Asymptotics in the random assignment problem

David Aldous *

Department of Statistics, University of California, Berkeley, CA 94720, USA

Summary. We show that, in the usual probabilistic model for the random assignment problem, the optimal cost tends to a limit constant in probability and in expectation. The method involves construction of an infinite limit structure, in terms of which the limit constant is defined. But we cannot improve on the known numerical bounds for the limit.

Mathematics Subject Classification: 60C05

1 Introduction

In the deterministic assignment problem, there are n jobs, n machines and a $n \times n$ non-negative matrix $(t_{j,m})$ representing the cost of performing job j on machine m. An *assignment* is a permutation π of $\{1, \ldots, n\}$, indicating that job j is assigned to machine $\pi(j)$. The optimal assignment has cost $\min_{\pi} \sum_{j} t_{j,\pi(j)}$, the minimum taken over permutations π. For the *random assignment problem* we define the $t_{j,m}$ to be independent r.v.'s with exponential distribution with mean n:

(1) $$P(t_{j,m} > x) = \exp(-x/n), \quad 0 \leq x < \infty$$

and we study the average cost per job in the optimal assignment, that is to say the random variable

(2) $$C(n) = \frac{1}{n} \min_{\pi} \sum_{j} t_{j,\pi(j)}.$$

It is traditional to study instead the total cost of all jobs, which corresponds to using exponential (mean 1) distributions. It is also traditional to use the uniform distribution on $[0, 1]$ but since it is only the density at 0 that ultimately matters, our set-up is asymptotically equivalent to the usual one.

Although modelling the costs as i.i.d. variables seems unrealistic in many contexts, the determination of asymptotic behavior of $EC(n)$ in this model has emerged as a challenging mathematical problem. The survey paper Steele [11] is devoted to this problem, and we refer the reader there for more background. The rigorous results known already can be summarized as

$$1 + e^{-1} + \imath \leq \liminf_{n \to \infty} EC(n) \leq \limsup_{n \to \infty} EC(n) \leq 2$$

where $\imath$ is a small explicit constant. The upper bound is due to Karp [7], who sharpened earlier results of Walkup [12] and whose technique is further developed in Dyer et al. [4]. The lower bound is due to Goemans and Kodialam [5]. The purpose of this paper is to prove

Theorem 1 *As $n \to \infty$*

$$E|C(n) - c^*| \to 0$$

for a certain constant c^* defined at (17).

Our argument uses the fact that $\limsup_{n \to \infty} EC(n) < \infty$ as well as ingredients of Walkup's [12] proof of this fact. Unfortunately our proof does not give any new numerical information about c^*.

To outline the argument, for a non-negative $n \times n$ matrix Q define a type of discriminant function

(3)
$$\chi(Q) = \frac{1}{n} \sum_j |1 - \sum_m q_{j,m}| + \frac{1}{n} \sum_m |1 - \sum_j q_{j,m}|.$$

Given the random cost matrix $T = (t_{j,m})$ defined at (1), we further define

$$c(n, \varepsilon) = \min_{Q:\, E\chi(Q) \leq \varepsilon} E\left(\frac{1}{n} \sum_j \sum_m q_{j,m} t_{j,m}\right)$$

where the minimum is over all random non-negative matrices Q. Note that the minimizing matrix Q will be strongly dependent on T.

The property $\chi(Q) = 0$ means Q is doubly-stochastic. So Birkhoff's theorem (doubly-stochastic matrices are mixtures of permutation matrices) says that $EC(n) = c(n, 0)$. In other words, replacing the assignment problem by its "continuous relaxation" makes no difference. For $\varepsilon > 0$ we have $c(n, \varepsilon) \leq EC(n)$, and the "discrete half" of Theorem 1 is the following result, which says roughly that one can use *almost* doubly-stochastic matrices.

Proposition 2 $\lim_{\varepsilon \to 0} \limsup_{n \to \infty} (EC(n) - c(n, \varepsilon)) = 0$.

This result is proved in Sect. 2, and though the ingredients are standard, combining them effectively is not entirely trivial.

The conceptually new development of this paper is that (for the purposes of the random assignment problem) there is a well-defined "limit random object" representing the $n \to \infty$ limit of the $n \times n$ cost matrices. The limit is an infinite tree with random costs associated with edges. The detailed description is in

Sect. 3, but here is an outline of the idea. The root of the tree represents a random job; the first generation offspring represent the machines to which that job could be assigned with small cost, with the costs marked on the edge; the second generation offspring represent other jobs which could be assigned to the first generation machines at small cost, and so on. Limits of assignments in the finite problem correspond to matchings of this tree which satisfy a certain compatibility condition. So we can define a constant c^* as the minimum, over all compatible matchings of the infinite tree, of the expected cost of the assignment of the "root" job to its machine. Once this structure is set up, "soft" arguments in Sect. 3 lead to the following result.

Proposition 3 (a) $c^* \leq \liminf_{n \to \infty} EC(n)$.

(b) *For fixed* $\varepsilon > 0$, $\limsup_{n \to \infty} c(n, \varepsilon) \leq c^*$.

The "convergence of expectations" part of Theorem 1 is immediate from Propositions 2 and 3, and in Sect. 3.5 we show that convergence in probability follows from an ergodicity property of the limit tree.

Remarks. 1. Proposition 2 might be useful in the context of analyzing some explicit algorithm for non-optimal assignment in order to improve the known upper bound on c^*.
2. The general technique of proving limit theorems for discrete structures by exhibiting a limit random object has been called the "objective method" by Mike Steele. See [2, 1, 3] for other recent examples. In many applications of this technique it is easy to see intuitively what the limit object is; in this application the existence and nature of a useful limit object seem less intuitively obvious.
3. Steel [11] observes that a natural "global greedy algorithm" produces a non-optimal matching with cost $\sim \log n$, although a naive analysis that ignores the effect of conditioning at each stage suggests (incorrectly) that the expected cost of this greedy matching is $\sum_{i=1}^{n} \frac{1}{(n+1-i)^2} \sim \pi^2/6$. Mezard and Parisi [8] give a non-rigorous argument based on group renormalization techniques and claim that $c^* = \pi^2/6$. Avram and Bertsimas have expanded these ideas to give a detailed outline of an argument for Theorem 1 which includes a plausible explicit expression for c^*, but a complete proof has not been presented at the time of writing.
4. It has been conjectured that $EC(n)$ is increasing in n, and an elementary proof of this fact would of course establish part of Theorem 1. No useful bounds on the variance of $C(n)$ are known, but perhaps the modern martingale methods which have proved useful in other combinatorial and algorithmic contexts (e.g. Rhee and Talagrand [10]) could be applied to the assignment problem.

2 Proof of Proposition 2

We collect in Lemmas 4–6 some tools from previous analyses of the random assignment problem. Then we give our analysis in Proposition 7 and 9, and at the end of the section we show how Proposition 2 follows from these.

The first tool is the matching lemma of Walkup [13].

Lemma 4 *In the uniform random bipartite digraph with k vertices in each partite class and with out-degree 2 at each vertex, the probability that there exists a perfect matching tends to 1 as $k \to \infty$.*

Second, by examining an existing proof (e.g. Walkup [12]) that $EC(n)$ is bounded in n, we see that the argument establishes the slightly stronger *uniform integrability* property.

Lemma 5 *There exist matchings $\hat{\pi}_n$ with costs $\hat{C}(n)$, and a function $\delta(\theta) \to 0$ as $\theta \to 0$ such that, for arbitrary events Ω_n,*

$$\limsup_{n \to \infty} E\hat{C}(n) \, 1_{\Omega_n} \leq \delta(\varepsilon)$$

where $\varepsilon = \limsup\limits_{n \to \infty} P(\Omega_n)$.

The third ingredient is elementary, but it forms the basis of "independent splitting" arguments.

Lemma 6 *If ξ_1 and ξ_2 are independent r.v.'s with exponential distributions with means $1/\lambda_1$ and $1/\lambda_2$, then $\min(\xi_1, \xi_2)$ has exponential distribution with mean $1/(\lambda_1 + \lambda_2)$.*

By a *partial assignment* π we mean an assignment of some subset $U(\pi)$ of jobs to different machines. If $|U(\pi)| \geq (1 - \theta)n$ we call π a $1 - \theta$ partial assignment.

Proposition 7 *Let Q and T be given non-random non-negative $n \times n$ matrices. Suppose $200 \, \chi(Q) \leq \theta < 1$, for $\chi(Q)$ defined at (3). Then there exists a $1 - \theta$ partial assignment π_0 such that*

$$\sum_{j \in U(\pi_0)} q_{j, \pi_0(j)} \, t_{j, \pi_0(j)} \leq (1 + 4\sqrt{\chi(Q)/\theta}) \sum_j \sum_m q_{j,m} \, t_{j,m}.$$

Proof. Define

$$q_{j,\cdot} = \sum_m q_{j,m}, \qquad q_{\cdot,m} = \sum_j q_{j,m}, \qquad a_{j,m} = \frac{q_{j,m}}{\max(1, q_{j,\cdot}) \max(1, q_{\cdot,m})}.$$

We then have

$$a_{j,\cdot} = \sum_m a_{j,m} \leq 1, \qquad a_{\cdot,m} = \sum_j a_{j,m} \leq 1$$

and we set, assuming $s > 0$,

$$b_{j,m} = \frac{(1 - a_{j,\cdot})(1 - a_{\cdot,m})}{s}, \qquad \text{for } s = \sum_j (1 - a_{j,\cdot}) = \sum_m (1 - a_{\cdot,m}) = n - \sum_j \sum_m a_{j,m}.$$

It is easy to check that A and B are non-negative matrices and $A + B$ is doubly stochastic. The next lemma shows the average entry of B is small when $\chi(Q)$ is small.

