Weighted tree games

Patrick Bennett* Alan Frieze†

Abstract

We consider a variation on Maker-Breaker games on graphs or digraphs where the edges have random costs. We assume that Maker wishes to choose the edges of a spanning tree, but wishes to minimise his cost. Meanwhile Breaker wants to make Maker's cost as large as possible.

1 Introduction

We consider a variation on Maker-Breaker games on graphs or digraphs where the edges have costs. Maker-Breaker games were intoduced in Erdős and Selfridge [3] and in Chvátal and Erdős [2]. In a traditional Maker-Breaker game, there are two players Maker and Breaker. There is a board X, typically the complete graph K_n and Maker and Breaker take turns in choosing elements of X. Maker has the goal of choosing a set of edges that contain some particular structure, e.g. a spanning tree and Breaker has the goal of preventing this. Typically, Breaker can choose b elements to each element chosen by Maker. If both players play optimally then there is a threshold b^* to the *bias* b such that if $b < b^*$ then Maker wins and if $b \geq b^*$ then Breaker wins. It is clear that there is a rich choice for X and structures in X. For an excellent introduction to this topic see Beck [1] or Hefetz, Krivelevich, Stojaković and Szabo [7]

We now assume that each edge has a random cost. These costs are known to both players before the game starts. Maker's goal is build some some structure, e.g. a spanning tree, but wishes to minimise her total cost. In a play of the game, Maker chooses an edge for her structure and then Breaker deletes b edges. If Maker is unable to build the desired structure at all we say the cost is infinity. Maker will have to avoid this situation to get a meaningful result. So, given any Maker-Breaker game in the traditional sense, one can make a weighted version where the object of Maker is to optimize the cost of the structure she can build. In this paper we use the standard asymptotic notation, i.e. $o(), O()$, and asymptotics are as $n \to \infty$ unless otherwise noted. We say an event $\mathcal{E}_n$ happens *with high probability (w.h.p.)* if $\mathbb{P}(\mathcal{E}_n) \to 1$.

We begin with a simple case that does not involve graphs: let $\mathbb{N} = \{1, 2, \dots, \}$ be the set of positive integers. Maker in her turn has to choose an $i \in [n]$ and irrevocably assign a value to $f(i) \in \mathbb{N}$. We let M_t denote the set of elements $i \in [n]$ that have been selected in this way by Maker after t rounds of play. Breaker in his turn selects $i_1, i_2, \dots, i_b \in [n] = \{1, 2, \dots, n\}$ and $j_1, j_2, \dots, j_b \in \mathbb{N}$ and makes j_r unavailable to Maker for the value of $f(i_r)$, $r = 1, 2, \dots, b$. We let B_t denote the set of pairs (i, j) for which Breaker has made $f(i) = j$ unavailable to Maker after t rounds. Thus $M_0 = B_0 = \emptyset$ and in round t, Maker adds one element to M_t to create M_{t+1} and Breaker adds b pairs to B_t to create B_{t+1}.

*Research supported in part by Simons Foundation Grant #426894.

†Research supported in part by NSF grant DMS1952285

Let $\phi(f) = \sum_{i=1}^n f(i)$. Maker's aim is to keep ϕ as small as possible and Breaker has the opposite intention. Our first result is the following.

Theorem 1. *If Breaker goes first then Maker can choose f such $\phi(f) \leq (b+1)n$ and this is optimal. If Maker goes first then Maker can choose f such $\phi(f) \leq (n-1)b+n$ and this is optimal*

In preparation for more complicated scenarios, suppose that we replace $\mathbb{N}$ by $[m]$ i.e. we insist that $f(i) \leq m$.

Theorem 2. *Suppose that $m > \sum_{i=1}^n 1/i$. Then Maker can choose f such $f(i) \leq m$ for $i \in [n]$ and $\phi(f) \leq (b+1)n$ and this is optimal.*

Suppose next that instead of paying $j = f(i)$, Maker pays $X(i,j)$ where the X-values are independent uniform $[0,1]$ random variables and j ranges over $[m]$ where $m = \rho^{-1}n$ and where $m > \sum_{i=1}^n 1/i = (1+o(1))\log n$. We assume that these values are known to the players at the start of the game.

