lambda < 4$. Hence, $U(x)$ monotonically increases to $+\infty$.

The conclusion in the case $\epsilon = 1$ simply follows from Lemma 5.6. $\square$

Now we are ready to prove a lemma, which completely characterizes the basins of attraction S_0 , S_1 , and S_2 .

LEMMA 5.8. *Suppose $\lambda < 4$. Then there exists a unique $\epsilon_0 \in (0, 1)$ such that*

$$(5.10) \quad S_0 = \{\epsilon_0\}, \quad S_1 = (\epsilon_0, 1], \quad S_2 = [0, \epsilon_0).$$

PROOF. Lemma 5.7 implies that $1 \in S_1$ and $0 \in S_2$. Note that $UV < 1$ by Lemma 5.2. Thus it follows from Lemma 5.6 that S_1 and S_2 are disjoint.

We first prove that S_1 is left open. Fix any $\epsilon \in S_1$. Since $V(x, \epsilon) \geq 1$ for some finite x , it follows from Lemma 5.6 that there exists an x_0 such that $V(x_0, \epsilon) > 1$. By the continuity of $V(x_0, \epsilon)$ in ϵ , there exists a $\delta > 0$ such that for all $\epsilon' \in [\epsilon - \delta, \epsilon]$, $V(x_0, \epsilon') > 1$, and thus $V(x, \epsilon') \rightarrow +\infty$ and $U(x, \epsilon') \rightarrow 0$ as $x \rightarrow +\infty$. Hence, $[\epsilon - \delta, \epsilon] \subset S_1$. Thus S_1 is left open.

Analogously, we can prove that S_2 is right open. Note that $S_0 = [0, 1] \setminus (S_1 \cup S_2)$, and S_1 and S_2 are disjoint. It follows that S_0 is nonempty. Let ϵ_0 be any point in S_0 . Next we prove (5.10).

We first fix any $\epsilon \in (\epsilon_0, 1)$. Since $\epsilon_0 \in S_0$, it follows that $U(x, \epsilon_0) < 1$ and $V(x, \epsilon_0) < 1$ for all $x \geq 0$. In view of Lemma 5.4, we have that $U(x, \epsilon) < 1$ for all $\epsilon \in (\epsilon_0, 1)$ and all $x > 0$. It follows from Lemma 5.3 that $\partial V(x, \epsilon)/\partial \epsilon \geq 1$ for all $x > 0$ and all $\epsilon \in (\epsilon_0, 1)$. Thus for all $x \geq 0$,

$$V(x, \epsilon) = V(x, \epsilon_0) + \int_{\epsilon_0}^{\epsilon} \frac{\partial V(x, \eta)}{\partial \eta} d\eta \geq V(x, \epsilon_0) + (\epsilon - \epsilon_0).$$

Since $V(x, \epsilon_0) \rightarrow 1$ as $x \rightarrow +\infty$, there exists an x_0 such that for all $x \geq x_0$,

$$V(x, \epsilon_0) \geq 1 - (\epsilon - \epsilon_0)/2.$$

Combining the last two displayed equation gives that for all $x \geq x_0$,

$$V(x, \epsilon) \geq 1 + (\epsilon - \epsilon_0)/2 > 1.$$

We conclude that $\epsilon \in S_1$ and thus $(\epsilon_0, 1] \subset S_1$.

Next we fix any $\epsilon \in (0, \epsilon_0)$ and show that $\epsilon \in S_2$. Suppose not. Then there exists an $\epsilon_1 \in (0, \epsilon_0)$ such that $U(x, \epsilon_1) < 1$ for all $x \geq 0$. By Lemma 5.4, we have that $U(x, \epsilon) < 1$ for all $\epsilon \in (\epsilon_1, 1)$ and all $x > 0$. In view of Lemma 5.3, it immediately follows that (5.9) holds for all $x > 0$ and all $\epsilon \in (\epsilon_1, 1)$. Thus,

$$V(x, \epsilon_0) = V(x, \epsilon_1) + \int_{\epsilon_1}^{\epsilon_0} \frac{\partial V(x, \epsilon)}{\partial \epsilon} d\epsilon \geq V(x, \epsilon_1) + (\epsilon_0 - \epsilon_1).$$

Note that since $\epsilon_0 \in S_0$, $V(x, \epsilon_0) < 1$ for all $x \geq 0$, it follows that for all $x \geq 0$,

$$V(x, \epsilon_1) < 1 - (\epsilon_0 - \epsilon_1),$$

which contradicts the conclusion of Lemma 5.5. Thus we conclude that $\epsilon \in S_2$ and thus $[0, \epsilon_0) \subset S_2$.

Since S_0 , S_1 , and S_2 are all disjoint, the desired (5.10) readily follows. $\square$

5.5. *Proof of Theorem 5.1.* We are now ready to prove Theorem 5.1. Let $S_0 = \{\epsilon_0\}$, and let $(U(x, \epsilon_0), V(x, \epsilon_0), W(x, \epsilon_0))$ be the unique solution to the system of ODEs (5.1) with the initial condition (5.2). For $x \in [0, +\infty)$, define

$$F(x) = U(x, \epsilon_0)V(x, \epsilon_0), \quad F(-x) = (1 - U(x, \epsilon_0))W(x, \epsilon_0),$$

$$V(x) = V(x, \epsilon_0), \quad V(-x) = W(x, \epsilon_0),$$

$$G(x) = F(-x), \quad G(-x) = F(x),$$

$$W(x) = V(-x), \quad W(-x) = V(x).$$

We show that (F, G, V, W) is a solution to the system of ODEs (3.8)–(3.11) with conditions (3.12)–(3.13). First, by construction (F, G, V, W) satisfy the system of ODEs (3.8)–(3.11). In particular, since $U(0, \epsilon_0) = 1/2$ and $V(0, \epsilon_0) = W(0, \epsilon_0) = \epsilon_0$, we have

$$\lim_{x \downarrow 0} \frac{dV}{dx}(x, \epsilon) = \lambda V(0, \epsilon_0)(1 - U(0, \epsilon_0)) = \frac{\lambda \epsilon_0}{2},$$

$$\lim_{x \uparrow 0} \frac{dV}{dx}(x, \epsilon) = -\lim_{x \downarrow 0} \frac{dW}{dx}(x, \epsilon) = \lambda W(0, \epsilon_0)U(0, \epsilon_0) = \frac{\lambda \epsilon_0}{2}.$$

Therefore, V is differentiable at 0. Analogously, we can verify that F, G , and W are differentiable at 0.

Second, since $\epsilon_0 \in S_0$, by definition $U(x, \epsilon_0) < 1$ and $V(x, \epsilon_0) < 1$ for all $x \geq 0$. Thus it follows from Lemma 5.5 that as $x \rightarrow +\infty$, $U(x, \epsilon_0) \rightarrow 1$, $V(x, \epsilon_0) \rightarrow 1$, and $W(x, \epsilon_0) \rightarrow 0$. Hence, (F, G, V, W) satisfy condition (3.12). Third, in view of Lemma 5.2, we have that $U(x, \epsilon_0), V(x, \epsilon_0), W(x, \epsilon_0) > 0$, $W(x, \epsilon_0) < 1$, $U(x, \epsilon_0)V(x, \epsilon_0) < 1$, and $(1 - U(x, \epsilon_0))W(x, \epsilon_0) < 1$. Therefore, $0 < V, W < 1$ and $0 < F, G < 1$, satisfying condition (3.13).

Next, we show that the solution (F, G, V, W) is unique. Let $(\tilde{F}, \tilde{G}, \tilde{V}, \tilde{W})$ denote another solution to system of ODEs (3.8)–(3.11) with conditions (3.12)–(3.13). Let $\tilde{U} = \tilde{F}/\tilde{V}$. Then $(\tilde{U}, \tilde{V}, \tilde{W})$ is a solution to the system of ODEs (5.1), satisfying the initial condition (5.2) with $\epsilon = \tilde{V}(0) = \tilde{W}(0)$. Moreover, $\tilde{U}(x) < 1$ and $\tilde{V}(x) < 1$ for all $x \geq 0$, because otherwise by Lemma 5.6, either $\tilde{U}(x) \rightarrow +\infty$ or $\tilde{V}(x) \rightarrow +\infty$, violating that $\tilde{F}(x), \tilde{V}(x) \rightarrow 1$. As a consequence, $\tilde{V}(0) \in S_0$. It follows from Lemma 5.8 that $\epsilon_0 = \tilde{V}(0)$. By the uniqueness of the solution to system of ODEs (5.1) with the initial condition (5.2), we have $\tilde{U}(x) \equiv U(x, \epsilon_0)$, $\tilde{V}(x) \equiv V(x, \epsilon_0)$, $\tilde{W}(x) \equiv W(x, \epsilon_0)$. Thus, $(\tilde{F}, \tilde{G}, \tilde{V}, \tilde{W}) = (F, G, V, W)$.

6. Planted networks and local weak convergence. In this section and the succeeding ones we define the planted Poisson Weighted Infinite Tree (planted PWIT), define a matching $\mathcal{M}_{\infty, \text{opt}}$ on it, prove that it is optimal and unique, and prove that the minimum weight matching $M_{n, \text{min}}$ on $K_{n, n}$ converges to it in the local weak sense. We follow the strategy of Aldous' celebrated proof of the $\pi^2/6$ conjecture in the unplanted model [2, 6], and in a few places the review article of Aldous and Steele [5]. There are some places where we can simply re-use the steps of that proof, and others where the planted model requires a nontrivial generalization or modification, but throughout we try to keep our proof as self-contained as possible.

In this section we lay out our notation, and formally define local weak convergence. We apologize to the reader in advance for the notational complications they are about to endure: there are far too many superscripts, subscripts, diacritical marks, and general doodads on these symbols. But some level of this seems to be unavoidable if we want to carefully define the various objects and spaces we need to work with.

First off, the exponential distribution with rate ν is denoted by $\exp(\nu)$. Its cumulative distribution function is $\mathbb{P}[x > t] = e^{-\nu t}$ and its mean is $1/\nu$. For a Borel space $(S, \mathcal{S})$ consisting of a set S and a σ -algebra $\mathcal{S}$, $\mathcal{P}(S)$ is the set of all Borel probability measures defined on S .

We use $\mathbb{Z}$ for the integers, $\mathbb{N}_0 = \{0, 1, 2, \dots\}$ and $\mathbb{N}_+ = \{1, 2, \dots\}$ for the natural numbers, and $\mathbb{R}_+$ for the set of nonnegative real numbers. The number of elements of a set A is denoted by $|A|$. Random variables are denoted by capital letters; when we need to refer to a specific realization we sometimes use small letters.

Our graphs will be simple and undirected unless otherwise specified. Given an undirected graph $G = (V, E)$, a (perfect) matching $M \subset E$ is a set of edges where every vertex $v \in V$ is incident to exactly one edge in M . For each v , we refer to the unique v' such that $\{v, v'\} \in M$ as the partner to v , and will sometimes denote it as $M(v)$; then $M(M(v)) = v$. In a bipartite

graph, a matching defines a one-to-one correspondence between the vertices on the left and those on the right. In a forgivable abuse of notation, will often write $M(v, v') = 1$ if $\{v, v'\} \in M$ and 0 otherwise.

A *rooted graph* $G_\circ = (V, E, \emptyset)$ is a graph $G = (V, E)$ with a distinguished vertex $\emptyset \in V$. The *height* of a vertex $v \in V$ in a rooted graph $G_\circ = (V, E, \emptyset)$ is the shortest-path distance from $\emptyset$ to v , that is, the minimum number of edges among all paths from $\emptyset$ to v .

A *planted graph* $G = (V, E, M^*)$ is a graph (V, E) together with a *planted matching* $M^* \subset E$. Similarly a *rooted planted graph* $G_\circ = (V, E, M^*, \emptyset)$ is a planted graph with a distinguished vertex $\emptyset$. We refer to the edges in M^* and $E \setminus M^*$ as the planted edges and unplanted edges respectively.

Two planted graphs $G = (V, E, M^*)$ and $G' = (V', E', M^{*'})$ are said to be isomorphic if there exists a bijection $\gamma : V \rightarrow V'$ such that $\{v_1, v_2\} \in E$ if and only if $\gamma(\{v_1, v_2\}) := \{\gamma(v_1), \gamma(v_2)\} \in E'$, and $\{v_1, v_2\} \in M^*$ if and only if $\gamma(\{v_1, v_2\}) = \{\gamma(v_1), \gamma(v_2)\} \in M^{*'}$. Thus the isomorphism γ preserves the planted and unplanted edges. A rooted isomorphism from $G_\circ = (V, E, M^*, \emptyset)$ to $G'_\circ = (V', E', M^{*'}, \emptyset')$ is an isomorphism between $G = (V, E, M^*)$ and $G' = (V', E', M^{*'})$ such that $\gamma(\emptyset) = \emptyset'$.

Next we endow a planted graph with a weight function. A *planted network* $N = (G, \ell)$ is a planted graph $G = (V, E, M^*)$ together with a function $\ell : E \rightarrow \mathbb{R}_+$ that assigns weights to the edges. For the sake of brevity, we write $\ell(v, w)$ instead of $\ell(\{v, w\})$.

Now let $K_{n,n} = (V_n, E_n, M_n^*)$ denote a complete bipartite graph together with a planted matching. We use $[n]$ to denote the set of integers $\{1, 2, \dots, n\}$. We label the vertices on the left-hand side of $K_{n,n}$ as $\{1, 2, \dots, n\}$, and the vertices on the right-hand side as $\{1', 2', \dots, n'\}$. In a slight abuse of notation, we denote these sets of labels $[n]$ and $[n']$ respectively. Thus $V_n = [n] \cup [n']$, $E_n = \{\{i, j'\} : i \in [n] \text{ and } j' \in [n']\}$, and $M_n^* = \{\{i, i'\} : i \in [n] \text{ and } i' \in [n']\}$.

Let ℓ_n denote a random function that assigns weights to the edges of $K_{n,n}$ as follows: if $e \in M^*$, then $\ell_n(e) \sim \exp(\lambda)$, and if $e \notin M_n^*$ then $\ell_n(e) \sim \exp(1/n)$. We denote the resulting planted network as $(K_{n,n}, \ell_n)$. We denote the minimum matching on $(K_{n,n}, \ell_n)$ as $M_{n,\min}$. Figure 5 illustrates a realization of the planted model.

We want to define a metric on planted networks, or rather on their isomorphism classes. Two planted networks $N = (G, \ell)$ and $N' = (G', \ell')$ are isomorphic if there is an isomorphism γ between G and G' that preserves the length of the edges, that is, if $\ell(v_1, v_2) =$




FIG. 5. A realization of $(K_{4,4}, \ell_4)$ for $n = 4$ and $\lambda = 1$. Red bold edges are in M_n^* (planted edges), and solid blue edges are in $E_n \setminus M_n^*$.

$\ell'(\gamma(v_1), \gamma(v_2))$. A rooted planted network $N_o = (G_o, \ell)$ is a rooted planted graph G_o together with a weight function ℓ , and we define rooted isomorphism as before. Let $[N_o]$ denote the class of rooted planted networks that are isomorphic to N_o . Henceforth, we use N_o to denote a typical member of $[N_o]$.

Next, we define a distance function $d_\ell(v, v')$ as the shortest-path weighted distance between vertices but treating planted edges as if they have zero weight. That is,

$$(6.1) \quad d_\ell(v, v') := \inf_{\text{paths } p \text{ from } v \text{ to } v'} \sum_{e \in p \setminus M^*} \ell(e).$$

For any vertex $v \in V$ and any $\rho \in \mathbb{R}_+$, we can consider the neighborhood $N_\rho(v) = \{v' : d_\ell(v, v') \leq \rho\}$. A network is *locally finite* if $|N_\rho(v)|$ is finite for all v and all ρ .

Now let $\mathcal{G}_*$ denote the set of all isomorphism classes $[N_o]$, where N_o ranges over all connected locally finite rooted planted networks. There is a natural way to equip $\mathcal{G}_*$ with a metric. Consider a connected locally finite rooted planted network $N_o = (G_o, \ell)$. Now, for $\rho \in \mathbb{R}_+$, we can turn the neighborhood $N_\rho(\emptyset)$ into a rooted subgraph $(G_o)_\rho$. To be precise, $(G_o)_\rho = (V_\rho, E_\rho, M_\rho^*, \emptyset)$ is given as follows:

1. *Vertex set*: $V_\rho = N_\rho(\emptyset) = \{v : d_\ell(v, \emptyset) \leq \rho\}$.
2. *Edge set*: $e \in E_\rho$ if $e \in p$ for some path p starting from $\emptyset$ such that $\sum_{e \in p \setminus M^*} \ell(e) \leq \rho$.
3. *Planted matching*: $M_\rho^* = M^* \cap E_\rho$.

Given this definition, for any $[N_o], [N'_o] \in \mathcal{G}_*$ a natural way to define a distance is

$$d([N_o], [N'_o]) = \frac{1}{R+1},$$

where R is the largest ρ at which the corresponding rooted subnetworks $((G_o)_\rho, \ell)$ and $((G'_o)_\rho, \ell')$ cease to be approximately isomorphic in the following sense:

$$(6.2) \quad R = \sup\{\rho \geq 0 : \text{there exists a rooted isomorphism } \gamma_\rho : (G_o)_\rho \rightarrow (G'_o)_\rho \text{ such that } \forall e \in E_\rho, |\ell(e) - \ell'(\gamma_\rho(e))| < 1/\rho\}.$$

(Note that this isomorphism is between the rooted subgraphs $(G_o)_\rho$ and $(G'_o)_\rho$, not the corresponding rooted networks, so it is not required to preserve the weights exactly.) In other words, N_o and N'_o are close whenever there is a large neighborhood around $\emptyset$ where the edge weights are approximately the same, up to isomorphism. In particular, a continuous function is one that we can approximate arbitrarily well by looking at larger and larger neighborhoods of the root.

Equipped with this distance, we say that a sequence $([N_{n,o} = (G_{n,o}, \ell_n)])_{n=1}^\infty$ converges locally to $[N_{\infty,o}]$, and write $[N_{n,o}] \xrightarrow{\text{loc}} [N_{\infty,o}]$, if the following holds: for all $\rho \in \mathbb{R}_+$ such that $N_{\infty,o}$ does not have a vertex at a distance exactly ρ from the root $\emptyset$, there is an $n_\rho \in \mathbb{N}_0$ such that for all $n > n_\rho$ there is a rooted isomorphism $\gamma_{n,\rho} : (G_{n,o})_\rho \rightarrow (G_{\infty,o})_\rho$ such that $\ell_n(\gamma_{n,\rho}^{-1}(e)) \rightarrow \ell(e)$ for all $e \in E_\rho$ where E_ρ is defined from $G_{\infty,o}$ as above. That is, as n increases, $N_{n,o}$ becomes arbitrarily close to $N_{\infty,o}$ on arbitrarily large neighborhoods.

It is easy to check that d defines a metric on $\mathcal{G}_*$. Moreover, $\mathcal{G}_*$ equipped with this metric is a Polish space: a complete metric space which is *separable*, that is, it has a countable dense subset. Hence, we can use the usual tools in the theory of weak convergence to study sequence of probability measures on $\mathcal{G}_*$. More precisely, define $\mathcal{P}(\mathcal{G}_*)$ as the set of all probability measures on $\mathcal{G}_*$ and endow this space with the topology of weak convergence: a sequence $\mu_n \in \mathcal{P}(\mathcal{G}_*)$ converges weakly to μ_∞ , denoted by $\mu_n \xrightarrow{w} \mu$, if for any continuous bounded function $f : \mathcal{G}_* \rightarrow \mathbb{R}$,

$$\int_{\mathcal{G}_*} f d\mu_n \rightarrow \int_{\mathcal{G}_*} f d\mu_\infty.$$

Since $\mathcal{G}_*$ is a Polish space, $\mathcal{P}(\mathcal{G}_*)$ is a Polish space as well with the Lévy–Prokhorov metric [14], pages 394–395, Theorem 11.3.1 and Theorem 11.3.3. Also, Skorokhod’s theorem ([19], p. 79, Theorem 4.30) implies that μ_n converges weakly to μ_∞ if and only if there are random variables $[N_{n,\circ}]$ and $[N_{\infty,\circ}]$ defined over $\mathcal{G}_*$ such that $[N_{n,\circ}] \sim \mu_n$, $[N_{\infty,\circ}] \sim \mu_\infty$, and $[N_{n,\circ}] \xrightarrow{\text{loc}} [N_{\infty,\circ}]$ almost surely.

This notion of convergence in $\mathcal{G}_*$ was first discussed by Aldous and Steele in [5]. It is called *local weak convergence* to emphasize the fact that this notion of convergence only informs us about the local properties of measure around the root. We are going to use this framework to study the asymptotics of a sequence of finite planted networks. This methodology is known as *the objective method* [5] and has been used to analyze combinatorial optimization problems in a variety of random structures (e.g., [5, 15–17, 21, 30]).

In order to apply this machinery to random finite planted networks, consider a finite planted network $N = (G, \ell)$. For a vertex $v \in V$, let $N_\circ(v)$ denote the planted network rooted at v consisting of v ’s connected component. Then we can define a measure $U(N) \in \mathcal{P}(\mathcal{G}_*)$ as follows:

(6.3)
$$U(N) = \frac{1}{|V|} \sum_{v \in V} \delta_{[N_\circ(v)]},$$

where $\delta_{[N_\circ(v)]} \in \mathcal{P}(\mathcal{G}_*)$ is the Dirac measure that assigns 1 to $[N_\circ(v)] \in \mathcal{G}_*$ and 0 to any other member of $\mathcal{G}_*$. In other words, $U(N)$ is the law of $[N_\circ(\emptyset)]$ where $\emptyset$ is picked uniformly from V . Now, to study the local behavior of a sequence of finite networks $(N_n)_n$, the objective method suggests studying the weak limit of the sequence of measures $(U(N_n))_n$.

DEFINITION 6.1 (Random weak limit). A sequence of finite planted networks $(N_n)_{n=1}^\infty$ has a random weak limit $\mu \in \mathcal{P}(\mathcal{G}_*)$ if $U(N_n) \xrightarrow{w} \mu$.

If N_n is a random planted network, we replace $U(N_n)$ in the above definition with $\mathbb{E}U(N_n)$, where

(6.4)
$$\mathbb{E}U(N)(A) := \mathbb{E}[U(N)(A)] \quad \text{for all Borel sets } A \subseteq \mathcal{G}_*,$$

and the expectation is taken with respect to the randomness of N . For us, in both $(K_{n,n}, \ell_n)$ and the weighted infinite tree $(N_{\infty,\circ}, \ell_\infty)$ we define below, the only source of randomness is the edge weights. It is easy to see that if N is vertex transitive, so that every vertex has the same distribution of neighborhoods, then $\mathbb{E}U(N)$ is the law of $[N(\emptyset)]$ (or of $[N(v)]$ for any vertex v). In many settings, for example, sparse Erdős–Rényi graphs, $U(N)$ converges in distribution to $N(\emptyset)$, since averaging over all possible root vertices effectively averages over N as well. But taking the expectation over N as we do here avoids having to prove this.

Not all probability measures $\mu \in \mathcal{P}(\mathcal{G}_*)$ can be random weak limits. The uniform rooting in the measure associated with finite networks implies a modest symmetry property on the asymptotic measure. One necessary condition for a probability measure to be a random weak limit is called *unimodularity* [3].

To define unimodularity, let $\mathcal{G}_{**}$ denote the set of all isomorphism classes $[N_{\circ\circ}]$, where $N_{\circ\circ}$ ranges over all connected locally finite *doubly-rooted* planted networks—that is, networks with an ordered pair of distinguished vertices. We define $\mathcal{G}_{**}$ as the set of equivalence classes under isomorphisms that preserve both roots, and equip it with a metric analogous to (6.2) to make it complete and separable. A continuous function $f([N_{\circ\circ}(\emptyset, v)])$ is then one which we can approximate arbitrarily well by looking at neighborhoods of increasing size that contain both $\emptyset$ and v .

