

Mean field stable matchings

Daniel Ahlberg*

Maria Deijfen*

Matteo Sfragara*

June 2024

Abstract

Consider the complete bipartite graph on $n + n$ vertices where the edges are equipped with i.i.d. exponential costs. A matching of the vertices is stable if it does not contain any pair of vertices where the connecting edge is cheaper than both matching costs. There exists a unique stable matching obtained by iteratively pairing vertices with small edge costs. We show that the total cost $C_{n,n}$ of this matching is of order $\log n$ with bounded variance, and that $C_{n,n} - \log n$ converges to a Gumbel distribution. We also show that the typical cost of an edge in the matching is of order $1/n$, with an explicit density on this scale, and analyze the rank of a typical edge. These results parallel those of Aldous for the minimal cost matching in the same setting. We then consider the sensitivity of the matching and the matching cost to perturbations of the underlying edge costs. The matching itself is shown to be robust in the sense that two matchings based on largely identical edge costs will have a substantial overlap. The matching cost however is shown to be noise sensitive, as a result of the fact that the most expensive edges will with high probability be replaced after resampling. Our proofs also apply to the complete (unipartite) graph and the results in this case are qualitatively similar.

Keywords: Stable matching, bipartite matching, matching cost, Poisson weighted infinite tree, chaos, noise sensitivity.

AMS 2020 Subject Classification: 60C05, 05C70.

1 Introduction

Consider a situation where a number of objects acting to maximize their own satisfaction are to be matched. Each object ranks the other objects and a matching is then said to be stable if there is no pair of objects that would prefer to be matched to each other rather than their current partners. The concept was introduced in the seminal paper [8] by David Gale and Lloyd Shapley in 1962 and has since received a lot of attention in many different research areas. In 2012, Lloyd Shapley and Alvin Roth received the Nobel Memorial Prize in Economic Sciences for their work on developing mathematical theory for stable matchings and for applications in economics, respectively.

The most basic situation described in [8] consists of matching n men and n women on the marriage market, with only matchings between men and women allowed. This is referred to

*Department of Mathematics, Stockholm University; {daniel.ahlberg}{mia}{matteo.sfragara}@math.su.se

as the stable marriage problem. It is shown that this problem always (that is, for all ranking lists) has at least one solution, and an algorithm for producing a stable matching is also given. The corresponding problem without the bipartite structure is known as the stable roommates problem, alluding to the problem of allocating a number of students to double rooms in a dormitory. In this case, a stable matching may not exist. A polynomial time algorithm that determines if a matching exists and, if so, outputs the matching is described in [12]. For more extensive accounts on general theory for stable matchings, we refer to the books [10, 14, 16] and references therein.

We will consider stable matchings on the complete bipartite graph $K_{n,n}$ and on the complete graph K_n , where the preferences are governed by i.i.d. random edge costs. Let us first focus on $K_{n,n}$, which consists of two disjoint vertex sets $V_n = \{v_1, \dots, v_n\}$ and $V'_n = \{v'_1, \dots, v'_n\}$, and edge set $E_n = \{(v, v') : v \in V_n, v' \in V'_n\}$. Each edge $e = (v, v')$ in the graph is independently assigned an exponential random variable $\omega(e)$ with mean 1. A **matching** is a subset $M \subset E_n$ of non-adjacent edges, and a vertex is **matched** in M if it is contained in an edge of M . The matching is **perfect** if all vertices are matched. The **partner** of v in M is given by

$$M(v) = \begin{cases} v' & \text{if } (v, v') \in M; \\ \emptyset & \text{if } v \text{ is not matched,} \end{cases}$$

and the **matching cost** of v in M is defined as

$$c(v) = \begin{cases} \omega((v, M(v))) & \text{if } v \text{ is matched;} \\ \infty & \text{if } v \text{ is not matched.} \end{cases}$$

A matching is **stable** if there do not exist any pair of vertices with an edge between them that is cheaper than both matching costs, that is, if

$$\forall v \in V_n, v' \in V'_n : (v, v') \notin M \Rightarrow \omega((v, v')) > \min\{c(v), c(v')\}. \quad (1)$$

Vertices hence rank potential partners based on the cost of the connecting edge, and prefer to be matched as cheaply as possible. A vertex pair violating (1) consists of vertices that would prefer to be matched to each other rather than to their current partners, and is therefore called an **unstable pair**. The following algorithm yields an almost surely unique stable matching on $K_{n,n}$:

Greedy algorithm. First select the cheapest edge (v, v') in the graph and include this in the matching. Erase all other edges incident to v and v' . Then select the cheapest edge (u, u') among the remaining edges and include this in the matching. Erase all other edges incident to u and u' . Repeat until all vertices have been matched.

It follows by induction over the steps in the algorithm that all edges created by the algorithm must be included in any stable matching, since omitting any of the edges would result in an unstable pair. Note that it is important that the edge costs are almost surely distinct. Also note that the matching is perfect.

The concept of a stable matching can be defined analogously on the complete graph K_n , and the algorithm then produces a unique stable matching which is perfect if and only if n is even (and otherwise has exactly one unmatched vertex). In our setting, a stable matching hence

always exists also in the non-bipartite case. This is because basing the preferences on random edge cost leads to heavily correlated ranking lists. Indeed, if v is highly ranked by v' , it means that the edge (v, v') has a small cost, which implies that v' is most likely also highly ranked by v .

Matchings on weighted graphs have previously been studied in connection with the so-called random assignment problem. The task is then to assign n jobs to n machines in such a way that the total cost of performing all jobs is minimized. The input consists of a complete bipartite graph with i.i.d. exponential edge weights, specifying the pairwise costs, and the goal is to find a perfect matching that minimizes the total cost

$$C(M) = \sum_{v \in V} c(v).$$

In the seminal paper [3], Aldous proved that the total cost of the minimal matching converges to $\pi^2/6$, which had been conjectured for quite some time. He also analyzed the cost and rank of a typical edge in the minimal matching, and showed that any matching differing from the minimal one in $O(1)$ edges is asymptotically significantly more costly; see Section 1.2 for further details. Background and results predating [3] can be found in [17, 19, 20], and later results e.g. in [21, 22].

In this paper, we derive results for the stable matching that parallel those of Aldous [3] for the minimum matching; see Theorems 1.1-1.4 below. The behaviour that we encounter differs from that of the minimum matching in that the greedy matching results in a heavier edges being added at the end of the process. We then proceed to study the sensitivity of the stable matching with respect to small perturbations of the edge costs. In analogy with Aldous' asymptotic essential uniqueness (AEU) property, we show that updating a small proportion of the edge costs has a limited effect on which edges are contained in the matching (Theorem 1.4). As highlight of the paper, however, we show that the most expensive edges (the 'tail') of the matching are very likely to be replaced by such a perturbation (Theorem 1.5). This is a consequence of the larger cost of the stable matching compared to the minimum matching, where the same behaviour should not occur. Moreover, although the bulk of the stable matching contributes with the lion part of its cost, most of the randomness in its total cost comes from its tail. As a consequence of the sensitivity of the most expensive edges in the matching it follows that the matching cost is highly sensitive to resampling a small proportion of the edge weights (Theorem 1.6). To the best of our knowledge, this is the first confirmed instance where chaotic behaviour of the minimising structure and noise sensitivity of the minimised function does not come hand in hand; however, see recent work of Israeli and Peled [13] for results of a similar flavour.

1.1 Results

Let $S_{n,n}$ denote the unique stable matching on $K_{n,n}$ based on i.i.d. exponential edge weights $\{\omega(e)\}_{e \in E_n}$ with mean 1, and write $C_{n,n} = C(S_{n,n})$ for the total cost of the matching. Our first result specifies the asymptotic behavior of $C_{n,n}$ and is the analogue of [3, Theorem 1]. In contrast to [3, Theorem 1], we also obtain a distributional limit of the centered total cost.

Theorem 1.1 (The total cost). *We have that*

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}[C_{n,n}]}{\log n} = 1 \quad \text{and} \quad \lim_{n \rightarrow \infty} \text{Var}(C_{n,n}) = \frac{\pi^2}{6}.$$

Furthermore, $C_{n,n} - \log n \xrightarrow{d} G$, where G is a Gumbel distributed random variable.

There are n^2 edges in $K_{n,n}$ and hence the cost of the cheapest edge, which is for sure part of the stable matching, is of the order $1/n^2$. The typical cost of an edge in the matching however is of the order $1/n$, as stated in the next theorem. Note that the vertices are exchangeable and hence the matching cost $c(v)$ of vertex v has the same distribution for all vertices $v \in V_n$. This is also the distribution of the cost of a randomly chosen edge contained in the matching. Scaling the cost by n turns out to give rise to a proper random variable with an explicit distribution in the limit. This is the analogue of [3, Theorem 2].

Theorem 1.2 (The typical matching cost). *For any vertex v , the cost $nc(v)$ in $S_{n,n}$ converges in distribution as $n \rightarrow \infty$ to a random variable W with density*

$$f_W(x) = \frac{1}{(1+x)^2}, \quad x \in [0, \infty). \quad (2)$$

Next consider the typical rank of an edge in the matching. Specifically, order the edges incident to vertex $v \in V_n$ in $K_{n,n}$ according to increasing edge cost and let R_n be a random variable indicating the rank of the edge that is used in the stable matching, that is, $R_n = r$ if the matching uses the r th cheapest edge of vertex v . The following result is the analogue of [3, Theorem 3].

Theorem 1.3 (The edge rank). *We have that $R_n \xrightarrow{d} R$ as $n \rightarrow \infty$, where*

- (i) $\mathbb{P}(R = 1) = e \int_1^\infty \frac{e^{-x}}{x} dx \approx 0.596$;
- (ii) $\mathbb{P}(R \geq r) \sim \frac{1}{r}$, as $r \rightarrow \infty$.

Some structures arising from i.i.d. configurations have recently been shown to exhibit a chaotic behavior with respect to perturbations of the underlying configuration. This direction of research first arose in the literature on disordered systems, to which combinatorial optimization problems such as minimal matchings are considered related. Specifically, it has been observed that resampling only a very small fraction of the underlying configuration can cause substantial changes to some structures; see e.g. [1, 6, 7, 9]. Our next result shows that this is not the case for $S_{n,n}$. Let $\omega = \{\omega(e)\}_{e \in E_n}$ and $\omega' = \{\omega'(e)\}_{e \in E_n}$ be two independent random configurations of i.i.d. mean 1 exponential edge costs, and let $\{U(e)\}_{e \in E_n}$ be i.i.d. uniform variables on $[0, 1]$ independent of ω and ω' . For $\varepsilon \in [0, 1]$, define $\omega_\varepsilon = \{\omega_\varepsilon(e)\}_{e \in E_n}$ to be a configuration where a fraction ε of the entries in ω are replaced by their counterparts in ω' , that is,

$$\omega_\varepsilon(e) := \begin{cases} \omega(e) & \text{if } U(e) > \varepsilon, \\ \omega'(e) & \text{if } U(e) \leq \varepsilon. \end{cases} \quad (3)$$

Let $S_{n,n}^\varepsilon$ denote the stable matching based on ω_ε . The following result shows that the fraction of edges in $S_{n,n}^0$ that are also part of $S_{n,n}^\varepsilon$ converges to 1 as $\varepsilon \rightarrow 0$.

Theorem 1.4 (Robustness of the matching). *For any $\varepsilon > 0$, there exists a constant $C > 7$ such that*

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E} [|S_{n,n}^0 \cap S_{n,n}^\varepsilon|]}{n} \geq 1 - C \frac{1}{\log(\frac{1}{\varepsilon})}.$$

While a small perturbation of the edge costs will leave the stable matching largely intact, it turns out that the most expensive edges of the matching, on the contrary, will be replaced with high probability. For $m \geq 1$ and $\varepsilon \in [0, 1]$, let $L_\varepsilon(m)$ denote the sets of vertices corresponding to the m most expensive edges in the matching $S_{n,n}^\varepsilon$ (that is, the last m edges to be picked by the greedy algorithm).