Theorem 3. *In this uniform random scenario, Maker can pay at most $\mu^{>1}(b)(b+1+o(1))\rho$, w.h.p., where $\mu^{>1}(b)$ is the solution to $\frac{\log(b+1)+1}{b+1} = \mu - 1 - \log\mu$ in $(1,\infty)$. (Note that $\mu^{>1}(b) = 1 + O(\frac{\log^{1/2} b}{b^{1/2}})$ as $b \to \infty$.) This is asymptotically optimal in the sense that w.h.p. for sufficiently large b Breaker can force Maker to pay at least $\mu^{<1}(b)(b+1-o(1))\rho$ where $\mu^{<1}(b)$ is the solution to $\frac{\log(b+1)+1}{b+1} = \mu - 1 - \log\mu$ in $(0,1)$. Note that the upper and lower bounds differ only in their $o(1)$ terms.*

Now we turn to a more complex problem. Here the *board* is the set of edges of the complete (loopless) digraph $\vec{K}_n$. There is an $n \times n$ cost matrix C. For each $i \in [n]$ we have $C(i,i) = \infty$ and the sequence of values $C(i,j), j \in [n] \setminus \{i\}$ is an independent uniform random permutation of $[n-1]$. Maker's aim is to construct a spanning arborescence T of low total cost, where the cost $C(T) = \sum_{(i,j)\in E(T)} C(i,j)$. (A spanning arborescence is a spanning tree whose edges have been oriented towards one vertex, the root.)

After t rounds, Maker will have selected a set of t edges M_t that induce a digraph where each vertex has out-degree at most one. Furthermore, these edges induce a forest when edge orientation is ignored. Similarly, Breaker will have selected a set B_t of bt edges where $M_t \cap B_t = \emptyset$.

Theorem 4. *W.h.p., over the random choice of C, Maker can construct an arborescence T of cost $C(T) \leq \big(b/\theta^* + b + 1 + o(1)\big)n$, where $\theta^* \approx 0.2938....$ (θ^* is the solution to $(1+\theta^*)\log(1+\theta^*) - \theta^*\log\theta^* = \log 2$.)*

We doubt that this is optimal for Maker.

Conjecture: W.h.p. Maker can construct an arborescence T of cost $C(T) \leq (b+1+o(1))n$.

We base this conjecture on the fact that if $f : [n] \to [n]$ is a uniform random mapping then the digraph with edge-set $\{(i, f(i))\}$ is almost a random arborescence, i.e. it differs from a arborescence by $O(\log n)$ edges. (This follows from Joyal's proof [9] of Cayley's formula for the number of spanning trees in K_n. Joyal describes a 1 to n^2 mapping from spanning trees to maps from $[n] \to [n]$.)

Theorem 5. *If the costs of the edges (i,j) are independent uniform $[0,1]$ random variables then Maker can pay at most $\mu^{>1}(b/\theta^* + b)(b/\theta^* + b + 1) + o(1)$, w.h.p.*

We now consider the undirected versions of Theorem 4 and Theorem 5. I.e., now we have a weighted complete graph and Maker wishes to build a low cost spanning tree. We note that Hefetz, Kupferman, Lellouche and Vardi [8] considered a worst-case version of this problem, in the context of finding a maximum weight spanning tree.

In the following theorem the edges of the complete graph are given independent uniform $[0,1]$ costs.

Theorem 6. *Maker can build a spanning tree of cost at most* $2\mu^{>1}(2b/\theta^* + 2b)(2b/\theta^* + 2b + 1) + o(1)$, *w.h.p.*

2 Mappings

2.1 Proof of Theorem 1

Proof. Maker's strategy is simply to choose an arbitrary $i \notin M_t$ and put $f(i) = r_i(t)$ and so that $M_{t+1} = M_t \cup \{i\}$. Note that $f(i)$ is at most 1 plus the number of pairs (i, j) that have been taken by Breaker. Since Breaker takes b pairs per turn, the total amount $f(1) + \cdots + f(n)$ paid by Maker is at most n plus b times the number of turns Breaker takes. If Breaker goes first then Breaker gets n turns and Maker pays at most $n + bn$. If Maker goes first then Breaker gets $n - 1$ turns and Maker pays at most $n + b(n - 1)$.