Then we can define unimodularity as follows.

DEFINITION 6.2 (Unimodularity). A probability measure $\mu \in \mathcal{P}(\mathcal{G}_*)$ is *unimodular* if for all Borel functions $f : \mathcal{G}_{**} \rightarrow \mathbb{R}_+$,

$$(6.5) \quad \int_{\mathcal{G}_*} \sum_{v \in V} f([N_{\circ\circ}(\emptyset, v)]) d\mu([N_{\circ}(\emptyset)]) = \int_{\mathcal{G}_*} \sum_{v \in V} f([N_{\circ\circ}(v, \emptyset)]) d\mu([N_{\circ}(\emptyset)]).$$

In other words, the expectation over μ of the sum (either finite or $+\infty$) over all v of $f([N_{\circ\circ}(\emptyset, v)])$ remains the same if we swap $\emptyset$ and v . Since in a connected graph we can swap any vertex v with $\emptyset$ by a sequence of swaps between $\emptyset$ and its neighbors, each of which moves v closer to the root, this definition is equivalent to one where we restrict f to Borel functions with support on $\{[N_{\circ\circ}(\emptyset, v)] \mid \emptyset \text{ and } v \text{ are neighbors}\}$. With this restriction, unimodularity is known as *involution invariance* [3], Proposition 2.2.

LEMMA 6.3 (Involution invariance). A probability measure $\mu \in \mathcal{P}(\mathcal{G}_*)$ is unimodular if and only if (6.5) holds for all Borel functions $f : \mathcal{G}_{**} \rightarrow \mathbb{R}_+$ such that $f([N_{\circ\circ}(\emptyset, v)]) = 0$ unless $\{\emptyset, v\} \in E$.

Aldous in [5] uses another characterization of involution invariance. Given a probability measure $\mu \in \mathcal{P}(\mathcal{G}_*)$, define a measure $\tilde{\mu}$ on $\mathcal{G}_{**}$ as the product measure of μ and the counting measure on the neighbors of the root, that is,

$$(6.6) \quad \tilde{\mu}(\cdot) := \int_{\mathcal{G}_*} \sum_{v: \{\emptyset, v\} \in E} \mathbf{1}([N_{\circ\circ}(\emptyset, v)] \in \cdot) d\mu([N_{\circ}(\emptyset)]),$$

where $\mathbf{1}$ is the indicator function. Like μ , $\tilde{\mu}$ is a σ -finite measure. Throughout the following sections, we use the tilde to distinguish a measure associated with doubly-rooted planted networks from the corresponding measure associated with singly-rooted ones.

Then Aldous' definition of involution invariance in [5] is as follows.

DEFINITION 6.4 (Involution invariance, again). A probability measure $\mu \in \mathcal{P}(\mathcal{G}_*)$ is said to be *involution invariant* if the induced measure $\tilde{\mu}$ on $\mathcal{G}_{**}$ is invariant under the involution map $\iota : \mathcal{G}_{**} \rightarrow \mathcal{G}_{**}$, that is,

$$\tilde{\mu}(A) = \tilde{\mu}(\iota^{-1}(A)) \quad \text{for all Borel sets } A \subseteq \mathcal{G}_{**},$$

where $\iota([N_{\circ\circ}(\emptyset, v)]) = [N_{\circ\circ}(v, \emptyset)]$.

Crucially, unimodularity and involution invariance are preserved under local weak convergence. Any random weak limit satisfies unimodularity and is involution invariant [3, 5] (although the converse is an open problem).

The theory of local weak convergence is a powerful tool for studying random combinatorial problems. In the succeeding sections we will prove a series of propositions analogous to [2, 6] showing local weak convergence between our planted model of randomly weighted graphs $K_{n,n}$ and a kind of infinite tree $N_{\infty, \circ}$. These propositions make a rigorous connection between the minimum matching on $K_{n,n}$ and the minimum involution invariant matching $\mathcal{M}_{\infty, \text{opt}}$ on $N_{\infty, \circ}$. Finally, we analyze $\mathcal{M}_{\infty, \text{opt}}$ using the RDEs that we solved with differential equations above.

7. The PWIT and planted PWIT. In this section we define the planted Poisson Weighted Infinite Tree, and show that it is the weak limit of the planted model $(K_{n,n}, \ell_n)$.

Let us ignore the planted matching for the moment and assume that $\ell_n(e) \sim \exp(n)$ for all $e \in E_n$. The problem of finding the minimum matching on this unplanted network is known as the random assignment problem. Kurtzberg [22] introduced this problem with i.i.d. uniform edge lengths on $[0, n]$, and Walkup [33] proved that the expected cost of the minimum matching is bounded and is independent of n . In the succeeding years, many researchers tightened the bound for $\mathbb{E}[X_n]$ (e.g., [18, 20, 23]). Using powerful but nonrigorous methods from statistical physics, Mezard and Parisi [25] conjectured that $\mathbb{E}[X_n]$ has the limiting value $\zeta(2) = \pi^2/6$ as $n \rightarrow \infty$. Aldous first proved [2] that $\mathbb{E}[X_n]$ indeed has a limit, and then [6] proved the $\pi^2/6$ conjecture, using the local weak convergence approach we follow here.

Other methods have been introduced to study this problem [24, 27, 35], including the marvelous fact that for finite n , the expected cost of the minimum matching is the sum of the first n terms of the Riemann series for $\zeta(2)$, namely $1 + 1/4 + 1/9 + \cdots + 1/n^2$. But these methods rely heavily on the specifics of the matching problem, and we will not discuss them here.

As the first step in applying local weak convergence to the planted problem, we are going to identify the weak limit of the planted model according to Definition 6.1: that is, the kind of infinite randomly weighted tree that corresponds to $K_{n,n}$ with weights drawn from our model. To be more precise, we are interested in a probability measure $\mu_\infty \in \mathcal{P}(\mathcal{G}_*)$ that $\mu_n = \mathbb{E}U(N_n) \in \mathcal{P}(\mathcal{G}_*)$ converges to in the local weak sense, where $N_n = (K_{n,n}, \ell_n)$ is the planted model, $U(N_n)$ is the random measure defined in (6.3) by rooting N_n at a uniformly random vertex, and $\mathbb{E}U(N_n)$ is the measure defined in (6.4). Since every neighborhood has the same distribution of neighborhoods in the planted model, the root might as well be at vertex 1, so μ_n is simply the distribution of $[N_{n,\circ}(1)]$. Thus

$$\begin{aligned} \mu_n(A) &= \mathbb{E}U(N_n)(A) \\ (7.1) \qquad &= \frac{1}{2n} \sum_{v \in V_n} \mathbb{E}[\delta_{[N_n(v)]}(A)] = \mathbb{P}[[N_{n,\circ}(1)] \in A] \quad \text{for all Borel sets } A \subseteq \mathcal{G}_*. \end{aligned}$$

In the unplanted model studied by Aldous and others, the weak limit of the random matching problem is the Poisson Weighted Infinite Tree (PWIT). The planted case is similar but more elaborate: the weights of the unplanted edges are Poisson arrivals, but the weights of the planted edges have to be treated separately. We call this the *planted PWIT*, and define it as follows.

We label the vertices V_∞ of the planted PWIT with sequences over $\mathbb{N}_0$, which we denote with bold letters. The root is labeled by the empty sequence $\emptyset$. The children of a vertex $\mathbf{i} = (i_1, i_2, \dots, i_t)$, are $\mathbf{i}j := (i_1, i_2, \dots, i_t, j)$ for some $j \in \mathbb{N}_0$, and if $t > 0$ its parent is $\text{parent}(\mathbf{i}) := (i_1, i_2, \dots, i_{t-1})$. We say that $\mathbf{i}$ belongs to the t th generation of the tree, and write $\text{gen}(\mathbf{i}) = t$.

Appending $j \in \{1, 2, \dots\}$ to $\mathbf{i}$ gives the j th nonplanted child, that is, the child with the j th smallest edge weight among the nonplanted edges descending from the parent $\mathbf{i}$. However, appending $j = 0$ indicates $\mathbf{i}$'s planted child if any, that is, $\mathbf{i}$'s partner in the planted matching if its partner is one of its children instead of its parent. Since the planted partner of a planted child is its parent, these sequences never have two consecutive zeroes. (Note that the root has a planted child, so the first entry in the sequence is allowed to be 0.) We denote the set of such sequences of length t as Σ^t , and the set of all finite such sequences as $\Sigma = \bigcup_{t \in \mathbb{N}_0} \Sigma^t$. Thus the edge set is $E_\infty = \{\{\mathbf{i}, \mathbf{i}j\} \mid j \in \mathbb{N}_0 \text{ and } \mathbf{i}, \mathbf{i}j \in \Sigma\}$, and the planted matching $M_\infty^* \subset E_\infty$ consists of the edges $\{\{\mathbf{i}, \mathbf{i}0\} \mid \mathbf{i}, \mathbf{i}0 \in \Sigma\}$. Let $T_\infty = (V_\infty = \Sigma, E_\infty, M_\infty^*)$ be the resulting planted tree.

Next we define the random edge weights $\ell_\infty : E_\infty \rightarrow \mathbb{R}_+$. The weights of the unplanted edges are distributed just as in the PWIT: that is, for each vertex $i \in V_\infty$, the sequence $(\ell_\infty(i, i_j))_{j=1,2,\dots}$ is distributed jointly as the arrivals $\zeta_1, \zeta_2, \dots$ of a Poisson process with rate 1. Then we have the planted edges: if $i0 \in \Sigma$, then $\ell_\infty(i, i0) \sim \exp(\lambda)$ independent of everything else. Note that these random weights are independent for different parents $i \in V_\infty$.

Finally, let $N_\infty = (T_\infty, \ell_\infty)$ denote the random planted tree and let $N_{\infty, \circ}$ denote the version of N_∞ rooted at $\emptyset$. We call $N_{\infty, \circ}$ the *planted Poisson Weighted Infinite Tree* or the planted PWIT for short. Its structure is shown in Figure 2.

As in Section 6, let $[N_{\infty, \circ}]$ denote the equivalence class of $N_{\infty, \circ}$ up to rooted isomorphisms, and denote by $\mu_\infty \in \mathcal{P}(\mathcal{G}_*)$ the probability distribution of $[N_{\infty, \circ}]$ in $\mathcal{G}_*$. The following theorem shows that μ_n converges weakly to μ_∞ .

THEOREM 7.1. *The planted PWIT is the random weak limit of the planted model on $K_{n,n}$, i.e., $\mu_n \xrightarrow{w} \mu_\infty$.*

SKETCH OF THE PROOF. Similar to the unplanted case [2], Lemmas 10 and 11, the proof follows from the following steps:

1. Recall that μ_n is the distribution of $[N_{n, \circ}(1)]$. We define an exploration process that explores the vertices of $N_{n, \circ}(1)$ starting from the root vertex 1 in a series of stages. At stage m , this process reveals a tree of depth $m + 1$ and maximum arity $m + 1$, where the children of each vertex are its m lightest unplanted neighbors (among the remaining vertices) and possibly its planted partner (if its planted partner is not its parent).

2. In the limit $n \rightarrow \infty$, the tree explored at each stage is asymptotically the same as a truncated version of the planted PWIT, that is, the analogous stage- m neighborhood of the root $\emptyset$.

3. For large enough m (independent of n), the ρ -neighborhood $(G_{n, \circ}(1))_\rho$ of vertex 1 in $N_{n, \circ}(1)$ is a subgraph of the explored tree at stage m of the process with high probability. This is due to the fact that, while $K_{n,n}$ has plenty of cycles that are topologically short, it is very unlikely that any short cycle containing vertex 1 consists entirely of low-weight edges.

4. Finally, the result follows by using the Portmanteau theorem, which enables us to extend the convergence of distributions on local neighborhoods in total variation distance to the desired local weak convergence.

The complete proof is presented in Appendix B. $\square$

Since the planted model on $K_{n,n}$ converges to the planted PWIT, we have every reason to believe that—just as Aldous showed for the unplanted problem—the minimum matching on the planted model converges locally weakly to the minimum involution invariant matching on the planted PWIT. We make this statement rigorous in the following sections, following and generalizing arguments in [2, 5, 6].

8. The optimal involution invariant matching on the planted PWIT. In this section we define the optimal involution invariant random matching $\mathcal{M}_{\infty, \text{opt}}$ on the planted PWIT—or more precisely, the joint distribution $(\ell_\infty, \mathcal{M}_{\infty, \text{opt}})$. We define it in terms of fixed points of a message-passing algorithm, construct it rigorously on the infinite tree, and prove that it is optimal and unique.

Since the planted PWIT is an infinite tree, the total weight of any matching is infinite. This makes it unclear whether there is a well-defined notion of a minimum-weight matching. But since we are ultimately interested in the cost per vertex of the minimum matching on $K_{n,n}$, we call a random matching $(\ell_\infty, \mathcal{M}_\infty)$ on the planted PWIT optimal if it minimizes the expected cost of the edge incident to the root, $\mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}_\infty(\emptyset))]$.

However, since μ_n is involution invariant and involution invariance is preserved under weak limit, we need to restrict our search for minimum matching to involution invariant matchings. This restriction is crucial. For instance, if we simply want to minimize the expected cost at the root, we could construct a matching as follows, akin to a greedy algorithm: first match the root to its lightest child, that is, the one with the lowest edge weight. Then match each of its other children with their lightest child, and so on. For this matching, $\mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}_\infty(\emptyset))] = \mathbb{E}[\min(\eta, \zeta)] = 1/(1 + \lambda)$ where $\eta \sim \exp(\lambda)$ is the weight of the root’s planted edge and $\zeta \sim \exp(1)$ is the weight of its lightest unplanted edge.

However, as pointed out by Aldous for the unplanted model [6], Section 5.1, this matching is not involution invariant. For instance, suppose 1 is $\emptyset$ ’s lightest child, but that 1 has a descending edge whose weight is even less. In this case, if we swap $\emptyset$ and 1, we won’t include the edge $\{1, \emptyset\}$ in the resulting matching. Indeed, in the unplanted case the optimal involution invariant matching has expected weight $\pi^2/6$ per vertex, while this greedy matching has expected weight 1. The lesson here is that the only matchings on the PWIT (or the planted PWIT) that correspond to genuine matchings on $K_{n,n}$ are those that are involution invariant.

Before we proceed, we make a small increment to our formalism. For a network N we define $\mathcal{M}[N]$ as the set of all matchings on N . Now, a *random matching* $(\ell, \mathcal{M})$ on N is a joint distribution of edge weights and matchings, that is, a probability measure on $\mathbb{R}_+^E \times \mathcal{M}[N]$ with marginal ℓ on $\mathbb{R}_+^E$. Intuitively, the reader would probably interpret the phrase “random matching” as a measurable function from $\mathbb{R}_+^E \rightarrow \mathcal{M}[N]$, assigning a distribution of matchings to each realization of the edge weights ℓ . However, here we follow Aldous by using it to mean a distribution over both ℓ and $\mathcal{M}$. Note that $\mathcal{M}$ may have additional randomness even after conditioning on ℓ ; we will eventually learn, however, that $\mathcal{M}_{\infty,\text{opt}}$ does not.

8.1. *The message-passing algorithm.* We start by describing a message-passing algorithm on the planted PWIT that we will use to define $\mathcal{M}_{\infty,\text{opt}}$. We have already discussed this, but we do it here in our notation for the infinite tree.

If $(\ell_\infty, \mathcal{M})$ is involution invariant, $\mathbb{E}[\ell_\infty(v, \mathcal{M}(v))]$ is independent of the choice of $v \in V_\infty$. Let us pretend for now that the total weight of the minimum involution invariant matching $\mathcal{M}_{\infty,\text{opt}}$ is finite, and minimize it with a kind of message-passing algorithm.

For a vertex $v \in V_\infty$, let $T_\infty(v)$ denote the subtree consisting of v and its descendants, rooted at v (in particular, $T_\infty(\emptyset) = T_\infty$). Let $\ell_\infty(T_\infty(v))$ and $\ell_\infty(T_\infty(v) \setminus \{v\})$ denote the total weight of the minimum involution invariant matching on $T_\infty(v)$ and $T_\infty(v) \setminus \{v\}$ respectively. The difference between these, which we denote

$$(8.1) \qquad X_v = \ell_\infty(T_\infty(v)) - \ell_\infty(T_\infty(v) \setminus \{v\}),$$

is the cost of matching v with one of its children, as opposed to leaving it unmatched (or rather matching it with its parent, without including the cost of that edge). This is the difference between two infinite quantities, but as Aldous and Steele say [5] we should “continue in the brave tradition of physical scientists” and see where it leads. While we have already seen the resulting RDEs in the proof of Theorem 3.1, it will be helpful to restate them here in this more precise notation.

Suppose that in a realization of $\mathcal{M}_{\infty,\text{opt}}$, $\emptyset$ is matched with its child i . Then we have

$$\begin{aligned} \ell_\infty(T_\infty(\emptyset)) &= \ell_\infty(\emptyset, i) + \ell_\infty(T_\infty(i) \setminus \{i\}) + \sum_{j \neq i} \ell_\infty(T_\infty(j)) \\ &= \ell_\infty(\emptyset, i) + \ell_\infty(T_\infty(i) \setminus \{i\}) + \ell_\infty(T_\infty(\emptyset) \setminus \{\emptyset\}) - \ell_\infty(T_\infty(i)). \end{aligned}$$

Rearranging and using (8.1), we have

$$X_\emptyset = \ell_\infty(\emptyset, i) - X_i.$$

We can read this as follows: by matching $\emptyset$ with its child i , we pay the weight $\ell_\infty(\emptyset, i)$ of the edge between them, but avoid the cost X_i of having i matching with one of its own children. But of course we want to match $\emptyset$ with whichever child minimizes this cost, giving

$$(8.2) \quad X_\emptyset = \min_{j \geq 0} (\ell_\infty(\emptyset, j) - X_j).$$

Using the same argument, this relation holds for any vertex $v \in V_\infty$. Recalling that the children of v are labeled vj (i.e., v 's label sequence with j appended) for $j \geq 0$, we have

$$(8.3) \quad X_v = \min_{j \geq 0, vj \in V_\infty} (\ell_\infty(v, vj) - X_{vj}).$$

Now recall that v 's planted partner is either its parent or its 0th child. If the former, then this minimization ranges over v 's unplanted children vj for $j \geq 1$. If the latter, then it also includes v 's planted child $v0$. Let us assume that X_v is drawn from one of two distributions over $\mathbb{R}$, and denote this random variable X in the first case and Y in the second case. We expect these distributions to be fixed if we draw X_{vj} independently for each j , and obtain X_v by applying (8.3). Since v 's unplanted children have planted children, but v 's planted child (if any) only has unplanted children, we get the following recursive distributional equations (RDEs):

$$(8.4) \quad X \stackrel{d}{=} \min(\{\zeta_i - Y_i\}_{i=0}^\infty),$$

$$(8.5) \quad Y \stackrel{d}{=} \min(\eta - X, \{\zeta_i - Y_i\}_{i=0}^\infty),$$

where X is independent of everything else, $\{Y_i\}_{i=1}^\infty$ and Y are i.i.d. and $\{\zeta_i\}_{i=1}^\infty$ are the arrivals of a Poisson process with rate 1, and $\eta \sim \exp(\lambda)$ is the weight $\ell(v, v0)$ of the planted edge—these are the edge weights of the planted PWIT described in Section 7.

As we saw in Section 3, the distributional equations (8.4)–(8.5) have a unique fixed point supported on $\mathbb{R}$ whenever $\lambda < 4$. Our next task is to turn this heuristic derivation into a rigorous construction of random variables on the planted PWIT, and use them to construct the minimum involution invariant random matching $\mathcal{M}_{\infty, \text{opt}}$.

8.2. A rigorous construction of $\mathcal{M}_{\infty, \text{opt}}$. The construction is similar to the one in the unplanted model (see [6], Section 4.3, and [5], Section 5.6). We draw random variables X from a fixed point of the system of recursive distributional equations (8.4)–(8.5). Then we show that these random variables generate an involution invariant random matching, by constructing it (randomly) on finite neighborhoods, and then extending it to the infinite tree. In the next subsection, we analyze this matching and show that it is optimal.

Define the set of directed edges $\overleftrightarrow{E}_\infty = \overrightarrow{E}_\infty \cup \overleftarrow{E}_\infty$ of T_∞ by assigning two directions to each edge $e \in E_\infty$: for an edge $e = \{i, ij\}$ let $\overrightarrow{e} = (i, ij) \in \overrightarrow{E}_\infty$ denote the edge directed downward, that is, away from the root, and let $\overleftarrow{e} = (ij, i) \in \overleftarrow{E}_\infty$ denote the edge directed upward toward the root. We use $\overleftrightarrow{e}$ to denote a typical member of $\overleftrightarrow{E}_\infty$. We extend the edge weights to $\overleftrightarrow{E}_\infty$, as $\ell_\infty(\overleftarrow{e}) = \ell_\infty(\overrightarrow{e}) = \ell_\infty(e)$.

The following lemma shows how to define “costs” or “messages” on $\overleftrightarrow{E}_\infty$. It is essentially identical to [6], Lemma 14, and [5], Lemma 5.8, except that we have different distributions of messages on the planted and unplanted edges.

LEMMA 8.1. *Let (X_0, Y_0) be a solution of the system of recursive distributional equations (8.4)–(8.5). Jointly with the edge weights ℓ_∞ , we can construct a random function $X : \overleftrightarrow{E}_\infty \rightarrow \mathbb{R}$ such that the following hold:*

- (i) For every edge $(u, v) \in \overleftrightarrow{E}_\infty$ we have
- (8.6)
$$X(u, v) = \min_{(v, w) \in E_\infty, w \neq u} (\ell_\infty(v, w) - X(v, w)).$$
- (ii) For every planted edge $e \in M_\infty^*$, $X(\overrightarrow{e})$ and $X(\overleftarrow{e})$ each have the same distribution as X_0 .
- (iii) For every unplanted edge $e \notin M_\infty^*$, $X(\overrightarrow{e})$ and $X(\overleftarrow{e})$ each have the same distribution as Y_0 .
- (iv) For every edge $e \in E_\infty$, $X(\overrightarrow{e})$ and $X(\overleftarrow{e})$ are independent.

PROOF. The idea is to construct these random variables on the subtree consisting of all edges up to a given depth h . We do this by initially “seeding” them on the downward-pointing edges at that depth, drawing their X independently from the appropriate fixed-point distribution. We then use the message-passing algorithm given by (8.6) to propagate them through this subtree. As with belief propagation on a tree, this propagation consists of one sweep upward to the root, and then one sweep back downward toward the leaves. Finally, we use the Kolmogorov consistency theorem [19], p. 115, Theorem 6.16, to take the limit $h \rightarrow \infty$, extending the distribution on these finite-depth subtrees to T_∞ .

Formally, let $h \in \mathbb{N}_+$. Let $\overrightarrow{E}_\infty(h)$ and $\overleftarrow{E}_\infty(h)$ respectively denote the set of downward- and upward-directed edges at depth $h - 1$, and let $\overleftrightarrow{E}_\infty(\leq h)$ denote the set of all directed edges up to depth h :

$$\begin{aligned}\overrightarrow{E}_\infty(h) &= \{\overrightarrow{e} = (v, vj) : \text{gen}(v) = h - 1, \{v, vj\} \in E_\infty\}, \\ \overleftarrow{E}_\infty(h) &= \{\overleftarrow{e} = (vj, v) : \text{gen}(v) = h - 1, \{v, vj\} \in E_\infty\}, \\ \overleftrightarrow{E}_\infty(\leq h) &= \{\overleftrightarrow{e} = (v, w) : \text{gen}(v), \text{gen}(w) \leq h, \{v, w\} \in E_\infty\}.\end{aligned}$$

In particular, $\overrightarrow{E}_\infty(h)$ is the set of downward-pointing edges at the leaves of the subtree of depth h , and $\overleftrightarrow{E}_\infty(\leq h)$ is the set of all edges, pointing in both directions, within that subtree (see Figure 6). Our goal is to define X on $\overleftrightarrow{E}_\infty(\leq h)$.

