

ONLINE PURCHASING UNDER UNCERTAINTY

ALAN FRIEZE AND WESLEY PEGDEN

ABSTRACT. Suppose there is a collection x_1, x_2, \dots, x_N of independent uniform $[0, 1]$ random variables, and a hypergraph \mathcal{F} of *target structures* on the vertex set $\{1, \dots, N\}$. We would like to purchase a target structure at small cost, but we do not know all the costs x_i ahead of time. Instead, we inspect the random variables x_i one at a time, and after each inspection, choose to either keep the vertex i at cost x_i , or reject vertex i forever.

In the present paper, we consider the case where $\{1, \dots, N\}$ is the edge-set of a complete graph (or digraph), and the target structures are the spanning trees of a graph, spanning arborescences of a digraph, the paths between a fixed pair of vertices, perfect matchings, Hamilton cycles or the cliques of some fixed size.

1. INTRODUCTION

Suppose we inspect independent and uniform $[0, 1]$ random variables x_1, x_2, \dots, x_N one at a time. After each i th inspection, we decide to pay the cost x_i to purchase the index i , or pass. If we *must* purchase some i , what is the minimum expected cost of an optimal strategy? More generally, given if we are given a *target hypergraph* \mathcal{F} on the vertex set $\{1, \dots, N\}$, what is the minimum expected cost (in terms of \mathcal{F}) to purchase elements covering at least one set $I \in \mathcal{F}$?

This problem is closely related to two well-studied and well-generalized problems. In the setting of the *prophet inequality*, we have a sequence x_1, x_2, \dots , of random variables and we wish to determine a stopping time T such that $V = x_T$ is as large as possible compared to the maximum M of the sequence. A result of Krengel and Sucheston and Garling [23] states that if x_1, x_2, \dots , are independent and non-negative then

$$V \geq M/2$$

and the bound 2 is sharp. For i.i.d. x_i , 2 is no longer the best constant. Hill and Kertz [16] proved a bound of 1.6 which was improved by Kertz [22]. Optimal stopping rules can be complex and difficult to implement. Samuel-Cahn [27] studied threshold rules and showed that the constant 2 can be achieved using one. See Hill and Kertz [17] for a survey on this problem.

Another related problem is the secretary problem. Here $x_1, x_2, \dots, x_N \geq 0$ are presented in random order; at each step we decide whether to accept the presented element. The goal is to maximise the probability that one accepts the item with the largest value. Ferguson

Research supported in part by NSF Grants DMS1362785, CCF1522984 and a grant(333329) from the Simons Foundation.

Research supported in part by NSF grant DMS1363136.

[11] surveys the literature on this problem. A related, but different, goal is to maximize the expected total weight of the accepted element. The problem is attributed by Ferguson to Cayley [5]. It was solved by Moser [26]. This has been generalised to the matroid secretary problem where weights of the elements of a matroid, presented in random order; at each step we decide whether to accept the presented element, subject to the condition that the set of accepted elements must be independent in the matroid. The goal is to maximize the expected total weight of the accepted set, see Babaioff, Immorlica and Kleinberg [2]. If the matroid has rank one then we have Cayley's problem. For more on these directions, see [9, 19, 20, 28].

Note that interestingly enough. Kleinberg and Weinberg [24] manage to combine the two problems into the discussion of Matroid Prophet Inequalities.

For our problem, the x_i are considered as costs, and the aim is minimize rather than maximise the weight of the selected items. We are interested in the minimum expected cost paid by an optimal algorithm which always succeeds at purchasing every element of at least one hyper-edge of the target hypergraph \mathcal{F} , which we call the *purchase price* of \mathcal{F} . In the bulk of the paper, the vertices of our hypergraph are the edges of K_n and so $N = \binom{n}{2}$. I.e. our target \mathcal{F} will correspond to some graph structure such a path or a tree or a cycle.

Presumably, the order in which the x_i 's are examined should have a large effect on the expected price paid for structures. We will in fact concern ourselves with three distinct models, in which we have progressively less control over the order of inspection.

Freely Ordered Online Model – FOM: In this model, at each step, we are allowed to inspect any uninspected x_i to see if we wish to purchase it. It is therefore an on-line model where we choose the order of the items.