Both results are optimal. For $i \notin M_t$ we let $r_i(t) = \min \{j : (i, j) \notin B_t\}$ be the minimum of the possible values for $f(i)$ available to Maker at the end of round t. Let $\delta_i(t)$ be the indicator for $i \notin M_t$. We use the following potential

$$\Phi(t) = \sum_{i=1}^{n} r_i(t)\delta_i(t).$$

We observe that Breaker on his turn can increase Φ by at least b by choosing the pairs (i, j) for some $i \notin M_t$ and for j ranging through the b smallest values such that (i, j) is not yet taken, making $r_i(t + 1) \geq r_i(t) + b$. Thus, if Breaker goes first we now have

$$-n = \sum_{i=0}^{n-1}(\Phi(t + 1) - \Phi(t)) \leq \sum_{i=0}^{n-1}(b - r_i(t)) = nb - \sum_{i=0}^{n-1} r_i(t) \geq nb - \sum_{i=1}^{n} f(i),$$

which proves our result is optimal in this case. The case where Maker goes first is similar. This completes the proof of Theorem 1. □

2.2 Proof of Theorem 2

Proof. Here Maker has to be more careful, since she must prevent Breaker from taking all the pairs (i, j) for any i. Maker will exploit a version of the Box Game of Chvátal and Erdős [2]. In this game there are n disjoint sets $A_1, A_2, \ldots, A_n$, the boxes. The elements of the A_i are called balls. There are two players, BoxMaker and BoxBreaker. In each round BoxMaker removes p balls from $A = \bigcup_{i=1}^{n} A_i$ and BoxBreaker removes q boxes. BoxMaker's goal is to remove all the balls from some box before BoxBreaker is able to remove that box. In the context of our mapping game, we let $A_i = \{(i, j), j \in [m]\}$. Maker takes the role of BoxBreaker with $q = 1$ and Breaker takes the role of BoxMaker with $p = b$. After t rounds there will be $n - t$ boxes remaining and their contents will have been reduced. We will assume that BoxMaker (Breaker) goes first and the BoxBreaker (Maker) always chooses a remaining box A_i of minimum size (and puts $f(i)$ equal to the smallest element left in A_i).

Theorem 3.4.1 of [7] shows that if $|A_i| = m > \sum_{i=1}^{n} 1/i$ then the above strategy for BoxBreaker (Maker) guarantees her a win the BoxGame, i.e. she prevents Breaker from taking all the pairs (i, j) for any i. Also, by the analysis of Section 2, she will end with a value of $\phi(f) \leq (b+1)n$. This completes the proof of Theorem 2. □

2.3 Proof of Theorem 3

Proof. Assume without loss of generality that $X(i,j+1) \geq X(i,j)$ for $1 \leq j \leq m$ for all i (in other words, to form the sequence $X(i,1),\dots,X(i,n)$ we generate m random values and then put them in decreasing order). For an upper bound we assume that Maker tries to choose f so as to minimise $\sum_{i=1}^n f(i)$, i.e. she essentially plays the game in Theorem 1. We need the following lemma from Frieze and Grimmett [5].

Lemma 7. *Suppose that $k_1+k_2+\cdots+k_M \leq aN$, and $Y_1,Y_2,\dots,Y_M$ are independent random variables with Y_i distributed as the k_ith smallest of N independent uniform [0,1] random variables. If $\mu > 1$ then*

$$\mathbb{P}\left(Y_1+\cdots+Y_M \geq \frac{\mu a N}{N+1}\right) \leq e^{aN(1+\log\mu-\mu)}. \tag{1}$$

(The lemma in [5] is given in terms of $a\log N$ instead of aN. We have replaced a by $aN/\log N$.)

Now naively, we could observe that $\mathbb{E}[X(i,j)] = j/(m+1)$ and then, at least in expectation, we could replace a cost of j in the model of Theorem 1 by $j/(m+1)$, giving us a bound that is asymptotic to $b+1$. We must however take account of the variability in $X(i,j)$ and so we are forced to take a union bound over Maker's possibilities. This leads to the claimed inflated constant.

There are at most $\binom{(b+1)n-1}{n-1} \leq \binom{(b+1)n}{n}$ choices for the $f(i) \geq 1$ that add up to $(b+1)n$. Let F denote this set of choices. For a fixed $\mu > 1$ we have

$$\begin{aligned}\mathbb{P}\left(\exists f \in F : \sum_{i=1}^n X(i,f(i)) \geq \mu(b+1)\rho\right) &\leq \binom{(b+1)n}{n} e^{-(b+1)n(1+\log\mu-\mu)} \\ &\leq (e(b+1))^n e^{-(b+1)n(1+\log\mu-\mu)} \\ &= \left((b+1)e^{1-(b+1)(1+\log\mu-\mu)}\right)^n = o(1),\end{aligned} \tag{2}$$

if $\mu > \mu^{>1}(b)$.