To initialize the process, for each $\overrightarrow{e} \in \overrightarrow{E}_\infty(h)$ we assign the random variable $X(\overrightarrow{e})$ by drawing independently from X_0 if $e \in M_\infty^*$ and from Y_0 if $e \notin M_\infty^*$. We then use (8.6) recursively to define $\{X(\overrightarrow{e}) : \overrightarrow{e} \in \overrightarrow{E}_\infty(k)\}$ for $k \in \{h - 1, h - 2, \dots, 1\}$. Once we have $X(\overrightarrow{e})$ for all edges incident to the root, we use (8.6) to obtain $X(\overleftarrow{e})$ for these edges, that is, for $\overleftarrow{E}_\infty(1)$. We then move back down the tree, using (8.6) at each level to define $\{X(\overleftarrow{e}) : \overleftarrow{e} \in \overleftarrow{E}_\infty(k)\}$ for $k \in \{1, 2, \dots, h\}$.

Parts (ii) and (iii) of the lemma follow from the fact that (X_0, Y_0) are fixed points of (8.4)–(8.5). Part (iv) follows from the fact that, for all $e \in E_\infty(\leq h)$, $X(\overrightarrow{e})$ and $X(\overleftarrow{e})$ are determined by disjoint subsets of $\{X(\overrightarrow{e}) : \overrightarrow{e} \in \overrightarrow{E}_\infty(h)\}$ and hence are independent.




FIG. 6. The sets $\overrightarrow{E}_\infty(2)$ and $\overleftrightarrow{E}_\infty(\leq 2)$. Bold red edges are in M_∞^* and solid blue edges are in $E_\infty \setminus M_\infty^*$.

Finally, we extend these random variables to the entire planted PWIT. For each finite depth h , the above construction gives a collection of random variables

$$\mathcal{X}_h = \{(\ell_\infty(\overrightarrow{e}), X(\overleftarrow{e})) : \overrightarrow{e} \in \overleftarrow{E}_\infty(\leq h)\},$$

that satisfies (i), (ii), (iii), and (iv). Moreover, the marginal distribution of $\mathcal{X}_{h+1}$ restricted to depth h is the same as the distribution of $\mathcal{X}_h$. Now, by the Kolmogorov consistency theorem, there exists a collection of random variables $\mathcal{X}_\infty$ that satisfies (i), (ii), (iii), and (iv), such that the marginal distribution of $\mathcal{X}_\infty$ restricted to depth h is the same as the distribution of $\mathcal{X}_h$. $\square$

One important implication of Lemma 8.1 is the following corollary.

COROLLARY 8.2. *Consider the collection of random variables $\mathcal{X}_\infty$ given by Lemma 8.1.*

- (i) *Let $e = \{\emptyset, 0\}$ denote the planted edge incident to the root. Then $X(\overrightarrow{e})$ and $X(\overleftarrow{e})$ are independent and identically distributed as X_0 , and are independent of $\ell_\infty(e)$.*
- (ii) *Suppose we condition on the existence of an unplanted edge $e = \{\emptyset, i\}$ incident to $\emptyset$ with $\ell_\infty(e) = \zeta$. Then $X(\overrightarrow{e})$ and $X(\overleftarrow{e})$ are independent and identically distributed as Y_0 .*

PROOF. Part (i) follows immediately from the construction in Lemma 8.1. Part (ii) follows from the fact that if we condition on the existence of a Poisson arrival at time ζ , the other arrivals are jointly distributed according to the same Poisson process. There is a subtlety here in that it is important to condition on ζ but not on i , since knowing where ζ is in the sorted order of the unplanted weights affects their distribution. On the other hand, if we fix an edge e before doing this sorting, then $X(\overrightarrow{e})$ and $X(\overleftarrow{e})$ are independent of $\ell_\infty(e)$ for both planted and unplanted edges, and we will use this fact below. $\square$

Our next task is to transform the above construction into a random matching $(\ell_\infty, \mathcal{M}_{\infty, \text{opt}})$. There are two ways we might do this. One would be to define a function on V_∞ that yields a proposed partner w for each vertex v . As in (8.6), matching v with w would cost the weight of the edge between them, but remove the cost of having w pair with one of its other neighbors. Minimizing this total cost over all neighbors w (rather than over all but one as in the message-passing algorithm) gives

$$(8.7) \quad \mathcal{M}_{\infty, \text{opt}}(v) = \arg \min_{w: \{v, w\} \in E_\infty} (\ell_\infty(v, w) - X(v, w)).$$

Since each edge weight $\ell_\infty(v, w)$ is drawn from a continuous distribution, and Corollary 8.2 implies that it is independent of $X(v, w)$, with probability 1 the elements of the set we are minimizing over are distinct and this arg min is well-defined.

Alternately, we could define a mark function on E_∞ as described above, namely the indicator function for the event that an edge e is in the matching. Including e in the matching makes sense if $\ell_\infty(e)$ is less than the cost of matching each of its endpoints to one of their other neighbors. So (abusing notation) this suggests

$$(8.8) \quad \mathcal{M}_{\infty, \text{opt}}(e) = \begin{cases} 1 & \text{if } \ell_\infty(e) < X(\overrightarrow{e}) + X(\overleftarrow{e}), \\ 0 & \text{otherwise.} \end{cases}$$

A priori, there is no guarantee that either of these functions is a matching, or that they agree with each other. The following lemma (which is a reformulation of [5], Lemma 5.9) gives the good news that they are, and they do.

LEMMA 8.3. *The following are equivalent:*

1. $u = \arg \min_{w:\{v,w\} \in E_\infty} (\ell_\infty(v, w) - X(v, w))$
2. $v = \arg \min_{w:\{u,w\} \in E_\infty} (\ell_\infty(u, w) - X(u, w))$
3. $\ell_\infty(u, v) < X(u, v) + X(v, u)$.

Therefore, $u = \mathcal{M}_{\infty, \text{opt}}(v)$ if and only if $v = \mathcal{M}_{\infty, \text{opt}}(u)$ (with $\mathcal{M}_{\infty, \text{opt}}$ defined as in (8.7)), and these are equivalent to $\mathcal{M}_{\infty, \text{opt}}(u, v) = 1$ (with $\mathcal{M}_{\infty, \text{opt}}$ defined as in (8.8)).

PROOF. By (8.6), condition (1) holds if and only if

$$\ell_\infty(v, u) - X(v, u) < \arg \min_{w:\{v,w\} \in E_\infty, w \neq u} (\ell_\infty(v, w) - X(v, w)) = X(u, v).$$

Rearranging gives (3), so (1) and (3) are equivalent. Since (3) is symmetric with respect to swapping u and v , (2) and (3) are also equivalent. $\square$

Finally, given the symmetric dependency of $\mathcal{M}_{\infty, \text{opt}}(e)$ on the values of $X(\vec{e})$ and $X(\overleftarrow{e})$, it is intuitive that the random matching $(\ell_\infty, \mathcal{M}_{\infty, \text{opt}})$ is involution invariant. The following lemma corresponds to [6], Lemma 24, in the unplanted case, but defining the involutions in a way that preserves the (un)planted edges takes a little more work. We give the proof in Appendix C.

PROPOSITION 8.4. *The random matching $(\ell_\infty, \mathcal{M}_{\infty, \text{opt}})$ is involution invariant.*

8.3. *Optimality of $\mathcal{M}_{\infty, \text{opt}}$.* Now that we have constructed $(\ell_\infty, \mathcal{M}_{\infty, \text{opt}})$, it is time to prove that $(\ell_\infty, \mathcal{M}_{\infty, \text{opt}})$ is the minimum involution invariant random matching. The steps we take to prove this claim are mostly the same as in [6], Sections 4.4 and 4.5, but a few details differ in the planted model, so for the sake of completeness and consistency with our notation we give a self-contained proof.

As the first step, we are going to prove that $(\ell_\infty, \mathcal{M}_{\infty, \text{opt}})$ is a minimum involution invariant matching: that is, it achieves the minimum expected length at the root. We follow the discussion at the beginning of Section 4.5 in [6].

PROPOSITION 8.5. *Let $(\ell_\infty, \mathcal{M}'_\infty)$ be an involution invariant random matching on the planted PWIT. Then $\mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}'_\infty(\emptyset))] \geq \mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}_{\infty, \text{opt}}(\emptyset))]$.*

PROOF. Note that in addition to depending on the edge weights ℓ_∞ , $\mathcal{M}'_\infty$ might also have additional randomness. However, we can always couple $(\ell_\infty, \mathcal{M}'_\infty)$ and $(\ell_\infty, \mathcal{M}_{\infty, \text{opt}})$ so that if we condition on ℓ_∞ then $\mathcal{M}'_\infty$ and $\mathcal{M}_{\infty, \text{opt}}$ are independent. Let A be the event that $\mathcal{M}'_\infty(\emptyset) \neq \mathcal{M}_{\infty, \text{opt}}(\emptyset)$, and assume without loss of generality that $\mathbb{P}[A] > 0$.

Conditioned on A , there is a doubly-infinite alternating path that passes through the root $\emptyset$, alternating between edges in $\mathcal{M}'_\infty$ and $\mathcal{M}_{\infty, \text{opt}}$. That is to say, there is a doubly-infinite sequence of distinct vertices $\dots, v_{-2}, v_{-1}, v_0, v_1, v_2, \dots$ where $v_0 = \emptyset$, $v_1 = \mathcal{M}_{\infty, \text{opt}}(\emptyset)$, and $v_{-1} = \mathcal{M}'_\infty(\emptyset)$, and where for all even integers m we have $\mathcal{M}_{\infty, \text{opt}}(v_m) = v_{m+1}$ and $\mathcal{M}'_\infty(v_m) = v_{m-1}$.

By the construction of $\mathcal{M}_{\infty, \text{opt}}$, we know that v_1 achieves the minimum in Equation (8.7):

$$(8.9) \quad \ell_\infty(v_0, v_1) - X(v_0, v_1) = \min_{w:\{v_0,w\} \in E_\infty} (\ell_\infty(v_0, w) - X(v_0, w)).$$

We also have the message-passing equation (8.6) for $X(v_{-1}, v_0)$,

$$(8.10) \quad X(v_{-1}, v_0) = \min_{w:\{v_0,w\} \in E_\infty, w \neq v_{-1}} (\ell_\infty(v_0, w) - X(v_0, w)).$$

The right-hand sides of (8.9) and (8.10) are the same except that v_{-1} is excluded in (8.10). But since the minimum is achieved by v_1 , excluding v_{-1} makes no difference, and the right-hand sides are equal. Rearranging gives

$$(8.11) \quad \ell_\infty(v_0, v_1) = X(v_0, v_1) + X(v_{-1}, v_0).$$

On the other hand, (8.6) also implies $X(u, v) \leq \ell_\infty(v, w) - X(v, w)$ for any u, v, w where u and w are distinct neighbors of v , and in particular

$$(8.12) \quad X(v_1, v_0) \leq \ell_\infty(v_0, v_{-1}) - X(v_0, v_{-1}).$$

Now, using (8.11), the expected difference in the length at the root is

$$(8.13) \quad \begin{aligned} & \mathbb{E}[(\ell_\infty(\emptyset, \mathcal{M}'_\infty(\emptyset)) - \ell_\infty(\emptyset, \mathcal{M}_{\infty, \text{opt}}(\emptyset)))] \\ &= \mathbb{E}[(\ell_\infty(v_0, v_{-1}) - \ell_\infty(v_0, v_1))\mathbf{1}_A] \\ &= \mathbb{E}[(\ell_\infty(v_0, v_{-1}) - X(v_0, v_1) - X(v_{-1}, v_0))\mathbf{1}_A]. \end{aligned}$$

Now we use the fact that $\mathcal{M}_{\infty, \text{opt}}$ and $\mathcal{M}'_\infty$ are both involution invariant. There is a subtlety here in that conditioning on A breaks involution invariance, since it requires $\mathcal{M}_{\infty, \text{opt}}$ and $\mathcal{M}'_\infty$ to differ at the root specifically. However, the involutions that swap v_0 with v_1 or with v_{-1} maintain this conditioning, since $\mathcal{M}_{\infty, \text{opt}}$ and $\mathcal{M}'_\infty$ differ at these vertices as well. It follows that $X(v_0, v_1)$ and $X(v_1, v_0)$ have the same conditional distribution and hence the same conditional expectation, and similarly for $X(v_0, v_{-1})$ and $X(v_{-1}, v_0)$. Then (8.13) becomes

$$(8.14) \quad \begin{aligned} & \mathbb{E}[(\ell_\infty(\emptyset, \mathcal{M}'_\infty(\emptyset)) - \ell_\infty(\emptyset, \mathcal{M}_{\infty, \text{opt}}(\emptyset)))] \\ &= \mathbb{E}[(\ell_\infty(v_0, v_{-1}) - X(v_1, v_0) - X(v_0, v_{-1}))\mathbf{1}_A], \end{aligned}$$

which is greater than or equal to zero by (8.12). $\square$

Even given Proposition 8.5, it is still possible a priori that there might be a random involution invariant matching $(\ell_\infty, \mathcal{M}'_\infty)$ with the same expected length at the root as $(\ell_\infty, \mathcal{M}_{\infty, \text{opt}})$. If we were simply trying to calculate the expected length of the minimum matching, this would not be an issue. But our object is the overlap, not the length. If there are two minimal matchings with the same length but different overlap, it would not be clear which is the weak limit of the minimum matching on $K_{n,n}$.

Happily, we can follow a path similar to [6], Section 4.4 and 4.5, to show that $\mathcal{M}_{\infty, \text{opt}}$ is unique, making the inequality in Proposition 8.5 strict. The following is essentially Proposition 18 of [6].

PROPOSITION 8.6. *Let $(\ell_\infty, \mathcal{M}'_\infty)$ be an involution invariant random matching on the planted PWIT. If $\mathbb{P}[\mathcal{M}'_\infty(\emptyset) \neq \mathcal{M}_{\infty, \text{opt}}(\emptyset)] > 0$ then $\mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}'_\infty(\emptyset))] > \mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}_{\infty, \text{opt}}(\emptyset))]$.*

PROOF. For sake of contradiction, assume there is an involution invariant random matching $(\ell_\infty, \mathcal{M}'_\infty)$ such that $\mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}'_\infty(\emptyset))] = \mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}_{\infty, \text{opt}}(\emptyset))]$. By the proof of Proposition 8.5, we have $\mathbb{E}[D\mathbf{1}_A] = 0$ where

$$D = \ell_\infty(v_0, v_{-1}) - X(v_1, v_0) - X(v_0, v_{-1}) \geq 0,$$

and where A is again the event $\{\mathcal{M}'_\infty(\emptyset) \neq \mathcal{M}_{\infty, \text{opt}}(\emptyset)\}$, and where the inequality $D \geq 0$ is given by Equation (8.12). Therefore, conditioned on A , almost surely

$$(8.15) \quad X(v_1, v_0) = \ell_\infty(v_0, v_{-1}) - X(v_0, v_{-1}).$$

Now recall that v_1 achieves the minimum, over all w in v_0 ’s neighborhood, of $\ell_\infty(v_0, w) - X(v_0, w)$. By Equation (8.6), $X(v_1, v_0)$ is the minimum of this same quantity over all $w \neq v_1$. But this is the second minimum, that is, the second-smallest value, and (8.15) implies

$$(8.16) \qquad v_{-1} = \arg \min^{[2]}_i (\ell_\infty(\emptyset, i) - X(\emptyset, i)),$$

where $\min^{[2]}$ denotes the second minimum. Thus the following holds almost surely: either $\mathcal{M}'_\infty$ agrees with $\mathcal{M}_{\infty, \text{opt}}$ at the root, or it matches the root with the second minimum of $\ell_\infty(\emptyset, i) - X(\emptyset, i)$ rather than the minimum. That is, without conditioning on A ,

$$\mathbb{P}\Big(\mathcal{M}'_\infty(\emptyset) \in \Big\{\arg \min_i (\ell_\infty(\emptyset, i) - X(\emptyset, i)) \text{ or } \arg \min^{[2]}_i (\ell_\infty(\emptyset, i) - X(\emptyset, i))\Big\}\Big) = 1.$$

Since $(\ell_\infty, \mathcal{M}'_\infty)$ is involution invariant, the same relation holds for each vertex $v \in V_\infty$, that is,

$$(8.17) \qquad \mathbb{P}\Big(\mathcal{M}'_\infty(v) \in \Big\{\arg \min_{w:\{w,v\} \in E_\infty} (\ell_\infty(v, w) - X(v, w)) \text{ or } \arg \min^{[2]}_{w:\{w,v\} \in E_\infty} (\ell_\infty(v, w) - X(v, w))\Big\}\Big) = 1.$$

Thus any matching with the same expected length as $\mathcal{M}_{\infty, \text{opt}}$ must, almost surely at almost all vertices v , match v with its best or second-best partner according to $\ell_\infty(v, w) - X(v, w)$.

Surprisingly, no involution-invariant matching can choose the second-best partner with nonzero probability. The following proposition shows that (8.17) cannot hold unless $\mathcal{M}'_\infty = \mathcal{M}_{\infty, \text{opt}}$ almost surely.

PROPOSITION 8.7 (Proposition 20 of [6]). *The only involution invariant random matching that satisfies (8.17) is $\mathcal{M}_{\infty, \text{opt}}$.*

PROOF. The reader might be wondering why we can’t simply assign everyone to their second-best partner. But recall the key fact from Lemma 8.3 that if

$$\mathcal{M}_{\infty, \text{opt}}(v) = \arg \min_{w:\{v,w\} \in E_\infty} (\ell_\infty(v, w) - X(v, w)),$$

then $\mathcal{M}_{\infty, \text{opt}}(\mathcal{M}_{\infty, \text{opt}}(v)) = v$ and $\{\{v, \mathcal{M}_{\infty, \text{opt}}(v)\} : v \in V_\infty\}$ is indeed a matching. The problem is that this fact does not generally hold if we replace $\arg \min$ with $\arg \min^{[2]}$.

If $\mathcal{M}'_\infty$ and $\mathcal{M}_{\infty, \text{opt}}$ differ anywhere with positive probability, then by involution invariance they differ at the root with positive probability. In that case, as before, there is a doubly-infinite alternating path from the root to infinity. Thus once $\mathcal{M}'_\infty$ matches the root with its second-best partner, it must keep doing this forever on that path. But in order for $\mathcal{M}'_\infty$ to be involution invariant, it must make the same choices if we follow the path in reverse, and so each vertex on this path must be the second-best partner of its second-best partner. We will see that the probability that this is true on every step of the path, all the way to infinity, is zero.

Let $\dots, v_{-2}, v_{-1}, v_0, v_1, v_2, \dots$ be the alternating path defined as follows. First let $v_0 = \emptyset$. To define v_t for $t > 0$, we extend the path by alternately apply the best and second-best rules,

$$v_{t+1} = \begin{cases} \arg \min_{u:\{u,v_t\} \in E_\infty} (\ell_\infty(v_t, u) - X(v_t, u)) & \text{if } t \text{ is even.} \\ \arg \min^{[2]}_{u:\{u,v_t\} \in E_\infty} (\ell_\infty(v_t, u) - X(v_t, u)) & \text{if } t \text{ is odd.} \end{cases}$$

Similarly, for $t < 0$ we extend the path backwards,

$$v_{t-1} = \begin{cases} \arg \min_{u: \{u, v_t\} \in E_\infty}^{[2]} (\ell_\infty(v_t, u) - X(v_t, u)) & \text{if } t \text{ is even,} \\ \arg \min_{u: \{u, v_t\} \in E_\infty} (\ell_\infty(v_t, u) - X(v_t, u)) & \text{if } t \text{ is odd.} \end{cases}$$

In particular, $v_1 = \mathcal{M}_{\infty, \text{opt}}(\emptyset)$ and (if A holds) $v_{-1} = \mathcal{M}'_\infty(\emptyset)$.

Now for each odd integer t , define the event B_t that v_t and v_{t+1} are the second-best partners of each other. For odd $t > 0$ we can write

$$B_t = \{v_t = \arg \min_{u: \{u, v_{t+1}\} \in E_\infty}^{[2]} (\ell_\infty(v_{t+1}, u) - X(v_{t+1}, u))\},$$

and for odd $t < 0$,

$$B_t = \{v_{t+1} = \arg \min_{u: \{u, v_t\} \in E_\infty}^{[2]} (\ell_\infty(v_t, u) - X(v_t, u))\}.$$

As discussed above, since $\mathcal{M}'_\infty$ is involution invariant A implies B_t , in particular, for all $t = 1, 3, 5, \dots$. Thus

$$A \subset \overline{B}_\infty := \bigcap_{t=1,3,5,\dots}^\infty B_t.$$

Writing $\overline{B}_t = \bigcap_{t'=1,3,5,\dots}^t B_{t'}$, this implies

$$\mathbb{P}[A] \leq \mathbb{P}[\overline{B}_\infty] = \prod_{t=1,3,5,\dots} \mathbb{P}[B_{t+2} \mid \overline{B}_t] = \prod_{t=1,3,5,\dots} \frac{\mathbb{P}[\overline{B}_{t+2}]}{\mathbb{P}[\overline{B}_t]}.$$

and so

$$(8.18) \quad \text{if } \mathbb{P}[A] > 0 \text{ then } \lim_{t \rightarrow \infty, t \text{ odd}} \frac{\mathbb{P}[\overline{B}_{t+2}]}{\mathbb{P}[\overline{B}_t]} = 1.$$

Now we use involution invariance again. If we root the planted PWIT at v_2 instead of v_0 , sliding the alternating path two steps to the left, the event $\overline{B}_{t+2}$ becomes the event $B_{-1} \cap \overline{B}_t$ (and A still holds). By involution invariance the probability of these two events is the same, so

$$\frac{\mathbb{P}[\overline{B}_{t+2}]}{\mathbb{P}[\overline{B}_t]} = \frac{\mathbb{P}[B_{-1} \cap \overline{B}_t]}{\mathbb{P}[\overline{B}_t]} = \mathbb{P}[B_{-1} \mid \overline{B}_t].$$

By continuity of probability measure, if $\mathbb{P}[\overline{B}_\infty] > 0$ —which holds if $\mathbb{P}[A] > 0$ —we also have

$$\lim_{t \rightarrow \infty} \mathbb{P}[B_{-1} \cap \overline{B}_t] = \mathbb{P}[B_{-1} \cap \overline{B}_\infty] \quad \text{and} \quad \lim_{t \rightarrow \infty} \mathbb{P}[\overline{B}_t] = \mathbb{P}[\overline{B}_\infty],$$

in which case

$$\lim_{t \rightarrow \infty} \mathbb{P}[B_{-1} \mid \overline{B}_t] = \mathbb{P}[B_{-1} \mid \overline{B}_\infty].$$

Thus (8.18) demands that this conditional probability is 1. But the following lemma, which generalizes Lemma 22 of [6] to the planted case, shows that this is not so.

LEMMA 8.8. *If $\mathbb{P}[\overline{B}_\infty] > 0$ then $\mathbb{P}[B_{-1} \mid \overline{B}_\infty] < 1$.*

PROOF. As in [6], Remark on page 402, the idea is that B_{-1} only depends on what happens on the “leftward” branch of the alternating path, $v_0, v_{-1}, v_{-2}, \dots$, while $\overline{B}_\infty$ depends only on the “rightward” branch $v_0, v_1, v_2, \dots$. For the details, see Appendix D. $\square$

Lemma 8.8 implies that $P(A) = 0$, and by the discussion above that $\mathcal{M}'_\infty = \mathcal{M}_{\infty,\text{opt}}$ almost surely. This completes the proof of Proposition 8.7... $\square$

... which completes the proof of Proposition 8.6. $\square$

An immediate corollary of Proposition 8.6 is the following.