Randomly Ordered Online Model – ROM: In this model the x_i are presented to us in random order x_{i_1}, \dots, x_{i_N} . At each step t we learn the index i_t and the associated cost x_{i_t} , and must decide immediately whether or not to purchase e_i .

Adversarially Ordered Online Model – AOM In this model, the order of inspections is determined by an *adaptive* adversary. (The adversary does not know the costs of uninspected x_i 's.)

Note that when evaluating the purchase price of \mathcal{F} , the expected cost is computed just over a probability space of weights for FOM and AOM, but in a product space of weights and edge-orderings, for ROM.

Surprisingly, for the problems we consider, we achieve upper bounds in ROM which are close to our FOM lower bound. In particular, the effect of optimum control over the order of inspection (versus random order) is small for these problems. We will also show that some of the upper bounds in ROM can also be achieved even in AOM. On the other hand, there will be cases where there is a substantial gap between what can be achieved in these models.

For most of the paper x_i will be the cost of the edge e_i of the complete graph (or digraph) on n vertices, in which case $N = \binom{n}{2}$ (or $n(n-1)$). Each hyper-edge of \mathcal{F} will correspond to

a particular desired graphical structure. As an example, suppose that the hyper-edges of \mathcal{F} correspond to the edge-sets of paths between vertices 1 and 2. A minimum cost hyper-edge of \mathcal{F} corresponds to a shortest path between vertices 1 and 2 in K_n . We know that w.h.p.¹ this is asymptotically $\frac{\log n}{n}$ [21]. We will prove (see Theorem 1.3) that even in the FOM framework, an optimal strategy to purchase such a path pays much more in expectation.

We will begin our discussion though with something closely related to the famous “secretary problem”. In section 3 we consider the particularly simple case of determining the purchase price $\rho_{N,k}$ where \mathcal{F} is the complete k -uniform hypergraph on N vertices; in other words, the task is simply to purchase any k items. (Note that the S_n symmetry of the hypergraph means that FOM, ROM and AOM are equivalent for this hypergraph.) This case will arise as a tool in many of our more complicated analyses. We call this the k -purchase problem.

Theorem 1.1. *For a fixed $k \geq 1$ we have*

$$(1) \quad \rho_{N,k} \approx \frac{c_k}{N} \text{ where } c_1 = 2 \text{ and } c_k = c_{k-1} + 1 + \sqrt{1 + 2c_{k-1}} \text{ for } k \geq 2.$$

Here we use the notation $a_n \approx b_n$ to indicate that $\lim_{n \rightarrow \infty} a_n/b_n = 1$.

The following corollary will be needed in the proof of Theorem 1.9. We only need it for $k = 1$, although we claim it for arbitrary k .

Corollary 1.2. *If we replace uniform $[0, 1]$ by a distribution that has a bounded support and a density $F(x) = D + \psi(x)$ where $D > 0$ is a constant and $\psi(x) = O(x)$ as $x \rightarrow 0$ then we simply replace $\frac{c_k}{N}$ by $\frac{c_k}{DN}$.*

With these tools available, we can evaluate the cost of finding a shortest path between two specified vertices.

Theorem 1.3. *The FOM and ROM purchase prices of a path between vertices 1 and n in K_n are both $n^{-2/3 \pm o(1)}$. The AOM price is $\Omega(n^{-1/2})$.*

This stands in stark contrast to the length $\approx \frac{\log n}{n}$ we could achieve if we were allowed to examine all the edges ahead of time.

A similarly stark increase in cost occurs when we aim for the seemingly simple target of a triangle. Although for arbitrary $\omega \rightarrow \infty$, a uniform $[0, 1]$ weighting of K_n has a triangle of cost $O\left(\frac{\omega}{n}\right)$, w.h.p., we prove:

Theorem 1.4. *The FOM and ROM purchase prices of a triangle in K_n are both $n^{-4/7 \pm o(1)}$. The AOM price is $\Omega(n^{-1/2})$.*

We can generalize the ROM upper bound in Theorem 1.4 to the problem of purchasing a copy of a clique K_r . We will prove that if $\kappa_{r,n}^{\text{ROM}}$ is the expected cost of purchasing a copy of $K_r, r \geq 3$ in ROM, then

¹A sequence of events $\mathcal{E}_n, n = 1, 2, \dots$ is said to occur with high probability (w.h.p.) if $\lim_{n \rightarrow \infty} \Pr(\mathcal{E}_n) = 1$.