We now go about proving that Breaker can force Maker to pay at least $\mu^{<1}(b)(b+1-o(1))\rho$. Of course Breaker can play so that $\sum_{i=1}^n f(i) \geq (b+1)n$. The argument of Lemma 4.2(b) of [5] can be used to prove that for $1/2 < \mu < 1$ and $k_1+k_2+\cdots+k_M \geq aN$,

$$\mathbb{P}\left(Y_1+\cdots+Y_M \leq \frac{\mu a N}{N+1}\right) \leq e^{aN(1+\log\mu-\mu)}. \tag{3}$$

Indeed, let $Y := Y_1+\cdots+Y_M$, $y = \frac{\mu a N}{N+1}$ and $t := (\mu^{-1}-1)(N+1)$. Note that $0 < t < N+1$. We have

$$\mathbb{P}(Y \leq y) = \mathbb{P}\left(e^{-tY} \geq e^{-ty}\right) \leq e^{ty}\mathbb{E}\left[e^{-tY}\right]. \tag{4}$$

Now, following the proof of Lemma 4.2(b) of [5], we have

$$\mathbb{E}\left[Y_i^j\right] = \int_0^1 y^{k_i+j-1}(1-y)^{N-k_i} = \binom{N}{k_i} k_i \frac{(j+k_i-1)!(N-k_i)!}{(N+j)!} \leq \frac{(k_i+j-1)_j}{(N+1)^j}.$$

Since $|t/(N+1)| < 1$, by Newton's binomial formula we have

$$\mathbb{E}\left[e^{-tY_i}\right] \leq \sum_{j=0}^{\infty} \frac{(-t)^j}{j!} \cdot \frac{(k_i+j-1)_j}{(N+1)^j} = \left(1+\frac{t}{N+1}\right)^{-k_i}.$$

Therefore

$$\mathbb{E}\left[e^{-tY}\right] \le \left(1 + \frac{t}{N+1}\right)^{-aN}.$$

Picking up from (4), we have

$$\mathbb{P}(Y \le y) \le e^{ty}\left(1 + \frac{t}{N+1}\right)^{-aN} = e^{aN(1+\log\mu-\mu)}.$$

Thus, just as in (2) we have for $\mu < \mu^{<1}(b)$

$$\mathbb{P}\left(\exists f \in F : \sum_{i=1}^{n} X(i, f(i)) \le \mu(b+1)\rho\right) \le \left((b+1)e^{1-(b+1)(1+\log\mu-\mu)}\right)^{n} = o(1)$$

This completes the proof of Theorem 3. □

3 Arborescences

3.1 Proof of Theorem 4

Proof. First we describe Maker's strategy. Let F_t denote the set of oriented trees induced by M_t. Each of these trees/components will have a root and Maker must choose an edge leaving one of these roots and going to a different component in F_t. This situation is almost like the setup for Theorem 2, where each vertex $i \in [n]$ (except for one) must choose an out-neighbor $f(i) \in [n]$. However our present situation is complicated by the graph structure, i.e. an arborescence cannot have any cycle. Maker's strategy for Theorem 2 was to always take the minimum available $f(i)$ for some i, but in our present situation it is possible that this minimum $f(i)$ would correspond to an edge that would create a cycle, forcing Maker to take some slightly more expensive edge. Our proof will focus on the "extra" amount paid by Maker on these steps where the cheapest edge cannot be chosen.

On Maker's turn, if there is no root of any component on at most $n/2$ vertices such that Breaker has taken at least n^{β} edges leaving the root, we call this a *normal turn* for Maker. Here $0 < \beta < 1$ is a constant to be determined later. On a normal turn, Maker always chooses the root i of the smallest component K (in case of a tie, say choose the least indexed root). Maker then chooses the edge $(i, j) \notin B_t$ that (i) minimises $C(i, k), k \notin B_t$ and (ii) does not point into K. We call this the *sensible* choice from root i.