COROLLARY 8.9. *In the minimum involution invariant random matching $(\ell_\infty, \mathcal{M}_{\infty,\text{opt}})$, $\mathcal{M}_{\infty,\text{opt}}$ is a function of the edge lengths ℓ_∞ . That is to say, given a realization of $(\ell_\infty(e), e \in E_\infty)$, $\mathcal{M}_{\infty,\text{opt}}$ is a fixed matching on the planted PWIT.*

PROOF. Consider a coupling $(\ell_\infty, \mathcal{M}_{\infty,\text{opt}}, \mathcal{M}'_{\infty,\text{opt}})$ such that conditioned on $(\ell_\infty(e), e \in E_\infty)$, $\mathcal{M}_{\infty,\text{opt}}$ and $\mathcal{M}'_{\infty,\text{opt}}$ are i.i.d. Then, by Proposition 8.6 we have $\mathcal{M}_{\infty,\text{opt}} = \mathcal{M}'_{\infty,\text{opt}}$ almost surely. $\square$

In other words, $\mathcal{M}_{\infty,\text{opt}}$ does not have any additional randomness besides its dependence on ℓ_∞ . This was left as an open question for the unplanted case in [6], Remark (d), although we claim that that paper in fact resolved it! As later stated in [5], this implies that if we use the construction of Section 8.2 to define random variables X on neighborhoods of depth h , then (conditioning on ℓ_∞) the random matching defined by these variables becomes concentrated around a single matching as $h \rightarrow \infty$.

This does not quite imply that the messages X on the directed edges of the planted PWIT are determined by ℓ_∞ . This was shown for the unplanted case by Bandyopadhyay using the concept of endogeny [9]. We believe endogeny holds for the planted case, but we leave this as an open question. In any case, as long as the system of recursive distributional equations (8.4)–(8.5) has a solution supported on $\mathbb{R}$, whether it is unique or not, the minimum involution invariant random matching $\mathcal{M}_{\infty,\text{opt}}$ is uniquely defined. Therefore, whenever we focus on a realization of ℓ_∞ , there is no need to call $\mathcal{M}_{\infty,\text{opt}}$ a random matching.

8.4. *Uniqueness of the solution of RDEs.* Recall from Section 8.2 that $\mathcal{M}_{\infty,\text{opt}}$ is defined by drawing messages at the boundary of neighborhoods of increasing size from a fixed point of the RDEs (8.4)–(8.5), propagating these messages throughout the neighborhood, and then including edges (u, v) whose weights $\ell_\infty(u, v)$ are less than the sum of their messages $X(u, v) + X(v, u)$.

However, Corollary 8.9 shows that $\mathcal{M}_{\infty,\text{opt}}$ is a function of the weights ℓ_∞ . As we commented there, this doesn’t quite imply that the messages X are also functions of ℓ_∞ . However, Corollary 8.9 imposes strong conditions on the possible solutions of the RDEs. Specifically, if the RDEs have more than one solution, then each one must somehow result in the same matching $\mathcal{M}_{\infty,\text{opt}}$ given the edge weights. In this section, we show that this implies that the fixed point is indeed unique. This provides an interesting counterpart to the dynamical proof of uniqueness given in Theorem 5.1.

First we show that any solution has a well-defined moment generating function in a neighborhood of the origin.

LEMMA 8.10. *Let (X, Y) be a solution of the system of recursive distributional equations (8.4)–(8.5) supported on $\mathbb{R}$. Then the random variable X has a finite moment generating function $\mathbb{E}[e^{\mu X}]$ for μ in an open neighborhood of 0.*

PROOF. Recall that F_X and F_Y denote the cumulative distribution functions of X and Y respectively, and $\bar{F}_X$ and $\bar{F}_Y$ denote their complements. On the one hand, by (3.6), for all $x > 0$ we have

$$\bar{F}_X(x) = \exp\left(-\int_{z=-x}^{\infty} \bar{F}_Y(z) dz\right) \leq \exp(-x\bar{F}_Y(0)).$$

On the other hand, for every $x_0 > 0$, Lemma 3.3 gives

$$f_X(x_0) = \bar{F}_X(x_0)\bar{F}_X(-x_0)\mathbb{E}[F_X(\eta + x_0)] \geq \bar{F}_X(-x_0)\bar{F}_X(x_0)\mathbb{E}[F_X(\eta - x_0)] = f_X(-x_0),$$

where the inequality follows by the fact that $F_X(\eta - x_0) \leq F_X(\eta + x_0)$ for all η . Hence,

$$\mathbb{P}[X < -x] \leq \mathbb{P}[X > x] \leq \exp(-x\bar{F}_Y(0)),$$

and $\bar{F}_X(0) \geq 1/2$. Then (3.5) implies

$$\bar{F}_Y(0) = \bar{F}_X(0)\mathbb{E}[F_X(\eta)] \geq \frac{1}{2}\mathbb{E}[F_X(\eta)] > 0.$$

where the last inequality holds because η can be arbitrarily large and X is supported on $\mathbb{R}$. The result now follows by simple algebra. If $0 \leq \mu < \bar{F}_Y(0)$ we have

$$\begin{aligned} \mathbb{E}[e^{\mu X}] &= \int_0^\infty \mathbb{P}[e^{\mu X} > s] ds \leq 1 + \int_1^\infty \mathbb{P}[e^{\mu X} > s] ds \\ &= 1 + \int_1^\infty \mathbb{P}\left[X > \frac{\ln s}{\mu}\right] ds \leq 1 + \int_1^\infty s^{-\bar{F}_Y(0)/\mu} ds < \infty, \end{aligned}$$

and the proof for $-\bar{F}_Y(0) < \mu \leq 0$ is similar. Hence $\mathbb{E}[e^{\mu X}] < \infty$ for $\mu \in (-\bar{F}_Y(0), \bar{F}_Y(0))$. $\square$

Now recall that by Lemmas 8.1 and Lemma 8.3, given $\ell(u, v) = x$, the probability that $(u, v) \in \mathcal{M}_{\infty, \text{opt}}$ equals $\mathbb{P}[X + X' > x]$ where $X = X(u, v)$ and $X' = X(v, u)$ are i.i.d. copies of the random variable X . If the RDEs have two distinct solutions (X_1, Y_1) and (X_2, Y_2) , Corollary 8.9 implies that $\mathbb{P}[X_1 + X'_1 > x] = \mathbb{P}[X_2 + X'_2 > x]$ for all x , so that $X_1 + X'_1$ and $X_2 + X'_2$ have the same distribution. But since

$$\mathbb{E}[e^{\mu(X_1 + X'_1)}] = (\mathbb{E}[e^{\mu X_1}])^2,$$

and similarly for X_2 , this implies that X_1 and X_2 have the same moment generating function, which by Lemma 8.10 is well-defined in a neighborhood of the origin. It follows that X_1 and X_2 have the same distribution [13], Theorem 1. Using (3.3), Y_1 and Y_2 are equidistributed as well, and we have proved the following theorem.

THEOREM 8.11. *Assume the system of recursive distributional equations (8.4)–(8.5) has a solution supported on $\mathbb{R}$. Then such solution is unique.*

9. Convergence of the minimum matching on $(K_{n,n}, \ell_n)$ to $\mathcal{M}_{\infty, \text{opt}}$. At this point we have constructed $(\ell_\infty, \mathcal{M}_{\infty, \text{opt}})$ and shown that it is the unique involution invariant matching on the planted PWIT that minimizes the weight at the root. It is finally time to show that the minimum matching $(\ell_n, M_{n, \text{min}})$ on our original planted model on $K_{n,n}$ converges to $(\ell_\infty, \mathcal{M}_{\infty, \text{opt}})$ in the local weak sense. This implies that these two objects have the same joint distribution of edge weights, and which edges they include in the matching, on neighborhoods of any finite radius. In particular, they have the same expected overlap—which is the overlap we computed in Section 3. Thus we finally complete the proof of Theorem 3.1.

To use the framework of local weak convergence to study minimum matchings, we append $\{0, 1\}$ to the edges of planted networks in $\mathcal{G}_*$. In a slight abuse of terminology, we add a hat and also call $\widehat{N} = (G, \ell, \mathcal{M})$ a planted network where ℓ is the weight function and $\mathcal{M} : E \rightarrow \{0, 1\}$ is a mark function (which may or may not be a matching). This lets us discuss the joint distribution of edge weights, the planted matching, and the minimum matching of vertex neighborhoods in either model.

In particular, let $\widehat{\mu}_n$ be this distribution in the vicinity of a uniformly random vertex in the finite model,

$$\widehat{\mu}_n = \mathbb{E}U(\widehat{N}_n) \quad \text{is the law of} \quad [\widehat{N}_{n,o}(1)] \quad \text{where} \quad \widehat{N}_n = (K_{n,n}, \ell_n, M_{n,\min}),$$

and let $\widehat{\mu}_\infty$ be the analogous distribution at the root of the planted PWIT,

$$\widehat{\mu}_\infty \quad \text{is the law of} \quad [\widehat{N}_{\infty,o}] \quad \text{where} \quad \widehat{N}_\infty = (T_\infty, \ell_\infty, \mathcal{M}_{\infty,\text{opt}}).$$

We will show that $\widehat{\mu}_\infty$ is the weak limit of $\widehat{\mu}_n$. Thus the two models have all the same local statistical properties, including their expected weight and overlap.

The proof consists of two main steps, namely, the easier half and the harder half. In the easier half, using a simple compactness argument we prove that any subsequence of probability measures $\widehat{\mu}_n$ has a subsequence that converges to an involution invariant random matching on planted PWIT. Using Skorokhod’s theorem this shows that the weight of the minimum matching on $K_{n,n}$ is at least that of $\mathcal{M}_{\infty,\text{opt}}$:

$$(9.1) \qquad \liminf_{n \rightarrow \infty} \mathbb{E}[\ell_n(1, M_{n,\min}(1))] \geq \mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}_{\infty,\text{opt}}(\emptyset))].$$

Informally, this follows by contradiction. If $\liminf_n \ell_n(1, M_{n,\min}(1))$ were smaller than this, then the subsequence of sizes n on which it converges to that smaller value would itself have a subsequence that converges to an involution invariant matching on the planted PWIT with that weight. . . but this would contradict the optimality of $\mathcal{M}_{\infty,\text{opt}}$.

In the harder half, using $(\ell_\infty, \mathcal{M}_{\infty,\text{opt}})$ we follow the strategy of [2]. First we construct an almost-perfect matching on $(K_{n,n}, \ell_n)$ with weight close to $\mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}_{\infty,\text{opt}}(\emptyset))]$. Then, we fix this almost-perfect matching to make a perfect matching on $(K_{n,n}, \ell_n)$ without changing the weight too much. This proves that

$$(9.2) \qquad \limsup_{n \rightarrow \infty} \mathbb{E}[\ell_n(1, M_{n,\min}(1))] \leq \mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}_{\infty,\text{opt}}(\emptyset))].$$

Combining (9.1) and (9.2), we have

$$(9.3) \qquad \lim_{n \rightarrow \infty} \mathbb{E}[\ell_n(1, M_{n,\min}(1))] = \mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}_{\infty,\text{opt}}(\emptyset))].$$

As in Aldous’ proof of the $\zeta(2)$ conjecture for the unplanted model, this establishes the expected weight of the minimum matching in the planted model. But much more is true. Since $\mathcal{M}_{\infty,\text{opt}}$ is unique, we get the following theorem.

THEOREM 9.1. *The random weak limit of $(K_{n,n}, \ell_n, M_{n,\min})$ is $(T_\infty, \ell_\infty, \mathcal{M}_{\infty,\text{opt}})$, that is, $\widehat{\mu}_n \xrightarrow{w} \widehat{\mu}_\infty$. In particular, their expected overlap is equal to*

$$\begin{aligned} \alpha(\lambda) &:= \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[|M_{n,\min} \cap M_n^*|] \\ &= \lim_{n \rightarrow \infty} \mathbb{P}(\{1, 1'\} \in M_n^*) = \mathbb{P}(\{\emptyset, 0\} \in \mathcal{M}_{\infty,\text{opt}}). \end{aligned}$$

PROOF. In the easy-half proof, we show that every subsequence of $\widehat{\mu}_n$ has a further subsequence that converges to an involution random matching on the planted PWIT (see next subsection). Now, by (9.3) and Proposition 8.6, every subsequence of $\widehat{\mu}_n$ converges to

$\hat{\mu}_\infty$; so does the whole sequence. Hence, the random weak limit of $(K_{n,n}, \ell_n, M_{n,\min})$ is $(T_\infty, \ell_\infty, \mathcal{M}_{\infty,\text{opt}})$. $\square$

Finally, using Corollary 8.2 we have for the expected overlap

$$(9.4) \quad \mathbb{P}(\{\emptyset, 0\} \in \mathcal{M}_{\infty,\text{opt}}) = \mathbb{P}(X + \hat{X} > \ell_\infty(\emptyset, 0)),$$

where X and $\hat{X}$ are independent copies of X_0 and (X_0, Y_0) is the unique solution of the system of recursive distributional equations (8.4)–(8.5). But this computation is exactly what we have done in Section 3 by transforming these distributional equations into a system of ordinary differential equations. This completes the proof of Theorem 3.1.

9.1. The easy half: A simple compactness argument. As the first step toward the proof of the local weak convergence of $\hat{\mu}_n$ to $\hat{\mu}_\infty$, we show that for any sequence of n that tends to infinity, there is a subsequence that converges weakly to some involution invariant random matching $\hat{\mu}'_\infty$ on the planted PWIT. Saying this again in symbols, for any sequence (n_k) there is an involution invariant $\hat{\mu}'_\infty$ and a subsequence $(n_j) \subseteq (n_k)$ such that $\hat{\mu}_{n_j} \xrightarrow{w} \hat{\mu}'_\infty$. Our argument is somewhat simplified from [5], Section 5.8, pages 53–54.

By Theorem 7.1, we already know $\mu_{n_k} \xrightarrow{w} \mu_\infty$: that is, the two models agree on their local distributions of weighted neighborhoods. Since $\mathcal{G}_*$ is a Polish space, the Prokhorov theorem ([19], page 309, Theorem 16.3) implies that the sequence μ_{n_k} is tight, that is, for every $\epsilon > 0$ there is a compact set $\mathcal{K} \subset \mathcal{G}_*$ such that $P([N_{n_k,\circ}(1)] \in \mathcal{K}) > 1 - \epsilon$ (recall that μ_{n_k} is the law of $[N_{n_k,\circ}(1)]$).

Define $\hat{\mathcal{K}}$ by appending $\{0, 1\}$ to the edges of each member of $\mathcal{K}$. Since $\mathcal{K}$ is compact, so is $\hat{\mathcal{K}}$. Moreover, since $P([N_{n_k,\circ}(1)] \in \mathcal{K}) > 1 - \epsilon$, so is $P([\hat{N}_{n_k,\circ}(1)] \in \hat{\mathcal{K}}) > 1 - \epsilon$. Hence, the sequence $\hat{\mu}_{n_k}$ is also tight, and by the Prokhorov theorem there is a further subsequence n_j such that $\hat{\mu}_{n_j} \xrightarrow{w} \hat{\mu}'_\infty$ where $\hat{\mu}'_\infty$ is some random matching on the planted PWIT. Since involution invariance passes through limit, $\hat{\mu}'_\infty$ is involution invariant. Note that we cannot invoke the Portmanteau theorem to show that $\mathbb{E}[\ell_{n_j}(1, M_{n_j,\min}(1))] \rightarrow \mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}'_\infty(\emptyset))]$, since the weight of the minimum matching is not a bounded continuous function with respect to the topology of (local) weak convergence.

By Skorokhod's theorem we can assume that $[\hat{N}_{n_j,\circ}(1)] \xrightarrow{\text{loc}} [\hat{N}'_{\infty,\circ}]$ almost surely, where $[\hat{N}'_{\infty,\circ}] \sim \hat{\mu}'_\infty$, and $\hat{N}'_\infty = (T_\infty, \ell_\infty, \mathcal{M}'_\infty)$. By the definition of local convergence,

$$\ell_{n_j}(1, M_{n_j,\min}(1)) \rightarrow \ell_\infty(\emptyset, \mathcal{M}'_\infty(\emptyset)) \quad \text{as } n \rightarrow \infty, \text{ almost surely.}$$

Using Fatou's lemma, we have

$$\liminf_{n_j \rightarrow \infty} \mathbb{E}[\ell_{n_j}(1, M_{n_j,\min}(1))] \geq \mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}'_\infty(\emptyset))].$$

The lower bound (9.1) follows by assuming n_k is a subsequence of n that achieves $\liminf_{n \rightarrow \infty} \mathbb{E}[\ell_n(1, M_{n,\min}(1))]$.

9.2. The harder half. In the “easy half” above we proved the inequality (9.1), namely that the average weight of the minimum matching on $(K_{n,n}, \ell_n)$ is bounded below by $\mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}_{\infty,\text{opt}}(\emptyset))]$. Now we are going to prove the inequality (9.2) in the opposite direction, and therefore that inequality can be replaced by equality within arbitrarily small ϵ . The key idea is to construct a low weight matching on $(K_{n,n}, \ell_n)$ using $(\ell_\infty, \mathcal{M}_{\infty,\text{opt}})$. In particular, we want the average weight of this matching to be arbitrarily close to $\mathbb{E}[\ell_\infty(\emptyset, \mathcal{M}_{\infty,\text{opt}}(\emptyset))]$ for large enough n .

Recall that the weight of planted and unplanted edges in $(K_{n,n}, \ell_n)$ are distributed as $\exp(\lambda)$ and $\exp(1/n)$ respectively, independent of everything else. Intuitively speaking, $K_{n,n}$

viewed from the planted directed edge $(1, 1')$ corresponds to the doubly-rooted planted PWIT viewed from the root’s planted edge $(\emptyset, 0)$, and $K_{n,n}$ viewed from an unplanted directed edge such as $(1, 2)$ corresponds to the doubly-rooted planted PWIT viewed from an unplanted edge incident to the root, namely $(\emptyset, i)$ where $i \in \mathbb{N}$ is arbitrary.

Now, using this “edge-centric” viewpoint, and following the approach of Aldous in [2, 6], we will assign possibly fractional values to the edges of $(K_{n,n}, \ell_n)$ such that the value assigned to the edge $e = \{i, j'\} \in E_n$ corresponds to the probability that $\mathcal{M}_{\infty,\text{opt}}(e) = 1$, assuming its local neighborhood in $(K_{n,n}, \ell_n)$ is a realization of the planted PWIT.

Consider the $n \times n$ matrix Q_n where $q_{i,j'}$ is the value assigned to the edge (i, j') . If these values corresponded exactly to probabilities, then Q_n would be doubly-stochastic. That is, Q_n would be a fractional matching, that is, an element of the matching polytope, with weight

$$\ell_n(Q_n) = \sum_{i \in [n], j \in [n']} q_{i,j'} \ell_n(i, j').$$

Since minimizing the weight is a linear programming problem, there is a vertex of this polytope, that is, an honest matching $M(\pi)_{i,j'} = 1$ if $j' = \pi(i)$ and 0 otherwise for some permutation π , whose weight is less than or equal to $\ell_n(Q_n)$. Alternatively, we could use the Birkhoff theorem to write Q as a convex combination of permutation matrices,

$$Q = \sum_{\pi} c_{\pi} M(\pi).$$

Then if we choose a random matching π with probability c_{π} , the expected weight would be $\ell_n(Q_n)$. Finally, if $(1/n)\ell_n(Q_n)$ also converges to the expected weight of $\mathcal{M}_{\infty,\text{opt}}$, we would be done.

All this is almost true. As we will see, we will define Q_n by looking at a bi-infinite version of the planted PWIT, extending it in either direction from an edge [6], Section 5.2. By looking at $(\ell_{\infty}, \mathcal{M}_{\infty,\text{opt}})$ on large neighborhoods of this edge, we will obtain probabilities that almost, but not quite, sum to 1, since the true partner of a vertex in $\mathcal{M}_{\infty,\text{opt}}$ might be outside this neighborhood. As a result, Q_n is almost doubly-stochastic in a certain sense. Following [2], we then build an almost-perfect matching with weight close to that of $\mathcal{M}_{\infty,\text{opt}}$, and then—by swapping a small fraction of edges—convert this into a perfect matching within increasing the weight very much. At this point in the proof, we can use lemmas in [2] virtually unchanged.

REMARK 1. For sake of notational simplicity, for the remainder of the section, we will drop the subscripts ∞ and $\infty\infty$. Thus objects without the subscript n live on the planted PWIT, while those with n live on the finite model.

9.2.1. *The bi-infinite planted PWIT.* As the first step toward the proof of the “harder half”, we change the viewpoint from a vertex to an edge. There are two doubly-rooted infinite versions of the planted PWIT that we need to study. One is rooted at a planted edge $(\emptyset, 0)$, and the other is rooted at an unplanted edge $(\emptyset, i)$ for some $i \in \mathbb{N}$. We illustrate these in Figure 7.

The measure on $i \in \mathbb{N}$ will be the uniform counting measure—that is, every position in the order of $\emptyset$ ’s unplanted edges has the same measure. Although this measure is not normalizable, we will speak informally of i as “uniformly chosen.” One intuition for this uniformity comes from the finite model $K_{n,n}$. If we choose uniformly from the $n(n - 1)$ unplanted edges, or (by symmetry) from the $n - 1$ unplanted edges of the root vertex 1, then its order in the sorted list of edge weights is uniform on the set $\{1, \dots, n - 1\}$.

The other intuition is as follows. The edge weight of an unplanted edge is distributed as $\exp(1/n)$, which for weights of constant size is asymptotically $1/n$ times the Lebesgue measure on $\mathbb{R}_+$. If the weight is x , the probability that the Poisson process of weight 1




FIG. 7. The structure of (a) $T_u^{\leftrightarrow}$ and (b) $T_p^{\leftrightarrow}$. The bold red edges are planted edges and the solid blue edges are unplanted.

generates $i - 1$ arrivals in the interval $[0, x]$ is $e^{-x} x^{i-1} / (i - 1)!$. The total Lebesgue measure of this event is then

$$(9.5) \quad \int_0^\infty \frac{e^{-x} x^{i-1}}{(i - 1)!} dx = \frac{\Gamma(i)}{(i - 1)!} = 1 \quad \text{for all } i,$$

so every i has equal measure.

We make this intuition rigorous below, showing how the appropriate measure on bi-infinite PWITs around both types of edge is related to the planted PWIT by extending the strategy of [6], Section 5.2, to the planted case.

In the case of the planted edge $\{\emptyset, 0\}$, the corresponding bi-infinite tree is just a relabeling of the vertices: relabel $\emptyset$ as $-\emptyset$ and relabel 0 as $+\emptyset$, and then relabel all the other vertices “below” these two roots as we did before (see Figure 7). However, for a planted edge $\{\emptyset, i\}$ for a “uniformly” selected $i \in \mathbb{N}$, things are a bit more complicated. Since $i \in \mathbb{N}$ is “uniformly” selected, the cost of the edge $\{\emptyset, i\}$ is “uniformly distributed” over $\mathbb{R}_+$, that is, Lebesgue measure on $\mathbb{R}_+$. On the other hand, as we pointed out in the proof of Corollary 8.2, the other unplanted children of $\emptyset$ are still arrivals of a Poisson process. Specifically, if we remove $\{\emptyset, i\}$, the remaining connected component of $\emptyset$ is still a planted PWIT. Hence, the corresponding bi-infinite tree is obtained by gluing two independent copies of the planted PWIT, using an edge with edge weight “distributed” as Lebesgue measure, and then relabeling the vertices: $\emptyset$ as $-\emptyset$, i as $+\emptyset$, and the others accordingly (see Figure 7). In the rest of the section, we are going to give detailed construction of the bi-infinite trees, and then we show that they are equivalent to the corresponding doubly-rooted planted PWIT.

REMARK 2. To distinguish the edge-centric viewpoint from the singly-rooted vertex viewpoint, we use the superscript $\leftrightarrow$.