Theorem 1.5 (Sensitivity of the tail). *Let $m \geq 1$ and $\varepsilon \in (0, 1]$ satisfy $m \ll \varepsilon \log n$ as $n \rightarrow \infty$. Then, with high probability as $n \rightarrow \infty$, none of the edges in $L_0(m)$ remain in the matching after perturbation and the two sets $L_0(m)$ and $L_\varepsilon(m)$ are hence disjoint.*

Let $C_{n,n}^\varepsilon$ denote the total cost of the stable matching based on ω_ε . It will turn out that the most expensive edges are responsible for most of the randomness in the matching cost. A consequence of the above result is hence that the total cost of the matching is sensitive to the perturbation of the edge costs, in the sense that the matching costs before and after resampling are asymptotically uncorrelated.

Theorem 1.6 (Noise sensitivity of the matching cost). *For $\varepsilon \log n \gg 1$ we have that*

$$\text{Corr}(C_{n,n}^0, C_{n,n}^\varepsilon) \rightarrow 0 \text{ as } n \rightarrow \infty.$$

The study of noise sensitivity was initiated by Benjamini, Kalai and Schramm [5] in the context of Boolean functions. The topic has since developed substantially, but results are still mainly restricted to Boolean functions. Theorem 1.6 is one of the first instances of noise sensitivity for a more general function (the matching cost).

1.2 Comparison with the minimal matching

As mentioned above, the asymptotic total cost of the minimal matching on $K_{n,n}$ is a constant $\pi^2/6$, while for the stable matching it grows logarithmically with n according to Theorem 1.1. Indeed, the stable matching arises from a greedy algorithm that selects cheap edges in the early stages, but will pay a price for this in the later stages when more expensive edges have to be selected. The typical cost of an edge in the matching however is of the order $1/n$ in both matchings. For the stable matching, the density of the limiting typical edge cost W on this scale is given by (2), and for the minimal matching it is shown in [3, Theorem 2] to equal

$$h(x) = \frac{e^{-x}(e^{-x} - 1 + x)}{(1 - e^{-x})^2}, \quad x \geq 0.$$

The distribution has an exponentially decaying tail for the minimal matching and a power law tail with infinite mean for the stable matching indicating that, also on the typical scale, the stable matching is more likely to produce edges with a large cost.

At the other end of the spectrum, the expected total number of edges in $K_{n,n}$ with cost at most x/n is given by $n^2\mathbb{P}(\omega(e) \leq x/n) \sim xn$ for small x , and the expected number of edges in the matching with cost at most x/n is given by $n\mathbb{P}(W \leq x)$, which according to the given densities of W scales as xn for the stable matching and as $xn/2$ for the minimal matching for small x . The fraction of edges with a small weight on the typical scale that will be a part of the matching hence equals 1 for the stable matching and 1/2 for the minimal matching, so that the stable matching hence includes all but a vanishing fraction of the cheap edges on the typical scale, while the minimal matching uses only half of those edges. Being less greedy in this regime turns out to be beneficial for the minimal matching, since it helps to avoid the expensive edges created at the end of the algorithm by the stable matching.

The edge rank R_n is shown in [3, Theorem 3] to converge to a random variable R with probability function $\mathbb{P}(R = r) = 2^{-r}$. In particular, the probability that the cheapest edge of a vertex is used is 1/2. In the stable matching, on the other hand, this probability is approximately 0.596 and the rank distribution has a power law tail with infinite mean. This again reflects the fact that the stable matching is more likely to use the very cheapest edges, but will in return include more edges with a large cost.

As for the robustness of the stable matching established in Theorem 1.4, a related property, referred to as an asymptotic essential uniqueness (AEU) property, is established for the minimal matching in [3, Theorem 4]. It is shown that, if a matching differs from the minimal one by a proportion at least δ , then its cost is at least $\varepsilon = \varepsilon(\delta)$ larger than the minimal one. The minimal matching is hence unique in the sense that a matching with a cost close to the optimal one must to a large extent coincide with the minimal matching.

1.3 Outline of proofs

Write Y_k for the cost of the edge selected in the k th step of the greedy algorithm. In the first step, the cost is the minimum of n^2 exponential variables with mean 1 and is hence $\text{Exp}(n^2)$ -distributed. With the convention that $Y_0 = 0$, by the memoryless property of the exponential distribution, we can for $k \geq 1$ write

$$Y_k = Y_{k-1} + X_k, \tag{4}$$

where X_k is the minimum of $(n - k + 1)^2$ exponential variables with mean 1 and hence $\text{Exp}((n - k + 1)^2)$ -distributed. The total cost is obtained as

$$C_{n,n} = \sum_{k=1}^n Y_k = \sum_{k=1}^n (n - k + 1)X_k = \sum_{k=1}^n Z_k \quad \text{where } Z_k \sim \text{Exp}(k). \tag{5}$$

Theorem 1.1 follows immediately from this expression, and Theorem 1.2 follows by analyzing the weight of the U th selected edge, where U is uniform on $[n] = \{1, 2, \dots, n\}$.

Theorem 1.3 is proved by transferring the problem to a limiting object known as the Poisson Weighted Infinite Tree (PWIT), which is also the strategy used in [3] (there also the first two results are obtained from computations on the PWIT, since the algorithm to obtain the minimal matching is less explicit). As a preparation for this, we extend the concept of stable matchings to general (possibly infinite) graphs and adapt the greedy algorithm. We also explore connections

between the stable matching and so-called descending paths, which are paths with strictly decreasing edge costs. These ideas have previously appeared in [11]. To obtain the PWIT as a local limit of the weighted version of $K_{n,n}$, we transform the edge costs to the typical scale by multiplying them by n . We then show that the rank on $K_{n,n}$ converges in distribution to the rank on the PWIT, where the latter can be explicitly computed.

Theorem 1.4 is proved by making use of the relation between the stable matching and descending paths. Specifically, changes in the stable matching arising from resampling a certain proportion of the edge costs can be estimated by aid of crude bounds on the set of descending paths emanating from a given vertex.

The results concerning sensitivity of the most expensive edges and the total matching cost are established by splitting the vertex set into two sets $L_0(m)$ and $L_0^c(m)$, corresponding to the m most expensive and the $n-m$ cheapest edges of the original matching, respectively. Theorem 1.5 is proved by showing that, after resampling, every vertex in $L_0(m)$ is desired by many vertices in $L_0^c(m)$ and, with high probability, the desire is reciprocated. This implies that the vertices in $L_0(m)$ are with high probability matched to vertices in $L_0^c(m)$ after resampling. As for Theorem 1.6, we observe that most of the matching cost is generated by the bulk of the matching, which turns out to be essentially deterministic, while most of the randomness comes from the last few, most expensive, edges. We then construct the original matching $S_{n,n}$ and the matching $S_{n,n}^e$ based on the perturbed configuration dynamically by adding edges at times prescribed by their costs. Most edges are the same in both matchings but, by Theorem 1.5, the last edges correspond to disjoint subgraphs and are therefore generated by independent times/costs. Since this phase is responsible for most of the randomness in the matching, the correlation of the matching costs will be small.

1.4 Results for the complete graph

Before proceeding with the proofs, we comment briefly on results for the stable matching on the complete graph K_n , where n is assumed to be even. All our proofs extend, with very minor adjustments, to this case. For the total weight C_n we obtain that

$$\frac{\mathbb{E}[C_n]}{\log n} \rightarrow 1/2 \quad \text{and} \quad \text{Var}(C_n) \rightarrow \pi^2/8.$$

The number of edges in the matching is $n/2$ on K_n while it is n on $K_{n,n}$, so the expected total matching cost is asymptotically the same in relation to the number of edges. This is proved by noting that, on K_n , the representation in (5) is replaced by

$$C_n = \sum_{k=1}^{n/2} (n/2 - k + 1) X'_k = \sum_{k=1}^{n/2} Z'_k \tag{6}$$

where $X'_k \sim \text{Exp}\left(\binom{n-2k+2}{2}\right)$ and $Z'_k \sim \text{Exp}(2k-1)$. The centered total matching cost $C_n - \log n/2$ converges in distribution to a proper random variable also on K_n . However, perhaps somewhat surprisingly, the limiting distribution is not a Gumbel. We explain this in more detail after the proof of Theorem 1.1. Our other results apply in identical formulations also on K_n ,

except that the normalization in Theorem 1.4 is $n/2$ (the number of edges) instead of n . As for Theorem 1.2, the proof is identical, except that we need to work with X'_k instead of X_k and recall that there are $n/2$ instead of n edges to choose from. Theorem 1.3 is proved by computing the rank on the limiting PWIT. It is well known that also the weighted graph K_n converges locally to the PWIT (see Section 3.2) and the distribution of the rank is therefore the same. The proof of Theorem 1.4 applies verbatim on K_n , and so do the proofs of Theorem 1.5 and 1.6, provided we again work with X'_k instead of X_k and recall that there are $n/2$ edges in total.

1.5 Further work

One natural question is to what extent our results generalize to other distributions of the edge costs. Some results will certainly be different, for instance the quantification of the matching cost and its fluctuations in Theorem 1.1 will be affected, as well as the explicit density of the typical matching cost in Theorem 1.2. Note however that the stable matching is defined only through the relative ordering of the edge costs. This implies that, if the edge costs are transformed by a strictly increasing continuous function, then the stable matching does not change. Transforming the costs by a strictly decreasing function, on the other hand, yields a stable matching where expensive edges in the original configuration are preferred. The edges can then be relabelled by inverting their order, so that in particular the rank R_n has the same meaning as above. Since Theorem 1.3 is about the relative ordering of the edges, it can hence be extended to all continuous cost distributions on $(0, \infty)$. Similarly, Theorem 1.4 is only concerned with the matching as a geometrical object and thus also extends to all continuous cost distributions on $(0, \infty)$. The proofs of Theorem 1.5 and Theorem 1.6 rely heavily on specific estimates for the exponential distribution and the memoryless property so would need to be revised for other distributions.

Another question is whether the decorrelation in Theorem 1.6 ceases to hold when instead $\varepsilon \log n \ll 1$. We conjecture that this is indeed the case, so that there is hence a transition at $\varepsilon \sim (\log n)^{-1}$: For $\varepsilon \log n \gg 1$, the matching costs decorrelate, while for $\varepsilon \log n \ll 1$ they do not.

2 Proofs of Theorems 1.1 and 1.2

In this section we give the short proofs of Theorem 1.1 and Theorem 1.2.