If it is not a normal turn we say it is an *emergency turn*. On an emergency turn, Maker picks some component on at most $n/2$ vertices from which Breaker has removed at least n^{β} edges, and Maker makes the sensible choice out of this root. We will argue that w.h.p. this strategy gets Maker an arborescence whose total cost is at most $(b/\theta^* + b + 1 + o(1))n$.

Consider the event to the contrary, i.e. that the total cost in the end is say $(b/\theta^* + b + 1 + 3\varepsilon)n$ for some fixed $\varepsilon > 0$, meaning that the extra cost paid for edges pointing within components is $(b/\theta^* + 3\varepsilon)n$, over and above the $(b+1)n$ achievable for Theorem 2. We bound the probability of this event as follows. First it is at most $\pi_1 + \pi_2 + \pi_3$, defined as follows. π_1 is the probability that we pay an extra εn on normal steps falling in the first n_0 steps, where $n_0 = n - n^{\alpha}$ for some $0 < \alpha < 1$. π_2 is the probability of paying εn on normal steps after step n_0. π_3 is the probability of paying $bn/\theta^* + \varepsilon n$ on emergency steps.

First we bound π_1. Let N_1 be the set of normal steps in the first n_0 steps. We take the union bound over choices for numbers $a_t, t \in N_1$ where the intended meaning of the a_t is as follows. At step t, Maker takes an edge from a root to another arborescense, paying an extra a_t because of edges pointing into its own component. So we have $\sum_t a_t = \varepsilon n$, and the number of choices for the a_t is $\binom{\varepsilon n+|N_1|-1}{|N_1|-1} = \exp\{O(n)\}$.

We will also take the union bound over sequences $x_t, t \in N_1$ which will mean the following. At step i, if v is the root that Maker is taking an edge from and Maker chooses the edge of cost r_t, $x_t = r_t - 1 - a_t$. In other words, the reason Maker has to choose the edge of cost $r_t = 1 + a_t + x_t$ is because among the edges of smaller cost, a_t of them point into v's component and x_t edges have been taken by Breaker. In particular, knowing x_t and a_t tells us the cost r_t of the edge that Maker will choose at step t. We have $\sum_t x_t \leq bn$, and the number of choices for the x_t is $\exp\{O(n)\}$.

Having fixed the a_t and x_t we will reveal the random digraph step by step as the game progresses. More specifically, on Maker's turn at step $t \leq n-1$ we see the current component structure which uniquely determines the root, say v from which Maker will take an edge. We reveal the costs of all edges coming from v, which determines which edge Maker will take (i.e. the lowest cost edge which is not taken by Breaker and which does not point into v's component). Recall that among all the edges from v not taken by Breaker, the a_t lowest in cost all point into v's component. The size of the smallest component is at most $\left\lfloor \frac{n}{n-t+1} \right\rfloor$ and so each out-edge has probability at most $\frac{\frac{n}{n-t+1}-1}{n} = \frac{t}{n(n-t+1)}$ of pointing into the component. There are at most $\binom{a_t+x_t}{a_t}$ choices for the costs of the a_t edges pointing within the component. Thus we bound

$$\begin{aligned}
\pi_1 &\leq \exp\{O(n)\} \prod_{t \in N_1} \binom{a_t+x_t}{a_t} \left(\frac{t}{n(n-t+1)}\right)^{a_t} \\
&\leq \exp\{O(n)\} \prod_{t \in N_1} \left(e\left(1+\frac{x_t}{a_t}\right)\right)^{a_t} \left(\frac{t}{n(n-t+1)}\right)^{a_t} \\
&\leq \exp\{O(n)\} \prod_{t \in N_1} \left(\frac{t}{n(n-t+1)}\right)^{a_t} \\
&\leq \exp\{O(n)\} \prod_{t \in N_1} \left(\frac{n-n^\alpha}{n^{1+\alpha}}\right)^{a_t} \leq n^{-\alpha\varepsilon n + o(n)} = o(1). \qquad (5)
\end{aligned}$$