Let T_u denote the planted PWIT rooted at $\emptyset$, and let T_p denote the subtree, rooted at 0 , obtained by removing the edge $\{\emptyset, 0\}$. Relabel the root of T_p to be $\emptyset$, and then relabel all the vertices of T_p using the same rule as in the planted PWIT. Let T_u^- and T_u^+ be two independent copies of T_u . Similarly, let T_p^- and T_p^+ be two independent copies of T_p . Relabel the vertices of T_u^+ and T_p^+ by adding “+” sign to the original labels, and relabel the vertices of T_u^- and T_p^- by adding “−” sign to the original labels.

Now, let $T_u^{\leftrightarrow}$ denote a bi-infinite tree, rooted at $(-\emptyset, +\emptyset)$, obtained by joining the roots of T_u^- and T_u^+ . Let $V_u^{\leftrightarrow}$ and $E_u^{\leftrightarrow}$ denote the vertices and the edges of $T_u^{\leftrightarrow}$ respectively. Let

$M_u^{\leftrightarrow}$ denote the set of planted edges of $T_u^{\leftrightarrow}$, which is the union of the planted edges in T_u^- and T_u^+ . Let $\ell_u^{\leftrightarrow}$ denote the function that assigns weight to the edges of $T_u^{\leftrightarrow}$ by using the weight of the edges in T_u^- and T_u^+ , and specifying the weight of $\{-\emptyset, +\emptyset\}$ to be uniformly “distributed” on $[0, \infty)$, that is, Lebesgue measure on $\mathbb{R}_+$, independent of everything else. Write $T_u^{\leftrightarrow} = (G_u^{\leftrightarrow}, \ell_u^{\leftrightarrow})$, where $G_u^{\leftrightarrow} = (V_u^{\leftrightarrow}, E_u^{\leftrightarrow}, M_u^{\leftrightarrow}, (-\emptyset, +\emptyset))$, and let $\mu_u^{\leftrightarrow}$ denote the σ -finite measure associated with $T_u^{\leftrightarrow}$.

Similarly, define $T_p^{\leftrightarrow}$ by joining the roots of T_p^- and T_p^+ . However, this time include $\{-\emptyset, +\emptyset\}$ as a planted edge in $M_p^{\leftrightarrow}$, and specify the weight of $\{-\emptyset, +\emptyset\}$ to be an exponentially distributed random variable with parameter λ , independent of everything else. Write $T_p^{\leftrightarrow} = (G_p^{\leftrightarrow}, \ell_p^{\leftrightarrow})$, where $G_p^{\leftrightarrow} = (V_p^{\leftrightarrow}, E_p^{\leftrightarrow}, M_p^{\leftrightarrow}, (-\emptyset, +\emptyset))$, and let $\mu_p^{\leftrightarrow}$ denote the probability distribution of $T_p^{\leftrightarrow}$. Figure 7 illustrates a realization of $T_u^{\leftrightarrow}$ and $T_p^{\leftrightarrow}$.

Recall that the doubly-rooted planted PWIT is the product measure $\mu \times \text{count}$ on $\{0, 1, 2, 3, \dots\}$, where count is the counting measure. We can think of it as a product measure on $[0, \infty)^E \times \{0, 1, 2, \dots\}$. Similarly, we can think of $\mu_u^{\leftrightarrow}$ and $\mu_p^{\leftrightarrow}$ as a σ -finite measure on $[0, \infty)^{E_u^{\leftrightarrow}}$ and a probability measure on $[0, \infty)^{E_p^{\leftrightarrow}}$, respectively. Now, depending on whether the second root is 0 or not, there is a natural map from the doubly-rooted planted PWIT to $T_p^{\leftrightarrow}$ or $T_u^{\leftrightarrow}$: (1) If the second root is 0, relabel the vertices as:

- i. relabel 0 as $+\emptyset$ and $\emptyset$ as $-\emptyset$,
- ii. relabel any sequence $0i_1i_2 \dots i_l$ for $l \geq 0$ as $+i_1i_2 \dots i_l$,
- iii. relabel any sequence $i_1i_2 \dots i_l$ for $l \geq 1$ such that $i_1 \neq 0$ as $-i_1i_2 \dots i_l$.

This relabeling induces a bijection $\psi|_p : [0, \infty)^E \times \{0\} \rightarrow [0, \infty)^{E_p^{\leftrightarrow}}$; (2) If the second root is $k \in \{1, 2, 3, \dots\}$, relabel the vertices as:

- i. relabel k as $+\emptyset$ and $\emptyset$ as $-\emptyset$,
- ii. relabel any sequence $ki_1i_2 \dots i_l$ for $l \geq 0$ as $+i_1i_2 \dots i_l$,
- iii. relabel any sequence $i_1i_2 \dots i_l$ for $l \geq 1$ such that $i_1 > k$ as $-(i_1 - 1)i_2 \dots i_l$,
- iv. relabel any sequence $i_1i_2 \dots i_l$ for $l \geq 1$ such that $i_1 < k$ as $-i_1i_2 \dots i_l$.

This relabeling induces a bijection $\psi|_u : [0, \infty)^E \times \{1, 2, 3, \dots\} \rightarrow [0, \infty)^{E_u^{\leftrightarrow}}$. Clearly, $\psi|_p$ maps $\mu \times \delta_0$ to $\mu_p^{\leftrightarrow}$, where δ_0 is the delta measure on the second root. It is also easy to see that $\psi|_u$ maps $\mu \times \text{count}$ on $\{1, 2, 3, \dots\}$ to $\mu_u^{\leftrightarrow}$ using the following lemma. Finally, note that

$$\mu \times \text{count on } \{0, 1, 2, 3, \dots\} = \mu \times \delta_0 + \mu \times \text{count on } \{1, 2, 3, \dots\}.$$

LEMMA 9.2 ([6], Lemma 25). Write $\Delta := \{(x_i) : 0 < x_1 < x_2 < \dots, x_i \rightarrow \infty\}$. Write Pois for the probability measure on Δ which is the distribution of the Poisson process of rate 1. Consider the map $\mathcal{X} : \Delta \times \{1, 2, 3, \dots\} \rightarrow \Delta \times [0, \infty)$ which takes $((x_i), k)$ to $((x_i, i \neq k), x_k)$. Then $\mathcal{X}$ maps $\text{Pois} \times \text{count}$ to $\text{Pois} \times \text{Leb}$, where Leb is the Lebesgue measure on $[0, \infty)$.

REMARK 3. For sake of brevity, we use $E^{\leftrightarrow}$ instead of $E_u^{\leftrightarrow}$ or $E_p^{\leftrightarrow}$ whenever the subscript does not affect the discussion. Similar changes applies to other symbols. More specifically, all the following arguments are true if we replace all the “.” with “ $_u$ ” or “ $_p$.”

We introduce two additional pieces of notation that we will use in the following subsections. Recall that $G^{\leftrightarrow} = (V^{\leftrightarrow}, E^{\leftrightarrow}, M^{\leftrightarrow}, (-\emptyset, +\emptyset))$ is a bi-infinite tree (without the edge weights). We define $G_{\leftarrow, B}^{\leftrightarrow}$ to be a subtree of $G^{\leftrightarrow}$ induced by $V_{\leftarrow, B}^{\leftrightarrow} := \{\pm\emptyset\} \cup \{\pm i_1i_2 \dots i_l \in V^{\leftrightarrow} : i_s \in \{0, 1, \dots, B\}\}$, that is, the subtree obtained by restricting the number of unplanted children of each vertex to B . We also define $G_{\leftarrow, B, H}^{\leftrightarrow}$ to be a subtree of $G_{\leftarrow, B}^{\leftrightarrow}$ induced by

$V_{\cdot, B, H}^{\leftrightarrow} := \{\pm \emptyset\} \cup \{\pm i_1 i_2 \cdots i_l \in V_{\cdot, B}^{\leftrightarrow} : l \leq H + 1 \text{ and } l = H + 1 \text{ iff } i_l = 0\}$, that is, the subtree obtained by restricting the depth of vertices to H or $H + 1$, depending on whether the vertex is a planted pair of its parent or not. Define $\ell_{\cdot, B}^{\leftrightarrow}$ to be the restriction of $\ell_{\cdot}^{\leftrightarrow}$ to $\ell_{\cdot, B}^{\leftrightarrow}$, and $\ell_{\cdot, B, H}^{\leftrightarrow}$ to be the restriction of $\ell_{\cdot, B}^{\leftrightarrow}$ to $\ell_{\cdot, B, H}^{\leftrightarrow}$. Now, define $T_{\cdot, B}^{\leftrightarrow} := (G_{\cdot, B}^{\leftrightarrow}, \ell_{\cdot, B}^{\leftrightarrow})$ and $T_{\cdot, B, H}^{\leftrightarrow} := (G_{\cdot, B, H}^{\leftrightarrow}, \ell_{\cdot, B, H}^{\leftrightarrow})$, and let $\mu_{\cdot, B}^{\leftrightarrow}$ and $\mu_{\cdot, B, H}^{\leftrightarrow}$ to be the associated measures. Note that there is a natural restriction $\rho_{\cdot, H}$ that maps $T_{\cdot, B}^{\leftrightarrow}$ to $T_{\cdot, B, H}^{\leftrightarrow}$.

Thus far, we show that the doubly-rooted planted PWITs and bi-infinite trees are equivalent. Next, we are going to establish the connection between the planted model $(K_{n, n}, \ell_n)$ and the bi-infinite trees.

9.2.2. The unfolding map. Now that we have discussed how to view the planted PWIT from an edge (a planted edge or a “uniformly” selected unplanted edge), we are going to discuss the similar viewpoint in $K_{n, n} = (V_n, E_n, M_n^*)$. This is done via an unfolding map that unfolds $(K_{n, n}, \ell_n)$ viewed from a planted or unplanted edge. This unfolding map is similar to the one discussed in [2], Section 3.2, with two additional properties: the number of unplanted children of every vertex is the same (for a fixed B when n is large enough), and the set of unplanted edges M_n^* is a matching on the unfolded graph. Eventually, we want to show that the local neighborhood of an edge in $K_{n, n}$ is the same as the local neighborhood of the edge in the bi-infinite trees. This should not come as a surprise, given Theorem 7.1.

Next, we are going to give detailed construction of the unfolding map. Fix $B \in \mathbb{N}_+$, and fix some edge $\{i, j'\}$ from $K_{n, n}$. We are going to unfold $(K_{n, n}, \ell_n)$ from the viewpoint of the edge $\{i, j'\}$ and construct the doubly-rooted tree $U_{B, n}^{\leftrightarrow}(i, j')$, rooted at (i, j') , with maximum $B + 1$ arity. This unfolding map resembles the exploration process discussed in the proof of Theorem 7.1, with one key difference: this map unfolds $(K_{n, n}, \ell_n)$ from an edge viewpoint.

Include the edge $\{i, j'\}$ in $U_{B, n}^{\leftrightarrow}(i, j')$. The unfolding map proceeds as follows. In this process, all vertices will be “live,” “dead,” or “neutral.” The live vertices will be contained in a queue. Initially, i and j' are live and the queue consists of only i, j' in order and all the other vertices are neutral. At each time step, a live vertex v is popped from the head of the queue. If v ’s planted neighbor $M_n^*(v)$ is neutral, then include the edge $\{v, M_n^*(v)\}$ in the doubly-rooted tree $U_{B, n}^{\leftrightarrow}(i, j')$, add $M_n^*(v)$ to the end of the queue and let $M_n^*(v)$ live. Let $v_1, v_2, \dots, v_B$ denote the B closest unplanted, neutral neighbors of vertex v in a nondecreasing order of distance to v . Include all the edges $\{v, v_k\}$ in the doubly-rooted tree $U_{B, n}^{\leftrightarrow}(i, j')$ for $1 \leq k \leq B$. Also, include all the edges $\{v_k, M_n^*(v_k)\}$ in $U_{B, n}^{\leftrightarrow}(i, j')$ for $1 \leq k \leq B$. The popped vertex v is dead. Add those neutral vertices $v_1, v_2, \dots, v_B$ to the end of the queue in order and they are live. Also, add those neutral vertices $M_n^*(v_1), M_n^*(v_2), \dots, M_n^*(v_B)$ to the end of the queue in order and they are live. The process ends when the queue is empty. According to this rule, the order of the vertices that are selected after j' is the planted neighbor i_0 of i (if $\{i, j'\}$ is not planted), $i_1, i_2, \dots, i_B, M_n^*(i_1), M_n^*(i_2), \dots, M_n^*(i_B)$, the planted neighbor j'_0 of j' (if $\{i, j'\}$ is not planted), $j'_1, j'_2, \dots, j'_B, M_n^*(j'_1), M_n^*(j'_2), \dots, M_n^*(j'_B)$, etc. This unfolding process stops when all vertices of $K_{n, n}$ are included in $U_{B, n}^{\leftrightarrow}(i, j')$. Figure 8 illustrates the $U_{1, 4}^{\leftrightarrow}(1, 2')$ and $U_{1, 4}^{\leftrightarrow}(1, 1')$ for the planted network given in Figure 5. Note that in both cases, M_n^* is a matching.

Let us define a relabeling bijection $\phi_{B, n}^{i, j'}$ (not to be confused with $\phi|_p$ and $\phi|_u$) from V_n to a subset of $+\Sigma \cup -\Sigma$, where $+\Sigma := \{+i : i \in \Sigma\}$ and $-\Sigma := \{-i : i \in \Sigma\}$ denote the set of vertex labelings of the bi-infinite planted PWIT. Define $\phi_{B, n}^{i, j'}(i) = -\emptyset$ and $\phi_{B, n}^{i, j'}(j') = +\emptyset$. At each step of the unfolding process, when we pick vertex v and add $\{v, v_k\}$ and $\{v_k, M_n^*(v_k)\}$ to $U_{B, n}^{\leftrightarrow}(i, j')$, set $\phi_{B, n}^{i, j'}(v_k) = \phi_{B, n}^{i, j'}(v)k$ and $\phi_{B, n}^{i, j'}(M_n^*(v_k)) = \phi_{B, n}^{i, j'}(v_k)0$.

Here, using the bijection $\phi_{B, n}^{i, j'}$, we first map $U_{B, n}^{\leftrightarrow}(i, j')$ to a subtree of $T_{\cdot, B}^{\leftrightarrow}$ and then apply the restriction map $\rho_{\cdot, H}$ on this subtree. For ease of notation, we abbreviate $\rho_{\cdot, H}(\phi_{B, n}^{i, j'}(U_{B, n}^{\leftrightarrow}(i, j')))$ as $\rho_{\cdot, H}(U_{B, n}^{\leftrightarrow}(i, j'))$ whenever the context is clear. It is easy




FIG. 8. The doubly-rooted trees (a) $U_{1,4}^{\leftrightarrow}(1, 2')$ and (b) $U_{1,4}^{\leftrightarrow}(1, 1')$ obtained by unfolding the planted network in Figure 5.

to see that the graph structure of $\rho_{u,H}(U_{B,n}^{\leftrightarrow}(i, j'))$ for $j' \neq i'$ and all sufficiently large n is isomorphic to $G_{u,B,H}^{\leftrightarrow}$, and the graph structure of $\rho_{p,H}(U_{B,n}^{\leftrightarrow}(i, i'))$ for all sufficiently large n is isomorphic to $G_{p,B,H}^{\leftrightarrow}$. Let $\mu_{n,u,B,H}^{\leftrightarrow}$ denote the probability measure associated with $\rho_{u,H}(U_{B,n}^{\leftrightarrow}(i, j'))$ for $j' \neq i'$, and let $\mu_{n,p,B,H}^{\leftrightarrow}$ denote the probability measure associated with $\rho_{p,H}(U_{B,n}^{\leftrightarrow}(i, i'))$. Note that if $D_{u,x} = \{T_{u,B,H}^{\leftrightarrow} : \ell_{u,B,H}^{\leftrightarrow}(\{-\emptyset, \emptyset\}) < x\}$, then $\mu_{n,u,B,H}^{\leftrightarrow}(D_{u,x}) = 1 - \exp(-x/n) \approx x/n$ as $n \rightarrow \infty$. Also, note that $\mu_{u,B,H}^{\leftrightarrow}(D_{u,x}) = x$. We now generalize Lemma 10 in [2] to the planted case.

LEMMA 9.3. For any fixed $B, H \in \mathbb{N}_+$, and $x > 0$ we have

$$\begin{aligned} n\mu_{n,u,B,H}^{\leftrightarrow}(D_{u,x} \cap \cdot) &\xrightarrow{\text{TV}} \mu_{u,B,H}^{\leftrightarrow}(D_{u,x} \cap \cdot), \\ \mu_{n,p,B,H}^{\leftrightarrow} &\xrightarrow{\text{TV}} \mu_{p,B,H}^{\leftrightarrow}, \end{aligned}$$

where the total variation convergence of positive measures is defined as:

$$\mu_n \xrightarrow{\text{TV}} \mu \quad \text{iff} \quad \sup_A |\mu_n(A) - \mu(A)| \rightarrow 0.$$

PROOF. The proof is almost identical to the proof of Lemma B.1, and has been omitted. The only subtle difference is the restriction to $D_{u,x}$, which causes no problem since $n \times \frac{1}{n}e^{-x/n} \rightarrow 1$ as $n \rightarrow \infty$. $\square$

Next, we are going to use the involution invariant random matching $(\ell, \mathcal{M}_{\text{opt}})$ on the planted PWIT to assign values to the edges of $(K_{n,n}, \ell_n)$. Ideally, the value assigned to an edge $e = \{i, j'\}$ corresponds to the probability of e being in the matching that we want to construct.

9.2.3. Assigning values to the edges of $(K_{n,n}, \ell_n)$. Now it's time to assign fractional values to the edges of $(K_{n,n}, \ell_n)$. This is done by pretending that the local neighborhood of an edge e in $(K_{n,n}, \ell_n)$, is a realization of the local neighborhood of the corresponding bi-infinite tree. So, as the first step toward assigning values to the edges of $(K_{n,n}, \ell_n)$, we need to know how to assign value to $\{-\emptyset, +\emptyset\}$ in the bi-infinite tree using the minimum matching on the planted PWIT. The idea is to use the inverse image of $\psi|_{\cdot}$ and map the edge $\{-\emptyset, +\emptyset\}$ to the corresponding edge on the planted PWIT. This gives us a function $g_{\cdot} : [0, \infty)^{E^{\leftrightarrow}} \rightarrow [0, 1]$, which we then use to assign fractional values to the edges of $(K_{n,n}, \ell_n)$, by conditioning on

its neighborhood. We are going to discuss these steps in detail. We follow the discussion in [6], Section 5.5, and then [2], Introduction.

Recall that $\tilde{\mu} = \mu \times \text{count}$ on $\{0, 1, 2, 3, \dots\}$ is the measure associated with the doubly-rooted planted PWIT. Also, recall that since μ is involution invariant, the product measure $\tilde{\mu}$ is invariant under the involution map ι that swaps the roots. Following the discussion of [6], Section 5.5, we define the function $\gamma_{\text{opt}} : [0, \infty)^E \times \{0, 1, 2, \dots\} \rightarrow [0, 1]$ by

$$\gamma_{\text{opt}}(\mathbf{w}, i) = \mathbb{P}[\mathcal{M}_{\text{opt}}(\emptyset) = i \mid \ell(e) = w(e) \forall e \in E].$$

As Aldous points out, the function γ_{opt} satisfies certain consistency properties:

- (i) $\sum_{i=0}^{\infty} \gamma_{\text{opt}}(\mathbf{w}, i) = 1$ since $\mathcal{M}_{\text{opt}}$ is a matching.
- (ii) $\gamma_{\text{opt}}(\iota(\mathbf{w}, i)) = \gamma_{\text{opt}}(\mathbf{w}, i)$ since the random matching $(\ell, \mathcal{M}_{\text{opt}})$ is involution invariance, where $\iota(\mathbf{w}, i)$ swaps i and $\emptyset$ given $\{\ell(e) = w(e) \forall e \in E\}$.

Also, we have (iii) $\mathbb{E}[\ell(\mathcal{M}_{\text{opt}}(\emptyset), \emptyset)] = \int_{\mathbf{w}} \sum_{i=0}^{\infty} \gamma_{\text{opt}}(\mathbf{w}, i) w(\emptyset, i) \mu(d\mathbf{w})$. Now, define $g : [0, \infty)^{E^{\leftrightarrow}} \rightarrow [0, 1]$ as follows:

$$g(\mathbf{w}^{\leftrightarrow}) = \gamma_{\text{opt}}(\psi| \cdot^{-1}(\mathbf{w}^{\leftrightarrow})),$$

where $\psi| \cdot$ is the bijection map that we defined earlier in Section 9.2.1, and $\psi| \cdot^{-1}$ is its inverse. We can think of the function g as the probability that the edge $\{-\emptyset, +\emptyset\}$ is in the matching. The function g satisfies similar consistency properties corresponding to (i) and (ii):

- (i') The function g assigns honest probabilities to the neighbors of $-\emptyset$ as well as $+\emptyset$. To be more specific, let $\iota_{p,i}^+, \iota_{p,i}^- : [0, \infty)^{E_p^{\leftrightarrow}} \rightarrow [0, \infty)^{E_u^{\leftrightarrow}}$ denote the root swapping maps, that change the root from $(-\emptyset, +\emptyset)$ to $(-\emptyset, -i)$, and from $(-\emptyset, +\emptyset)$ to $(+i, +\emptyset)$ respectively (and then relabel all the vertices). Similarly, define $\iota_{u,i}^+$ and $\iota_{u,i}^-$. We have:

- (a) In $T_p^{\leftrightarrow}$, the values assigned to the neighbors of $-\emptyset$ as well as $+\emptyset$ sums to one:

$$g_p(\mathbf{w}_p^{\leftrightarrow}) + \sum_{i=1}^{\infty} g_u(\iota_{p,i}^{\pm}(\mathbf{w}_p^{\leftrightarrow})) = 1.$$

- (b) In $T_u^{\leftrightarrow}$, the values assigned to the neighbors of $-\emptyset$ as well as $+\emptyset$ sums to one:

$$g_u(\mathbf{w}_u^{\leftrightarrow}) + g_p(\iota_{u,0}^{\pm}(\mathbf{w}_u^{\leftrightarrow})) + \sum_{i=1}^{\infty} g_u(\iota_{u,i}^{\pm}(\mathbf{w}_u^{\leftrightarrow})) = 1.$$

- (ii') Let $\iota^{\leftrightarrow}$ denote the root swamping map, that swaps the root $(-\emptyset, +\emptyset)$ to $(+\emptyset, -\emptyset)$ (and then relabels all the vertices). We have $g(\mathbf{w}^{\leftrightarrow}) = g(\iota^{\leftrightarrow}(\mathbf{w}^{\leftrightarrow}))$.

Also note that the function g is measurable with respect to the product σ -algebra on $[0, \infty)^{E^{\leftrightarrow}}$. Hence, (iii) becomes (iii')

$$\begin{aligned} \mathbb{E}[\ell(\mathcal{M}_{\text{opt}}(\emptyset), \emptyset)] &= \int_{\mathbf{w}_p^{\leftrightarrow}} \mathbf{w}_p^{\leftrightarrow}(-\emptyset, +\emptyset) g_p(\mathbf{w}_p^{\leftrightarrow}) \mu_p^{\leftrightarrow}(d\mathbf{w}_p^{\leftrightarrow}) \\ &\quad + \int_{\mathbf{w}_u^{\leftrightarrow}} \mathbf{w}_u^{\leftrightarrow}(-\emptyset, +\emptyset) g_u(\mathbf{w}_u^{\leftrightarrow}) \mu_u^{\leftrightarrow}(d\mathbf{w}_u^{\leftrightarrow}). \end{aligned}$$

Now we use the function g to assign fractional values to the edges of the planted model. We store the values assigned to the edges of $(K_{n,n}, \ell_n)$ in an $n \times n$ matrix $\mathcal{Q}_n = [q_{i,j'}]_{i,j'}$, where $q_{i,j'}$ is the value assigned to the edge $\{i, j'\}$. Let $g_{\cdot, B, H}$ to be the conditional expectation of g given $\sigma(\ell^{\leftrightarrow}(e), e \in E_{\cdot, B, H}^{\leftrightarrow})$, with respect to the measure $\mu_{\cdot}^{\leftrightarrow}$. Also, let $g_{u,x,B,H} = g_{u,B,H} \mathbf{1}_{D_{u,x}}$. Now, $q_{i,j'}$ is defined as follows:

$$(9.6) \quad q_{i,j'} := \begin{cases} g_{u,x,B,H}(\rho_{u,H}(U_{B,n}^{\leftrightarrow}(i, j'))) & \text{if } j' \neq i', \\ g_{p,B,H}(\rho_{p,H}(U_{B,n}^{\leftrightarrow}(i, i'))) & \text{if } j' = i'. \end{cases}$$

As we mentioned before, if Q_n were doubly-stochastic, we could have used it to construct a matching on $K_{n,n}$. However, it is not hard to see that $\mathbb{P}[Q_n \text{ is doubly-stochastic}] = 0$. Nevertheless, we expect Q_n to become almost doubly-stochastic, as a larger neighborhood is revealed.