Proof of Theorem 1.1. Recall the expression (5) for the total cost and note that the variables $\{Z_k\}_{k \geq 1}$ are independent. The expectation is given by

$$\mathbb{E}[C_{n,n}] = \sum_{k=1}^n \mathbb{E}[Z_k] = \sum_{k=1}^n \frac{1}{k} \sim \log n$$

and the variance by

$$\text{Var}(C_{n,n}) = \sum_{k=1}^n \text{Var}(Z_k) = \sum_{k=1}^n \frac{1}{k^2} \rightarrow \frac{\pi^2}{6}.$$

To obtain the distributional limit, define $\tilde{Z}_k = Z_k - \mathbb{E}[Z_k] = Z_k - 1/k$ and $\tilde{C}_{n,n} = \sum_{k=1}^n \tilde{Z}_k$. The moment generating function of $\tilde{C}_{n,n}$ is given by

$$\Psi_{\tilde{C}_{n,n}}(t) = \prod_{k=1}^n \frac{1}{1-t/k} e^{-t/k} \rightarrow \Gamma(1-t)e^{-\gamma t} \quad \text{as } n \rightarrow \infty,$$

with $\gamma \approx 0.577$ denoting the Euler Mascheroni constant, where the convergence follows from the expansion $\Gamma(t) = \frac{e^{\gamma t}}{t} \prod_{k=1}^{\infty} \frac{1}{1+t/k} e^{t/k}$ and the relation $\Gamma(1-t) = -t\Gamma(-t)$ for the Gamma function. The limit is recognized as the generating function of a Gumbel variable with location parameter $-\gamma$ and scale parameter 1. Finally, note that $\tilde{C}_{n,n} = C_{n,n} - \sum_{k=1}^n 1/k$, where $\sum_{k=1}^n 1/k \sim \log n$. \square

Before proceeding with the proof of Theorem 1.2, we comment on the distributional limit for the stable matching on K_n . The total cost C_n is then given by (6). Centering Z'_k , as in the proof of Theorem 1.1, we obtain \tilde{Z}'_k and the corresponding sum \tilde{C}_n with moment generating function

$$\Psi_{\tilde{C}_n}(t) = \prod_{k=1}^{n/2} \frac{1}{1-t/(2k-1)} e^{-t/(2k-1)} = \prod_{k=1}^n \frac{1}{1-t/k} e^{-t/k} \left[\prod_{k=1}^{n/2} \frac{1}{1-t/2k} e^{-t/2k} \right]^{-1}.$$

Using the same results for the Gamma function as in the proof of Theorem 1.1, we obtain that the first product on the right-hand side converges to $\Gamma(1-t)e^{-\gamma t}$ while the second product converges to $\Gamma(1-t/2)e^{-\gamma t/2}$. Hence

$$\Psi_{\tilde{C}_n}(t) \rightarrow \frac{\Gamma(1-t)}{\Gamma(1-t/2)} e^{\gamma t/2} \quad \text{as } n \rightarrow \infty.$$

We conclude, as in the proof of Theorem 1.1, that $C_n - \log n/2$ converges in distribution to a random variable with this generating function. However, the generating function does not correspond to a Gumbel distribution, and hence the limiting distribution is not Gumbel.

Proof of Theorem 1.2. Recall from Section 1.3 that the cost of the edge selected in the k th step is given by $Y_k = \sum_{i=1}^k X_k$, where $X_k \sim \text{Exp}((n-k+1)^2)$. Let U be uniform on $[0, 1]$. Note that the matching cost $c(v)$ of vertex v has the same distribution as the cost $Y_{\lceil Un \rceil}$ of a randomly chosen edge in the matching. To analyze the latter, note that, for $\alpha \in [0, 1]$, we have that

$$\mathbb{E}[Y_{\lceil \alpha n \rceil}] = \sum_{i=1}^{\lceil \alpha n \rceil} \frac{1}{(n-i+1)^2} \sim \int_1^{\alpha n} \frac{1}{(n-x)^2} dx \sim \frac{\alpha}{n(1-\alpha)}$$

and

$$\text{Var}[Y_{\lceil \alpha n \rceil}] = \sum_{i=1}^{\lceil \alpha n \rceil} \frac{1}{(n-i+1)^2} \sim \int_1^{\alpha n} \frac{1}{(n-x)^4} dx = O(1/n^3).$$

Hence $nY_{\lceil \alpha n \rceil}$ converges in probability to $\alpha/(1-\alpha)$. Now fix $\varepsilon > 0$ and decompose

$$\begin{aligned} \mathbb{P}\left(nY_{\lceil Un \rceil} \leq \frac{\alpha}{1-\alpha}\right) &= \mathbb{P}\left(nY_{\lceil Un \rceil} \leq \frac{\alpha}{1-\alpha}, U \leq \alpha - \varepsilon\right) + \mathbb{P}\left(nY_{\lceil Un \rceil} \leq \frac{\alpha}{1-\alpha}, U \geq \alpha + \varepsilon\right) \\ &\quad + \mathbb{P}\left(nY_{\lceil Un \rceil} \leq \frac{\alpha}{1-\alpha}, U \in (\alpha - \varepsilon, \alpha + \varepsilon)\right). \end{aligned}$$

The first term converges to $\mathbb{P}(U \leq \alpha - \varepsilon) = \alpha - \varepsilon$, the second term converges to 0 and the last term is bounded from above by $\mathbb{P}(U \in (\alpha - \varepsilon, \alpha + \varepsilon)) = 2\varepsilon$. Sending $n \rightarrow \infty$ and $\varepsilon \rightarrow 0$ yields that the limit equals α . Hence $nY_{[Un]}$ converges to a random variable W with a distribution function satisfying $F_W(\frac{\alpha}{1-\alpha}) = \alpha$. The latter can be inverted to $F_W(x) = \frac{x}{1+x}$, which corresponds to the stated density. \square

3 Stable matchings, descending paths and the PWIT

In this section we extend the definition of stable matchings to general weighted graphs, introduce the notion of descending paths and describe how stable matchings are related to such paths. We then define the PWIT, which is a well-known infinite tree arising as local limit of $K_{n,n}$ with exponential weights. This will be useful in the next section, where Theorem 1.3 is proved by transferring the computations to the PWIT and Theorem 1.4 by exploiting the connection between the stable matching and descending paths. Some of these auxiliary results can be found in similar form in [11], and we present them here for completeness.

3.1 Stable matchings and descending paths

Consider a weighted graph $G = (V, E)$, with finite or countably infinite vertex set V , edge set E and edge costs $\{\tau(e)\}_{e \in E}$ (random or deterministic). A matching on G is a subset $M \subset E$ of non-adjacent edges. The concepts of matched vertices, perfect matching, the partner $M(v)$ and matching cost $c(v)$ of a vertex v are defined analogously as in Section 1. A matching is stable if

$$\forall u, v \in V \text{ with } (u, v) \in E : (u, v) \notin M \Rightarrow \tau(u, v) > \min\{c(u), c(v)\}.$$

Note that, if u and v are neighbors in G and $\tau(u, v) < \infty$, then u and v cannot both be unmatched in a stable matching. A stable matching may not exist and, if it does, it may not be unique. Sufficient conditions for existence and uniqueness involve the concept of descending paths. For a weighted graph $G = (V, E)$, a descending path is a weighted subgraph consisting of a sequence of adjacent edges e_1, e_2, \dots such that $\tau(e_1) > \tau(e_2) > \tau(e_3) > \dots$. The set of descending paths emanating from a given vertex $v \in V$ is denoted by $D_v(G)$.

The following proposition from [11] gives conditions that guarantee the existence of a unique stable matching. We include a proof for completeness.

Proposition 3.1 (Holroyd, Martin, Peres (2020)). *Given a weighted graph $G = (V, E)$, there exists a unique stable matching $S(G)$ if*

- (i) *the edge costs are finite and all distinct;*
- (ii) *for each vertex $v \in V$ and all finite $s > 0$, the set of vertices connected to v by an edge with weight less than s is finite;*
- (iii) *there are no infinite descending paths.*

Proof. We prove the proposition by giving an algorithm that produces the matching. The greedy algorithm described in Section 1 works only on finite graphs, but the following algorithm is well-defined on any graph satisfying (i) and (ii):

General greedy algorithm. Two vertices u and v are called potential partners if $(u, v) \in E$, and two potential partners u and v are called mutual favourites if (u, v) is the cheapest among all edges of u and also the cheapest among all edges of v . Note that (i) and (ii) guarantee that any vertex has a unique cheapest edge. Match all mutual favourites and remove them from the graph. Then match all mutual favorites in the remaining graph. Repeat (possibly indefinitely) until no unmatched potential partners remain.

We claim that this produces a unique stable matching. As for the algorithm in Section 1, it follows from induction over the stages in the algorithm that all edges created must be included in any stable matching, since otherwise there would be an unstable pair. We also need to show that all vertices that are left unmatched by the algorithm are unmatched in any stable matching. To this end, let v be a vertex that is unmatched in the matching S arising from the algorithm, and assume there is another stable matching S' where v is matched, say to u_1 . The fact that v is not matched to u_1 in S means that u_1 must be matched to a vertex u_2 with $\tau(u_1, u_2) < \tau(v, u_1)$ in S , since v and u_1 would otherwise constitute an unstable pair in S . Similarly, the fact that u_1 is not matched to u_2 in S' means that u_2 must be matched to a vertex u_3 with $\tau(u_2, u_3) < \tau(u_1, u_2)$, since u_1 and u_2 would otherwise constitute an unstable pair in S' . Iterating this leads to the conclusion that the graph must contain an infinite descending path. If no such path exists, there can hence not exist stable matchings where v is matched. \square

Descending paths turn out to have further importance for the stable matching. In essence, in order to find out if a vertex is matched in the stable matching and, if so, identify its partner, it is sufficient to investigate the set of descending paths emanating from the vertex. To formulate this, write $S(G)$ for the unique stable matching of a graph satisfying the assumptions of Proposition 3.1. Also, denote the set of descending paths including only edges with weight at most $s > 0$ by $D_v(G, s)$.

Proposition 3.2. *Consider a weighted graph $G = (V, E)$ satisfying conditions (i)-(iii) of Proposition 3.1 and fix a vertex $v \in V$. For any $s > 0$, we have that*

$$S(D_v(G, s)) \subseteq S(D_v(G)) \subseteq S(G). \quad (7)$$

Furthermore, if v is unmatched in $S(D_v(G, s))$ for all $s > 0$, then v is unmatched in $S(G)$.

Proof. Note that $D_v(G, s) \subseteq D_v(G)$. Assume that $S(D_v(G, s)) \not\subseteq S(D_v(G))$. This means that there exists a vertex $v \in D_v(G, s)$ that is matched to a vertex u_2 in $S(D_v(G, s))$, but that is matched to another vertex u_1 (or possibly unmatched) in $S(D_v(G))$. We consider the case when $\tau(v, u_2) < \tau(v, u_1)$, so that v has a higher matching cost in $S(D_v(G))$ (including also the possibility that v is unmatched in $S(D_v(G))$), but the opposite case can be handled analogously. By definition of $D_v(G, s)$, no vertex that is matched in $S(D_v(G, s))$ prefers a vertex in $D_v(G, s)^c$ before its partner in $S(D_v(G, s))$, since edges to vertices in $D_v(G, s)^c$ are more expensive than edges to vertices in $D_v(G, s)$. It follows that the vertex u_2 must be matched in $S(D_v(G))$ to a vertex $u_3 \in D_v(G, s)$ that is matched in $S(D_v(G, s))$ and with $\tau(u_2, u_3) < \tau(v, u_2)$, since v and u_2 would otherwise constitute an unstable pair in $S(D_v(G))$. Let u_4 denote the partner of u_3 in $S(D_v(G, s))$. Then $\tau(u_3, u_4) < \tau(u_3, u_2)$, since otherwise u_3 and u_4 would be unstable in $S(D_v(G, s))$. Furthermore, as with u_2 , the vertex u_4 must be matched in $S(D_v(G))$ to a

vertex $u_5 \in D_v(G, s)$ that is matched in $S(D_v(G, s))$ and with $\tau(u_4, u_5) < \tau(u_3, u_4)$, since u_3 and u_4 would otherwise constitute an unstable pair in $S(D_v(G))$. Iterating this leads to an infinite descending chain, which by assumption does not exist, and therefore a contradiction. We conclude that all vertices in $D_v(G, s)$ that are matched in $S(D_v(G, s))$ must be matched to the same partner in $S(D_v(G))$, that is, $S(D_v(G, s)) \subseteq S(D_v(G))$. The other inclusion in (7) follows from an analogous argument, noting that no vertex in $D_v(G)$ prefers a vertex in $D_v(G)^c$ before its partner in $S(D_v(G))$.