We bound π_2 similarly. Let N_2 be the set of normal steps after step n_0. We choose numbers $a_t, t \in N_2$ adding up to εn and there are at most $\binom{\varepsilon n+n^\alpha}{n^\alpha} = \exp\{o(n)\}$ choices here. Likewise there are $\exp\{o(n)\}$ choices for the $x_t, i \in N_2$ where $\sum_{t \in N_2} x_t \leq n^{\alpha+\beta}$. To bound the number of choices for the interior edges we can instead choose Breaker's edges and there are at most n^{x_t} ways to do that. So

$$\begin{aligned}
\pi_2 &\leq \exp\{o(n)\} \prod_{t \in N_2} n^{x_t} \left(\frac{t}{n(n-t+1)}\right)^{a_t} \\
&\leq \exp\{o(n)\} n^{\sum_{t \in N_2} x_t} \frac{1}{2^{\sum_{t \in N_2} a_t}} \\
&\leq \exp\{o(n)\} \times n^{n^{\alpha+\beta}} \times 2^{-en} = o(1)
\end{aligned}$$

assuming only that $\alpha + \beta < 1$.

Finally we bound π_3. Let N_3 be the set of emergency steps. Note first that there are at most $bn^{1-\beta}$ emergency steps, so we take the union bound over at most $n^{bn^{1-\beta}} = \exp\{o(n)\}$ choices for N_3. We fix numbers $a_t \geq n^\beta, t \in N_3$ adding up to say an where $a = b/\theta^* + \varepsilon$, there being $\exp\{o(n)\}$ choices. We also fix numbers $x_t, t \in N_3$ adding up to say $X \leq bn$, there being $\exp\{o(n)\}$ choices. Since on an emergency step we always have a root of a component on at most $n/2$ vertices, the probability of an edge landing in the same component is at most $1/2$. Thus we bound π_3 by

$$\exp\{o(n)\} \prod_{t\in N_3} \binom{a_t+x_t}{a_t} \left(\frac{1}{2}\right)^{a_t} \leq 2^{-an+o(n)} \prod_{t\in N_3} e^{(a_t+x_t)\log(a_t+x_t)-a_t\log(a_t)-x_t\log(x_t)}. \tag{6}$$

Suppose now that we fix the values for $a_t, t\in N_3$. Let

$$\phi = \sum_{t\in N_3} f(a_t, x_t) \text{ where } f(a,x) = (a+x)\log(a+x) - a\log(a) - x\log(x).$$

We argue next that to maximise ϕ we must have x_t/a_t taking the same value for all $t \in N_3$. Consider the function $g(x) = f(a,x) + f(b, L-x)$ for some $a, b, L > 0$. Then

$$g'(x) = \log(a+x) - \log(x) - \log(b+L-x) + \log(L-x).$$
$$g''(x) = -\frac{a}{x(a+x)} - \frac{b}{(L-x)(b+L-x)} < 0.$$

So, g is strictly concave and its derivative vanishes when $(a+x)(L-x) = x(b+L-x)$ equivalently when $\frac{a}{x} = \frac{b}{L-x}$. Thus, if we fix $a_t, t\in N_3$ and maximise over $x_t, t \in N_3$ then we have $x_t = \theta a_t$ for $t \in N_3$ where $\theta = X/an \leq b/a$. We can therefore bound the product in (6) by

$$\prod_{t\in N_3} e^{(a_t(1+\theta))\log(a_t(1+\theta))-a_t\log(a_t)-\theta a_t\log(\theta a_t)} = \prod_{t\in N_3} e^{a_t((1+\theta)\log(1+\theta)-\theta(\log\theta))} = e^{an((1+\theta)\log(1+\theta)-\theta\log\theta)}. \tag{7}$$

Now recall that $\theta^* \approx 0.2938...$ is defined as the root of $(1+\theta^*)\log(1+\theta^*) - \theta^*\log\theta^* = \log 2$. Since $a > b/\theta^*$ we have $\theta^* > b/a \geq \theta$ and so $(1+\theta)\log(1+\theta) - \theta\log\theta < \log 2$. Now by (6) and (7) we have $\pi_3 = o(1)$.

We finally note that with the above Maker strategy, if a component reaches size greater than $n/2$ then its root will become the root of the final arborescence. This completes the proof of Theorem 4. $\square$