Let us define a discrimination factor, as Aldous does in [2], Introduction:

$$\mathcal{X}(Q_n) = \frac{1}{n} \sum_{i=1}^n \left| 1 - \sum_{j'=1'}^{n'} q_{i,j'} \right| + \frac{1}{n} \sum_{j'=1'}^{n'} \left| 1 - \sum_{i=1}^n q_{i,j'} \right|.$$

Note that if Q_n is doubly-stochastic, then $\mathcal{X}(Q_n) = 0$. Naturally, we should expect that $\mathbb{E}[\mathcal{X}(Q_n)] \approx 0$, for large values x , B , H , and n . We should also expect the average expected cost of Q_n to be close to the expected cost of $\mathcal{M}_{\text{opt}}$, that is,

$$\frac{1}{n} \mathbb{E} \left[\sum_{i,j'} q_{i,j'} \ell_n(i, j') \right] \approx \mathbb{E}[\ell(\mathcal{M}_{\text{opt}}(\phi), \phi)],$$

for large enough values of x , B , H , and n . Using the same set of inequalities as in [2], Section 3.4, it follows that the both intuitions are correct.

LEMMA 9.4.

(i) *For any $\epsilon > 0$, there is an x_0 , B_0 , H_0 and n_0 such that for all $x > x_0$, $B > B_0$, $H > H_0$, and $n > n_0$ we have $\mathbb{E}[\mathcal{X}(Q_n)] < \epsilon$. In the other words,*

$$\lim_{x \rightarrow \infty} \limsup_{B \rightarrow \infty} \limsup_{H \rightarrow \infty} \limsup_{n \rightarrow \infty} \mathbb{E}[\mathcal{X}(Q_n)] = 0.$$

(ii) *For any $\delta > 0$, there is an x_0 and n_0 such that for all $x > x_0$ and $n > n_0$, the average expected cost of Q_n is in the δ neighborhood of the cost of $\mathcal{M}_{\text{opt}}$ on the planted PWIT, for all B and H . In the other words,*

$$\lim_{x \rightarrow \infty} \limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\sum_{i,j'} q_{i,j'} \ell_n(i, j') \right] = \mathbb{E}[\ell(\mathcal{M}_{\text{opt}}(\phi), \phi)] \quad \text{for all } B \text{ and } H.$$

PROOF. The proof is almost identical to the proof presented in [2], Section 3.4. The proof of part (ii) is a direct consequence of Lemma 9.3 and linearity of expectation. The proof of part (i) needs more work, but the inequalities are the same as the one presented in [2], Section 3.4. The key factor is the consistency properties of functions g_u and g_p , Lemma 9.3, and the fact that there is no short cycle containing the root consists entirely of low-weight edges (as mentioned in the proof of Theorem 7.1). $\square$

Now that we know Q_n eventually becomes a doubly-stochastic matrix with weight close to that of $\mathcal{M}_{\text{opt}}$, we will construct a perfect matching on $(K_{n,n}, \ell_n)$ with (within ϵ) the same weight. First, by invoking [2], Proposition 7, we construct a partial matching with the cost close to the cost of Q_n . Next, using [2], Proposition 9, we construct a perfect matching by swapping operation while keeping the cost almost the same. The changes required to extend this analysis to the planted case are very minor, but we present the strategy in the next section for completeness.

9.2.4. *Construction of the matching.* Finally, we are going to give the precise construction of the low cost matching on $(K_{n,n}, \ell_n)$ and prove (9.2). By Lemma 9.4, we know that Q_n becomes arbitrary close to a doubly-stochastic matrix, that is, $\mathcal{X}(Q_n)$ becomes arbitrary close to 0 with high probability. Now, using Q_n for sufficiently large n , we construct a low

cost partial matching which matches most of the vertices and leaves a small fraction of vertices isolated. This is done by invoking Proposition 7 in [2]. Then we use Proposition 9 in [2] and swap some of the edges to obtain a perfect matching on $(K_{n,n}, \ell_n)$ while keeping the cost almost the same. On the rare occasions that we fail to either construct the partial matching or swap the edges, we use M_n^* . Since these events are rare, the use of planted matching won't affect the total cost. We follow the discussion in [2], Section 2.

Let us begin with describing how to construct a partial matching using an almost doubly-stochastic matrix. We say Q_n is an almost doubly-stochastic matrix if its discrimination factor is close enough to 0, or more precisely, if $\mathcal{X}(Q_n) < 1/200$. We say $v_n : E_n \rightarrow \{0, 1, \emptyset\}$ is a $(1 - \theta)$ partial matching if vertices in $U(v_n) = \{i \in [n] : v_n(i, j') = 1 \text{ for some } j' \in [n']\}$ are matched to different vertices in $[n']$ and $|U(v_n)| \geq (1 - \theta)n$.

The first step is to convert Q_n to a doubly-stochastic matrix. Define an $n \times n$ matrix $A_n = [a_{i,j'}]$ as follows:

$$a_{i,j'} := \frac{q_{i,j'}}{\max(1, q_{i,:}) \max(1, q_{:,j'})},$$

where $q_{i,:} := \sum_{j' \in [n']} q_{i,j'}$ and $q_{:,j'} := \sum_{i \in [n]} q_{i,j'}$. Similarly, define $a_{i,:}$ and $a_{:,j'}$ for all $i \in [n]$ and $j' \in [n']$. Note that $a_{i,:} \leq 1$ and $a_{:,j'} \leq 1$. Define an $n \times n$ matrix $B_n = [b_{i,j'}]$ as follows:

$$b_{i,j'} = \frac{(1 - a_{i,:})(1 - a_{:,j'})}{n - \sum_{i \in [n]} \sum_{j' \in [n']} a_{i,j'}}.$$

It is easy to check that $A_n + B_n$ is doubly-stochastic. Hence, by Birkhoff–von Neumann theorem $A_n + B_n$ can be written as a convex combination of permutations. Hence, there exists a random matching $\mathcal{M}_n$ on (K_n, ℓ_n) such that $\forall i \in [n], j' \in [n'] : P(\mathcal{M}_n(i) = j') = a_{i,j'} + b_{i,j'}$. Note that for all $i \in [n]$ and $j' \in [n']$ we have $a_{i,j'} \leq q_{i,j'}$, however, there is no such a bound for $b_{i,j'}$. As a result, we may end up assigning high probabilities to undesired edges which can affect the expected cost of the matching $\mathcal{M}_n$. Now, the idea is to use some part of the matching $\mathcal{M}_n$ that is behaving well enough.

PROPOSITION 9.5 ([2], Proposition 7). *Let $Q_n = [q_{i,j'}]$ and $\ell_n = [\ell_n(i, j')]$ be given nonrandom $n \times n$ matrices. Suppose $200\mathcal{X}(Q_n) \leq \theta < 1$. Consider the random matching $\mathcal{M}_n$ given as above and define the random set $D(\mathcal{M}_n) := \{i \in [n] : b_{i,\mathcal{M}_n(i)} \leq \eta a_{i,\mathcal{M}_n(i)}\}$ where $\eta = \sqrt{3\mathcal{X}(Q_n)/\theta}$. Then*

$$\mathbb{E} \left[\sum_{i \in D(\mathcal{M}_n)} q_{i,\mathcal{M}_n(i)} \ell_n(i, \mathcal{M}_n(i)) \right] \leq (1 + \eta) \sum_{i \in [n]} \sum_{j' \in [n']} q_{i,j'} \ell_n(i, j'),$$

$$\mathbb{P}[|D(\mathcal{M}_n)| \geq (1 - \theta)n] \geq 1 - \frac{3(1 + \eta^{-1})\mathcal{X}(Q_n)}{\theta},$$

that is, the random matching $\mathcal{M}_n$ restricted to the random set $D(\mathcal{M}_n)$ is a $(1 - \theta)$ matching with high probability with cost close to the cost of Q_n . Specifically, there is a $(1 - \theta)$ partial matching v_n (nonrandom) such that

$$\sum_{i \in U(v_n)} q_{i,v_n(i)} \ell_n(i, v_n(i)) \leq (1 + 4\sqrt{\mathcal{X}(Q_n)/\theta}) \sum_{i \in [n]} \sum_{j' \in [n']} q_{i,j'} \ell_n(i, j').$$

Next, we patch the partial matching v_n given by the Proposition 9.5 (which exists for almost all realization of the edge weights and Q_n), to construct a perfect matching without distorting the expected cost and the given partial matching too much. Aldous suggests using the greedy algorithm to do that. In the planted setting, the idea is to simply remove all the planted edges and then use the greedy algorithm.

PROPOSITION 9.6 ([2], Proposition 9). *Fix $0 < \theta < 1/10$, and let $k = \lfloor \theta n \rfloor$. Let $\ell_n = [\ell_n(i, j')]$ denote the matrix of the edge weights. Let v_n denote a $1 - \theta$ partial matching. Let $N_0 = [n] \setminus U(v_n)$ and $N'_0 = [n'] \setminus v_n(U(v_n))$ denote the set of unpaired vertices in either side. (Here, N_0 , N'_0 and v_n may depend on ℓ_n in an arbitrary way.) Then there exists a random subset $S \subset [n] \setminus N_0$ of size k , random bijections $v_1 : S \rightarrow N'_0$, $v_2 : N_0 \rightarrow v_n(S)$, and events Ω_n with $P(\Omega_n) \rightarrow 0$ such that*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\mathbf{1}_{\Omega_n^c} \left(\sum_{i \in N_0} \ell_n(i, v_2(i)) + \sum_{i \in S} \ell_n(i, v_1(i)) \right) \right] \leq 24\theta^{1/2}.$$

It remains to combine Proposition 9.5 and Proposition 9.6 to construct a matching on $K_{n,n}$. The key idea is to rewrite the edge weights of $K_{n,n}$ as the minimum of two independent exponential random variables.

LEMMA 9.7. *If $X \sim \exp(\mu_1)$ and $Y \sim \exp(\mu_2)$ are independent, then $\min(X, Y) \sim \exp(\mu_1 + \mu_2)$.*

Now, for any fixed $0 < \alpha < 1$, we can write

$$(9.7) \quad \ell_n(i, j') = \min \left(\frac{\ell_n^1(i, j')}{1 - \alpha}, \frac{\ell_n^2(i, j')}{\alpha} \right),$$

where $\ell_n^1(i, j')$ and $\ell_n^2(i, j')$ are independent copies of $\ell_n(i, j')$. We use $\ell_n^1(i, j')$ to construct the partial matching, and then $\ell_n^2(i, j')$ to patch the partial matching and construct a complete matching. On the event Ω_n^* that the construction is not possible, that is either Proposition 9.5 or Proposition 9.6 failed, we can always use the planted matching to construct a matching on $K_{n,n}$. Since the probability of this failure goes to zero, this does not affect the cost of the matching that much. Specifically, we need to show that $L_n := \frac{1}{n} \sum_{i \in [n]} \ell_n(i, i')$ is uniformly integrable.

LEMMA 9.8. *There exists a function $\delta(\cdot)$ with $\delta(x) \rightarrow 0$ as $x \rightarrow 0$ such that for arbitrary events Ω_n^* ,*

$$(9.8) \quad \limsup_{n \rightarrow \infty} \mathbb{E}[L_n \mathbf{1}_{\Omega_n^*}] \leq \delta(\epsilon),$$

where $\epsilon = \limsup_{n \rightarrow \infty} \mathbb{P}[\Omega_n^*]$.

PROOF. Since $\mathbb{E}[(L_n)^2] < \infty$, L_n is uniformly integrable. In particular, by Cauchy–Schwarz inequality, we have $\mathbb{E}[L_n \mathbf{1}_{\Omega_n^*}] \leq \sqrt{\mathbb{E}[(L_n)^2] \mathbb{P}[\Omega_n^*]} = \sqrt{\mathbb{E}[(L_n)^2]} \epsilon$. $\square$

Now we are ready to present the formal proof of (9.2), which closely follows the proof of Proposition 2 in [2]. Fix $0 < \theta < 1$ and let $\epsilon = \theta^3/200$. Construct $Q_n = [q_{i,j'}]_{i,j'}$ as per (9.6) using the edge cost ℓ_n^1 . It follows from Lemma 9.4 that for all sufficiently large x , B , H , n , $\mathbb{E}[\mathcal{X}(Q_n)] < \epsilon$ and

$$(9.9) \quad \frac{1}{n} \mathbb{E} \left[\sum_{i,j'} q_{i,j'} \ell_n^1(i, j') \right] \leq \mathbb{E}[\ell(\mathcal{M}_{\text{opt}}(\emptyset), \emptyset)] + \epsilon.$$

Define event $\Omega_n^1 = \{200\mathcal{X}(Q_n) > \theta^2\}$. Then by Markov’s inequality,

$$\mathbb{P}[\Omega_n^1] = \mathbb{P}[200\mathcal{X}(Q_n) > \theta^2] \leq \frac{200\mathbb{E}[\mathcal{X}(Q_n)]}{\theta^2} \leq \frac{200\epsilon}{\theta^2} = \theta.$$

Outside event Ω_n^1 , we have $200\mathcal{X}(Q_n) \leq \theta^2 < \theta$. It follows from Proposition 9.5 that outside event Ω_n^1 , there exists a $(1 - \theta)$ partial matching v_n such that

$$(9.10) \quad \sum_{i \in U(v_n)} q_{i, v_n(i)} \ell_n^1(i, v_n(i)) \leq (1 + \theta) \sum_{i \in [n]} \sum_{j' \in [n']} q_{i, j'} \ell_n^1(i, j'),$$

where $U(v_n)$ is the set of vertices i matched under v_n . Now, condition on a realization (outside of Ω_n^1) of ℓ_n^1 , and apply Proposition 9.6 to ℓ_n^2 . Since ℓ_n^1 and ℓ_n^2 are independent, it follows from Proposition 9.6 that there exist an event Ω_n^2 and bijections v_1 and v_2 such that for all sufficiently large n , $\mathbb{P}[\Omega_n^2] \leq \theta$ and

$$(9.11) \quad \frac{1}{n} \mathbb{E} \left[1_{(\Omega_n^1)^c} \left(\sum_{i \in [n] \setminus U(v_n)} \ell_n^2(i, v_2(i)) + \sum_{i \in v_n^{-1}(v_2([n] \setminus U(v_n)))} \ell_n^2(i, v_1(i)) \right) \right] \leq 24\theta^{1/2}.$$

Outside event $\Omega_n^* = \Omega_n^1 \cup \Omega_n^2$, we can construct a complete matching $\pi_n : [n] \rightarrow [n']$ such that $\pi_n(i) = v_2(i)$ if $i \in [n] \setminus U(v_n)$; and $\pi_n(i) = v_1(i)$ if $i \in v_n^{-1}(v_2([n] \setminus U(v_n)))$; and $\pi_n(i) = v_n(i)$ otherwise. On event Ω_n^* , we just let π_n to be the planted matching. Combining (9.7), (9.9), (9.10), (9.11), and (9.8) yields that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\sum_{i, j'} \pi_n(i, j') \ell_n(i, j') \right] \leq \frac{(1 + \theta)(\mathbb{E}[\ell(\mathcal{M}_{\text{opt}}(\emptyset), \emptyset)] + \epsilon)}{1 - \alpha} + \frac{24\theta^{1/2}}{\alpha} + \delta(2\theta).$$

Letting θ (and hence ϵ) $\rightarrow 0$, then letting $\alpha \rightarrow 0$, we establish (9.2).

APPENDIX A: PROOF OF LEMMA 2.2

The moment generating function for an exponential random variable Y with rate λ is

$$\mathbb{E}[e^{\mu Y}] = \frac{\lambda}{\lambda - \mu}.$$

Since X_1 and X_2 are independent, the exponential generating function for $X_1 - X_2$ is

$$(A.1) \quad \mathbb{E}[e^{\mu(X_1 - X_2)}] = \mathbb{E}[e^{\mu X_1}] \mathbb{E}[e^{-\mu X_2}] = \left(\frac{\lambda_1 \lambda_2}{(\lambda_1 - \mu)(\lambda_2 + \mu)} \right)^t.$$

By Markov's inequality

$$\mathbb{P}[X_1 > X_2] = \mathbb{P}[e^{\mu(X_1 - X_2)} > 1] \leq \mathbb{E}[e^{\mu(X_1 - X_2)}]$$

for any $\mu > 0$. The right-hand side of (A.1) is minimized when $\mu = (\lambda_1 - \lambda_2)/2$, giving the desired result.

APPENDIX B: PROOF OF THEOREM 7.1

Step 1: The exploration process

As we pointed out in the sketch of the proof, the first step is to define an exploration process that explores vertices of $N_{n, \circ}(1)$ in a series of stages. The stage m of the exploration process reveals a rooted subtree of $(K_{n, n}, \ell_n)$, denoted by $N_{n, \circ}[m]$. The root of $N_{n, \circ}[m]$ is vertex 1, the number of unplanted children of every vertex is m , and the set of planted edges restricted to $N_{n, \circ}[m]$ is a matching. Next, we provide a formal construction of $N_{n, \circ}[m]$.

The construction begins with vertex 1. Include the edge $\{1, M_n^*(1) = 1'\}$ in $N_{n, \circ}[m]$. Let $\{v_1, v_2, \dots, v_m\}$ denote the m closest unplanted neighbors of vertex 1. Add all the edges $\{1, v_k\}$ and then $\{v_k, M_n^*(v_k)\}$ to $N_{n, \circ}[m]$. Next, continue with the vertex $1'$. Let $\{w_1, w_2, \dots, w_m\}$ denote the m closest unplanted neighbors of vertex $1'$, among all the vertices that has not been added to $N_{n, \circ}[m]$. Include all the edges $\{1', w_k\}$ and then $\{w_k, M_n^*(w_k)\}$




FIG. 9. Stage 2 of the exploration process on $(K_{4,4}, \ell_4)$ given by Figure 5.

to $N_{n,o}[m]$. At each step of the construction, we follow two simple rules: (1) the next vertex to pick is the oldest one in $N_{n,o}[m]$; (2) when we add m closest unplanted children of this vertex and their planted pairs, we avoid all the vertices that has already been added to $N_{n,o}[m]$.

The construction continues until we are about to pick a vertex at depth m , at which point it stops. Note that the only vertices at depth $m + 1$ are the planted partners of the vertices at depth m . Let $V_n^{(m)}$ and $E_n^{(m)}$ denote the set of the vertices and the edges of $N_{n,o}[m]$ respectively. Note that for all sufficiently large n , $|V_n^{(m)}|$ and $|E_n^{(m)}|$ are independent of n . Let $\mu_n^{(m)} \in \mathcal{P}(\mathcal{G}_*)$ denote the law of $[N_{n,o}[m]]$. Figure 9 demonstrates the construction of $N_{n,o}[2]$ for the graph given by Figure 5.

Step 2: A total variation convergence

The rooted planted tree $N_{n,o}[m]$ has the same graph structure as a truncated version of the planted PWIT: remove all vertices $\mathbf{i} = (i_1, i_2, \dots, i_l)$ such that either (1) $i_s > m$ for some $s \in \{1, 2, \dots, l\}$, or (2) $l > m + 1$, or (3) $l = m + 1$ and $i_l \neq 0$. In particular, the number of unplanted children of every vertex in the truncated version is m , the depth of vertices are bounded by $m + 1$, and the only vertices at depth $m + 1$ are the planted pairs of the vertices at depth m . Let $N_\infty[m]$ denote the truncated planted PWIT, and let $\mu_\infty^{(m)}$ denote the law of $[N_\infty[m]]$. Now, using the same approach as in [2], Lemma 10, we show that $\mu_n^{(m)}$ converges to $\mu_\infty^{(m)}$ in total variation norm.

LEMMA B.1. For any fixed m , $\mu_n^{(m)} \xrightarrow{\text{TV}} \mu_\infty^{(m)}$ where the total variation convergence of positive measures is defined as follows:

$$\mu_n \xrightarrow{\text{TV}} \mu \quad \text{iff} \quad \sup_A |\mu_n(A) - \mu(A)| \rightarrow 0.$$

PROOF. It is easy to see that $\mu_n^{(m)}$ is absolutely continuous with respect to $\mu_\infty^{(m)}$. Moreover, the Radon-Nikodym derivative of $\mu_n^{(m)}$ with respect to $\mu_\infty^{(m)}$ equals the likelihood ratio.

Consider similar steps on the planted PWIT to construct $N_\infty[m]$. Conditioned on the first $t - 1$ steps of the construction of $N_\infty[m]$ and $N_{n,o}[m]$, we will calculate the ratio of the conditional densities for the next step of the construction. Since planted edges have the same $\exp(\lambda)$ distribution in both cases, we are only interested in the corresponding ratio of unplanted edges.

At t th step of the construction of $N_\infty[m]$, the conditional density of $(x_1, x_1 + x_2, \dots, x_1 + \dots + x_m)$ is $\exp(-(x_1 + x_2 + \dots + x_m))$. At t th step of the construction of $N_{n,o}[m]$, using the

memoryless property of exponential random variables, the conditional density of $(x_1, x_1 + x_2, \dots, x_1 + \dots + x_m)$ is

$$\prod_{i=0}^{m-1} \frac{|I_{t-1}| - i}{n} \exp\left(-\frac{x_{i+1}(|I_{t-1}| - i)}{n}\right),$$

where I_{t-1} is the set of vertices that has not been added to $N_\infty[m]$ yet up to the t th step. Hence, the ratio of conditional densities is at least

$$\prod_{i=0}^{m-1} \frac{|I_{k-1}| - i}{n} \geq \left(\frac{n - |V^{(m)}|}{n}\right)^m,$$

and we have

$$\frac{d\mu_n^{(m)}}{d\mu_\infty^{(m)}} \geq \left(1 - \frac{|V^{(m)}|}{n}\right)^{|E^{(m)}| - |V^{(m)}|/2},$$

where the exponent $|E^{(m)}| - |V^{(m)}|/2$ is the number of unplanted edges explored.

Note that $|V^{(m)}|$ and $|E^{(m)}|$ do not depend on n , for all sufficiently large n . Hence, as $n \rightarrow \infty$ the right-hand side of the above inequality goes to 1. Now, the result follows by the fact that $\mu_n^{(m)}$ and $\mu_\infty^{(m)}$ are probability measures. $\square$

Step 3: Locally tree-like property

Fix some $\rho > 0$. Recall that $(G_{n,\circ}(1))_\rho$ denotes the ρ -neighborhood of vertex 1 in $N_{n,\circ}(1)$ as is defined in Section 6. Similarly, $(G_{n,\circ}[m])_\rho$ denotes the neighborhood ρ of node 1 in $N_{n,\circ}[m]$. The question is, whether these two neighborhoods are the same. Note that $(G_{n,\circ}[m])_\rho$ is a tree but $(G_{n,\circ}(1))_\rho$ is not necessary. However, it becomes a tree with high probability.