To show the last statement, assume that v is unmatched in $S(D_v(G, s))$ for all $s > 0$, but that v is matched in $S(G)$, say to u . The matching cost of v in $S(G)$ is $c(v) = \tau(v, u)$. By (7), the vertex u cannot be matched to a different vertex in $S(D_v(G, c(v)))$, since it would then be matched to this other vertex also in $S(G)$. Hence both u and v are unmatched in $S(D_v(G, c(v)))$ and thus constitute an unstable pair. We conclude that v cannot be matched in $S(G)$. \square

3.2 The PWIT

The Poisson Weighted Infinite Tree (PWIT) was first introduced in [2]. To describe it, consider first a root vertex with an infinite number of children. The edges from the root to the children are assigned weights according to a Poisson process with rate 1. Recursively, each child is then given an infinite number of new children and the edges to these new children are again assigned weights according to the arrival times of independent Poisson processes with rate 1. Continuing this procedure, leads to a rooted infinite tree \mathcal{T} known as the PWIT. Formally, the PWIT is a rooted weighted graph with vertex set

$$\mathcal{V} = \bigcup_{k=0}^{\infty} \mathbb{N}^k = \{0, 1, 2, \dots, 11, 12, \dots, 21, 22, \dots, 111, 112, \dots\},$$

where 0 is the root, and edges (v, vj) , for each $v \in \mathcal{V}$ and $j \in \mathbb{N}$, where vj is referred to as a child of v . For $v \in \mathcal{V}$, let $(T_j^{(v)})_{j \in \mathbb{N}}$ be the points (in increasing order) of a Poisson process on \mathbb{R}_+ with rate 1. The cost of an edge (v, vj) is given by $T_{vj} = T_j^{(v)}$, where we write $T_{0j} = T_j$; see Figure 1.

Now consider $K_{n,n}$ with edge costs $\{n\omega(e)\}_{e \in E_n}$, where $\{\omega(e)\}_{e \in E_n}$ are i.i.d. exponential with mean 1. With this scaling of the weights, the cheapest edge of a given vertex is $\text{Exp}(1)$, the second cheapest is $\text{Exp}(1) + \text{Exp}(1)$ etc, that is, the ordered weights are described by the arrival times of a rate 1 Poisson process. It is well known that $K_{n,n}$ with costs $\{n\omega(e)\}_{e \in E_n}$ converges to the PWIT in a certain sense. Specifically, write \mathcal{G}_* for the set of rooted weighted graphs satisfying the assumption (ii) of Proposition 3.1. It can be shown that \mathcal{G}_* is a complete separable metric space, and a notion of local weak convergence can be defined for probability measures on \mathcal{G}_* . A sequence of weighted graphs $\{G_n\}_{n \geq 1}$ converges locally to the PWIT if the following holds: Fix a radius $\rho > 0$ and, given a vertex v of G_n , consider the subgraph consisting of all paths from v with total cost at most ρ . Similarly, consider the subtree of the PWIT consisting of all paths from the root with total cost at most ρ . Then, for any given ρ , the graph G_n can be coupled with the PWIT so that, with high probability as $n \rightarrow \infty$, there is an isomorphism between the two subgraphs which identifies v with the root of the PWIT and which preserves the edge costs. In particular, this means that it is unlikely to encounter short cycles in G_n . We refer to [2, 4] for further details and a general framework for local weak convergence. Note that also the complete graph K_n with exponential edge weights converges to the PWIT.




Figure 1: The PWIT.

Proposition 3.3 (Aldous (1992)). *The complete bipartite graph $K_{n,n}$ with i.i.d. exponential edge costs with mean n converges locally to the PWIT:*

$$K_{n,n} \xrightarrow{d} \mathcal{T} \quad \text{as } n \rightarrow \infty.$$

Next, we want to apply Proposition 3.1 to establish the existence of a unique stable matching on the PWIT. To this end, we first recall from [11, Lemma 4.8] that the PWIT does not contain infinite descending paths. Here, $|G|$ denotes the number of vertices in a graph G .

Proposition 3.4 (Holroyd, Peres, Martin (2020)). *Consider \mathcal{T} and its root 0. For all $s > 0$, we have that*

$$\mathbb{E}[|D_0(\mathcal{T}, s)|] = e^s. \quad (8)$$

In particular, there are almost surely no infinite descending paths in \mathcal{T} .

Proof. For $k \geq 0$, consider descending paths from 0 of length k and with edge costs less than s . Each such path consists of k edges with decreasing costs, where the first edge has cost $s_1 < s$, the second edge has cost $s_2 < s_1$, and the j th edge has cost $s_j < s_{j-1}$, for $j = 3, \dots, k$. The costs along paths of length k can be represented by the points of a unit rate Poisson process on \mathbb{R}^k and, integrating over the region $0 < s_k < \dots < s_1 < s$, we obtain that the expected number of descending paths of length k with costs less than s is $\int_{0 < s_k < \dots < s_1 < s} ds_1 \cdots ds_k = \frac{s^k}{k!}$. Each vertex is the endpoint of at most one such path and thus the expression for $\mathbb{E}[|D_0(\mathcal{T}, s)|]$ follows by summing over k .

Recall that T_n is the cost of the edge from the root of \mathcal{T} to its n th child. It follows from (8) that $|D_0(\mathcal{T}, T_n)|$ is finite almost surely for any n , implying that \mathcal{T} does not contain infinite descending paths. \square

Given this, it is clear that \mathcal{T} satisfies the assumptions of Proposition 3.1 and we can therefore conclude that it has a unique stable matching.

Proposition 3.5. *There exists almost surely a unique stable matching $S(\mathcal{T})$ on the PWIT.*

Note that we do not yet know that $S(\mathcal{T})$ is perfect. This will follow from Proposition 3.7 below. First we note that the set of descending paths in $K_{n,n}$ can be coupled to the set of descending paths in \mathcal{T} . This will allow us to derive results for $S(K_{n,n})$ from results for $S(\mathcal{T})$ since, by Proposition 3.2, the stable matching on a graph is determined by descending paths.

Proposition 3.6. *Consider $K_{n,n}$ with exponential edge costs with mean n , and fix a vertex v . For all $s > 0$, there exists a coupling of $D_v(K_{n,n}, s)$ and $D_0(\mathcal{T}, s)$ such that the weighted graphs coincide with high probability as $n \rightarrow \infty$.*

Proof. By Proposition 3.4, the set of descending paths $D_0(\mathcal{T}, s)$ is contained in the set of paths from 0 with total weight at most ρ for some value of $\rho < \infty$. The claim hence follows from Proposition 3.3. \square

Write W_0 for the matching cost of the root in the stable matching on the PWIT. We end this section by determining the distribution of W_0 . Since W_0 is finite almost surely, it follows that the stable matching on the PWIT is perfect almost surely. This is proved in [11, Section 3.2.1], but we give a different argument based on Theorem 1.2 and the connection between the stable matching and descending paths.

Proposition 3.7. *We have that $W_0 \stackrel{d}{=} W$, where the density of W is given by (2).*

Proof. Recall that $c(v)$ denotes the matching cost of vertex v in $K_{n,n}$ equipped with i.i.d. exponential edge weights with mean 1. Write $\tilde{c}_n(v) := nc(v)$ for the cost when the weights are scaled to have mean n . By Theorem 1.2, the cost $\tilde{c}_n(v)$ converges in distribution to a proper random variable W with density (2). The claim hence follows from the uniqueness of the limiting distribution if we show that $\tilde{c}_n(v)$ converges in distribution to W_0 . To this end, let $\tilde{c}_n^{(s)}(v)$ denote the analogue of $\tilde{c}_n(v)$ in the stable matching on $D_v(K_{n,n}, s)$ (based on exponential weights with mean n) and, similarly, let $W_0^{(s)}$ be the analogue of W_0 on $D_0(\mathcal{T}, s)$. By Proposition 3.2, the root is matched to a vertex u in $S(\mathcal{T})$ if and only if it is matched to u in $S(D_0(\mathcal{T}, s))$ for large s . Furthermore, by Proposition 3.6, the graphs $D_0(\mathcal{T}, s)$ and $D_v(K_{n,n}, s)$ can be coupled so that they coincide with high probability as $n \rightarrow \infty$. Hence

$$\mathbb{P}(W_0 > x) = \lim_{s \rightarrow \infty} \mathbb{P}(W_0^{(s)} > x) = \lim_{s \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{P}(\tilde{c}_n^{(s)}(v) > x). \quad (9)$$

It follows from Proposition 3.2 applied to $K_{n,n}$ that $\tilde{c}_n^{(s)}(v) \geq \tilde{c}_n(v)$ (with equality if $\tilde{c}_n^{(s)}(v) < \infty$, that is, if v is matched in $S(D_v(K_{n,n}, s))$). Since $\tilde{c}_n(v)$ does not depend on s , we obtain that

$$\lim_{s \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{P}(\tilde{c}_n^{(s)}(v) > x) \geq \lim_{n \rightarrow \infty} \mathbb{P}(\tilde{c}_n(v) > x). \quad (10)$$

To get the reverse inequality, note that, on the event $\{\tilde{c}_n(v) \leq s\}$, we have that $\tilde{c}_n^{(s)}(v) = \tilde{c}_n(v)$, since v is then matched in $S(D_v(K_{n,n}, s))$. We can thus bound

$$\begin{aligned} \mathbb{P}(\tilde{c}_n^{(s)}(v) > x) &= \mathbb{P}(\tilde{c}_n^{(s)}(v) > x \cap \tilde{c}_n(v) \leq s) + \mathbb{P}(\tilde{c}_n^{(s)}(v) > x \cap \tilde{c}_n(v) > s) \\ &\leq \mathbb{P}(\tilde{c}_n(v) > x) + \mathbb{P}(\tilde{c}_n(v) > s). \end{aligned}$$

It follows from Theorem 1.2 that $\lim_{s \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{P}(\tilde{c}_n(v) > s) = \lim_{s \rightarrow \infty} \mathbb{P}(W > s) = 0$ and hence

$$\lim_{s \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{P}(\tilde{c}_n^{(s)}(v) > x) \leq \lim_{n \rightarrow \infty} \mathbb{P}(\tilde{c}_n(v) > x). \quad (11)$$

Combining (9)-(11) we conclude that $\tilde{c}_n^{(s)}(v) \xrightarrow{d} W_0$, as desired. \square

4 Proofs of Theorems 1.3 and Theorem 1.4

In this section, we prove Theorem 1.3 and Theorem 1.4. Consider a vertex $v \in V_n$ in $K_{n,n}$ and order the edges emanating from v according to cost, so that e_1 is the cheapest edge and e_n the most expensive one. Recall that R_n denotes the rank of the edge used by v in the stable matching, that is, $R_n = m$ if $e_m \in S_{n,n}$. Write R for the analogous quantity on the PWIT:

$$R = \begin{cases} m & \text{if } (0, m) \in S(\mathcal{T}); \\ \infty & \text{if } 0 \text{ is not matched in } S(\mathcal{T}). \end{cases}$$

Theorem 1.3 is a consequence of the following two propositions.