3.2 Proof of Theorem 5

Proof. As in Theorem 3, Maker just tries to minimise the sum of the ranks in the order statistics of the $X(i,j)$'s. By Theorem 4, w.h.p. Maker can achieve a rank sum of at most $s := \sigma n$ where $\sigma = b/\theta^* + b + 1 + o(1)$. The number of choices of ranks which have sum s is $\binom{s-1}{n-1} \leq \binom{s}{n}$. Using Lemma 7 and the union bound, the probability that the total cost of our arborescence exceeds $\frac{\mu s}{n} = \mu\sigma$ is at most

$$\binom{s}{n} e^{-s(1+\log\mu-\mu)} \leq \left(\frac{es}{n}\right)^n e^{-s(1+\log\mu-\mu)}$$
$$= \left(\sigma e^{1-\sigma(1+\log\mu-\mu))}\right)^n = o(1) \tag{8}$$

assuming that $\mu > \mu^{>1}(b/\theta^* + b) = (1+o(1))\mu^{>1}(\sigma - 1)$. Thus, w.h.p. the cost of our arborescence is at most

$$\mu\sigma = (1+o(1))\mu^{>1}(\sigma-1)\sigma = \mu^{>1}(b/\theta^* + b)(b/\theta^* + b + 1) + o(1),$$

completing the proof.

$\square$

4 Spanning Trees

4.1 Proof of Theorem 6

Proof. We reduce this to Theorem 5. We replace each edge $\{i, j\}$ with a pair of directed edges $(i, j), (j, i)$. Each directed edge (i, j) is given a random cost $\widehat{C}(i, j)$, which is an independent copy of the $[0, 1]$ random variable Z where $\mathbb{P}(Z > x) = (1 - x)^{1/2} \leq 1 - \frac{x}{2}$ for $0 \leq x \leq 1$. If Z_1, Z_2 are two independent copies of Z, then $\min\{Z_1, Z_2\}$ is distributed as a uniform $[0, 1]$ random variable. This is a nice idea, employed by Walkup [10] in bounding the expected value of a random assignment problem. Note that Z is dominated by $2U[0, 1]$ where $U[0, 1]$ is uniform on $[0, 1]$.

Given the above construction, Maker builds a spanning arborescence. Thus we look to Theorem 5, but we must make some adjustments. Using Z in place of $U[0, 1]$ at most doubles the cost of each selected edge. Also, if Breaker deletes (i, j) then he must also delete (j, i). So we double Breaker's power by replacing b by $2b$. Thus, our upper bound here is twice what we get from Theorem 5 by replacing b with $2b$. $\square$

5 Final thoughts

We have studied an interesting class of Maker-Breaker games where Maker's goal is build something cheaply. Our results are not all tight and we believe that there is a general meta theorem that states for many such games, the existence of Breaker increases the cost of the optimum solution by a factor of $(b + 1)$ on average. More precisely, suppose that if there is no Breaker and that Maker choosing optimally, can w.h.p. build a *structure* of cost Z^*. Then in the presence of Breaker, Maker can build a structure of cost $(b + 1 + o(1))Z^*$. As open problems we could list, in preceived order of difficulty, the following structures: mappings; spanning arborescences; spanning trees; perfect matchings in the bipartite graph $K_{n,n}$; perfect matchings in K_n; Hamilton cycles in K_n i.e. a version of the TSP; hypergraph versions of these.

References

[1] J. Beck, Combinatorial Games: Tic-Tac-Toe Theory, Encyclopedia of Mathematics and its Applications 114, Cambridge University Press, 2008.

[2] V. Chvátal and P. Erdős, Biased Positional Games, *Annals of Discrete Mathematics* 2 (1978) 221-229.

[3] P. Erdős and J. Selfridge, Ona combinatorial game, *Journal of Combinatorial Theory A* 14 (1973) 298-301.

[4] A.M. Frieze, On the value of a random minimum spanning tree problem, *Discrete Applied Mathematics* 10 (1985) 47-56.

[5] A.M. Frieze and G. Grimmett, The shortest path problem for graphs with random arc-lengths, *Discrete Applied Mathematics* 10 (1985) 57-77.

[6] A.M. Frieze and M. Karoński, Introduction to random Graphs, Cambridge University Press, 2015

[7] D. Hefetz, M. Krivelevich, M. Stojaković and T. Szabo, Positional Games, Birkhauser, 2014.

[8] D. Hefetz, O. Kupferman, A. Lellouche and G. Vardi, Spanning-tree games, *43rd International Symposium on Mathematical Foundations of Computer Science* (MFCS18).

[9] A. Joyal, Une théorie combinatoire des séries formelles, *Advances in Mathematics* 42 (1981) 1-82.

[10] D.W. Walkup, On the expected value of a random asignment problem, *SIAM Journal on Computing* 8 (1979) 440-442.