Lemma 8 $\bar{b} \equiv \dfrac{1}{n} \sum_j \sum_m b_{j,m} \le 3\chi(Q).$

Proof. First note $1/\max(1, x) \ge 1 - H(x)$, where

$$
\begin{aligned}
H(x) &= 0, \quad x \le 1 \\
&= x - 1, \quad 1 \le x \le 2 \\
&= 1, \quad x \ge 2.
\end{aligned}
$$

From the definition of $a_{j,m}$, we obtain

$$
a_{j,m} \ge q_{j,m}(1 - H(q_{j,.}) - H(q_{.,m})),
$$

so

$$
\begin{aligned}
\bar{b} &= \frac{1}{n} \sum_j \sum_m b_{j,m} \\
&= \frac{1}{n} \left(n - \sum_j \sum_m a_{j,m} \right) \\
&\le \frac{1}{n} \left(n - \sum_j \sum_m q_{j,m} \right) + \frac{1}{n} \sum_j \sum_m q_{j,m}(H(q_{j,.}) + H(q_{.,m})) \\
&= \frac{1}{n} \sum_j \left(1 - \sum_m q_{j,m} \right) + \frac{1}{n} \sum_j q_{j,.} H(q_{j,.}) + \frac{1}{n} \sum_m q_{.,m} H(q_{.,m}).
\end{aligned}
$$

Writing the definition of $\chi(Q)$ as $\chi(Q) = \chi_1 + \chi_2$, the first term above is bounded by χ_1. Since $xH(x) \le 2(x-1)^+$, the second term is bounded by

$$
\frac{2}{n} \sum_j (q_{j,.} - 1)^+ \le 2\chi_1.
$$

Similarly the third term is bounded by $2\chi_2$, establishing the lemma.

Returning to the proof of Proposition 7, choose $\eta = \sqrt{3\chi/\theta}$ and define

$$
\begin{aligned}
\Delta(j, m) &= 1 \quad \text{if } b_{j,m} \le \eta\, a_{j,m} \\
&= 0 \quad \text{if not.}
\end{aligned}
$$

By Birkhoff's theorem there exists a random permutation π such that

$$
P(\pi(j) = m) = a_{j,m} + b_{j,m}.
$$

Let

$$
U(\pi) = \{ j : \Delta(j, \pi(j)) = 1 \} = \{ j : b_{j,\pi(j)} \le \eta\, a_{j,\pi(j)} \}.
$$

Then for fixed j

$$E q_{j,\pi(j)} \Delta(j,\pi(j)) t_{j,\pi(j)} = \sum_m (a_{j,m} + b_{j,m}) \Delta(j,m) t_{j,m}$$

$$\leq (1+\eta) \sum_m a_{j,m} \Delta(j,m) t_{j,m}$$

$$\leq (1+\eta) \sum_m q_{j,m} t_{j,m}.$$

Summing over j,

$$E \sum_{j \in U(\pi)} q_{j,\pi(j)} t_{j,\pi(j)} \leq (1+\eta) \bar{t}, \quad \text{where } \bar{t} = \sum_j \sum_m q_{j,m} t_{j,m}.$$

Then by Markov's inequality, for any $0 < \delta < 1$,

(4) $$P\left(\sum_{j \in U(\pi)} q_{j,\pi(j)} t_{j,\pi(j)} \leq \frac{(1+\eta)\bar{t}}{1-\delta}\right) \geq \delta.$$

On the other hand, for fixed j

$$E(1 - \Delta(j,\pi(j))) = \sum_m (a_{j,m} + b_{j,m})(1 - \Delta(j,m))$$

$$\leq \sum_m (1+\eta^{-1}) b_{j,m}$$

and so, averaging over j and using Lemma 8,

$$E \frac{1}{n} \sum_j (1 - \Delta(j,\pi(j))) \leq (1+\eta^{-1}) \bar{b} \leq 3(1+\eta^{-1}) \chi(Q).$$

Now the left side equals $E(1 - |U(\pi)|/n)$ and so using Markov's inequality

(5) $$P(|U(\pi)| \geq (1-\theta)n) \geq 1 - \frac{3(1+\eta^{-1}) \chi(Q)}{\theta}.$$

Choosing $\delta = 3(2+\eta^{-1}) \chi(Q)/\theta$, the right sides of (4) and (5) sum to more than 1. So there must exist some permutation π_0 such that

(6) $$|U(\pi_0)| \geq (1-\theta)n$$

$$\sum_{j \in U(\pi_0)} q_{j,\pi_0(j)} t_{j,\pi_0(j)} \leq \frac{(1+\eta)\bar{t}}{1-\delta}.$$

The former says that π_0 is a $1-\theta$ partial assignment. To estimate the bound in (6) note

$$\frac{1+\eta}{1-\delta} \leq (1-\eta-\delta)^{-1} = (1 - 6\chi(Q)/\theta - 2\sqrt{3\chi(Q)/\theta})^{-1}$$

by choice of η and δ. It is straightforward to bound this by $(1+4\sqrt{\chi(Q)/\theta})$ under the assumption of the proposition, so the proof of Proposition 7 is complete.

Proposition 9 *Fix $0<\theta<1/10$. Let $\mathbf{J}$ be the set of n jobs and let $\mathbf{M}$ be the set of n machines. Let T be the random cost matrix* (1). *Let $k=[\theta n]$. Let $\mathbf{J}_0$ be a subset of jobs of size k, let $\mathbf{M}_0$ be a subset of machines of size k, and let π_0 be a bijection $\mathbf{J}\backslash\mathbf{J}_0 \rightarrow \mathbf{M}\backslash\mathbf{M}_0$. (Here $\mathbf{J}_0$, $\mathbf{M}_0$ and π_0 may depend on T in an arbitrary way.) Then there exists a random subset $S \subset \mathbf{J}\backslash\mathbf{J}_0$ of size k, random bijections*

$$\pi_1 : S \rightarrow \mathbf{M}_0, \qquad \pi_2 : \mathbf{J}_0 \rightarrow \pi_0(S)$$

and events Ω_n with $P(\Omega_n) \rightarrow 0$ such that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} E 1_{\Omega_n^c} \left(\sum_{j \in \mathbf{J}_0} t_{j, \pi_2(j)} + \sum_{j \in S} t_{j, \pi_1(j)} \right) \leq 20 \theta^{1/2}.$$
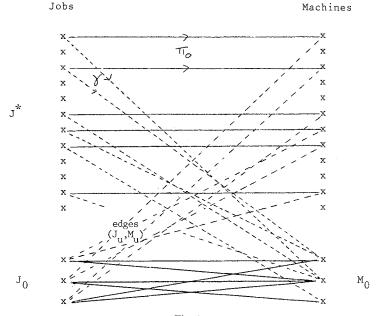

Fig. 1

Proof. The reader will find it helpful to consult Fig. 1. For each $j \in \mathbf{J}\backslash\mathbf{J}_0$ let

$$(7) \qquad X_j = \min_{m \in \mathbf{M}_0} t_{j,m}, \qquad y(j) = \arg\min_{m \in \mathbf{M}_0} t_{j,m}.$$

Let $\mathbf{J}^*$ be the set of j such that X_j is amongst the $[2\theta^{1/2}n]$ smallest of the X's. Since the X_j are independent in j and have exponential, mean n/k, distribution, it is easy to see

$$\lim_n E \frac{1}{|\mathbf{J}^*|} \sum_{j \in \mathbf{J}^*} X_j = \theta^{-1} E \xi 1_{(\xi \leq c)}$$

where ξ has exponential (1) distribution and c satisfies $P(\xi \leq c) = 2\theta^{1/2}$

$$(8) \qquad = \theta^{-1} \int_0^{2\theta^{1/2}} \log \frac{1}{1-z} \, dz \leq 3\theta^{-1/2}$$

the last inequality because θ is assumed small.

Now consider the array $\{t_{j,m} : j \in \mathbf{J}_0, \, m \in \pi_0(\mathbf{J}^*)\}$ which (conditional on $\mathbf{J}^*$) is independent of the X's above. We want to pick pairs $((J_u, M_u) : 1 \leq u \leq 2k)$ such that the t_{J_u, M_u} are small and such that each $j \in \mathbf{J}_0$ is picked exactly twice. We do this via the obvious greedy algorithm:

$$(J_u, M_u) = \arg \min_{j, m} t_{j,m}$$

where $(j, m) \in \mathbf{J}_0 \times \pi_0(\mathbf{J}^*)$ is constrained by

$$m \notin \{M_1, \ldots, M_{u-1}\}$$
$$|\{v < u : J_v = j\}| \leq 1.$$

Let $Y_u = t_{J_u, M_u}$ be the associated cost. We shall show

$$(9) \qquad \limsup_{n \to \infty} \frac{1}{n} E \sum_{u=1}^{2k} Y_u \leq 3\theta^{1/2}.$$

Consider the increment $Y_{2u+1} - Y_{2u}$. Conditional on the operation of the algorithm through the $2u$'th choice, there are at least $k - u$ j's and at least $[2\theta^{1/2}n] - 2k$ m's satisfying the constraints imposed upon the next step of the greedy algorithm. Using the memoryless property of the exponential distribution, we see that $Y_{2u+1} - Y_{2u}$ is stochastically smaller than the minimum of $(k - u)([2\theta^{1/2}n] - 2k)$ independent exponential (mean n) r.v.'s, and hence

$$E(Y_{2u+1} - Y_{2u}) \leq \frac{n}{(k-u)([2\theta^{1/2}n] - 2k)}.$$

The same inequality holds for $Y_{2u+2} - Y_{2u+1}$. Writing

$$\sum_{u=1}^{2k} Y_u \leq \sum_{u=0}^{k-1} 2(k-u)(Y_{2u+2} - Y_{2u})$$

we see that the left side of (9) is bounded by

$$\frac{4k}{[2\theta^{1/2}n] - 2k}$$

and the bound in (9) follows.