Proposition 4.1. *We have that $R_n \xrightarrow{d} R$ as $n \rightarrow \infty$.*

Proposition 4.2. *The rank R on the PWIT satisfies (i) and (ii) of Theorem 1.3.*

Proof of Proposition 4.1. This follows from the same arguments that were used to show that $\tilde{c}_n(v) \rightarrow W_0$ in the proof of Proposition 3.7. To see this, first note that scaling the edge costs does not affect the ranking of the edges. We can thus use the scaled edge weights $\{n\omega(e)\}_{e \in E}$, where $\{\omega(e)\}_{e \in E}$ are the original i.i.d. edge weights. Let $R_n^{(s)}$ and $R^{(s)}$ denote the analogues of R_n and R in the stable matchings on $D_v(K_{n,n}, s)$ and $D_v(\mathcal{T}, s)$ respectively, that is, $R_n^{(s)} = m$ if $e_m \in S(D_v(K_{n,n}, s))$ and $R^{(s)} = m$ if $e_m \in S(D_v(\mathcal{T}, s))$. The proof that $\tilde{c}_n(v) \rightarrow W_0$ in the proof of Proposition 3.7 can now be applied verbatim with $\tilde{c}_n(v)$ and W_0 replaced by R_n and R , and with $\tilde{c}_n^{(s)}(v)$ and $W_0^{(s)}$ replaced by $R_n^{(s)}$ and $R^{(s)}$. \square

Proof of Proposition 4.2. For $j = 1, 2, \dots$, let W_j denote the matching cost of vertex j in the PWIT in the stable matching on the subgraph consisting of j and its descendants, that is, the edge $(0, j)$ is removed and a stable matching is then constructed on the connected component of vertex j ; see Figure 2. These components have the same structure as the PWIT, implying that $\{W_j\}_{j=0}^\infty$ are i.i.d. random variables. By Proposition 3.7, the density is given by (2). Recall that the cost of the edge $(0, j)$ is T_j and note that $R = \min\{j \geq 1 : T_j \leq W_j\}$. It follows that

$$\mathbb{P}(R = 1) = \mathbb{P}(U_1 \leq Z_1) = \int_0^\infty F_{T_1}(w) f_W(w) dw = \int_0^\infty \frac{1 - e^{-w}}{(1+w)^2} dw = e \int_1^\infty \frac{e^{-t}}{t} dt,$$

where the last integral can be recognized as $-\text{Ei}(1)$ with $\text{Ei}(x) = -\int_{-x}^\infty \frac{e^{-t}}{t} dt$ denoting the exponential integral. This proves (i).

As for (ii), note that $\mathbb{P}(R > r) = \mathbb{P}(T_j > W_j, \forall j \leq r)$. We can compute this probability by considering an inhomogeneous Poisson process with rate $\lambda(t) = 1 - F_W(t)$. Indeed, first consider




Figure 2: The PWIT with vertices labeled by their matching cost in their respective subgraphs.

a standard Poisson process with rate 1 where the event times represent the variables $\{T_j\}_{j \in \mathbb{N}}$, and then generate an inhomogeneous process by accepting an event at time t independently with probability $1 - F_W(t)$. The first accepted event is by construction the R th event of the original process and $\mathbb{P}(R > r)$ is then the probability that the inhomogeneous Poisson process has no events before time T_r . Hence

$$\mathbb{P}(R > r) = \mathbb{E}[\mathbb{P}(\text{no events before } T_r \mid T_r)] = \mathbb{E}\left[e^{-\int_0^{T_r} (1 - F_W(t)) dt}\right] = \mathbb{E}\left[\frac{1}{1 + T_r}\right].$$

Since $\mathbb{E}[T_r] = \text{Var}(T_r) = r$, we have that $T_r \sim r$ as $r \rightarrow \infty$ with deviations of order \sqrt{r} , and hence $\mathbb{P}(R \geq r) \sim r^{-1}$. \square

It remains to prove Theorem 1.4. To this end, a bound on $|D_v(K_{n,n}, s)|$ uniformly in n is needed. This can be obtained from the bound on $|D_0(\mathcal{T}, s)|$ in Proposition 3.4.

Lemma 4.3. *For $K_{n,n}$ with exponential edge costs with mean n , we have that*

$$\mathbb{P}(|D_v(K_{n,n}, s)| > e^{2s}) \leq e^{-s}$$

uniformly in n .

Proof. First note that $D_v(K_{n,n}, s)$ and $D_v(K_{n+1,n+1}, s)$ can be coupled so that $D_v(K_{n,n}, s) \subseteq D_v(K_{n+1,n+1}, s)$. Indeed, if $K_{n+1,n+1}$ is constructed from $K_{n,n}$ by adding one vertex to each of the two vertex sets and equipping the edges of these vertices with i.i.d. weights, while the weights of existing edges in $K_{n,n}$ remain the same, then the set of descending paths is non-decreasing. Given this, we obtain that

$$\mathbb{P}(|D_v(K_{n,n}, s)| > e^{2s}) \leq \lim_{n \rightarrow \infty} \mathbb{P}(|D_v(K_{n,n}, s)| > e^{2s}) = \mathbb{P}(|D_0(\mathcal{T}, s)| > e^{2s}) \leq e^{-s},$$

where the equality follows from Proposition 3.6 and the last inequality follows from (8). \square

Proof of Theorem 1.4. Recall that $S_{n,n}^\varepsilon$ denotes the stable matching based on edge costs (3), that is, a proportion $\varepsilon > 0$ of the edge costs $\{\omega(e)\}_{e \in E_n}$ is resampled. Also, for a subgraph $G \subset K_{n,n}$, write $S^\varepsilon(G)$ for the stable matching of G based on the resampled set of edge weights. We will again work with scaled edge costs $\{n\omega_\varepsilon(e)\}_{e \in E_n}$, since this will allow us to make use of Lemma 4.3. Fix a vertex $v \in V_n$ and let $nc(v)$ refer to its matching cost in the initial configuration (with $\varepsilon = 0$). We will show that, if s is large, it is unlikely that an edge in $D_v(K_{n,n}, s)$ or its boundary is resampled in such a way that the stable matching on $D_v(K_{n,n}, s)$ is changed. This will prove the claim since, by Proposition 3.2, vertex v will be matched to the same partner in $S_{n,n}^\varepsilon$ as in $S^\varepsilon(D_v(K_{n,n}, s))$ in the limit.

Fix $\varepsilon > 0$ and $s > 0$, where s will later be chosen as a function of ε . Let $A_{v,n,s}$ be the event that at least one edge in $D_v(K_{n,n}, s)$ is resampled. By Lemma 4.3, we have that

$$\begin{aligned} \mathbb{P}(A_{v,n,s}) &= \mathbb{P}(A_{v,n,s} \mid |D_v(K_{n,n}, s)| > e^{2s}) \mathbb{P}(|D_v(K_{n,n}, s)| > e^{2s}) \\ &\quad + \mathbb{P}(A_{v,n,s} \mid |D_v(K_{n,n}, s)| \leq e^{2s}) \mathbb{P}(|D_v(K_{n,n}, s)| \leq e^{2s}) \\ &\leq e^{-s} + \varepsilon e^{2s}, \end{aligned} \tag{12}$$

uniformly in n . Define the edge boundary of $D_v(K_{n,n}, s)$ to be the set of edges in $K_{n,n}$ with exactly one endpoint in $D_v(K_{n,n}, s)$. Similarly, let $B_{v,n,s}$ be the event that an edge in the boundary of $D_v(K_{n,n}, s)$ is resampled and, in addition, that its new (scaled) cost is less than s . Using Lemma 4.3, the fact that there are at most ne^{2s} edges on the boundary if $|D_v(K_{n,n}, s)| \leq e^{2s}$ and a similar split as in (12), we obtain that

$$\mathbb{P}(B_{v,n,s}) = e^{-s} + \varepsilon ne^{2s}(1 - e^{-s/n}) \leq e^{-s} + \varepsilon se^{2s}, \tag{13}$$

uniformly in n . Write $M_\varepsilon(v)$ for the matching partner of v in $S_{n,n}^\varepsilon$. If $nc(v) < s$, then, by Proposition 3.2, vertex v is matched in $S^0(D_v(K_{n,n}, s))$ and the partner is the same as in $S_{n,n}^0$. Furthermore, on the event $A_{v,n,s}^c \cap B_{v,n,s}^c$, the resampled and the non-resampled configurations on $D_v(K_{n,n}, s)$ coincide and, in addition, no vertex that is matched in $S^\varepsilon(D_v(K_{n,n}, s))$ prefers a vertex outside of $D_v(K_{n,n}, s)$ before its partner in $D_v(K_{n,n}, s)$. It follows from the same argument as in the proof of Proposition 3.2 that v is matched to the same vertex in $S_{n,n}^\varepsilon$ as in $S^\varepsilon(D_v(K_{n,n}, s))$. Also, by the above, v is matched to the same vertex in $S^\varepsilon(D_v(K_{n,n}, s))$ as in $S_{n,n}^0$. Hence

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{P}(M_0(v) \neq M_\varepsilon(v)) &\leq \lim_{n \rightarrow \infty} [\mathbb{P}(nc(v) \geq s) + \mathbb{P}(A_{v,n,s}) + \mathbb{P}(B_{v,n,s})] \\ &\leq \frac{1}{1+s} + \varepsilon e^{2s} + \varepsilon se^{2s} \\ &\leq \frac{3}{s} + 2\varepsilon se^{2s}, \end{aligned}$$

where in the second inequality we have used Theorem 1.2 and (12)–(13), while in the last inequality we have used the fact that $e^{-s} \leq \frac{1}{s}$. Letting $s = C' \log(\frac{1}{\varepsilon})$, with $C' < 1/2$, we obtain for some $\delta > 0$ that

$$2\varepsilon se^{2s} = 2C' \varepsilon^{1-2C'} \log\left(\frac{1}{\varepsilon}\right) \leq \varepsilon^\delta \leq \frac{1}{\log\left(\frac{1}{\varepsilon}\right)}$$

so that hence

$$\lim_{n \rightarrow \infty} \mathbb{P}(M_0(v) \neq M_\varepsilon(v)) \leq \frac{C}{\log\left(\frac{1}{\varepsilon}\right)},$$

for any $C > 7$. Consequently, the probability that the edge $(v, M(v))$ is present both in $S_{n,n}^0$ and in $S_{n,n}^\varepsilon$ is given by

$$\lim_{n \rightarrow \infty} \mathbb{P}(M_0(v) = M_\varepsilon(v)) \geq 1 - \frac{C}{\log\left(\frac{1}{\varepsilon}\right)}.$$

Summing over all n edges in the stable matching gives the desired result

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}[|S_{n,n}^0 \cap S_{n,n}^\varepsilon|]}{n/2} \geq 1 - C \frac{1}{\log\left(\frac{1}{\varepsilon}\right)}.$$

□

5 Sensitivity of the tail of the matching

In this section, we prove Theorem 1.5, stating that the most expensive edges in $S_{n,n}^0$ and $S_{n,n}^\varepsilon$ are with high probability different. Recall that $L_\varepsilon(m)$ denotes the sets of the m most expensive edges in $S_{n,n}^\varepsilon$. To ease notation, we write $L_\varepsilon(m) = L_\varepsilon$ and abbreviate $L_0 = L$. Note that, before perturbing the costs, no edge connecting a vertex in L to a vertex in L^c is included in the matching. To establish the theorem, we will show that, for every vertex $u \in L$ there will be many vertices $v \in L^c$ for which the edge cost of (u, v) is resampled in such a way that v prefers to be rematched to u . In order to make sure that u also desires v , and that the cost of the new match is not among the m most expensive edges in the new matching, we require that resampled edges have costs below a certain threshold δ , which it is unlikely that any edge in L falls below. The core of the proof will be to establish two key lemmas, formalising this outline. First however we will require some information regarding the magnitude and concentration of the edge weights of the stable matching.