LEMMA B.2. *Fix $\epsilon > 0$ and $\rho > 0$. Then there exists large enough $m_0(\epsilon, \rho)$ such that for all fixed $m > m_0(\epsilon, \rho)$,*

$$\mathbb{P}[(G_{n,\circ}(1))_\rho \neq (G_{n,\circ}[m])_\rho] \leq \epsilon \quad \text{as } n \rightarrow \infty.$$

PROOF. Let m_0 to be large enough such that

$$\mathbb{P}[\text{number of vertices in } (G_{n,\circ}(1))_\rho > m_0] < \epsilon/2.$$

Fix $m > m_0$. Consider the event $\{(G_{n,\circ}(1))_\rho \neq (G_{n,\circ}[m])_\rho\}$. This event may happens if either the number of vertices in $(G_{n,\circ}(1))_\rho$ is greater than m or there are two vertices $v, w \in V_n^{(m)}$ such that $\{v, w\} \notin E_n^{(m)}$ but $\ell_n(v, w) \leq \rho$. The probability of the first event is bounded by $\epsilon/2$. For the other event, note that if $v, w \in V_n^{(m)}$ and $\{v, w\} \in E_n \setminus E_n^{(m)}$, then $\ell_n(v, w)$ is dominated by an exponentially distributed random variable with mean n . (To see this, assume that v was revealed earlier than w . Then we know that the cost of the edge $\{v, w\}$ is larger than the cost c_m of the edge $\{v, v_m\}$, where v_m is the m th closest unplanted neighbor of v . Hence, the probability distribution of the weight of $\{v, w\}$, using the memoryless property, is $1/n \exp((x - c_m)/n)$ which is stochastically larger than a random variable distributed as $\exp(1/n)$. Hence,

$$\begin{aligned} & \mathbb{P}[\exists v, w \in V_n^{(m)} \text{ such that } \{v, w\} \notin E_n^{(m)} \text{ and } \ell_n(v, w) \leq \rho] \\ & \leq \binom{|V_n^{(m)}|}{2} (1 - \exp(-\rho/n)) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \end{aligned}$$

As we mentioned before, $|V_n^{(m)}|$ is independent of n for all sufficiently large n . Now, the result follows by combining the last two displayed inequalities. $\square$

Now, combining Lemma B.1 and Lemma B.2, we get the following corollary.

COROLLARY B.3. *Fix $\rho > 0$. Let $\mu_{n,\rho}$ denote the law of $[(G_{n,\circ}(1))_\rho, \ell_n]$, and let $\mu_{\infty,\rho}$ denote the law of $[(G_\infty)_\rho, \ell_\infty]$. Then $\mu_{n,\rho} \xrightarrow{\text{TV}} \mu_{\infty,\rho}$.*

Step 4: Portmanteau theorem

For a fixed $R > 0$, since the condition $d([N_\circ], [T_\circ]) < (R + 1)^{-1}$ is equivalent to $d([(G_\circ)_R, \ell], [T_\circ]) < (R + 1)^{-1}$, Corollary B.3 implies that, for all finite rooted planted trees $[T_\circ] \in \mathcal{G}_*$, we have

$$(B.1) \quad |\mu_n(A_{T_\circ}) - \mu_\infty(A_{T_\circ})| \rightarrow 0 \quad \text{as } n \rightarrow \infty,$$

where $A_{T_\circ}$ is defined as

$$A_{T_\circ} := \{[N_\circ] \in \mathcal{G}_* : d([N_\circ], [T_\circ]) < (R + 1)^{-1}\}.$$

Note that the support of μ_∞ is rooted planted trees. Moreover, recall that $\mathcal{G}_*$ is separable, hence, the restriction of $\mathcal{G}_*$ to the rooted planted trees is also separable. Since μ_∞ is a probability measure, for any $R > 0$ and any $\epsilon > 0$, there exists a finite set $S(R, \epsilon)$ consisting of rooted planted trees $T_\circ = (G_\circ, \ell)$ with $(G_\circ)_R = G_\circ$ such that

$$\mu_\infty\left(\bigcup_{T_\circ \in S(R, \epsilon)} A_{T_\circ}\right) > 1 - \epsilon.$$

Using Corollary B.3, there exists $n_0(\epsilon) \in \mathbb{N}_+$ such that for all $n > n_0(\epsilon)$,

$$\mu_n\left(\bigcup_{T_\circ \in S(R, \epsilon)} A_{T_\circ}\right) > 1 - 2\epsilon.$$

Now we are going to prove that $\mu_n \xrightarrow{w} \mu_\infty$. By definition $\mu_n \xrightarrow{w} \mu_\infty$, if for any continuous bounded function $f : \mathcal{G}_* \rightarrow \mathbb{R}$,

$$\int_{\mathcal{G}_*} f d\mu_n \rightarrow \int_{\mathcal{G}_*} f d\mu_\infty.$$

Using the Portmanteau theorem, we can restrict our attention to the uniformly continuous bounded functions. Let $f : \mathcal{G}_* \rightarrow \mathbb{R}$ be a uniformly continuous bounded function. Now, for any $\epsilon > 0$, there is a $\delta > 0$ such that if $d([N_\circ], [N'_\circ]) < \delta$ then $|f([N_\circ]) - f([N'_\circ])| \leq \epsilon$. Fix the value of ϵ and let R to be large enough such that $(R + 1)^{-1} < \delta$. We have

$$\left| \int_{\mathcal{G}_*} f d\mu_n - \int_{\mathcal{G}_*} f d\mu_\infty \right| \leq 3\epsilon |f|_\infty + \sum_{T_\circ \in S(R, \epsilon)} f(T_\circ) |\mu_n(A_{T_\circ}) - \mu_\infty(A_{T_\circ})| + 2\epsilon,$$

where $|f|_\infty := \sup_{N \in \mathcal{G}_*} |f(N)|$. The result follows by arbitrary choice of ϵ , the fact that $|S(R, \epsilon)| < \infty$, and (B.1).

APPENDIX C: PROOF OF PROPOSITION 8.4

By Lemma 8.3, $\mathcal{M}_{\infty, \text{opt}}$ is a deterministic function of the collection of random variables

$$\mathcal{C}_\infty = \{\ell_\infty(e), X(\overleftrightarrow{e}); e \in E_\infty, \text{ and } \overleftrightarrow{e} \text{ is directed}\}.$$

Also note that $X(\overleftarrow{e})$ satisfies (8.6), which does not depend on the relabeling of the vertices. Now, by construction, we only need to show that the distribution of $(X(\overrightarrow{e}); e \in \overrightarrow{E}_\infty(h))$ for $h \geq 1$ is invariant with respect to the involution map ι .

Abusing the notation, let $\hat{\mu}_{\infty, \text{opt}}$ denote the law of $\mathcal{C}_\infty$ and define $\tilde{\mu}_{\infty, \text{opt}}$ to be $\hat{\mu}_{\infty, \text{opt}} \times \text{count on } \mathbb{N}_+$, similar to (6.6). Fix the second root $k \geq 0$ and let

$$(C.1) \quad B = \{(X(\overrightarrow{e}), \overrightarrow{e} \in \overrightarrow{E}_\infty(h)) \in \cdot, k \text{ is distinguished}\}$$

denote a measurable subset on $\hat{\mathcal{G}}_{**}$, where $\hat{\mathcal{G}}_{**}$ is defined similar to $\mathcal{G}_{**}$ (the set all isomorphism classes of connected locally finite doubly-rooted planted networks) with an additional mark on the edges $\overleftarrow{e}$ representing $X(\overleftarrow{e})$. Let (X_0, Y_0) denote a solution of the system of recursive distributional equations (8.4)–(8.5) as in Lemma 8.3. Recall that if $\overrightarrow{e}$ is a planted edge then $X(\overrightarrow{e})$ and X_0 have the same distribution; otherwise, $X(\overrightarrow{e})$ and Y_0 have the same distribution. Note that the collection of random variables $(X(\overrightarrow{e}), \overrightarrow{e} \in \overrightarrow{E}_\infty(h))$ are independent since they depend on messages received from disjoint subtrees.

We need to show that

$$\tilde{\mu}_{\infty, \text{opt}}(\iota^{-1}(B)) = \tilde{\mu}_{\infty, \text{opt}}(B).$$

We treat the cases $k = 0$ or $k > 0$ separately.

(i) If $k = 0$, then we have

$$\iota^{-1}(B) = \{\text{the double root is } (\emptyset, 0)\} \cap \{(X(\overrightarrow{e}), \overrightarrow{e} \in \overrightarrow{E}_\infty(h, 0)) \in \cdot\},$$

where $\overrightarrow{E}_\infty(h, 0) := \{(v, vj) : \{v, vj\} \in E_\infty \text{ s.t. } v = i_1 i_2 i_3 \cdots i_{h-2} \text{ with } i_1 \neq 0 \text{ or } v = 0 i_2 i_3 \cdots i_h\}$. Now to complete the proof, it suffices to show that $(X(\overrightarrow{e}), \overrightarrow{e} \in \overrightarrow{E}_\infty(h, 0))$ has the same distribution as $(X(\overrightarrow{e}), \overrightarrow{e} \in \overrightarrow{E}_\infty(h))$. Clearly, the collection of random variables $(X(\overrightarrow{e}), \overrightarrow{e} \in \overrightarrow{E}_\infty(h, 0))$ are independent, and $X(\overrightarrow{e})$ has the same distribution as X_0 or Y_0 depending on whether $\overrightarrow{e}$ is planted or unplanted. Thus it remains to prove that there is a one-to-one map from $\overrightarrow{E}_\infty(h, 0)$ to $\overrightarrow{E}_\infty(h)$ that maps (un)planted edges to (un)planted ones. Consider the relabeling function ϕ defined as follows:

$$\begin{aligned} \phi(0 i_2 i_3 \cdots i_h) &= i_2 i_3 \cdots i_h \quad \forall 0 i_2 i_3 \cdots i_h \in V_\infty, \\ \phi(i_1 i_2 \cdots i_{h-2}) &= 0 i_1 i_2 \cdots i_{h-2} \quad \forall i_1 i_2 \cdots i_{h-2} \text{ with } i_1 \neq 0, \end{aligned}$$

and define $\gamma : \overrightarrow{E}_\infty(h, 0) \rightarrow \overrightarrow{E}_\infty(h)$ by $\gamma(\{v, vj\}) := \{\phi(v), \phi(v)j\}$.

(ii) If $k > 0$, then we have

$$\iota^{-1}(B) = \bigcup_{l>0} \{\{\text{the double root is } (\emptyset, l)\} \cap A_l \cap \{(X(\overrightarrow{e}), \overrightarrow{e} \in \overrightarrow{E}_\infty(h, l)) \in \cdot\}\},$$

where

$$A_l := \{\ell_\infty(l, l(k-1)) < \ell_\infty(\emptyset, l) < \ell_\infty(l, lk)\},$$

and $\overrightarrow{E}_\infty(h, l)$ is defined similar to $\overrightarrow{E}_\infty(h, 0)$. Note that the events $\{(X(\overrightarrow{e}), \overrightarrow{e} \in \overrightarrow{E}_\infty(h, l)) \in \cdot\}$ and A_l are independent, and the distribution of $(X(\overrightarrow{e}), \overrightarrow{e} \in \overrightarrow{E}_\infty(h, l))$ does not depend on l . Also, note that by (9.5), we have

$$\sum_{l>0} \mathbb{P}[A_l] = \int_0^\infty \mathbb{P}[\text{exactly } k-1 \text{ arrivals before } x] dx = 1.$$

The result then follows using the same argument as in the previous case.

APPENDIX D: PROOF OF LEMMA 8.8

There are two branches that we are interested in: the alternating path from $\emptyset$ through v_{-1} , and the alternating path from $\emptyset$ through v_1 . It is more convenient to study these two branches on the doubly-rooted planted PWIT, rooted at $(\emptyset, v_1)$. The proof uses the discussion of bi-infinite planted PWITs $T_u^{\leftrightarrow}$ and $T_p^{\leftrightarrow}$ in Section 9.2.1. We follow the same notation and simplification (Remark 1 in Section 9.2.1) here.

Using the relabeling maps $\psi|_p$ and $\psi|_u$, we already know that $\mu_\infty \times \delta_0$ is equivalent to $\mu_p^{\leftrightarrow}$ and $\mu_\infty \times \text{count on } \{1, 2, 3, \dots\}$ is equivalent to $\mu_u^{\leftrightarrow}$. We can use the relabeling map $\psi|$ to define $\{X(\vec{e}), \vec{e} \text{ is a directed edge in } E^{\leftrightarrow}\}$ jointly with $\{\ell^{\leftrightarrow}(e), e \in E^{\leftrightarrow}\}$. Note that the joint distribution of $\{\ell^{\leftrightarrow}(e), X(\vec{e}); e \in E^{\leftrightarrow} \text{ and } \vec{e} \text{ is directed}\}$ is exactly the same as if we use the construction of Lemma 8.1 by redefining $\vec{E}(h)$ as

$$(D.1) \quad \begin{aligned} \vec{bE}^{\leftrightarrow}(h) := & \{ \vec{e} = (-v, -vj) : \text{gen}(v) = h - 1 \} \\ & \cup \{ \vec{e} = (+v, +vj) : \text{gen}(v) = h - 1 \}. \end{aligned}$$

Now, given $\{X(\vec{e}), \vec{e} \text{ is a directed edge in } E^{\leftrightarrow}\}$ we can define a minimum matching $\mathcal{M}_{\cdot, \text{opt}}^{\leftrightarrow}$ on $T^{\leftrightarrow}$, same as in Lemma 8.3, that is,

$$(D.2) \quad \forall v \in V^{\leftrightarrow} : \mathcal{M}_{\cdot, \text{opt}}^{\leftrightarrow}(v) = \arg \min_{w: \{v, w\} \in E^{\leftrightarrow}} (\ell^{\leftrightarrow}(v, w) - X(v, w)),$$

$$(D.3) \quad \forall e \in E^{\leftrightarrow} : \mathcal{M}_{\cdot, \text{opt}}^{\leftrightarrow}(e) = 1 \quad \text{if and only if} \quad \ell^{\leftrightarrow}(e) < X(\vec{e}) + X(\overleftarrow{e}).$$

Next, we are going to show that the bi-infinite tree $T_u^{\leftrightarrow}$ ($T_p^{\leftrightarrow}$) restricted to $\mathcal{M}_{u, \text{opt}}^{\leftrightarrow}(-\emptyset, +\emptyset) = 1$ ($\mathcal{M}_{p, \text{opt}}^{\leftrightarrow}(-\emptyset, +\emptyset) = 1$) is equivalent to the doubly-rooted planted PWIT, rooted at $\{\emptyset, \mathcal{M}_{\text{opt}}(\emptyset)\}$, restricted to $\mathcal{M}_{\text{opt}}(\emptyset) \neq 0$ ($\mathcal{M}_{\text{opt}}(\emptyset) = 0$).

On the planted PWIT T_u (recall that $T_u = T(\emptyset)$ as defined in Section 9.2.1), define $X_u^\downarrow = \min_{i \geq 0} (\ell(\emptyset, i) - X(\emptyset, i))$. Let $v_u(x)$ denote the conditional distribution of the set

$$\{\ell(e), X(\vec{e}); e \in E_h \text{ and } \vec{e} \text{ is directed away from } \emptyset\}$$

given $X_u^\downarrow = x$. Similarly, on the subtree T_p (recall that T_p is a relabeling of $T(0)$ as defined in Section 9.2.1), define $X_p^\downarrow = \min_{i \geq 1} (\ell(\emptyset, i) - X(\emptyset, i))$. Let $v_p(x)$ denote the conditional distribution of the set

$$\{\ell(e), X(\vec{e}); e \in E_p \text{ and } \vec{e} \text{ is directed away from } \emptyset\}$$

given $X_p^\downarrow = x$. On the bi-infinite tree $T^{\leftrightarrow}$, define μ^1 to be the measure obtained by restricting $\mu^{\leftrightarrow}$ to the set $\{\ell^{\leftrightarrow}(-\emptyset, +\emptyset) < X(-\emptyset, +\emptyset) + X(+\emptyset, -\emptyset)\}$, that is, $\mathcal{M}_{\cdot, \text{opt}}^{\leftrightarrow}(-\emptyset, +\emptyset) = 1$. Let $E^{\leftrightarrow+}$ and $E^{\leftrightarrow-}$ denote all edges of form $\{+v, +vj\}$ and $\{-v, -vj\}$ respectively. Clearly, $E^{\leftrightarrow} = E^{\leftrightarrow+} \cup E^{\leftrightarrow-} \cup \{-\emptyset, +\emptyset\}$. Let (X_0, Y_0) be a solution of the system of recursive distributional equations (8.4)–(8.5).

LEMMA D.1. *The measures μ_u^1 and μ_p^1 are finite positive measures. The total mass of μ_u^1 equals $\mathbb{P}[\mathcal{M}_{\text{opt}}(\emptyset) \neq 0]$ and the total mass of μ_p^1 is $\mathbb{P}[\mathcal{M}_{\text{opt}}(\emptyset) = 0]$. Under μ^1 we have:*

(i) *The joint density of $(\ell_u^{\leftrightarrow}(-\emptyset, +\emptyset), X_u(+\emptyset, -\emptyset), X_u(-\emptyset, +\emptyset))$ at point (l, x_1, x_2) is $f_u(x_1)f_u(x_2)\mathbf{1}_{(0 < l < x_1 + x_2)}$, where $f_u(\cdot)$ is the density of Y_0 and $\mathbf{1}_{(0 < l < x_1 + x_2)}$ is the indicator function; the joint density of $(\ell_p^{\leftrightarrow}(-\emptyset, +\emptyset), X_p(+\emptyset, -\emptyset), X_p(-\emptyset, +\emptyset))$ at point (l, x_1, x_2) is $f_p(x_1)f_p(x_2)\lambda \exp(-\lambda l)$, where $f_p(\cdot)$ is the density of X_0 .*

(ii) Conditioned on $(\ell^{\leftrightarrow}(-\phi, +\phi), X.(+\phi, -\phi), X.(-\phi, +\phi)) = (l, x_1, x_2)$ with $x_1 + x_2 > l$, the distribution of the family

$$\{\ell^{\leftrightarrow}(e), X.(\vec{e}); e \in E^{\leftrightarrow-} \text{ and } \vec{e} \text{ is directed away from } -\phi\}$$

is the image of $v.(x_1)$ under the natural embedding $T. \rightarrow T.^- \subset T.^{\leftrightarrow}$; the distribution of the family

$$\{\ell^{\leftrightarrow}(e), X.(\vec{e}); e \in E^{\leftrightarrow+} \text{ and } \vec{e} \text{ is directed away from } +\phi\}$$

is the image of $v.(x_2)$ under the natural embedding $T. \rightarrow T.^+ \subset T.^{\leftrightarrow}$ and these two families are conditionally independent.

REMARK 4. Conditioned on $(\ell^{\leftrightarrow}(-\phi, +\phi), X.(+\phi, -\phi), X.(-\phi, +\phi)) = (l, x_1, x_2)$ with $x_1 + x_2 > l$, we have $\mathcal{M}_{\cdot, \text{opt}}^{\leftrightarrow}(-\phi, +\phi) = 1$. Now, by (D.2) and the construction of $X.$ on $T.^{\leftrightarrow}$, we have $X.(-i, -\phi) = \ell^{\leftrightarrow}(-\phi, +\phi) - X.(-\phi, +\phi)$ for all i . Similarly, $X.(+i, +\phi) = \ell^{\leftrightarrow}(+\phi, -\phi) - X.(+\phi, -\phi)$ for all i . This combined with the families in part (ii) of Lemma D.1, specifies $X.$ on $T.^{\leftrightarrow}$ under μ^1 .

PROOF. By construction of $X.$ on $T.^{\leftrightarrow}$, we already know that $X.(+\phi, -\phi)$ and $X.(-\phi, +\phi)$ are independent with density $f(\cdot)$. Moreover, $\ell_u^{\leftrightarrow}(-\phi, +\phi)$ has uniform “distribution” on $[0, \infty)$, and $\ell_p^{\leftrightarrow}(-\phi, +\phi)$ is an exponentially distributed random variable with parameter λ . Hence, the joint density has the form mentioned in (i). Moreover, the total mass of μ_u^1 is

$$(D.4) \quad \int_{x_1=-\infty}^{\infty} \int_{x_2=-\infty}^{\infty} (x_1 + x_2)^+ f_u(x_1) f_u(x_2) dx_2 dx_1,$$

and the total mass of μ_p^1 is

$$(D.5) \quad \int_{x_1=-\infty}^{\infty} \int_{x_2=-\infty}^{\infty} (1 - \exp(-\lambda(x_1 + x_2)^+)) f_p(x_1) f_p(x_2) dx_2 dx_1,$$

where $(x_1 + x_2)^+ = \max(x_1 + x_2, 0)$. Now, using the joint density above, we can calculate the total mass of μ_u^1 and μ_p^1 as shown by the following lemma.

LEMMA D.2. The equation (D.4) equals $\mathbb{P}[\mathcal{M}_{\text{opt}}(\phi) \neq 0]$, and the equation (D.5) equals $\mathbb{P}[\mathcal{M}_{\text{opt}}(\phi) = 0]$.

PROOF OF LEMMA D.2. Let X_1 and X_2 denote two independent copies of X_0 , and Y_1 and Y_2 denote two independent copies of Y_0 . Let η denote an exponentially distributed random variable with parameter λ .

Using Corollary 8.2, we have

$$\begin{aligned} \text{Equation (D.4)} &= \mathbb{E}[(Y_1 + Y_2)^+] \\ &= \int_{x=0}^{\infty} \mathbb{P}[Y_1 + Y_2 > x] dx \\ &= \int_{\zeta=0}^{\infty} \mathbb{P}[X(\vec{e}) + X(\overleftarrow{e}) > \ell(e) \mid \exists e = \{\phi, i \geq 1\}, \ell(e) = \zeta] d\zeta \\ &= \mathbb{P}[\mathcal{M}_{\text{opt}}(\phi) \neq 0], \end{aligned}$$

$$\begin{aligned} \text{Equation (D.5)} &= \mathbb{P}[X_1 + X_2 > \eta] = \mathbb{P}[X(\phi, 0) + X(0, \phi) > \ell(\phi, 0)] \\ &= \mathbb{P}[\mathcal{M}_{\text{opt}}(\phi) = 0]. \end{aligned}$$

□

Next, by construction based on $(X.(\overrightarrow{e}), \overrightarrow{e} \in \overrightarrow{bE}^{\leftrightarrow}(h))$ as in Lemma 8.1, under $\mu^{\leftrightarrow}$ the families

$$\{\ell^{\leftrightarrow}(e), X.(\overrightarrow{e}); e \in E^{\leftrightarrow-} \text{ and } \overrightarrow{e} \text{ is directed away from } -\emptyset\} \cup X.(+\emptyset, -\emptyset),$$

and

$$\{\ell^{\leftrightarrow}(e), X.(\overrightarrow{e}); e \in E^{\leftrightarrow+} \text{ and } \overrightarrow{e} \text{ is directed away from } +\emptyset\} \cup X.(-\emptyset, +\emptyset),$$

are independent of each other and $\ell^{\leftrightarrow}(-\emptyset, +\emptyset)$. Therefore, the desired conditional independence in part (ii) follows, when conditioned on $(\ell^{\leftrightarrow}(-\emptyset, +\emptyset), X.(+\emptyset, -\emptyset), X.(-\emptyset, +\emptyset)) = (l, x_1, x_2)$.