We now specify a random diagraph with edges from $\mathbf{J}_0$ to $\mathbf{M}_0$. Recall the definition (7) of $\gamma(j)$. For each $1 \leq u \leq 2k$, create a directed edge $(J_u, \gamma(\pi_0^{-1}(M_u)))$, and give the edge "weight"

$$W_u = t_{J_u, M_u} + t_{\pi_0^{-1}(M_u), \, \gamma(\pi_0^{-1}(M_u))} = Y_u + X_{\pi_0^{-1}(M_u)}.$$

These are the edges at the bottom in Fig. 1. There are exactly two edges leading out of each vertex $j \in \mathbf{J}_0$. Using (8) and (9), the total weight satisfies

$$(10) \qquad \qquad \limsup_{n \to \infty} \frac{1}{n} E \sum_{u=1}^{2k} W_u \leqq 5 \theta^{1/2}.$$

By Lemma 6 the matrix T may be represented as $t_{j,m} = \min(2t_{j,m}^1, 2t_{j,m}^2)$ where T^1 and T^2 are independent cost matrices as at (1). Carry out the above construction of a graph directed from $\mathbf{J}_0$ to $\mathbf{M}_0$ using the cost matrix T^1, then repeat the construction using T^2 to define edges directed from $\mathbf{M}_0$ to $\mathbf{J}_0$. This defines a random bipartite digraph G_n. By (10) the sum $\mathbf{W}_n$ of the weights of the edges of G_n satisfies

$$(11) \qquad \qquad \limsup_{n \to \infty} \frac{1}{n} E \mathbf{W}_n \leqq 20 \theta^{1/2}.$$

To be slightly dishonest for a moment, pretend that G_n is distributed as the random graph in Lemma 4. Then that lemma (with k in place of n) implies that, outside an event whose probability tends to 0, there is a perfect matching, that is a bijection $\mu : \mathbf{J}_0 \to \mathbf{M}_0$ such that each $(j, \mu(j))$ is an edge of G_n, directed one way or the other. Note that if the matching connects an edge directed $j \to \mu(j)$, this edge is of the form $(J_u, \gamma(\pi_0^{-1}(M_u)))$ for some u, and thus specifies two correspondences

$$J_u \leftrightarrow M_u, \qquad \pi_0^{-1}(M_u) \leftrightarrow \gamma(\pi_0^{-1}(M_u))$$

which are the diagonal maps in Fig. 1. And similarly, an edge of the matching which is directed as $j \leftarrow \mu(j)$ specifies two such correspondences. Thus the k edges in the matching μ specify $2k$ such maps from jobs to machines, which define bijections π_1 and π_2 with the properties asserted in the proposition.

    To be honest, G_n is the following bipartite digraph on classes $\mathbf{J}_0$ and $\mathbf{M}_0$. The out-edges from vertex v go to v^* and v^{**}, say. The choice of (v^*, v^{**}) is independent as v varies, and for fixed v we choose v^* and v^{**} by making two uniform random draws *with* replacement. In the model of Lemma 4 the two draws are made *without* replacement. It seems plausible that Lemma 4 remains true under our model, but instead let us indicate a simple patching operation to complete the proof of Proposition 7. Note first that the mean number of v for which $v^{**} = v^*$ tends to 2 as $k \to \infty$. In the construction we have given, set aside at the beginning some subset $\mathbf{J}'$ of jobs and the subset $\pi(\mathbf{J}')$ of machines, where $\mathbf{J}'$ has size a_n such that $a_n \to \infty$ but $a_n/n \to 0$. Not using these jobs and machines does not affect (11). For each $v \in \mathbf{J}_0$ (and similarly in $\mathbf{M}_0$) with $v^{**} = v^*$, choose $m(v) \in \pi_0(\mathbf{J}')$ and $g(v) \in \mathbf{M}_0$ to minimize $W_v' = t_{v, m(v)} + t_{\pi_0(m(v)), g(v)}$, and add to G_n the edge $(v, g(v))$. It is straightforward to check that, outside an event of probability tending to zero, the resulting graph G_n' has the distribution specified in Lemma 4; and that $E \sum_v W_v' = O(n/a_n)$. So applying the previous argument to G_n' completes the proof.

Proof of Proposition 2. Here's the idea. Suppose we have a partial assignment π_0, and let $\mathbf{J}_0$ and $\mathbf{M}_0$ be the unassigned jobs and machines. Suppose also that we have a subset S of assigned jobs and bijections $\pi_1 : S \to \mathbf{M}_0$ and $\pi_2 :$

$\mathbf{J}_0 \to \pi_0(S)$. Then we can construct a complete assignment by assigning, for each $j \in \mathbf{J}_0$, job j to machine $\pi_2(j)$ and job $j' = \pi_0^{-1}(\pi_2(j)) \in S$ to machine $\pi_1(j') \in \mathbf{M}_0$. The other jobs j remain assigned to machine $\pi_0(j)$. The cost of this complete matching is less than the cost of the original partial matching plus the costs of the partial matching defined by the bijections π_1 and π_2. We shall apply this idea with the partial matching given by Proposition 7 and the bijections given by Proposition 9. Here are the details.

Fix $0 < \alpha < 1$ and represent the cost matrix T as

$$(12) \qquad t_{j,m} = \min \left(\frac{t_{j,m}^1}{1-\alpha}, \frac{t_{j,m}^2}{\alpha} \right)$$

where T^1 and T^2 are independent copies of the random cost matrix (1).

Fix $0 < \theta < 1$ and let $\varepsilon = \theta^3/200$. Let Q achieve the minimum in the definition of $c(n, \varepsilon)$, where we use the cost matrix T^1. Then

$$P(200\chi(Q) > \theta^2) \leqq \frac{200 \, E\chi(Q)}{\theta^2} \leqq \frac{200\varepsilon}{\theta^2} = \theta$$

and so, outside some event Ω_n^* of probability at most θ, we have $200\chi(Q) \leqq \theta^2 < \theta$. By Proposition 7, outside Ω_n^* there exists a random $1-\theta$ partial assignment π_0 such that

$$(13) \qquad \frac{1}{n} E 1_{\Omega_n^{*c}} \sum_{j \in \mathbf{J} \setminus \mathbf{J}_0} q_{j,\pi_0(j)} \, t_{j,\pi_0(j)}^1 \leqq (1+\theta^{1/2}) \, c(n, \varepsilon)$$

where $\mathbf{J}_0$ are the unmatched jobs. Now condition on a realization (outside Ω_n^*) of T^1 and π_0, and apply Proposition 9 to T^2. Proposition 9 says there exist asymptotically null sets Ω_n^{**} and random bijections π_1, π_2 such that

$$(14) \qquad \limsup_{n \to \infty} \frac{1}{n} E 1_{\Omega_n^{**c}} \left(\sum_{j \in \mathbf{J}_0} t_{j,\pi_2(j)}^2 + \sum_{j \in \pi_0^{-1}(\pi_2(\mathbf{J}_0))} t_{j,\pi_1(j)}^2 \right) \leqq 20\theta^{1/2}.$$

Outside $\Omega_n = \Omega_n^* \cup \Omega_n^{**}$ we can construct a complete matching π as described at the beginning of the proof, and on Ω_n we use the matching of Lemma 5. The average cost per machine C_n^+ of this matching satisfies, using (12, 13, 14),

$$\limsup_{n \to \infty} EC_n^+ \leqq \frac{(1+\theta^{1/2}) \, c(n, \varepsilon)}{1-\alpha} + \frac{20\theta^{1/2}}{\alpha} + \delta(\theta)$$

for $\delta(\theta)$ as in Lemma 5. Letting θ (and hence ε) $\to 0$, then letting $\alpha \to 0$, we establish Proposition 2.

3 The limit tree

3.1 Construction

We need to study a certain type of infinite (non-random) tree with random "costs" on its edges. We start with an informal description in terms of the

well-known *Yule process* or pure birth process: that is, the continuous-time branching process in which each individual has children at rate 1. Draw the family tree of such a branching process, started with a single progenitor, and to each parent-child edge let the "cost" on that edge be the age of the parent at the time of birth of the child. Next take an independent copy of this process. Finally, join the two progenitors by a special edge ϕ, and specify that this edge has a random cost chosen according to Lebesgue measure on $(0, \infty)$. Write c for a realization of this process, and λ for the "distribution" of the process, which is a σ-finite measure because of the specification above. Figure 2 shows part of a realization, with costs written above edges, and the edge labels (defined below) written beneath edges.
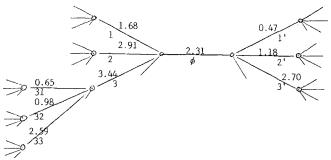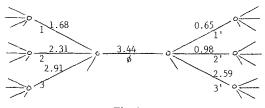

Fig. 2

Fig. 3

We now give a more precise construction, and introduce notation. Take the complete B-ary tree of height H. Each vertex can be labelled as a string $i_1 i_2 \ldots i_h$ where $0 \leq h \leq H$ and each $i_u \in \{1, 2, \ldots, B\}$. Thus the root is labelled as the null string ϕ, its children are labelled as vertices $1, 2, \ldots, B$, and so on. We can also label the edges with strings; the edge from vertex $i_1 \ldots i_{h-1}$ to vertex $i_1 \ldots i_h$ is labelled as edge $i_1 \ldots i_h$. In particular, the edges at the root ϕ are labelled as edges $1, 2, \ldots B$. Now take another copy of this tree, and label its vertices and edges using strings $i'_1 \ldots i'_h$. Finally join the two roots ϕ and ϕ' via an edge, which we label as edge ϕ. This defines a tree we shall call $\tau_{B,H}$. Taking the union over all H gives an infinite B-ary tree we shall

call τ_B. Then taking the union over all B gives an infinite infinitary tree we call τ. Regard these trees as sets of edges e. We study objects of the type $\mathbf{c} = (c(e); e \in \tau)$ where $0 < c(e) \leq \infty$ is a "cost" on edge e. These costs are required to satisfy a monotonicity condition: for each vertex $\mathbf{i} = i_1 \ldots i_h$ the costs $(c(\mathbf{i}1), c(\mathbf{i}2), c(\mathbf{i}3), \ldots)$ on the outward edges must be non-decreasing (and similarly for $\mathbf{i}'$). Call such a $\mathbf{c}$ a *cost tree* and write C for the set of cost trees. Mathematically, C is just a subset of countable product of $[0, \infty)$'s and so inherits the natural topology (coordinatewise convergence, i.e. convergence of costs on each edge) and σ-field. Write C_B for the set of cost trees $\mathbf{c} = (c(e); e \in \tau_B)$ and write $C_{B,H}$ for the set of cost trees $\mathbf{c} = (c(e); e \in \tau_{B,H})$.

Given $x > 0$, we define a random cost tree as follows. Let $(S_1, S_2, S_3, \ldots)$ be the times of a Poisson(1) process, i.e. distributed as $(\xi_1, \xi_1 + \xi_2, \xi_1 + \xi_2 + \xi_3, \ldots)$ where the (ξ_i) are independent with exponential (1) distribution. Let $c(\phi) = x$. For each vertex $\mathbf{i} = i_1 \ldots i_h$ let the costs $(c(\mathbf{i}1), c(\mathbf{i}2), c(\mathbf{i}3), \ldots)$ on the outward edges be distributed as $(S_i; i \geq 1)$, independently as $\mathbf{i}$ varies. Similarly for the $\mathbf{i}'$. Let λ_x be this distribution of the random cost tree. So λ_x is a probability measure on C. We shall need to use the σ-finite measure λ on C defined by

$$\lambda(\cdot) = \int\limits_0^\infty \lambda_x(\cdot) \, dx.$$

Needing to work with a σ-finite measure is a minor annoyance, which we often handle as follows. Let $D_x = \{\mathbf{c} : c(\phi) \leq x\}$. So $\lambda(D_x) = x$ by definition. So $x^{-1}\lambda(\cdot)$ restricted to D_x is a probability measure, and in the sequel computations with λ are often done (or implicitly justified) by appeal to these probability measures. In particular, conditional expectations w.r.t. λ can be defined this way.