5.1 Concentration of the matching costs

Recall that Y_1, Y_2, \dots, Y_n denote the costs of the edges in the stable matching $S_{n,n}^0$, ordered from cheapest to most expensive. By the representation in (4), the cost of the k th cheapest edge is

$$Y_k = \sum_{i=1}^k X_i, \tag{14}$$

where X_1, X_2, \dots, X_n are independent and exponentially distributed random variables where the parameter of X_i is $(n - i + 1)^2$. It follows, in particular, that

$$\mathbb{E}[Y_{n-\ell}] = \sum_{i=1}^{n-\ell} \mathbb{E}[X_i] = \sum_{i=\ell+1}^n \frac{1}{i^2}$$

and

$$\text{Var}(Y_{n-\ell}) = \sum_{i=1}^{n-\ell} \text{Var}(X_i) = \sum_{i=\ell+1}^n \frac{1}{i^4}.$$

Approximating the sum with an integral leads to

$$\frac{1}{\ell+1} - \frac{1}{n} \leq \mathbb{E}[Y_{n-\ell}] \leq \frac{1}{\ell} \quad \text{and} \quad \text{Var}(Y_{n-\ell}) \leq \frac{1}{3\ell^3}. \quad (15)$$

This yields the following concentration bounds on the final most expensive edges of the matching.

Lemma 5.1. *For $\ell \geq 1$ and $n \geq 6\ell$ we have*

$$\mathbb{P}\left(Y_{n-\ell} < \frac{1}{6\ell}\right) \leq \frac{12}{\ell} \quad \text{and} \quad \mathbb{P}\left(Y_{n-\ell} > \frac{7}{6\ell}\right) \leq \frac{12}{\ell}.$$

Proof. For $\ell \geq 1$ and $n \geq 6\ell$ we have from (15) that $\frac{1}{3\ell} \leq \mathbb{E}[Y_{n-\ell}] \leq \frac{1}{\ell}$, so the result follows from (15) and Chebyshev's inequality. \square

Summing over k in (14) gives the accumulated cost of the edges in the matching. In particular, $C_{n,n} = \sum_{k=1}^n Y_k$. The accumulated cost of the first $n - \ell$ edges is

$$\sum_{k=1}^{n-\ell} Y_k = \sum_{i=1}^{n-\ell} (n - \ell - i + 1) X_i.$$

Hence,

$$\sum_{k=1}^{n-\ell} \mathbb{E}[Y_k] = \sum_{i=1}^{n-\ell} (n - \ell - i + 1) \mathbb{E}[X_i] = \sum_{i=\ell+1}^n \frac{i - \ell}{i^2}.$$

Similarly,

$$\text{Var}\left(\sum_{k=1}^{n-\ell} Y_k\right) = \sum_{i=1}^{n-\ell} (n - \ell - i + 1)^2 \text{Var}(X_i) = \sum_{i=\ell+1}^n \frac{(i - \ell)^2}{i^4}.$$

Comparing the sums to integrals, for $\ell \ll n$, leads to the bounds

$$\log\left(\frac{n}{\ell}\right) - 2 \leq \sum_{k=1}^{n-\ell} \mathbb{E}[Y_k] \leq \log\left(\frac{n}{\ell}\right) \quad \text{and} \quad \text{Var}\left(\sum_{k=1}^{n-\ell} Y_k\right) \sim \frac{1}{3\ell}. \quad (16)$$

Finally, reversing the sum and using (15), we get that for all $n \geq 1$

$$\sum_{k=1}^n \mathbb{E}[Y_k^2] = \sum_{\ell=1}^n (\text{Var}(Y_{n-\ell}) + \mathbb{E}[Y_{n-\ell}]^2) \leq \sum_{\ell=1}^n \left(\frac{1}{3\ell^3} + \frac{1}{\ell^2}\right) \leq \frac{\zeta(3)}{3} + \frac{\pi^2}{6} \leq 3, \quad (17)$$

where $\zeta(s)$ is the Riemann zeta function.

5.2 Key lemmas

Set $\delta := (\log n)^{-3}$. By Lemma 5.1, it is unlikely for the last m edges of the matching to have cost below δ . As we shall see, the threshold δ is chosen so that it remains unlikely for the cost of edges between vertices in the set L of the m most expensive edges in the original matching to fall below δ even after the costs have been resampled.

Recall that ω_ε denotes the configuration of edge costs after an ε -perturbation. Given a vertex u , we denote by $c_\varepsilon^u(v)$ the cost of the vertex v in the stable matching of $K_{n,n}$ with respect to ω_ε where u has been removed (and one node is necessarily left unmatched). For $u \in L$ let

$$J_u := \{v \in L^c : (u, v) \text{ resampled}\},$$

$$N_u := \#\{v \in J_u : \omega'(u, v) < c_\varepsilon^u(v) \wedge \delta\},$$

where $\#$ denotes the cardinality of the set.

The key step towards Theorem 1.5 is to show that, with high probability $N_u \geq 1$ for all $u \in L$. In order to do that, we compare the costs of the edges in $S_{n,n}$ to the costs of the stable matching of $K_{n,n}$ when a vertex u has been removed. Note that the matching of $K_{n,n}$ with a vertex removed will contain $n - 1$ edges, and we denote their weights by $Y_1^u, Y_2^u, \dots, Y_{n-1}^u$ in increasing order.

Lemma 5.2. *Almost surely, we have for every vertex u and all $k = 1, 2, \dots, n - 1$ that*

$$Y_k \leq Y_k^u \leq Y_{k+1}.$$

Proof. Fix a vertex u in $K_{n,n}$. We will take a dynamic perspective on the construction of the matching, where we think of the weights ω as the times of the first rings of independent Poisson clocks associated with the edges of $K_{n,n}$. We may then construct the matching dynamically in time, by adding an edge e at time $\omega(e)$ unless either of its endpoints has already been matched at an earlier time.

In order to address the discrepancy between $S_{n,n}$ and the matching of $K_{n,n}$ with u removed, with respect to the same configuration ω , we colour ‘red’ the edges in $S_{n,n}$ that are not part of the matching with u removed, and ‘blue’ the edges in the matching with u removed, which are not in $S_{n,n}$. The edges that are in both matchings are not coloured, that is, they remain ‘black’. Note that the first time a coloured edge is added to either of the two matchings is when u is added to $S_{n,n}$, since this is the only discrepancy between the two weighted graphs. This edge is red, and we denote by r_1 its weight, and by v_1 its endpoint not equal to u . The discrepancy at time r_1 is now moved to the vertex v_1 . Either v_1 is left unmatched, or the next coloured edge added to the matching comes when v_1 is matched in $K_{n,n}$ with u removed. This edge is thus blue, and we let b_1 denote its cost and u_2 its endpoint other than v_1 . Since v_1 is added after u , we have $r_1 < b_1$, and the discrepancy in the two constructions is now transferred to u_2 . Repeating the above argument we find an alternating sequence of red and blue edges being added to the graph, starting and ending with a red edge, whose weights are similarly alternating

$$r_1 < b_1 < r_2 < b_2 \dots < r_\ell$$

for some $\ell \geq 1$. Since coloured edges are added alternately, there is at any time at most one more red than blue edge present, and never more blue than red. Consider the k th edge added

to the matching of $K_{n,n}$ with u removed, which happens at time Y_k^u . Either this edge is black, in which case it is either the k th or $(k+1)$ st edge added to $S_{n,n}$, and hence Y_k^u equals either Y_k or Y_{k+1} . Or the edge is blue, in which case $S_{n,n}$ already consists of k but not $k+1$ edges, and so $Y_k < Y_k^u < Y_{k+1}$. This holds for every u and $k = 1, 2, \dots, n-1$. \square

Our main step towards Theorem 1.5 is a moment analysis of N_u for $u \in L$. For ease of notation, we shall let $\mathbb{P}' := \mathbb{P}(\cdot | (L, L^c))$, and write \mathbb{E}' and Var' for expectation and variance with respect to \mathbb{P}' . Note that the law of the cost of a vertex, under \mathbb{P}' , depends on whether the vertex belongs to L or L^c , whereas the law of (Y_1, Y_2, \dots, Y_n) is equal under \mathbb{P} and \mathbb{P}' .

Lemma 5.3. *For $m \geq 1$ and $\varepsilon \in (0, 1]$ satisfying $2m \leq \varepsilon \log n$, we have that for every vertex $u \in L$ the two following statements hold:*

- (i) $\mathbb{E}'[N_u] = (1 + o(1))\varepsilon \log n$;
- (ii) $\text{Var}'(N_u) \leq 6\varepsilon \log n$.

Proof. We prove the two statements separately.

- (i) Fix a vertex $u \in L$ and let $F(x) = 1 - e^{-x}$ denote the distribution function of the exponential distribution. By conditioning on everything but the update variables U_e for the edges that connect u to L^c , we find that

$$\mathbb{E}'[N_u] = \varepsilon \sum_{v \in B} \mathbb{P}'(\omega'(u, v) < c_\varepsilon^u(v) \wedge \delta).$$

Conditioning, this time on the weight configuration ω_ε for edges not incident to u , we find that

$$\mathbb{E}'[N_u] = \varepsilon \sum_{v \in B} \mathbb{E}'[F(c_\varepsilon^u(v) \wedge \delta)]$$

Define $H := \sum_{v \in L^c} c_\varepsilon^u(v) \wedge \delta$. Using that $x - \frac{1}{2}x^2 \leq F(x) \leq x$, Lemma 5.2 and (17), we obtain that

$$\left| \mathbb{E}'[N_u] - \varepsilon \mathbb{E}'[H] \right| \leq \frac{\varepsilon}{2} \sum_{v \in L^c} \mathbb{E}'[(c_\varepsilon^u(v) \wedge \delta)^2] \leq \frac{\varepsilon}{2} \sum_{k=1}^n \mathbb{E}[Y_k^2] \leq 2\varepsilon. \quad (18)$$

Another application of Lemma 5.2 gives that

$$\sum_{k=1}^{n-m} \mathbb{E}[Y_k \wedge \delta] \leq \mathbb{E}'[H] \leq \sum_{k=1}^n \mathbb{E}[Y_k]. \quad (19)$$

Let $\ell_n = \lceil 2(\log n)^3 \rceil$ and set $E = \{Y_{n-\ell_n} \leq \delta\}$, where $\delta = (\log n)^{-3}$. Then, Lemma 5.1 gives $\mathbb{P}(E^c) \leq 6/(\log n)^3$. Consequently, using Cauchy-Schwartz' inequality and Theorem 1.1, we obtain that

$$\mathbb{E} \left[\sum_{k=1}^n Y_k \mathbf{1}_{E^c} \right] \leq \sqrt{\mathbb{E}[C_{n,n}^2] \mathbb{P}(E^c)} \leq \sqrt{(\text{Var}(C_{n,n}) + \mathbb{E}[C_{n,n}]^2) \mathbb{P}(E^c)} \leq \frac{4}{\sqrt{\log n}}.$$

Hence, for large n ,

$$\sum_{k=1}^{n-m} \mathbb{E}[Y_k \wedge \delta] \geq \mathbb{E} \left[\sum_{k=1}^{n-\ell_n} Y_k \mathbf{1}_E \right] \geq \mathbb{E} \left[\sum_{k=1}^{n-\ell_n} Y_k \right] - \mathbb{E} \left[\sum_{k=1}^{n-\ell_n} Y_k \mathbf{1}_{E^c} \right] \geq \sum_{k=1}^{n-\ell_n} \mathbb{E}[Y_k] - 1,$$

which, combined with (16) and (19), gives

$$\log n - 4 \log \log n \leq \sum_{k=1}^{n-\ell_n} \mathbb{E}[Y_k] - 1 \leq \mathbb{E}'[H] \leq \sum_{k=1}^n \mathbb{E}[Y_k] \leq \log n. \quad (20)$$

Together with (18), this shows that $\mathbb{E}'[N_u] = (1 + o(1))\varepsilon \log n$.