Finally, note that each families under $\mu^{\leftrightarrow}$ is distributed as the image of corresponding family on T , where $X.(+\emptyset, -\emptyset)$ (or $X.(-\emptyset, +\emptyset)$) corresponds to $X^\downarrow$. Now, the independence of these two families under $\mu^{\leftrightarrow}$ implies that the conditional distribution of families under μ^1 depends only on the corresponding value of $X^\downarrow$, that is, x_1 for the first family and x_2 for the second one. $\square$

Recall that $\psi|_p : [0, \infty)^E \times \{0\} \rightarrow [0, \infty)^{E_p^{\leftrightarrow}}$ maps $\mu \times \delta_0$ to $\mu_p^{\leftrightarrow}$, and $\psi|_u : [0, \infty)^E \times \{1, 2, 3, \dots\} \rightarrow [0, \infty)^{E_u^{\leftrightarrow}}$ maps $\mu \times \text{count}$ on $\{1, 2, 3, \dots\}$ to $\mu_u^{\leftrightarrow}$. Note that the inverse image of the event $\{\ell^{\leftrightarrow}(-\emptyset, +\emptyset) < X.(+\emptyset, -\emptyset) + X.(-\emptyset, +\emptyset)\}$ under $\psi|$ is the event

$$\{(\emptyset, \mathcal{M}_{\text{opt}}(\emptyset)) \text{ is the double root}\}.$$

Hence, $\psi|_p^{-1}$ maps the measure μ_p^1 to $\mu \times \delta_0$ restricted to $\{\text{the second root is } \mathcal{M}_{\text{opt}}(\emptyset) = 0\}$, and $\psi|_u^{-1}$ maps the measure μ_u^1 to $\mu \times \text{count}$ on $\{1, 2, 3, \dots\}$ restricted to

$$\{\text{the second root is } \mathcal{M}_{\text{opt}}(\emptyset) \text{ and } \mathcal{M}_{\text{opt}}(\emptyset) \neq 0\}.$$

Hence, to study the events $\overline{B}_\infty$ and B_{-1} on the planted PWIT, we can relabel the vertices by setting $\emptyset$ to be $-\emptyset$, $\mathcal{M}_{\text{opt}}(\emptyset)$ to be $+\emptyset$, map the doubly-rooted planted PWIT, rooted at $(\emptyset, \mathcal{M}_{\text{opt}}(\emptyset))$, to the corresponding bi-infinite tree $(T_u^{\leftrightarrow}$ or $T_p^{\leftrightarrow}$ depending on whether $\mathcal{M}_{\text{opt}}(\emptyset) = 0$ or not), and study the image of these events under μ_p^1 and μ_u^1 .

D.1. A lemma on T_u and T_p . Before analyzing the image of B_{-1} on $T_u^{\leftrightarrow}$ (or $T_p^{\leftrightarrow}$ depending on whether $\mathcal{M}_{\text{opt}}(\emptyset) = 0$ or not), let us present a technical lemma which generalizes Lemma 23 in [2] to the planted case.

LEMMA D.3. *On the planted PWIT, define*

$$\begin{aligned} X_u^\downarrow &= \min_{i \geq 0}(\ell(\emptyset, i) - X(\emptyset, i)), \\ X_p^\downarrow &= \min_{i \geq 1}(\ell(\emptyset, i) - X(\emptyset, i)), \\ I_u &= \arg \min_{i \geq 0}(\ell(\emptyset, i) - X(\emptyset, i)), \\ I_p &= \arg \min_{i \geq 1}(\ell(\emptyset, i) - X(\emptyset, i)). \end{aligned}$$

For $-\infty < b < a < \infty$ define

$$\begin{aligned} g_u(a, b) &:= \mathbb{P}\Big[\ell(\emptyset, I_u) - b > \min_{j: \{I_u, I_u j\} \in E}^{[2]} (\ell(I_u, I_u j) - X(I_u, I_u j)) \mid X_u^\downarrow = a\Big], \\ g_p(a, b) &:= \mathbb{P}\Big[\ell(\emptyset, I_p) - b > \min_{j: \{I_p, I_p j\} \in E}^{[2]} (\ell(I_p, I_p j) - X(I_p, I_p j)) \mid X_p^\downarrow = a\Big]. \end{aligned}$$

Then $g_u(a, b), g_p(a, b) > 0$.

PROOF. Let us begin with an observation, which is the continuous analogous of the splitting property of a Poisson process.

OBSERVATION. Let $\{X_i\}_{i=1}^\infty$ be independent real-valued continuous random variables with common distribution α_X . Let $\{\zeta_i\}_{i=1}^\infty$ denote the arrivals of a Poisson process with parameter 1. Then $\{(\zeta_i, X_i)\}_{i=1}^\infty$ forms a Poisson point process on $[0, \infty) \times (-\infty, +\infty)$ with mean intensity $\beta(z, x) dz dx = dz \alpha_X(dx)$.

Now, let $Y_i = \zeta_i - X_i$. The set of points $\{Y_i\}_{i=1}^\infty$ forms a certain inhomogeneous Poisson process on $(-\infty, \infty)$ with mean intensity $\gamma(y) dy = \alpha_X([-y, \infty)) dy$. Finally, it is easy to see that conditioned on the time of the first arrival to be y_0 , the other points in $\{Y_i\}_{i=1}^\infty$ are the points of a certain inhomogeneous Poisson process on (y_0, ∞) . Similar statement holds if we condition on no arrival before y_0 .

Note that by (8.6) we have

$$X(\emptyset, I) = \min_{j: \{I, I.j\} \in E} (\ell(I, I.j) - X(I, I.j)).$$

Now, by the above observation, the set of points $\{\ell(I, I.j) - X(I, I.j), j \geq 1\}$ conditioned on $X(\emptyset, I) = x$, are the points of a certain inhomogeneous Poisson process on (x, ∞) (note that the claim is true regardless of whether $\arg \min_{j: \{I, I.j\} \in E} (\ell(I, I.j) - X(I, I.j)) = 0$ or not). Hence,

$$\mathbb{P}\left[\min_{j: \{I, I.j\} \in E}^{[2]} (\ell(I, I.j) - X(I, I.j)) \in [y, y + dy] \mid X(\emptyset, I) = x\right] \geq \tilde{\beta}_x(y) dy,$$

where $\tilde{\beta}_x(y) > 0$ for all $y > x$. Since the above term does not depend on the value of $\ell(\emptyset, I)$, we have

$$\tilde{g} \cdot (a, b, x)$$

$$\begin{aligned} &:= \mathbb{P}\left[\ell(\emptyset, I) - b > \min_{j: \{I, I.j\} \in E}^{[2]} (\ell(I, I.j) - X(I, I.j)) \mid X(\emptyset, I) = x, \ell(\emptyset, I) = a + x\right] \\ &> 0, \end{aligned}$$

for all $-\infty < b < a < \infty$, and $-\infty < x < \infty$. Now, the result follows by

$$g \cdot (a, b) = \mathbb{E}[\tilde{g} \cdot (a, b, X(\emptyset, I)) \mid X^\downarrow = a],$$

since $X^\downarrow = \ell(\emptyset, I) - X(\emptyset, I)$. $\square$

D.2. Calculating with the bi-infinite tree. On the bi-infinite tree $T_{\cdot}^{\leftrightarrow}$, define the event $C_{\cdot, -1}$ as

$$C_{\cdot, -1} := \left\{ -\emptyset = \arg \min_{y: \{y, -I\} \in E_{\cdot}^{\leftrightarrow}}^{[2]} (\ell_{\cdot}^{\leftrightarrow}(-I, y) - X_{\cdot}(-I, y)) \right\},$$

where

$$-I = \arg \min_{-i: \{-\emptyset, -i\} \in E_{\cdot}^{\leftrightarrow}} (\ell_{\cdot}^{\leftrightarrow}(-\emptyset, -i) - X_{\cdot}(-\emptyset, -i)).$$

The event $C_{\cdot, -1}$ under μ^1 corresponds to the event B_{-1} on the doubly-rooted planted PWIT, rooted at $(\emptyset, v_1)$, where $v_1 = \mathcal{M}_{\text{opt}}(\emptyset)$. Define the following σ -algebras:

$$\mathcal{F}_{\cdot}^{-} = \sigma(X_{\cdot}(\vec{e}), \ell_{\cdot}^{\leftrightarrow}(e) : e \in E_{\cdot}^{\leftrightarrow-} \text{ and } \vec{e} \text{ is directed}),$$

$$\mathcal{F}_{\cdot}^{+} = \sigma(X_{\cdot}(\vec{e}), \ell_{\cdot}^{\leftrightarrow}(e) : e \in E_{\cdot}^{\leftrightarrow+} \text{ and } \vec{e} \text{ is directed}),$$

$$\mathcal{F}_{\cdot}^{\emptyset} = \sigma(\ell_{\cdot}^{\leftrightarrow}(-\emptyset, +\emptyset), X_{\cdot}(+\emptyset, -\emptyset), X_{\cdot}(-\emptyset, +\emptyset)).$$

LEMMA D.4. $\mu^1(C_{\cdot,-1}^c | \mathcal{F}^+, \mathcal{F}^\emptyset) = g.(X.(+\emptyset, -\emptyset), \ell^{\leftrightarrow}(-\emptyset, +\emptyset) - X.(-\emptyset, +\emptyset))$, where $C_{\cdot,-1}^c$ is the complement of the event $C_{\cdot,-1}$.

PROOF. By Remark 4 right after Lemma D.1,

$$\mathcal{F}^- \cap \mathcal{F}^\emptyset = \mathcal{F}^\emptyset \cap \sigma(X.(\vec{e}), \ell^{\leftrightarrow}(e) : e \in E^{\leftrightarrow-} \text{ and } \vec{e} \text{ is directed away from } -\emptyset),$$

$$\mathcal{F}^+ \cap \mathcal{F}^\emptyset = \mathcal{F}^\emptyset \cap \sigma(X.(\vec{e}), \ell^{\leftrightarrow}(e) : e \in E^{\leftrightarrow+} \text{ and } \vec{e} \text{ is directed away from } +\emptyset).$$

Now, since $C_{\cdot,-1}$ is $\mathcal{F}^-$ measurable, by conditional independence of Lemma D.1 part (ii) we have

$$\mu^1(C_{\cdot,-1}^c | \mathcal{F}^+, \mathcal{F}^\emptyset) = \mu^1(C_{\cdot,-1}^c | \mathcal{F}^\emptyset).$$

Hence, we need to show that for all (l, x_1, x_2) ,

$$\mu^1\{C_{\cdot,-1}^c | (\ell^{\leftrightarrow}(-\emptyset, +\emptyset), X.(+\emptyset, -\emptyset), X.(-\emptyset, +\emptyset)) = (l, x_1, x_2)\} = g.(x_1, l - x_2).$$

By Lemma D.1, conditioned on $(\ell^{\leftrightarrow}(-\emptyset, +\emptyset), X.(+\emptyset, -\emptyset), X.(-\emptyset, +\emptyset)) = (l, x_1, x_2)$, the distribution of the family

$$\{\ell^{\leftrightarrow}(e), X.(\vec{e}); e \in E^{\leftrightarrow-} \text{ and } \vec{e} \text{ is directed away from } -\emptyset\},$$

is the image of $v.(x_1)$ under the natural embedding $T. \rightarrow T_-^- \subset T_-^{\leftrightarrow}$. Recall that $v.(x_1)$ is the distribution of $\{\ell(e), X(\vec{e}); e \in E. \text{ and } \vec{e} \text{ is directed away from } \emptyset\}$ on $T.$, given $X^\downarrow = x_1$. Hence,

$$\begin{aligned} & g.(x_1, l - x_2) \\ &= \mu^1\left\{\ell^{\leftrightarrow}(-\emptyset, -I.) - (l - x_2) > \min_{j: \{-I., -I.j\} \in E^{\leftrightarrow-}}^{[2]} (\ell^{\leftrightarrow}(-I., -I.j) - X.(-I., -I.j)) \mid \right. \\ & \quad \left. (\ell^{\leftrightarrow}(-\emptyset, +\emptyset), X.(+\emptyset, -\emptyset), X.(-\emptyset, +\emptyset)) = (l, x_1, x_2)\right\}. \end{aligned}$$

However, under this conditioning

$$\begin{aligned} \ell^{\leftrightarrow}(-\emptyset, -I.) - (l - x_2) &= \ell^{\leftrightarrow}(-\emptyset, -I.) - (\ell^{\leftrightarrow}(-\emptyset, +\emptyset) - X.(-\emptyset, +\emptyset)) \\ &= \ell^{\leftrightarrow}(-I., -\emptyset) - X.(-I., -\emptyset), \end{aligned}$$

where the last equality follows by (8.6) and the fact that under μ^1 , $\mathcal{M}_{\cdot, \text{opt}}^{\leftrightarrow}(-\emptyset, \emptyset) = 1$ (see Remark 4). Finally, note that under μ^1 ,

$$C_{\cdot,-1}^c = \left\{ \ell^{\leftrightarrow}(-I., -\emptyset) - X.(-I., -\emptyset) > \min_{j: \{-I., -I.j\} \in E^{\leftrightarrow-}}^{[2]} (\ell^{\leftrightarrow}(-I., -I.j) - X.(-I., -I.j)) \right\}.$$

□

Now we have all the machinery to finish the proof of Lemma 8.8. By using the relabeling bijections, $\mathbb{P}[B_{-1} | \bar{B}_\infty, \{\mathcal{M}_{\text{opt}}(\emptyset) = 0\}]$ equals $\mu_p^1(C_{p,-1} | C_p)$ where C_p is a certain event which is measurable with respect to $\mathcal{F}_p^+ \cap \mathcal{F}_p^\emptyset$ such that $\mu_p^1(C_p) = \mathbb{P}[\bar{B}_\infty \cap \{\mathcal{M}_{\text{opt}}(\emptyset) = 0\}]$. Similarly, $\mathbb{P}[B_{-1} | \bar{B}_\infty, \{\mathcal{M}_{\text{opt}}(\emptyset) \neq 0\}]$ equals $\mu_u^1(C_{u,-1} | C_u)$ for a certain event C_u that is defined similar to C_p . Now, by Lemma D.3 and Lemma D.4, if $\mu^1(C.) > 0$, then we have

$$\begin{aligned} \mu^1(C_{\cdot,-1}^c \cap C.) &= \mathbb{E}_{\mu^1}[\mathbf{1}_{C.} \mu^1(C_{\cdot,-1}^c | \mathcal{F}^+, \mathcal{F}^\emptyset)] \\ &= \mathbb{E}_{\mu^1}[\mathbf{1}_{C.} g.(X.(+\emptyset, -\emptyset), \ell^{\leftrightarrow}(-\emptyset, +\emptyset) - X.(-\emptyset, +\emptyset))] > 0. \end{aligned}$$

Acknowledgments. We are very grateful to Venkat Anantharam, Charles Bordenave, Jian Ding, David Gamarnik, Christopher Jones, Vijay Subramanian, Yihong Wu and Lenka Zdeborová for helpful conversations. C.M. is also grateful to Microsoft Research New England for their hospitality. We also thank an anonymous reviewer for helpful comments.

Funding. The first author was supported by the Rackham Predoctoral Fellowship, a departmental Graduate Student Instructor appointment and by NSF Grant AST-1443972/AST-1516075. The majority of the work was done while the first author was at the University of Michigan.

The second author was supported in part by NSF Grant IIS-1838251.

The third author was supported by NSF Grants IIS-1838124, CCF-1850743, and CCF-1856424.

REFERENCES

- [1] ABBE, E. (2017). Community detection and stochastic block models: Recent developments. *J. Mach. Learn. Res.* **18** Paper No. 177, 86. [MR3827065](#)
- [2] ALDOUS, D. (1992). Asymptotics in the random assignment problem. *Probab. Theory Related Fields* **93** 507–534. [MR1183889](#) <https://doi.org/10.1007/BF01192719>
- [3] ALDOUS, D. and LYONS, R. (2007). Processes on unimodular random networks. *Electron. J. Probab.* **12** 1454–1508. [MR2354165](#) <https://doi.org/10.1214/EJP.v12-463>
- [4] ALDOUS, D. and PERCUS, A. G. (2003). Scaling and universality in continuous length combinatorial optimization. *Proc. Natl. Acad. Sci. USA* **100** 11211–11215. [MR2007286](#) <https://doi.org/10.1073/pnas.1635191100>
- [5] ALDOUS, D. and STEELE, J. M. (2004). The objective method: Probabilistic combinatorial optimization and local weak convergence. In *Probability on Discrete Structures. Encyclopaedia Math. Sci.* **110** 1–72. Springer, Berlin. [MR2023650](#) https://doi.org/10.1007/978-3-662-09444-0_1
- [6] ALDOUS, D. J. (2001). The $\zeta(2)$ limit in the random assignment problem. *Random Structures Algorithms* **18** 381–418. [MR1839499](#) <https://doi.org/10.1002/rsa.1015>
- [7] ALDOUS, D. J., BORDENAVE, C. and LELARGE, M. (2008). Near-minimal spanning trees: A scaling exponent in probability models. *Ann. Inst. Henri Poincaré Probab. Stat.* **44** 962–976. [MR2453778](#) <https://doi.org/10.1214/07-AIHP138>
- [8] ALDOUS, D. J., BORDENAVE, C. and LELARGE, M. (2009). Dynamic programming optimization over random data: The scaling exponent for near optimal solutions. *SIAM J. Comput.* **38** 2382–2410. [MR2506294](#) <https://doi.org/10.1137/070709037>
- [9] BANDYOPADHYAY, A. (2002). Bivariate uniqueness and endogeneity for the logistic recursive distributional equation. Technical Report #629, Department of Statistics, UC Berkeley, arXiv preprint [arXiv:math/0401389](https://arxiv.org/abs/math/0401389).
- [10] BARAK, B., HOPKINS, S. B., KELNER, J., KOTHARI, P., MOITRA, A. and POTECHIN, A. (2016). A nearly tight sum-of-squares lower bound for the planted clique problem. In *57th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2016* 428–437. IEEE Computer Soc., Los Alamitos, CA. [MR3631005](#)
- [11] BOLLOBÁS, B. (2001). *Random Graphs*, 2nd ed. *Cambridge Studies in Advanced Mathematics* **73**. Cambridge Univ. Press, Cambridge. [MR1864966](#) <https://doi.org/10.1017/CBO9780511814068>
- [12] CHERTKOV, M., KROC, L., KRZAKALA, F., VERGASSOLA, M. and ZDEBOROVÁ, L. (2010). Inference in particle tracking experiments by passing messages between images. *Proc. Natl. Acad. Sci. USA* **107** 7663–7668. <https://doi.org/10.1073/pnas.0910994107>
- [13] CURTISS, J. H. (1942). A note on the theory of moment generating functions. *Ann. Math. Stat.* **13** 430–433. [MR0007577](#) <https://doi.org/10.1214/aoms/1177731541>
- [14] DUDLEY, R. M. (2002). *Real Analysis and Probability. Cambridge Studies in Advanced Mathematics* **74**. Cambridge Univ. Press, Cambridge. Revised reprint of the 1989 original. [MR1932358](#) <https://doi.org/10.1017/CBO9780511755347>
- [15] GAMARNIK, D. (2004). Linear phase transition in random linear constraint satisfaction problems. *Probab. Theory Related Fields* **129** 410–440. [MR2128240](#) <https://doi.org/10.1007/s00440-004-0345-z>
- [16] GAMARNIK, D., NOWICKI, T. and SWIRSZCZ, G. (2006). Maximum weight independent sets and matchings in sparse random graphs. Exact results using the local weak convergence method. *Random Structures Algorithms* **28** 76–106. [MR2187483](#) <https://doi.org/10.1002/rsa.20072>

- [17] GAMARNIK, D., NOWICKI, T. and ŚWIRSZCZ, G. (2008). Invariant probability measures and dynamics of exponential linear type maps. *Ergodic Theory Dynam. Systems* **28** 1479–1495. [MR2449539](#) <https://doi.org/10.1017/S014338570700106X>
- [18] GOEMANS, M. X. and KODIALAM, M. S. (1993). A lower bound on the expected cost of an optimal assignment. *Math. Oper. Res.* **18** 267–274. [MR1250118](#) <https://doi.org/10.1287/moor.18.2.267>
- [19] KALLENBERG, O. (2002). *Foundations of Modern Probability*, 2nd ed. *Probability and Its Applications (New York)*. Springer, New York. [MR1876169](#) <https://doi.org/10.1007/978-1-4757-4015-8>
- [20] KARP, R. M. (1987). An upper bound on the expected cost of an optimal assignment. In *Discrete Algorithms and Complexity* (D. S. Johnson, T. Nishizeki, A. Nozaki and H. S. Wilf, eds.) 1–4. Academic Press. <https://doi.org/10.1016/B978-0-12-386870-1.50006-X>
- [21] KHANDAWALA, M. and SUNDARESAN, R. (2014). Belief propagation for optimal edge cover in the random complete graph. *Ann. Appl. Probab.* **24** 2414–2454. [MR3262507](#) <https://doi.org/10.1214/13-AAP981>
- [22] KURTZBERG, J. M. (1962). On approximation methods for the assignment problem. *J. Assoc. Comput. Mach.* **9** 419–439. [MR0172703](#) <https://doi.org/10.1145/321138.321140>
- [23] LAZARUS, A. J. (1993). Certain expected values in the random assignment problem. *Oper. Res. Lett.* **14** 207–214. [MR1250142](#) [https://doi.org/10.1016/0167-6377\(93\)90071-N](https://doi.org/10.1016/0167-6377(93)90071-N)
- [24] LINUSSON, S. and WÄSTLUND, J. (2004). A proof of Parisi’s conjecture on the random assignment problem. *Probab. Theory Related Fields* **128** 419–440. [MR2036492](#) <https://doi.org/10.1007/s00440-003-0308-9>
- [25] MÉZARD, M. and PARISI, G. (1987). On the solution of the random link matching problems. *J. Phys. France* **48** 1451–1459. <https://doi.org/10.1051/jphys:019870048090145100>
- [26] MOORE, C. (2017). The computer science and physics of community detection: Landscapes, phase transitions, and hardness. *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS* **121** 26–61. [MR3699594](#)
- [27] NAIR, C., PRABHAKAR, B. and SHARMA, M. (2003). Proofs of the Parisi and Coppersmith–Sorkin conjectures for the finite random assignment problem. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings* 168–178. <https://doi.org/10.1109/SFCS.2003.1238191>
- [28] PITTEL, B., SPENCER, J. and WORMALD, N. (1996). Sudden emergence of a giant k -core in a random graph. *J. Combin. Theory Ser. B* **67** 111–151. [MR1385386](#) <https://doi.org/10.1006/jctb.1996.0036>
- [29] PONTRYAGIN, L. S. (1962). *Ordinary Differential Equations*. *Adiw Wes International Series in Mathematics*. Addison-Wesley, Reading, Mass.-Palo Alto, Calif.-London. Translated from the Russian by Leonas Kacinskas and Walter B. Counts. [MR0140742](#)
- [30] SALEZ, J. and SHAH, D. (2009). Belief propagation: An asymptotically optimal algorithm for the random assignment problem. *Math. Oper. Res.* **34** 468–480. [MR2554069](#) <https://doi.org/10.1287/moor.1090.0380>
- [31] SEMERJIAN, G., SICURO, G. and ZDEBOROVÁ, L. (2020). Recovery thresholds in the sparse planted matching problem. *Phys. Rev. E* **102** 022304, 18. [MR4147946](#) <https://doi.org/10.1103/physreve.102.022304>
- [32] TALAGRAND, M. (1995). Concentration of measure and isoperimetric inequalities in product spaces. *Publ. Math. Inst. Hautes Études Sci.* **81** 73–205. [MR1361756](#)
- [33] WALKUP, D. W. (1979). On the expected value of a random assignment problem. *SIAM J. Comput.* **8** 440–442. [MR0539262](#) <https://doi.org/10.1137/0208036>
- [34] WÄSTLUND, J. (2005). The variance and higher moments in the random assignment problem. *Linköping Studies in Mathematics* 8.
- [35] WÄSTLUND, J. (2009). An easy proof of the $\zeta(2)$ limit in the random assignment problem. *Electron. Commun. Probab.* **14** 261–269. [MR2516261](#) <https://doi.org/10.1214/ECP.v14-1475>
- [36] WÄSTLUND, J. (2010). The mean field traveling salesman and related problems. *Acta Math.* **204** 91–150. [MR2600434](#) <https://doi.org/10.1007/s11511-010-0046-7>