For each $i \geq 1$ there is a map $\theta_i : C \to C$ which takes edge i to edge ϕ and re-labels edges to preserve monotonicity of edge-costs. A picture being worth a thousand words, Fig. 3 illustrates the action of θ_3 on the tree of Fig. 2.

Let θ_i' be the analogous map which takes edge i' to edge ϕ. Now let $\mathcal{G}$ denote the set of measurable functions $g : C \to [0, 1]$ such that

(15) $$g(\mathbf{c}) + \sum_{i=1}^\infty g(\theta_i \mathbf{c}) = 1 \quad \text{for } \lambda\text{-almost all } \mathbf{c} \in C.$$

(16) $$g(\mathbf{c}) + \sum_{i=1}^\infty g(\theta_i' \mathbf{c}) = 1 \quad \text{for } \lambda\text{-almost all } \mathbf{c} \in C.$$

It is not obvious that *any* such function g exists, though this fact emerges from the proof in Sect. 3.3 and the known bound in the random assignment problem. In Sect. 4 we sketch how to construct one $g \in \mathcal{G}$ in a comparatively explicit way.

Regard $g(\mathbf{c})$ as a "probability" associated with edge ϕ. Associate with g the number

$$\gamma(g) = \int\limits_C g(\mathbf{c}) \, c(\phi) \, \lambda(d\mathbf{c}).$$

Finally define

$$(17) \qquad\qquad\qquad c^* = \inf_{g \in \mathcal{G}} \gamma(g).$$

This is the constant which will be the asymptotic mean cost in the random assignment problem.

Remarks. 1. Our first "Yule process" description involves times, but for our purposes it is more helpful to regard the trees as "spatial" objects, and to regard λ as the natural "uniform measure" on C. For instance, λ has an "invariance" property w.r.t. the maps θ_i. To state this property, define

$$D_{(k)} = \{\mathbf{c} : c(k-1) < c(\phi) < c(k)\}$$

where $c(0) = 0$ (and $c(k)$ is the cost on edge k). Then θ_i maps $D_{(k)}$ into $D_{(i*k)}$, where

$$i * k = i + 1_{(i \geq k)}.$$

Let $\lambda_{(k)}(\cdot) = \lambda(D_{(k)} \cap \cdot)$. Then $\lambda_{(k)}$ is a probability measure. The invariance property is

$$(18) \qquad\qquad\qquad \theta_i \lambda_{(k)} = \lambda_{(i*k)} \qquad \text{for all } i, k$$

where $(\theta_i \lambda_{(k)})(\cdot) = \lambda_{(k)}(\theta_i^{-1} \cdot)$; and similarly for θ'. Property (18) could be checked directly from properties of the Poisson process. Alternatively, it is a consequence of Lemma 12 below, proved by a limit argument from the finite assignment problem.

2. When g takes value 0 and 1 only, we can use g to define a matching on τ. Informally, edge e is in the matching if g takes the value 1 when applied to the cost tree $\theta_e \mathbf{c}$ in which edge e is moved to edge ϕ. Conditions (15, 16) are exactly the conditions needed for this set of edges to be a matching. We are interested in the "average cost per edge" of this matching, and instead of formalizing that idea by taking limits of averages over finite subtrees it is more convenient to use the "ergodic" definition above (17).

3.2 The unfolding map

To relate the infinite tree with the finite assignment problem, the central idea is an "unfolding map". Fix B, n and $1 \leq j_0, m_0 \leq n$. The "unfolding map" takes a non-negative $n \times n$ matrix $(t_{j,m})$ to a B-ary cost tree $\mathbf{c} = (c(e); e \in \tau_B)$, as follows. Identify (j_0, m_0) with edge ϕ of τ_B, and let $c(\phi) = t_{j_0, m_0}$. Then let $m(1), \ldots, m(B)$ be the first B arguments in the increasing rearrangement of $(t_{j_0, m}; m \neq m_0)$. Identify $m(i)$ with edge $\phi(i)$ of τ_B and let $c(i) = t_{j_0, m(i)}$. Then let $j(1), \ldots, j(B)$ be the first B arguments in the increasing rearrangement of $(t_{j, m_0}; j \neq j_0)$. Identify $j(i)$ with edge $(\phi' i)$ of τ_B and let $c(i') = t_{j(i), m_0}$. Then let $j(1, 1,), \ldots, j(1, B)$ be the first B arguments in the increasing rearrangement of $(t_{j, m(1)}; j \in \{j_0, j(1), \ldots, j(B)\})$. Identify $(j(1, i), m(1))$ with edge $1i$ of τ_B and let $c(1i) = t_{j(1, i), m(1)}$. Continue with

$j(2, 1), \ldots, j(B, B)$ and then $m(1, 1), \ldots, m(B, B)$ and so on. When no matrix entries remain, assign cost $+\infty$ to the remaining edges of τ_B.

Applying the unfolding map (with $B = 3$) to the matrix in Fig. 4 gives the cost tree in Fig. 2.

Machines

m_0

..	..	..	..	.	0.98	.	.	.
.	.	.	.	.	.	.	.	.
.	.	.	1.18	..	.	.	.	.
.	.	.	.	.	.	.	.	.
j_0	2.91	.	.	2.31	.	3.44	.	.	1.68
.	.	.	.	.	.	.	.	.
.	.	.	.	.	0.65	.	.	.
.	.	.	.	.	.	.	.	.
.	.	.	2.70	.	.	.	.	.
.	.	.	.	.	2.59	.	.	.
.	.	.	.	.	.	.	.	.
.	.	.	0.47	.	.	.	.	.

Jobs

Fig. 4

Let $U_{B,n}^{j_0, m_0}$ be the random cost tree obtained by applying the unfolding map to the random cost matrix T of (1). Let $\lambda_{B,n}$ be the distribution of $U_{B,n}^{j_0, m_0}$. Then $\lambda_{B,n}$ is a probability measure on C_B which by symmetry does not depend on (j_0, m_0). Let $\lambda_{B,H,n}$ be the projection $\lambda_{B,n}$ onto $C_{B,H}$, i.e. the distribution of $\rho_H U_{B,n}^{j_0, m_0}$ where $\rho_H : C_B \to C_{B,H}$ is the natural restriction map.

Returning to the measure λ defined in Sect. 3.1, let λ_B be the projection of λ onto C_B, i.e. the measure on C_B induced by the restriction map $C \to C_B$. Similarly let $\lambda_{B,H}$ be the projection of λ onto $C_{B,H}$.

The connection between the $n \times n$ random cost matrices and the infinite random cost trees is provided by the following lemma. Recall $D_x = \{\mathbf{c} : c(\phi) \leq x\}$, so $\lambda_{B,H,n}(D_x) = P(t_{j_0, m_0} \leq x) \sim x/n$ as $n \to \infty$.

Lemma 10 *For fixed B, H, x,*

$$n \lambda_{B,H,n}(D_x \cap \cdot) \xrightarrow{TV} \lambda_{B,H}(D_x \cap \cdot) \quad as \ n \to \infty.$$

This involves *total variation* convergence of positive measures:

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

Proof. Write $|\tau_{B,H}|$ for the number of edges of $\tau_{B,H}$. Let $n > |\tau_{B,H}|$. The lemma is an immediate consequence of the following bound on the likelihood ratio.

$$(19) \qquad \frac{d\lambda_{B,H,n}}{d\lambda_{B,H}}(\mathbf{c}) \geq \frac{1}{n} \exp\left(-\frac{c(\phi)}{n}\right)\left(1 - \frac{|\tau_{B,H}|}{n}\right)^{|\tau_{B,H}|-1}.$$

To prove (19), consider the construction of $U_{B,n}^{1,1}$. At a typical stage of construction we have defined costs for certain edges of $\tau_{B,H}$, up through edge $\mathbf{i}a$, say, which was given cost $c(\mathbf{i}a) = t_{j^*,m^*}$ for some job j^* and machine m^*. If $a < B$ (the case $a = B$ leads to a similar bound) the next edge $\mathbf{i}(a+1)$ will be given cost $\min_{m \in S} t_{j^*,m}$ where S is the set of machines not yet identified with vertices of τ_B. Conditional on the construction thus far, the $(t_{j^*,m}; m \in S)$ are distributed as independent exponential r.v.'s conditioned on their minimum being greater than $c(\mathbf{i}a)$. Using the memoryless property of exponential distributions, the conditional density of $c(\mathbf{i}(a+1)) - c(\mathbf{i}a)$ at x is $\frac{|S|}{n}\exp(-x|S|/n)$. Under λ the conditional density is $\exp(-x)$, so the ratio of conditional densities is at least $|S|/n$, and hence at least $1 - \frac{|\tau_{B,H}|}{n}$. Now (19) holds because the likelihood ratio equals the ratio for edge ϕ, which is $\frac{1}{n}\exp\left(-\frac{c(\phi)}{n}\right)$, times the product of the conditional likelihood ratios.

We shall need a technical result which relates the effect of altering the initial pair (j_0, m_0) in the unfolding map to the maps θ_i defined in Sect. 3.1, which for $i \leq B$ we may regard as maps $C_B \to C_B$. Applying the unfolding map to the matrix in Fig. 4, but taking as initial pair the (j_0, m_0^*) with $t_{j_0, m_0^*} = 3.44$, we obtain the cost tree in Fig. 3. In general, the cost tree obtained by unfolding around (j_0, m_0^*) is essentially like the map θ_i applied to the cost tree obtained by unfolding about (j_0, m_0), where i is the (increasing) rank of (j_0, m_0^*) amongst $\{t_{j_0,m} : m \neq m_0\}$, with two provisos. First, both t_{j_0,m_0} and t_{j_0,m_0^*} must be among the $B+1$ smallest entries of $\{t_{j_0,m} : 1 \leq m \leq n\}$. Second, we only expect the cost trees to coincide "locally" around the special edge.