(ii) First note that we can express N_u as

$$N_u = \sum_{v \in L^c} \mathbf{1}_{\{v \in J_u\}} \mathbf{1}_{\{\omega'(u,v) < c_\varepsilon^u(v) \wedge \delta\}}.$$

Expanding the square and conditioning, first on everything but the update variables U_e for edges that connect u to L^c , and then on the weight configuration ω_ε for edges not incident to u , we obtain that

$$\begin{aligned} \mathbb{E}'[N_u^2] &= \varepsilon \sum_{v \in L^c} \mathbb{E}'[F(c_\varepsilon^u(v) \wedge \delta)] + \varepsilon^2 \sum_{\substack{v, v' \in L^c \\ v \neq v'}} \mathbb{E}'[F(c_\varepsilon^u(v) \wedge \delta) F(c_\varepsilon^u(v') \wedge \delta)] \\ &\leq \mathbb{E}'[N_u] + \varepsilon^2 \mathbb{E}' \left[\left(\sum_{v \in L^c} F(c_\varepsilon^u(v) \wedge \delta) \right)^2 \right] \\ &\leq \mathbb{E}'[N_u] + \varepsilon^2 \mathbb{E}' \left[\left(\sum_{v \in L^c} c_\varepsilon^u(v) \wedge \delta \right)^2 \right] \\ &\leq \mathbb{E}'[N_u] + \varepsilon^2 \mathbb{E}'[H^2]. \end{aligned}$$

Combining the above with (18) and (20), we get that

$$\text{Var}'(N_u) = \mathbb{E}'[N_u^2] - \mathbb{E}'[N_u]^2 \leq \mathbb{E}'[N_u] + \varepsilon^2 \left(\text{Var}'(H) + 4 \log n \right). \quad (21)$$

Next, we introduce three events $D_1 = \{Z > b\}$, $D_2 = \{Z \in [a, b]\}$ and $D_3 = \{Z < a\}$, where $a = \sum_{k=1}^{n-\ell_n} \mathbb{E}[Y_k] - 1$ and $b = \sum_{k=1}^n \mathbb{E}[Y_k]$. We bound the variance of H by estimating its contribution restricted to each of the events D_1 , D_2 and D_3 , that is,

$$\text{Var}'(H) = \mathbb{E}'[(H - \mathbb{E}'[H])^2 \mathbf{1}_{D_1}] + \mathbb{E}'[(H - \mathbb{E}'[H])^2 \mathbf{1}_{D_2}] + \mathbb{E}'[(H - \mathbb{E}'[H])^2 \mathbf{1}_{D_3}].$$

First, we note from (20) that $a \leq \mathbb{E}'[H] \leq b$ and that $b - a \leq 4 \log \log n$, which immediately gives

$$\mathbb{E}'[(H - \mathbb{E}'[H])^2 \mathbf{1}_{D_2}] \leq (b - a)^2 \leq 16(\log \log n)^2.$$

Second, by adding and subtracting b , and using that $(x + y)^2 \leq 4x^2 + 4y^2$, we find that

$$\mathbb{E}'[(H - \mathbb{E}'[H])^2 \mathbf{1}_{D_1}] \leq 4\mathbb{E}'[(H - b)^2 \mathbf{1}_{D_1}] + 4(\mathbb{E}'[H] - b)^2.$$

Using Lemma 5.2, and that we have restricted to the event D_1 , we obtain the further upper bound

$$4\mathbb{E}'\left[\left(\sum_{k=1}^n Y_k - b\right)^2 \mathbf{1}_{D_1}\right] + 4(b - a)^2 \leq 4\text{Var}(C_{n,n}) + 64(\log \log n)^2.$$

Third, by adding and subtracting a , and using that $(x + y)^2 \leq 4x^2 + 4y^2$, we find that

$$\mathbb{E}'[(H - \mathbb{E}'[H])^2 \mathbf{1}_{D_3}] \leq 4\mathbb{E}'[(H - a)^2 \mathbf{1}_{D_3}] + 4(\mathbb{E}'[H] - a)^2.$$

Recall that $E = \{Y_{n-\ell_n} \leq \delta\}$. Using Lemma 5.2 and (16), and that we have restricted to the event D_3 , we obtain that

$$\begin{aligned} \mathbb{E}'[(H - a)^2 \mathbf{1}_{D_3}] &\leq \mathbb{E}\left[\left(\sum_{k=1}^{n-\ell_n} Y_k \wedge \delta - a\right)^2 \mathbf{1}_{D_3}\right] \leq \mathbb{E}\left[\left(\sum_{k=1}^{n-\ell_n} Y_k - a\right)^2 \mathbf{1}_{D_3 \cap E}\right] + \mathbb{E}[a^2 \mathbf{1}_{D_3 \cap E^c}] \\ &\leq \text{Var}\left(\sum_{k=1}^{n-\ell_n} Y_k\right) + a^2 \mathbb{P}(E^c) = o(1). \end{aligned}$$

Hence, for large n ,

$$\mathbb{E}'[(H - \mathbb{E}'[H])^2 \mathbf{1}_{D_3}] \leq 1 + 4(b - a)^2 \leq 1 + 64(\log \log n)^2.$$

In conclusion, we get that $\text{Var}'(H) \leq 200(\log \log n)^2$ and hence, via (21) that for n large enough

$$\text{Var}'(N_u) \leq \varepsilon \log n + 5\varepsilon^2 \log n \leq 6\varepsilon \log n.$$

□

5.3 Proof of Theorem 1.5

Recall that for $\varepsilon \in (0, 1]$, we assume that $m \ll \varepsilon \log n$ and $\delta = (\log n)^{-3}$. Let $G_1 = \{Y_{n-m} > \delta, Y_{n-m}^\varepsilon > \delta\}$, where Y_k^ε denotes the k th cheapest edge in the matching $S_{n,n}^\varepsilon$, $G_2 = \{\omega'(u, u') > \delta \text{ for all } u, u' \in L\}$, and $G_3 = \bigcap_{u \in L} \{N_u \geq 1\}$. Finally, set $G = G_1 \cap G_2 \cap G_3$.

We start by bounding the probability that G fails. First, by Lemma 5.1 we have

$$\mathbb{P}(G_1^c) \leq 2\mathbb{P}(Y_{n-m} \leq \delta) \leq 2\mathbb{P}(Y_{n-\lfloor \log n \rfloor} \leq \delta) \leq \frac{24}{\log n}.$$

Second, conditioning on the division (L, L^c) and using the union bound, we have that

$$\mathbb{P}(G_2^c) \leq \mathbb{E}\left[\sum_{u, u' \in L} \mathbb{P}'(\omega'(u, u') \leq \delta)\right] \leq m^2 F(\delta) \leq m^2 \frac{1}{(\log n)^3},$$

where $F(x) = 1 - e^{-x} \leq x$ again denotes the distribution function of the exponential distribution. Third, the union bound, Chebyshev's inequality and Lemma 5.3 give that

$$\mathbb{P}(G_3^c) \leq \mathbb{E} \left[\sum_{u \in L} \mathbb{P}'(N_u = 0) \right] \leq \mathbb{E} \left[\sum_{u \in L} 4 \frac{\text{Var}'(N_u)}{\mathbb{E}'[N_u]^2} \right] \leq 2m \frac{24}{\varepsilon \log n}.$$

In conclusion,

$$\mathbb{P}(G^c) \leq \mathbb{P}(G_1^c) + \mathbb{P}(G_2^c) + \mathbb{P}(G_3^c) \leq \frac{24}{\log n} + \frac{m^2}{(\log n)^3} + \frac{48m}{\varepsilon \log n}, \quad (22)$$

which is $o(1)$ since $m \ll \varepsilon \log n$. Hence G occurs with high probability as $n \rightarrow \infty$.

Note that, on G_1 , we have $\omega(e) > \delta$ for every $e \in L$ and $\omega_\varepsilon(e) > \delta$ for every $e \in L_\varepsilon$. Moreover, on $G_1 \cap G_2$ we have $\omega_\varepsilon(e) > \delta$ for every $e \in L$, that is, all edges originally in L still have cost exceeding δ after perturbation. We claim that, on G_3 , every node in L is matched with cost at most δ after perturbation. Once the claim is proved, we conclude that, on the event G , we have that after perturbation:

- every node in L is rematched with cost at most δ ;
- every edge in L has cost exceeding δ and hence does not belong to the matching $S_{n,n}^\varepsilon$;
- the m most expensive edges have cost exceeding δ , implying that $L_\varepsilon \subseteq L^c$.

It remains to prove the claim that, on G_3 , every node in L is matched with cost at most δ after perturbation. We again argue using a dynamic construction of the matching in which an edge is added to the matching at the 'time' indicated by its cost, unless either of its endpoints has already been matched before that time. It is straightforward to verify that the matching obtained is indeed the stable matching $S_{n,n}^\varepsilon$. In the dynamic construction, a vertex being unmatched at time δ is equivalent to the cost of the vertex exceeding δ . Consequently, if a vertex $u \in L$ is left unmatched at time δ in the perturbed configuration, we have $c_\varepsilon(u) > \delta$. Assume that a vertex u is unmatched at time δ , which implies that the matching obtained at time δ coincides with the matching obtained until time δ when u is removed. In particular, it follows that

$$c_\varepsilon(v) \wedge \delta = c_\varepsilon^u(v) \quad \text{for every } v \in L^c. \quad (23)$$

On G_3 , we have that $N_u \geq 1$ for every $u \in L$. Hence there exists a vertex $v \in L^c$ such that

$$\omega_\varepsilon(u, v) < c_\varepsilon^u(v) \wedge \delta.$$

This contradicts (23), since it implies the existence of a vertex $v \in L^c$ which is unmatched at time $\omega_\varepsilon(u, v) < \delta$, to which u would therefore be matched to, unless it has already been matched before. In conclusion, on G_3 every node in L is matched with cost at most δ after perturbation, as required. This ends the proof of the theorem. \square

6 Noise sensitivity of the stable matching

In this section, we prove Theorem 1.6. The proof will roughly go as follows. We first observe that the bulk of the matching is responsible for most of the matching cost, whereas the cost of the bulk of the bulk of the matching is highly concentrated, so that most of the randomness comes from the tail of the last edges. We then dynamically construct the matchings $S_{n,n}$ and $S_{n,n}^\varepsilon$, by equipping each edge with a Poisson clock and adding it to the corresponding matching when its clock rings, if adding the edge is allowed. By concentration of the bulk of the matching, most edges added are the same in both matchings. However, by Theorem 1.5, we will reach a point in time when the remaining sets of unmatched vertices correspond to disjoint subgraphs. From this point on, we are waiting for independent sets of clocks to ring. The contributions to the matchings obtained from this phase will therefore be independent and, since this phase is responsible for most of the randomness in the construction of the matching, the correlation of the matching costs $C_{n,n}^0$ and $C_{n,n}^\varepsilon$ will be small.

Given $m \geq 1$, denote by $W_m^-(\omega)$ and $W_m^+(\omega)$ the cost of the matching that is detected in the matching of the first $n - m$ and last m edges, respectively. In the notation of (4), we have

$$W_m^-(\omega) = \sum_{k=1}^{n-m} (n - k + 1)X_k \quad \text{and} \quad W_m^+(\omega) = \sum_{k=n-m+1}^n (n - k + 1)X_k.$$

Note that $C_{n,n} = W_m^- + W_m^+$. In particular we find that

$$\mathbb{E}[W_m^-] = \sum_{k=m+1}^n \frac{1}{k} \quad \text{and} \quad \mathbb{E}[W_m^+] = \sum_{k=1}^m \frac{1}{k},$$

and hence that $\mathbb{E}[W_m^-] \sim \log(n/m)$ and $\mathbb{E}[W_m^+] \sim \log m$ for $1 \ll m \ll n$. In addition, we have

$$\text{Var}(W_m^-) = \sum_{k=m+1}^n \frac{1}{k^2} \leq \frac{1}{m} \quad \text{and} \quad \text{Var}(W_m^+) = \sum_{k=1}^m \frac{1}{k^2} \leq \frac{\pi^2}{6}.$$

That is, while little weight remains to be picked up at the end of the matching, most of the randomness comes from that part.

Proof of Theorem 1.6. Fix $m \geq 1$. We first decompose the covariance according to

$$\begin{aligned} \text{Cov}(C_{n,n}^0, C_{n,n}^\varepsilon) &= \text{Cov}(W_m^-(\omega), W_m^-(\omega_\varepsilon)) + \text{Cov}(W_m^-(\omega), W_m^+(\omega_\varepsilon)) \\ &\quad + \text{Cov}(W_m^+(\omega), W_m^-(\omega_\varepsilon)) + \text{Cov}(W_m^+(\omega), W_m^+(\omega_\varepsilon)). \end{aligned} \tag{24}$$

Since $\text{Var}(W_m^-) \leq \frac{1}{m}$, an application of Cauchy-Schwartz gives

$$|\text{Cov}(W_m^-(\omega), W_m^-(\omega_\varepsilon))| \leq \text{Var}(W_m^-) \leq \frac{1}{m},$$

and

$$|\text{Cov}(W_m^-(\omega), W_m^+(\omega_\varepsilon))| \leq \sqrt{\text{Var}(W_m^-)\text{Var}(W_m^+)} \leq \frac{2}{\sqrt{m}}.$$

This gives

$$|\text{Cov}(C_{n,n}^0, C_{n,n}^\varepsilon)| \leq |\text{Cov}(W_m^+(\omega), W_m^+(\omega_\varepsilon))| + \frac{5}{\sqrt{m}}. \quad (25)$$

Write T for the time at which the two subgraphs induced by the unmatched nodes in the two configurations ω and ω_ε become disjoint. Let

$$Q := \{T \leq \min\{Y_{n-m}, Y_{n-m}^\varepsilon\}\}$$

and note that, on the event Q , when matching the last m edges, we are waiting for disjoint sets of Poisson clocks to ring. We next show that Q occurs with high probability. Set $\ell = \sqrt{\varepsilon \log n}$ and let Q_1 denote the event that the subgraphs induced by the vertices of the last 7ℓ edges of the matching in ω and ω_ε are disjoint. In addition, let

$$Q_2 = \left\{ \max\{Y_{n-7\ell}, Y_{n-7\ell}^\varepsilon\} \leq \frac{1}{6\ell} \right\},$$

$$Q_3 = \left\{ \min\{Y_{n-\ell}, Y_{n-\ell}^\varepsilon\} \geq \frac{1}{6\ell} \right\}.$$

By assumption, we have $\varepsilon \log n \gg 1$, so that $m \leq \ell$ when n is large. It follows that, on $Q_1 \cap Q_2 \cap Q_3$, we have for large n that

$$T \leq \max\{Y_{n-7\ell}, Y_{n-7\ell}^\varepsilon\} \leq \min\{Y_{n-\ell}, Y_{n-\ell}^\varepsilon\} \leq \min\{Y_{n-m}, Y_{n-m}^\varepsilon\},$$

so that $Q_1 \cap Q_2 \cap Q_3 \subseteq Q$. Note that the probability of Q_1^c can be upper bounded using the quantitative bound (22) leading to Theorem 1.5, while the probabilities of Q_2^c and Q_3^c can be bounded using Lemma 5.1. Hence

$$\mathbb{P}(Q^c) \leq \mathbb{P}(Q_1^c) + \mathbb{P}(Q_2^c) + \mathbb{P}(Q_3^c) \leq \frac{350}{\sqrt{\varepsilon \log n}} + 2\frac{12}{7\ell} + 2\frac{12}{\ell} \leq \frac{400}{\sqrt{\varepsilon \log n}} \quad (26)$$

for sufficiently large n .

Decomposing the covariance depending on the event Q gives

$$\begin{aligned} \text{Cov}(W_m^+(\omega), W_m^+(\omega_\varepsilon)) &= \mathbb{E}[(W_m^+(\omega) - \mathbb{E}[W_m^+]) (W_m^+(\omega_\varepsilon) - \mathbb{E}[W_m^+]) \mathbf{1}_Q] \\ &\quad + \mathbb{E}[(W_m^+(\omega) - \mathbb{E}[W_m^+]) (W_m^+(\omega_\varepsilon) - \mathbb{E}[W_m^+]) \mathbf{1}_{Q^c}]. \end{aligned} \quad (27)$$

Note that Q depends on the Poisson clocks in ω and ω_ε before time $\min\{Y_{n-m}, Y_{n-m}^\varepsilon\}$, whereas $W_m^+(\omega)$ and $W_m^+(\omega_\varepsilon)$ depend on the clocks after time $\min\{Y_{n-m}, Y_{n-m}^\varepsilon\}$. Moreover, on Q , we have that $W_m^+(\omega)$ and $W_m^+(\omega_\varepsilon)$ are functions of disjoint sets of clocks. It follows that, on Q , we have

$$(W_m^+(\omega), W_m^+(\omega_\varepsilon)) \stackrel{d}{=} (W_m^+(\omega), W_m^+(\tilde{\omega})),$$

and hence that

$$\mathbb{E}[(W_m^+(\omega) - \mathbb{E}[W_m^+]) (W_m^+(\omega_\varepsilon) - \mathbb{E}[W_m^+]) \mathbf{1}_Q] = \mathbb{E}[(W_m^+(\omega) - \mathbb{E}[W_m^+]) (W_m^+(\tilde{\omega}) - \mathbb{E}[W_m^+]) \mathbf{1}_Q],$$

where $\tilde{\omega}$ indicates a cost configuration independent from ω . Using that $xy \leq x^2 + y^2$, on Q^c we get

$$\mathbb{E}[(W_m^+(\omega) - \mathbb{E}[W_m^+]) (W_m^+(\omega_\varepsilon) - \mathbb{E}[W_m^+]) \mathbf{1}_{Q^c}] \leq 2 \mathbb{E}[(W_m^+(\omega) - \mathbb{E}[W_m^+])^2 \mathbf{1}_{Q^c}]$$

and similarly

$$\mathbb{E}[(W_m^+(\omega) - \mathbb{E}[W_m^+])(W_m^+(\tilde{\omega}) - \mathbb{E}[W_m^+])\mathbf{1}_{Q^c}] \leq 2\mathbb{E}[(W_m^+(\omega) - \mathbb{E}[W_m^+])^2\mathbf{1}_{Q^c}].$$

From (27) we obtain

$$|\text{Cov}(W_m^+(\omega), W_m^+(\omega_\varepsilon))| \leq |\text{Cov}(W_m^+(\omega), W_m^+(\tilde{\omega}))| + 4\mathbb{E}[(W_m^+(\omega) - \mathbb{E}[W_m^+])^2\mathbf{1}_{Q^c}].$$

Note that, for fixed m and for $\varepsilon = \varepsilon(n)$ such that $\varepsilon \log n \rightarrow \infty$ as $n \rightarrow \infty$, we have from (26) that $\mathbb{P}(Q^c) = o(1)$. Moreover, $W_m^+(\omega)$ and $W_m^+(\tilde{\omega})$ are independent and $W_m^+(\omega)$ equals $C_{m,m}$ in distribution. It hence follows from dominated convergence (or the reverse Fatou lemma) that

$$\text{Cov}(W_m^+(\omega), W_m^+(\omega_\varepsilon)) \rightarrow 0.$$

Since m was arbitrary, we conclude from (25) that

$$\text{Cov}(C_{n,n}^0, C_{n,n}^\varepsilon) \rightarrow 0.$$

This completes the proof, since $\text{Var}(C_{n,n}) \rightarrow \pi^2/6$ as $n \rightarrow \infty$. \square

Acknowledgement.

The authors thank Svante Janson for discussions during the early part of this project, and in particular for indicating the correct limit law for the matching cost. The first author further thanks Marcelo Campos, Simon Griffiths and Rob Morris for discussions regarding the possible sensitivity of the tail of the matching. This work was in part supported by the Swedish Research Council (VR) through grant 2021-03964 (DA) and 2020-04479 (MD and MS).

References

- [1] D. Ahlberg, M. Deijfen, M. Sfragara. Chaos, concentration and multiple valleys in first-passage percolation, [arXiv:2302.11367], 2023.
- [2] D.J. Aldous. Asymptotics in the random assignment problem. *Probability Theory and Related Fields* 93, 507–534, 1992.
- [3] D.J. Aldous. The $\zeta(2)$ limit in the random assignment problem. *Random Structures and Algorithms* 18, 381–418, 2001.
- [4] D. Aldous, J.M. Steele. The objective method: probabilistic combinatorial optimization and local weak convergence. In *Probability on discrete structures*, volume 110 of *Encyclopaedia of Mathematical Sciences*, 1–72. Springer, Berlin, 2004.
- [5] I. Benjamini, G. Kalai, O. Schramm. Noise sensitivity of Boolean function and applications to percolation. In *Publication Mathématiques de l’Institut des Hautes Études Scientifiques* 90, 5–43, 1999.

- [6] C. Bordenave, G. Lugosi, N. Zhivotovskiy. Noise sensitivity of the top eigenvector of a Wigner matrix. *Probability Theory and Related Fields* 177, 1103–1135, 2020.
- [7] S. Chatterjee. *Superconcentration and related topics*. Springer Monographs in Mathematics. Springer, Cham, 2014.
- [8] Gale, D. and Shapely, L. College admissions and stability of marriage. *American Mathematical Monthly* 69, 9–15, 1962.
- [9] S. Ganguly, A. Hammond. Stability and chaos in dynamical last passage percolation, [arXiv:2010.05837], 2020.
- [10] D. Gusfield and R. W. Irving. *The stable marriage problem: Structure and algorithms*. MIT Press, 1989.
- [11] A.E. Holroyd, J.B. Martin, and Y. Peres. Stable matchings in high dimensions via the Poisson-weighted infinite tree. *Annales Institut Henri Poincaré (B) Probability and Statistics* 56, 826–846, 2020.
- [12] R. W. Irving. An efficient algorithm for the “stable roommates” problem. *Journal of Algorithms* 6, 577–595, 1985.
- [13] O. Israeli, Y. Peled. Noise sensitivity of the minimal spanning tree of the complete graph, [arXiv:2306.07357], 2023.
- [14] D. Knuth. *Stable marriage and its relation to other combinatorial problems: An introduction to the mathematical analysis of algorithms* CMR Lecture Notes vol **10**, American Mathematical Society, 1996.
- [15] S. Linusson, and J. Wästlund. A proof of Parisi’s conjecture on the random assignment problem. *Probability Theory and Related Fields* 128, 419–440, 2004.
- [16] D. Manlove. *Algorithmics of matching under preferences* World Scientific, 2013.
- [17] M. Mézard, and G. Parisi. On the solution of the random link matching problems. *Journal de Physique* 48, 1451–1459, 1987.
- [18] G. Parisi, and J. Wästlund. Mean field matching and traveling salesman in pseudo-dimension 1. [arXiv:2302.11367], 2017.
- [19] G. Parisi. A conjecture on random bipartite matching. [arxiv.org/abs/cond-mat/9801176], 1998.
- [20] J.M. Steele *Probability and combinatorial optimization*. No 69 in CBMS-NSF Regional Conference Series in Applied Math. SIAM, 1997.
- [21] J. Wästlund. An easy proof of the $\zeta(2)$ limit in the random assignment problem. *Electronic Communication in Probability* 14, 261–269, 2009.
- [22] J. Wästlund. Replica symmetry of the minimum matching. *Annals of Mathematics* 175, 1061–1091, 2012.