These ideas are formalized in the lemma below, in which we take (j_0, m_0) to be $(1, 1)$. Define

$$(20) \qquad R_n = \text{increasing rank of } t_{1,1} \text{ amongst } (t_{1,m}; 1 \leq m \leq n).$$

Lemma 11 *For fixed B, H and $i \leq B$,*

$$nP(R_n \leq B+1, \rho_H U_{B,n}^{1,m(i)} \neq \rho_H \theta_i U_{B,n}^{1,1}) \to 0 \qquad as \ n \to \infty$$

where $m(i)$ is the i'th ranked argument of $\{t_{1,m}; m \neq 1\}$.

Proof. For a cost tree $\mathbf{c} = (c(e): e \in T_B)$ write

$$\max_{(H)} \mathbf{c} = \max_{e \in \tau_{B,H}} c(e).$$

The key fact is: if

(21) $$R_n \leq B+1 \quad \text{and} \quad \rho_H U_{B,n}^{1,m(i)} \neq \rho_H \theta_i U_{B,n}^{1,1}$$
$$\text{and} \quad \max_{(H)} U_{B,n}^{1,m(i)} \leq y \quad \text{and} \quad \max_{(H)} \theta_i U_{B,n}^{1,1} \leq y$$

then for some $h \leq H$ there exists a "h-cycle"

$$j_0 = 1, m_0, j_1, m_1, j_2, m_2, \ldots, j_h, m_h, j_{h+1} = j_0 = 1$$

such that the j's are distinct, the m's are distinct, $m_0 \in \{m(1), \ldots, m(B)\}$ and

$$t_{j_i, m_i} \leq y, \quad t_{m_i, j_{i+1}} \leq y \quad \text{for all } i.$$

To prove this, imagine we have constructed $U_{B,n}^{1,1}$ and are now starting the construction of $U_{B,n}^{1,m(i)}$. Provided $R_n \leq B+1$, the job-machine pairs which $U^{1,m(i)}$ identifies with edges $\phi, 1, \ldots, B$ will be the same pairs (in different order) that $U^{1,1}$ identified with those edges. Continuing the construction, consider the first time that $U^{1,m(i)}$ identifies a job j^* (or machine, similarly) with a vertex $v \in \tau_{B,H}$ which is different from the job j^{**} identified with $\theta_i v$ by $U^{1,1}$. This can happen in one of two similar ways. Either j^* was previously identified by $U^{1,1}$ with some other vertex $v^* \in \tau_{B,H}$, or j^{**} was previously identified by $U^{1,m(i)}$ with some other vertex $v^{**} \in \tau_{B,H}$. In the first case (the second is similar) consider the path in τ from vertex ϕ to v and the sequence of jobs and machines associated by $U^{1,m(i)}$ with that path; then consider the path in τ from vertex ϕ to v^* and the sequence of jobs and machines associated by $U^{1,1}$ with that path. Concatenating paths gives a h-cycle.

By counting the number of possible h-cycles,

$$P(\text{there exists some } h\text{-cycle}) \leq B n^{2h} (y/n)^{2h+2}$$

and so for fixed y the probability of the event in (21) is $O(n^{-2})$. Thus the quantity in the lemma is asymptotically at most

$$\lim_{y \to \infty} \limsup_{n \to \infty} n P(R_n \leq B+1, \max_{(H+1)} U_{B,n}^{1,1} > y)$$

plus a similar term involving $U^{1,m(i)}$. By considering whether $t_{1,1} \leq x$ we can bound this as the sum of the following two quantities.

$$\lim_{y \to \infty} \limsup_{n \to \infty} n P(t_{1,1} \leq x, \max_{(H+1)} U_{B,n}^{1,1} > y)$$

$$\lim_{x \to \infty} \limsup_{n \to \infty} n P(t_{1,1} > x, R_n \leq B+1).$$

In the first quantity, the $n \to \infty$ limit is $\lambda\{D_x, \max_{(H+1)} \mathbf{c} > y\}$ by Lemma 10, and this $\to 0$ as $y \to \infty$. The second quantity is, for fixed n, bounded by $n \dfrac{B+1}{n} P(Y_n > x)$ where Y_n is the $B+1$'st smallest of $(t_{1,m}; 1 \leq m \leq n)$. So the $n \to \infty$ limit is at most $(B+1) P(S_{B+1} \geq x)$ and this $\to 0$ as $x \to \infty$.

3.3 Proof of Proposition 3 (a)

The optimal assignment π_n can be regarded as a random $\{0,1\}$-valued matrix $q_{j,m}=1_{(\pi_n(j)=m)}$. There is no loss of generality in assuming the symmetry property

(22) $\text{dist}((t_{j,m},q_{j,m}); 1\leqq j, m\leqq n)$ is invariant under
         permutations of jobs and permutations of machines

since this may be achieved by randomly permuting the labels of jobs and machines. Then

(23) $$EC(n)=\frac{1}{n}\sum_j\sum_m E q_{j,m}\, t_{j,m}=nE q_{1,1}\, t_{1,1}.$$

For fixed (B, n) consider the unfolding map with j_0 and m_0 uniform random on $\{1, ..., n\}$. When we construct the unfolding map, we associate edges e of C_B with entries (j, m) and set $c(e)=t_{j,m}$; now also set $p(e)=q_{j,m}$, with the convention that $p(e)=0$ when $c(e)=\infty$. This constructs a random element of the set C_B^2 whose generic element is $((c(e),p(e)); e\in\tau_B)$, that is the set of B-ary trees with two numbers associated with each edge. Let $\psi_{B,n}$ be the distribution of this random element. So

(24) $$E q_{1,1}\, t_{1,1}=\int_{C_B^2} p(\phi)\, c(\phi)\, d\psi_{B,n}.$$

Let $\psi_{B,H,n}$ be the projection of $\psi_{B,n}$ onto the set $C_{B,H}^2$ of $((c(e),p(e)); e\in\tau_{B,H})$. Then the sequence $(n\psi_{B,H,n}; n\geqq 1)$, restricted to D_x, is relatively compact w.r.t. weak convergence on $C_{B,H}^2$, because the first coordinates are convergent by Lemma 10 and because $0\leqq p(e)\leqq 1$. So take a subsequence of n's for which the lim inf in Proposition 3(a) is a limit; then we can take a further subsequence such that, for fixed B, H, x,

(25) $$n\psi_{B,H,n}(D_x\cap\,\cdot\,)\to\psi_{B,H}(D_x\cap\,\cdot\,),\text{say},\qquad\text{as } n\to\infty$$

where $\to$ denotes weak convergence; then by a diagonal argument we may suppose (25) holds for all B, H, x. It is easy to see that $\psi_{B,H}$ are consistent as B, H increase and so are projections of some measure ψ on the set C^2 of $((c(e),p(e)); e\in\tau)$. We then have

(26) $$\int_{C^2} p(\phi)\, c(\phi)\, d\psi=\int_{C_B^2} p(\phi)\, c(\phi)\, d\psi_{B,H}\leqq\liminf_{n\to\infty} EC(n),$$

the inequality by (23, 24, 25) and Fatou's lemma.

We next use Lemma 11 to show that ψ inherits the "invariance" property (18) we stated earlier for λ. Recall

$$D_{(k)}=\{\mathbf{c}: c(k-1)<c(\phi)<c(k)\}.$$

Write $\psi_{(k)}$ for ψ restricted to $D_{(k)}$.

Lemma 12 $\theta_i\,\psi_{(k)}=\psi_{(i*k)}$ for all i, k, where $i*k=i+1_{(i\geqq k)}$.

Proof. The assertion that convergence in (25) holds for all H and x can be written as

(27)
$$n\psi_{B,n} \to \rho_B \psi$$

where "convergence" is weak convergence of σ-finite measures on C_B^2. Now fix B and i, $k \leq B$. Regard $\psi_{B,n}$ as the distribution of the unfolding map applied to $(j_0, m_0) = (1, 1)$. Let $\psi_{B,n}^{(k)}$ be the distribution obtained when we condition on $t_{1,1}$ being the k'th smallest of $(t_{1,m}: 1 \leq m \leq n)$, i.e. condition on $\{R_n = k\}$. Using (27),

(28)
$$\psi_{B,n}^{(k)} \to \tilde{\rho}_B \psi_{(k)}$$

where $\tilde{\rho}_B$ is the restriction map $C \to C_B$, and hence

(29)
$$\theta_i \psi_{B,n}^{(k)} \to \tilde{\rho}_B \theta_1 \psi_{(k)}.$$

Now let $\psi_{B,n}^{(k,i)}$ be the distribution of $U_{B,n}^{1;m(i)}$ conditioned on $\{R_n = k\}$. Lemma 11 implies

$$P(\rho_H U_{B,n}^{1;m(i)} \neq \rho_H \theta_i U_{B,n}^{1,1} \mid R_n = k) \to 0 \quad \text{as } n \to \infty.$$

This and (29) show

(30)
$$\rho_H \psi_{B,n}^{(k,i)} \to \rho_H \tilde{\rho}_B \theta_i \psi_{(k)}.$$

If we condition on $m(i) = m^*$ as well as on $R_n = k$, then t_{1,m^*} is the $i*k$'th smallest of $\{t_{1,m}: 1 \leq m \leq n\}$. But by the symmetry property (22), the effect of these two conditionings on $U_{B,n}^{1;m^*}$ is the same as conditioning on t_{1,m^*} being the $i*k$'th smallest of $\{t_{1,m}: 1 \leq m \leq n\}$. In other words

$$\psi_{B,n}^{(k,i)} = \psi_{B,n}^{(i*k)}.$$

Then by (30) and (28)

$$\rho_H \tilde{\rho}_B \theta_i \psi_{(k)} = \rho_H \tilde{\rho}_B \psi_{(i*k)}$$

establishing the lemma.

Lemma 13

$$p(\phi) + \sum_{i=1}^{\infty} p(i) = 1 \quad \psi\text{-a.e.}$$

and similarly for $p(i')$.

Here the "i" in $p(i)$ refers to edge i, the edge from vertex ϕ to vertex i.

Proof. One side is easy. The fact that π_n is a matching says that $\sum_m q_{1,m} = 1$, and then from the description of the unfolding map

(31)
$$p(\phi) + \sum_{i=1}^{B} p(i) \leq 1 \quad \psi_{B,n}\text{-a.e.}$$

Letting $n \to \infty$, (25) shows that the inequality in (31) holds $\psi_{B,H}(D_x \cap \cdot)$-a.e. and hence ψ-a.e. Letting $B \to \infty$ gives " ≤ 1 " in the lemma.

Now define

$$\Omega_B\{\mathbf{c}: c(\phi) < \min(c(B), c(B'))\}$$

where $c(B)$ is the cost attached to edge B. Consider

$$\int_{\Omega_B \cap D_x} \left(1 - p(\phi) - \sum_{i=1}^{B} p(i)\right) d\psi_{B,n} = a(B, n, x) \quad \text{say.}$$

To interpret this, we may regard $\psi_{B,n}$ as the distribution of $U_{B;n}^{1;M}$ where M is random, uniform on $\{1, 2, \ldots, n\}$. Then $a(B, n, x)$ is the chance that

$t_{1,M} \leq x$; and
$t_{1,M}$ is one of the B smallest values of $(t_{1,m}; 1 \leq m \leq n)$; and
$t_{1,M}$ is one of the B smallest values of $(t_{j,M}; 1 \leq j \leq n)$; and
$t_{1,\pi_n(1)}$ is not of the $B+1$ smallest values of $(t_{1,m}; 1 \leq m \leq n)$.

To upper bound this chance, ignore the second and third constraints. This leads to the bound

$$a(B, n, x) \leq E n^{-1} V_{n,x} 1_{(R_n^\pi > B + 1)}$$

where

(32) $$V_{n,x} = |\{m: t_{1,m} \leq x\}|$$

R_n^π is the rank of $t_{1,\pi_n(1)}$ amongst $(t_{1,m}; 1 \leq m \leq n)$.

Letting $n \to \infty$ and using (25)

$$\int_{\Omega_B \cap D_x} \left(1 - p(\phi) - \sum_{i=1}^{B} p(i)\right) d\psi_B \leq \limsup_{n \to \infty} E V_{n,x} 1_{(R_n^\pi > B + 1)}.$$

Now the left side is unchanged by replacing ψ_B with ψ, and is made smaller by replacing Ω_B with Ω_b for $b < B$. So, letting $B \to \infty$ for fixed b,

(33) $$\int_{\Omega_b \cap D_x} \left(1 - p(\phi) - \sum_{i=1}^{\infty} p(i)\right) d\psi \leq \lim_B \limsup_{n \to \infty} E V_{n,x} 1_{(R_n^\pi > B + 1)}.$$

Write Z_x for a r.v. with Poisson(x) distribution. Then $V_{n,x}$ is stochastically smaller than Z_x, and hence $(V_{n,x}; n \geq 1)$ is uniformly integrable. Granted that $(R_n^\pi; n \geq 1)$ is tight, the double limit in (33) equals zero. Then letting $x, b \to \infty$ in the left side of (33) establishes Lemma 13. To verify tightness of $(R_n^\pi; n \geq 1)$,

$$P(R_n^\pi \geq 2r) \leq P(t_{1,\pi_n(1)} > r) + P(V_{n,r} \geq 2r)$$
$$\leq \frac{EC(n)}{r} + P(V_{n,r} \geq r)$$

since $Et_{1,\pi_n(1)}=EC(n)$ by symmetrization (23). So

$$\limsup_{n\to\infty} P(R_n^\pi\geqq 2r)\leqq\frac{\limsup EC(n)}{r}+P(Z_r\geqq 2r).$$

This bound $\to 0$ as $r\to\infty$.

Now regard $((c(e),p(e));e\in\tau)$ as random variables w.r.t. the (σ-finite) measure ψ. There is a function $g:C\to[0,1]$ such that $g(\mathbf{c})$ is the conditional expectation of $p(\phi)$ given $\mathscr{F}\equiv\sigma(c(e):e\in\tau)$. Using Lemma 12, $g(\theta_i\mathbf{c})$ is the conditional expectation of $p(i)$ given $\mathscr{F}$. Then taking conditional expectations in Lemma 13 shows that $g(\mathbf{c})+\sum\limits_{i=1}^{\infty} g(\theta_i\mathbf{c})=1$ for λ-almost all $\mathbf{c}\in C$; that is to say, $g\in\mathscr{G}$. And linearity of conditional expectations implies

$$\int\limits_C g(\phi)\,c(\phi)\,d\lambda=\int\limits_{C^2} p(\phi)\,c(\phi)\,d\psi$$

and so part (a) of Proposition 3 follows from (26).

3.4 Proof of Proposition 3 (b)

Fix a function $g:C\to[0,1]$ with $g\in\mathscr{G}$ and with $\gamma(g)<\infty$. Let $g_{B,H}:C_{B,H}\to[0,1]$ be the conditional expectation of g given $\mathscr{F}_{B,H}\equiv\sigma(c(e):e\in\tau_{B,H})$ w.r.t. the measure λ, and let $g_B:C_B\to[0,1]$ be the conditional expectation of g given $\mathscr{F}_B\equiv\sigma(c(e):e\in\tau_B)$. Define

$$g_x=g\,1_{D_x},\qquad g_{B,x}=g_B\,1_{D_x},\qquad g_{B,H,x}=g_{B,H}\,1_{D_x}.$$

In other words $g_{B,H,x}(\mathbf{c})=g_{B,H}(\mathbf{c})\,1_{(c(\phi)\leqq x)}$, etc. Recall the definition of $U_{B,n}^{j,m}$, the random cost tree obtained by applying the unfolding map to the random cost matrix T. For each (j,m) define

$$q_{j,m}=g_{B,H,x}(U_{B,n}^{j,m}),$$

This defines a random matrix Q, which depends on (B,H,x,n) although this is suppressed in our notation. We shall prove

(34) $$\lim_{x\to\infty}\limsup_{n\to\infty} E\frac{1}{n}\sum_j\sum_m q_{j,m}t_{j,m}=\gamma(g)\qquad\text{for all }(B,H).$$

(35) $$\lim_{x\to\infty}\limsup_{B\to\infty}\limsup_{H\to\infty}\limsup_{n\to\infty} E\chi(Q)=0.$$

Then given $\delta,\varepsilon>0$ we can first choose x_0 such that the lim sup in (34) is less than $\gamma(g)+\delta$. Then by (35) choose $H,B,x>x_0$ such that $E\chi(Q)<\varepsilon$ for all sufficiently large n. Then (34) shows $c(n,\varepsilon)\leqq\gamma(g)+\delta$ for all sufficiently large n. Part (b) of the Proposition follows.

To prove (34), recall that $\lambda_{B,H,n}$ denotes the distribution of $\rho_H U^{j,m}_{B,n}$ where ρ_H is the restriction map $C_B \to C_{B,H}$. So

$$E q_{1,1} t_{1,1} = \int c(\phi) g_{B,H,x}(\mathbf{c}) \, d\lambda_{B,H,n}.$$

Note we automatically have the symmetry property (22), and so

$$E \frac{1}{n} \sum_j \sum_m q_{j,m} t_{j,m} = E q_{1,1} t_{1,1}.$$

So letting $n \to \infty$ and using Lemma 10

$$\lim_{x \to \infty} E \frac{1}{n} \sum_j \sum_m q_{j,m} t_{j,m} = \int_{C_{B,H}} c(\phi) g_{B,H,x}(\mathbf{c}) \, d\lambda_{B,H}.$$

By linearity of conditional expectations, this equals $\int_C c(\phi) g_x(\mathbf{c}) \, d\lambda$. Then letting $x \to \infty$ this converges to $\int_C c(\phi) g(\mathbf{c}) \, d\lambda = \gamma(g)$, establishing (34).

To start the proof of (35), note that the symmetry property (23) implies

$$(36) \qquad E \chi(Q) = E \left| \sum_{m=1}^n q_{1,m} - 1 \right| + E \left| \sum_{j=1}^n q_{j,1} - 1 \right|.$$

Conditioning on $V_{n,x} \equiv |\{m : t_{1,m} \leq x\}|$, gives the crude inequality

$$E \left| \sum_{m=1}^n q_{1,m} - 1 \right| \leq P(V_{n,x} = 0) + E \left| \sum_{m=1}^n q_{1,m} - 1 \right| V_{n,x}.$$

By symmetry, the final term above equals

$$(37) \qquad nE \left| \sum_{m=1}^n q_{1,m} - 1 \right| 1_{(t_{1,1} \leq x)}.$$

Now let $m(1), m(2), \ldots, m(n-1)$ be the re-ordering of $\{2, 3, \ldots, n\}$ such that $t_{1,m(i)}$ is increasing. Then, since $q_{1,m} \leq 1_{(t_{i,m} \leq x)}$,

$$\sum_{i=B+1}^{n-1} q_{1,m(i)} \leq (V_{n,x} - B - 1)^+ \qquad \text{on } \{t_{1,1} \leq x\}$$

and so (37) is bounded by

$$nE(V_{n,x} - B - 1)^+ 1_{(t_{1,1} \leq x)} + nE \left| q_{1,1} + \sum_{i=1}^B q_{1,m(i)} - 1 \right| 1_{(t_{1,1} \leq x)}.$$

Collecting these bounds gives

(38)
$$E\left|\sum_{m=1}^{n} q_{1,m} - 1\right| \leq P(V_{n,x}=0) + nE(V_{n,x}-B-1)^{+} 1_{(t_{1,1}\leq x)}$$

$$+ nE\left|q_{1,1} + \sum_{i=1}^{B} q_{1,m(i)} - 1\right| 1_{(t_{1,1}\leq x)}.$$

Now the first term on the right equals e^{-x}. As for the second,

$$nE(V_{n,x}-B-1)^{+} 1_{(t_{1,1}\leq x)} = nP(t_{1,1}\leq x) E((V_{n,x}-1-B)^{+}|t_{1,1}\leq x)$$
$$\rightarrow xE(Z_{x}-B)^{+} \quad \text{as} \quad n\rightarrow\infty$$

where Z_x has Poisson(x) distribution. Thus under the limit procedure of (35) these two terms tend to 0, and using (36) the proof of (35) reduces to the proof of

(39)
$$\lim_{x\rightarrow\infty} \limsup_{B\rightarrow\infty} \limsup_{H\rightarrow\infty} \limsup_{n\rightarrow\infty} nE\left|q_{1,1} + \sum_{i=1}^{B} q_{1,m(i)} - 1\right| 1_{(t_{1,1}\leq x)} = 0$$

together with the reflected version involving $q_{j(i),1}$ whose proof is identical.

Fix B, H, x. Summing over i in Lemma 11 shows that, as $n \rightarrow \infty$,

$$nP\left(R_{n}\leq B+1, \sum_{i=1}^{B} g_{B,H,x}(U_{B,n}^{1,m(i)}) + \sum_{i=1}^{B} g_{B,H,x}(\theta_{i}(U_{B,n}^{1,1})) \right) \rightarrow 0$$

or in other words

(40) $nP\left(R_{n}\leq B+1, q_{1,1} + \sum_{i=1}^{B} q_{1,m(i)} \neq g_{B,H,x}(U_{B,n}^{1,1}) + \sum_{i=1}^{B} g_{B,H,x}(\theta_{i}(U_{B,n}^{1,1})) \right) \rightarrow 0$

Then

$$\limsup_{n\rightarrow\infty} nE\left|q_{1,1} + \sum_{i=1}^{B} q_{1,m(i)} - 1\right| 1_{(t_{1,1}\leq x)}$$

$$\leq \limsup_{n\rightarrow\infty} nE\left|q_{1,1} + \sum_{i=1}^{B} q_{1,m(i)} - 1\right| 1_{(t_{1,1}\leq x)} 1_{(R_{n}\leq B+1)}$$

$$+ B \limsup_{n\rightarrow\infty} nP(R_{n} > B+1, t_{1,1}\leq x)$$

because the quantity $|\cdot|$ is at most B. The second limit is bounded by $BxP(Z_x > B+1)$ which becomes 0 in the limit procedure of (39). The first limit equals

$$\limsup_{n\rightarrow\infty} nE\left|g_{B,H,x}(U_{B,n}^{1,1}) + \sum_{i=1}^{B} g_{B,H,x}(\theta_{i}(U_{B,n}^{1,1})) - 1\right| 1_{(t_{1,1}\leq x)} 1_{(R_{n}\leq B+1)}$$

using (40) and the fact that the quantity $|\cdot|$ is at most B;

$$\leq \lim_{n \to \infty} \int_{D_x} \left| g_{B,H,x}(\mathbf{c}) + \sum_{i=1}^{B} g_{B,H,x}(\theta_i \mathbf{c}) - 1 \right| n \, d\lambda_{B,H,n}$$

where we simply neglect the constraint on R_n;

$$= \int_{D_x} \left| g_{B,H,x}(\mathbf{c}) + \sum_{i=1}^{B} g_{B,H,x}(\theta_i \mathbf{c}) - 1 \right| d\lambda \quad \text{by Lemma 10.}$$

Thus (39) will follow if we can show

(41) $$\qquad \lim_{B \to \infty} \limsup_{H \to \infty} \int_{D_x} \left| g_{B,H,x}(\mathbf{c}) + \sum_{i=1}^{B} g_{B,H,x}(\theta_i \mathbf{c}) - 1 \right| d\lambda = 0 \qquad \text{for all } x.$$

To prove this, note first that the martingale convergence theorem implies

$$\lim_{H \to \infty} |g_{B,H,x}(\mathbf{c}) - g_{B,x}(\mathbf{c})| \, d\lambda = 0.$$

Next, let $\mathscr{F}_{B,H}^{i}$ denote $\theta_i \mathscr{F}_{B,H}$, that is to say the σ-field generated by $(c(e); e \in \tau_B, e$ within distance H of edge i). By the invariance property (18) we see that for $i \leq B$

$g_{B,H,x}(\theta_i \mathbf{c})$ is the cond. expectation of $g(\theta_i \mathbf{c})$ w.r.t. $\mathscr{F}_{B,H}^{i}$, on $\{c(\phi) < c(B+1)\}$.

Then, since $\sigma(\bigcup_H \mathscr{F}_{B,H}^{i}) = \sigma(c(e); e \in \tau_B)$, the martingale convergence theorem implies

$$\lim_{H \to \infty} \int_{\{c(\phi) < c(B+1)\} \cap D_x} |g_{B,H,x}(\theta_i \mathbf{c}) - g_{B,x}(\theta_i \mathbf{c})| \, d\lambda = 0.$$

On $D_x \setminus \{c(\phi) < c(B+1)\}$, the integrand above is bounded by 1. Thus the $\limsup\limits_{H \to \infty}$ term in (41) is bounded by

$$B\lambda \{D_x, c(\phi) \geq c(B+1)\} + \int_{D_x} \left| g_{B,x}(\mathbf{c}) + \sum_{i=1}^{B} g_{B,x}(\theta_i \mathbf{c}) - 1 \right| d\lambda.$$

The first term is at most $B \times P(Z_x > B)$ where Z_x has Poisson(x) distribution, and this quantity $\to 0$ as $B \to \infty$. So we have reduced (41) to the proof of

(42) $$\qquad \lim_{B \to \infty} \int_{D_x} \left| g_{B,x}(\mathbf{c}) + \sum_{i=1}^{B} g_{B,x}(\theta_i \mathbf{c}) - 1 \right| d\lambda = 0.$$

To prove this, first use the martingale convergence theorem and repeat arguments above to show

(43) $$\qquad \lim_{B \to \infty} \int |g_{B,x}(\mathbf{c}) - g_x(\mathbf{c})| \, d\lambda = 0$$

(44) $$\qquad \lim_{B \to \infty} \int_{D_x} |g_{B,x}(\theta_i \mathbf{c}) - g_x(\theta_i \mathbf{c})| \, d\lambda = 0.$$

Recall that $g \in \mathscr{G}$, that is

(45) $$g(\mathbf{c}) + \sum_{i=1}^{\infty} g(\theta_i \mathbf{c}) = 1 \quad \lambda\text{-a.e.}$$

Next let us show

(46) $$\lim_{b \to \infty} \limsup_{B \to \infty} \int_{D_x} \left| \sum_{i=b+1}^{B} g_{B,x}(\theta_i \mathbf{c}) \right| + \left| \sum_{i=b+1}^{\infty} g_x(\theta_i \mathbf{c}) \right| d\lambda = 0.$$

This holds because $g_{B,x}(\theta_i \mathbf{c})$ and $g_x(\theta_i \mathbf{c})$ are bounded by $1_{(c(i) \leq x)}$, and so the integral is bounded by

$$2 \sum_{i=b+1}^{\infty} \lambda\{D_x, c(i) \leq x\} = 2x \sum_{i=b+1}^{\infty} P(Z_x \geq i)$$

and the sum is convergent.

Combining (43–46) establishes (42).

3.5 Convergence in probability

Proposition 14 below states an ergodicity property for the measure λ on C. Let us first show how the ergodicity property implies $C(n) \to c^*$ in probability, thus completing the proof of Theorem 1. Knowing $EC(n) \to c^*$, it suffices to show

(47) $$P(C(n) \leq c) \to 0, \quad \text{each } c < c^*.$$

To argue by contradiction, suppose (47) fails for some fixed $c < c^*$. By Proposition 2 we cannot have $\limsup_{n \to \infty} P(C(n) \leq c) = 1$, and so we can pass to a subsequence in which

$$\lim_{n'} P(C(n') \leq c) = a \in (0, 1).$$

Now reconsider the proof (Sect. 3.3) of Proposition 3 (a). Recall $\psi_{B,n}$ is the distribution of a certain random element of the set C_B^2 with generic element $((c(e), p(e)); e \in \tau_B)$. We can write

$$\psi_{B,n} = P(C(n) \leq c) \, \psi_{\bar{B},n}^{\leq} + P(C(n) > c) \, \psi_{B,n}^{>}$$

where $\psi_{\bar{B},n}^{\leq}$ denotes the conditional distribution given $C(n) \leq c$. Taking subsequential limits through n', we can write the ψ of (25) as

$$\psi = a\psi^{\leq} + (1-a)\psi^{>}.$$

Copying the proof of Lemma 12 shows that $\psi^{\leq}$ has the invariance property stated in Lemma 12 for ψ. Recall that λ is the marginal distribution (on C) of ψ. So

$$\lambda = a\lambda^{\leq} + (1-a)\lambda^{>}$$

where $\lambda^{\leqq}$ and $\lambda^{>}$ are marginals of $\psi^{\leqq}$ and $\psi^{>}$. But by invariance the Radon-Nikodym density $F = \dfrac{d\lambda^{\leqq}}{d\lambda}$ has the invariance property (48) below, and so Proposition 14 applied to $a^{-1}f$ shows that $\lambda^{\leqq} = \lambda$. But now the argument in the final paragraph of Sect. 3.3, applied to $\psi^{\leqq}$ instead of ψ, shows

$$c^* \leqq \liminf_{n \to \infty} E(C(n') \,|\, C(n') \leqq c).$$

But the right side is at most c, and since we assumed $c < c^*$ we have reached a contradiction, and thereby established (47).

Proposition 14 *Let $f : C \to [0, 1]$ be measurable and have the invariance property*

(48) $$f(\theta_i \, \mathbf{c}) = f(\theta'_i \, c) = f(\mathbf{c}) \quad \lambda\text{-a.e.,} \quad \text{for each } i.$$

Then f is constant λ-a.e.

This is a variation on standard facts about branching processes (Lemma 15 below) and standard ergodic-theory ideas (mixing implies ergodicity). We first assemble the required ingredients. For $s > 0$ and $\mathbf{c} \in C$ let $\Gamma_s(\mathbf{c})$ be the set of edges e of τ such that, for the path $\phi = e^1, e^2, e^3, \ldots, e^d = e$ we have $c(\phi) + c(e^1) + \ldots + c(e) \leqq s$. Let $|\Gamma_s(\mathbf{c})|$ be the size of $\Gamma_s(\mathbf{c})$ so $|\Gamma_s(\mathbf{c})| = 0$ if $c(\phi) > s$. Recall from Sect. 3.1 the Yule process description of λ. It is well known that the population size N_t at time t in the Yule process (started with one individual at time 0) satisfies $EN_t = e^t$ and

(49) $$N_t / e^t \to W > 0 \text{ a.s.,} \quad (N_t / e^t; t \geqq 0) \quad \text{is uniformly integrable.}$$

Given $c(\phi) = x$, the edges in $\Gamma_s(\mathbf{c})$ correspond to individuals born before time $s - x$, and so there are an expected number e^{s-x} of such individuals on each side of τ; this leads to the calculation

(50) $$\int |\Gamma_s(\mathbf{c})| \, \lambda(d\mathbf{c}) = \int_0^s (2e^{s-x} - 1) \, dx = 2e^s - s - 2.$$

Appealing to (49),

(51) $$\frac{|\Gamma_s(\mathbf{c})|}{2e^s - s - 2} \to r(\mathbf{c}), \quad \text{say, } \lambda\text{-a.e.} \quad \text{as} \quad s \to \infty.$$

(52) $$r(\mathbf{c}) > 0 \ \lambda\text{-a.e.}$$

(53) $$\left(\frac{|\Gamma_s(\mathbf{c})|}{2e^s - s - 2}; s \geqq 0 \right) \quad \text{is uniformly } \lambda\text{-integrable.}$$

Next, recall that the map θ_i (resp. θ'_i) from C to C take edge i (resp. i') to edge ϕ. We can similarly define for each $e \in \tau$ the map $\theta_e : C \to C$ which takes edge e to edge ϕ.

Lemma 15 *There exists a probability distribution $\tilde{\lambda}$ on C such that, for continuous*
f_B: $C_{B,B} \to [0, 1]$

$$\lim_{s \to \infty} \frac{1}{|\Gamma_s(\mathbf{c})|} \sum_{e \in \Gamma_s(\mathbf{c})} f_B(\theta_e \mathbf{c}) \to \int_C f_B(\mathbf{c}) \tilde{\lambda}(\mathbf{c}) \quad \lambda\text{-a.e.}$$

($C_{B,B}$ is $C_{B,H}$ with $H = B$.) This follows, by specialization to the Yule process, from results about general supercritical branching processes: see Jagers and Nerman [6, 9]. In their terminology $\tilde{\lambda}$ is in general the "stable doubly infinite pedigree process". We will see below that in our setting, $\tilde{\lambda}$ can be specified via $\dfrac{d\tilde{\lambda}}{d\lambda} = r(\cdot)$, this being a very special property of the Yule process.

For the final ingredient, let V_s be a random cost tree whose density w.r.t. λ is $\dfrac{|\Gamma_s(\mathbf{c})|}{2e^s - s - 2}$. (This is a probability density, by (50)). Then define V_s^* by: given $V_s = \mathbf{c}$, let V_s^* be uniform on $\{\theta_e \mathbf{c}: e \in \Gamma_s(\mathbf{c})\}$. We now assert we have a symmetry property:

$$(54) \qquad\qquad (V_s, V_s^*) \overset{d}{=} (V_s^*, V_s).$$

The point is that for a pair $(\mathbf{c}, \mathbf{c}^*)$ the property "$\mathbf{c}^* = \theta_e \mathbf{c}$ for some $e \in \Gamma_s(\mathbf{c})$" in equivalent to the property "$\mathbf{c} = \theta_e \mathbf{c}^*$ for some $e \in \Gamma_s(\mathbf{c}^*)$". So for such a pair we have (abusing notation) $P(V_s = \mathbf{c}, V_s^* = \mathbf{c}^*) = \dfrac{\lambda(d\mathbf{c})}{2e^s - s - 2}$. So in showing (54) the issue is to show that $\lambda(d\mathbf{c}) = \lambda(d\mathbf{c}^*)$ for such a pair. This is the assertion, analogous to (18), that λ is invariant under θ_e (on the appropriate domain and range), and this fact follows from (18) iterated along the path to e.

Proof of Proposition 14 Let V_∞ be a random cost tree with density $r(\mathbf{c})$ w.r.t. λ. (This is a probability density, by (50, 51, 53)). Let f_B: $C_{B,B} \to [0, 1]$ be continuous. By (51)

$$(55) \qquad\qquad \lim_{s \to \infty} E f_B(V_s) = E f_B(V_\infty).$$

By Lemma 15, for λ-a.e. $\mathbf{c}$

$$E(f_B(V_s^*) | V_s = \mathbf{c}) \to \int_C f_B(\mathbf{c}^*) \tilde{\lambda}(d\mathbf{c}^*)$$

and then using (53)

$$(56) \qquad\qquad \lim_{s \to \infty} E f_B(V_s^*) = \int_C f_B(\mathbf{c}^*) \tilde{\lambda}(d\mathbf{c}^*).$$

But by (54) the limits in (55, 56) must be the same. In other words, $\tilde{\lambda}$ is the distribution of V_∞ and we can write (56) as

$$\lim_{s \to \infty} E f_B(V_s^*) = E f_B(V_\infty).$$

The point is that similarly by Lemma 15, for λ-a.e. $\mathbf{c}$

$$E(f_B(V_s^*) f_B(V_s) \mid V_s = \mathbf{c}) = f_B(\mathbf{c}) \, E(f_B(V_s^*) \mid V_s = \mathbf{c}) \to f_B(\mathbf{c}) \, E f_B(V_\infty)$$

and then by (53) and (51)

$$(57) \qquad\qquad \lim_{s \to \infty} E f_B(V_s^*) f_B(V_s) = (E f_B(V_\infty))^2.$$

This is a "mixing" property, and now the following routine argument establishes ergodicity. Let $f: C \to [0, 1]$ be invariant. By martingale convergence and approximation of measurable functions by continuous functions there exist continuous $f_B: C_{B,B} \to [0, 1]$ such that

$$(58) \qquad\qquad E \mid f(V_\infty) - f_B(V_\infty) \mid \to 0 \quad \text{as} \quad B \to \infty.$$

Now

$$(59) \quad |E f(V_s) f(V_s^*) - E f_B(V_s) f_B(V_s^*)| \leq E |f(V_s^*) - f_B(V_s^*)| + E |f(V_s) - f_B(V_s)|$$
$$= 2 E |f(V_s) - f_s(V_s)|$$

the inequality because $0 \leq f, f_B \leq 1$ and the equality by (54). And by (51)

$$\lim_{s \to \infty} E |f(V_s) - f_B(V_s)| = E |f(V_\infty) - f_B(V_\infty)|.$$

Thus letting $B \to \infty$ in (59), and using (58, 57)

$$\lim_{s \to \infty} E f(V_s) f(V_s^*) = (E f(V_\infty))^2.$$

But invariance of f implies that $f(V_s^*) = f(V_s)$, and then by (53, 51)

$$E f^2(V_\infty) = (E f(V_\infty))^2.$$

Thus $f(V_\infty)$ is a.s. constant, and so (52) establishes the proposition.

4 Remarks

The definition $c^* = \inf_{g \in \mathscr{G}} \gamma(g)$ at (17) is hard to intuit because it is not clear how to produce *any* example of a function $g \in \mathscr{G}$. The existing proofs [12, 4, 7] that $\lim\sup_{n \to \infty} EC(n) < \infty$ proceed via the marriage lemma or via linear programming and show that assignments exist without explicitly defining them in terms of the cost matrix. Let us briefly mention a different algorithm for a non-optimal assignment which, as $n \to \infty$, does lead to an expression for a $g \in \mathscr{G}$.

The algorithm is simple: start with all edges (j, m) of the bipartite graph; at each stage delete the most expensive edge which can be deleted while still guaranteeing that some matching exists in the remaining graph.

This leads to an "algorithm" for constructing a matching on the infinite tree τ with given costs $(c(e); e \in \tau)$. Regard time t as decreasing from ∞ to 0.

At each time t there is a subtree of τ which contains every vertex and which consists of isolated edges and of infinite components with no leaves. As t decreases the isolated edges are not changed but the infinite components break up as follows. If an edge e remains in an infinite component as $t \downarrow c(e)$ then when $t = c(e)$ the edge e is deleted. This may create one or two leaves l; if so, the remaining edge (l, v) containing the leaf l is made an isolated edge by deleting the other edges at v. This in turn may create another leaf, in which case we continue the "chain reaction" of deleting edges, all this happening instantaneously at time $t = c(e)$. At time $t = 0$ we have a matching on τ. The construction respects the symmetry of τ, and so the associated function

$$g(\mathbf{c}) = 1_{(\text{edge}\,\phi\,\text{in matching})}$$

is in $\mathcal{G}$.

Unfortunately the discrete algorithm, as well as being slow, yields assignments more expensive than the other known algorithms: simulations suggest mean cost around 2.7.

Note added in proof. Use of martingale concentration inequalities to prove Theorem 1 has been sketched by Bing Zhao (Study on the Limit Laws of a Class of Combinatorial Optimization Problems under Independent Model: I.E.O.R. Dept, Columbia University).

References

1. Aldous, D.J.: The continuum random tree II: an overview. In: Barlow, M.T., Bingham, N.H. (eds.) Proc. Durham Symp. Stochastic Analysis 1990. Cambridge: University Press 1992
2. Aldous, D.J.: A random tree model associated with random graphs. Random Structures Algorithms **1**, 383–402 (1990)
3. Aldous, D.J., Steele, J.M.: Asymptotics for Euclidean minimal spanning trees on random points. Probab. Theory Relat. Fields **92**, 247–258 (1992)
4. Dyer, M.E., Frieze, A.M., Mac Diarmid, C.J.H.: On linear programs with random costs. Math. Program. **35**, 3–16 (1986)
5. Goemans, M.X., Kodialam, M.S.: A lower bound on the expected cost of an optimal assignment. Technical Report, Operations Res. MIT (1989)
6. Jagers, P., Nerman, O.: The growth and composition of branching processes. Adv. Appl. Probab. **16**, 221–259 (1984)
7. Karp, R.M.: An upper bound on the expected cost of an optimal assignment. In: Johnson, D.S. et al. (eds.) Discrete algorithms and complexity: Proceedings of the Japan-U.S. Joint Seminar, pp. 1–4. New York: Academic Press 1987
8. Mezard, M., Parisi, G.: On the solution of the random link matching problem. J. Phys. (Paris) **48**, 1451–1459 (1987)
9. Nerman, O., Jagers, P.: The stable doubly infinite pedigree process of supercritical branching populations. Z. Wahrscheinlichkeitstheor. Verw. Geb. **65**, 445–460 (1984)
10. Rhee, W.T., Talagrand, M.: Martingale inequalities, interpolation and NP-complete problems. Math. Oper. Res. **14**, 91–96 (1989)
11. Steele, J.M.: Probability and statistics in the service of computer science: illustrations using the assignment problem. Commun. Stat., Theory Methods **19**, 4315–4329 (1990)
12. Walkup, D.W.: Matchings in random regular bipartite digraphs. Discrete Math. **31**, 59–64 (1980)
13. Walkup, D.W.: On the expected value of a random assignment problem. SIAM J. Comput. **8**, 440–442 (1979)