Theorem 1.5.

$$\kappa_{r,n}^{\text{ROM}} \leq O\left(\frac{1}{n^{d_r+o(1)}}\right)$$

$$d_3 = \frac{4}{7}, d_4 = \frac{2}{9} \text{ and } d_r = \frac{1}{11 \cdot 2^{r-5} - 1}, r \geq 5.$$

We note that the distribution of the minimum cost of a clique when we are allowed to examine all the edges is of the order $n^{-2/(k-1)}$, see [15].

Despite these examples, purchasing a structure in our model is not always so prohibitively costly, compared with the minimum weight structure available. In particular, recall that the minimum cost spanning tree for a uniform $[0, 1]$ weighting of K_n has asymptotic length $\zeta(3)$ w.h.p. [12]

Theorem 1.6. *Let $\beta_n^{\text{FOM}}, \beta_n^{\text{ROM}}$ be the purchase price of a spanning tree in the FOM and ROM models, respectively. Then, for n large, we have*

$$\zeta(3) < 1.38 \leq \beta_n^{\text{FOM}} \leq \beta_n^{\text{ROM}} < 2\zeta(3).$$

Furthermore, the AOM price is $O(1)$.

We will give a short proof that $\beta_n^{\text{FOM}} \geq 1.25$ in Section 7.1.2. The bound $\beta_n^{\text{FOM}} \geq 1.38$ is due to unpublished prior work of Aldous, Angel and Berestycki [1]. (Although their paper considers the ROM problem, the proof of their lower bound is valid for β_n^{FOM} also.) We have included a sketch of their proof in Section 7.1.3.

For our next result, let us consider the setting where \mathcal{F} is the hypergraph of perfect matchings on the edge-set of the complete bipartite graph $K_{n,n}$.

Theorem 1.7. *Letting $\mu_n^{\text{FOM}}, \mu_n^{\text{ROM}}, \mu_n^{\text{AOM}}$ be the FOM, ROM and AOM purchase prices of perfect matchings on $K_{n,n}$, we have*

$$2 \lesssim \mu_n^{\text{FOM}} \leq \mu_n^{\text{ROM}} \lesssim \mu_n^{\text{AOM}} \lesssim 4c_3,$$

where c_3 is as in Theorem 1.1. Here $a \lesssim b$ denotes $a \leq (1 + o(1))b$.

After this we turn to the non-bipartite case.

Theorem 1.8. *If $\nu_n^{\text{FOM}}, \nu_n^{\text{ROM}}, \nu_n^{\text{AOM}}$ are the FOM, ROM and AOM purchase prices of a perfect matching in K_n , we have*

$$2 \lesssim \nu_n^{\text{FOM}} \leq \nu_n^{\text{ROM}} \leq \nu_n^{\text{AOM}} \lesssim 4c_3.$$

We can also purchase a Traveling Salesperson tour at constant expected cost:

Theorem 1.9. *If $h_n^{\text{FOM}}, h_n^{\text{ROM}}, h_n^{\text{AOM}}$ are the FOM, ROM and AOM purchase prices of a Hamilton cycle in K_n , then*

$$c_2 \lesssim h_n^{\text{FOM}} \leq h_n^{\text{ROM}} \leq h_n^{\text{AOM}} \lesssim 200.$$

In Theorems 1.6, 1.7, 1.8, 1.9, we prove the purchase prices are $\Theta(1)$, but we do not determine an asymptotic constant. In our next example, we can determine such a constant.

We consider the case where the underlying set of \mathcal{F} is the set of $n(n-1)$ directed edges in the complete digraph \vec{K}_n , and the target sets are the *arborescences*; i.e., directed rooted spanning trees, such that every edge is oriented towards the root.

Theorem 1.10. *If $\alpha_n^{FOM}, \alpha_n^{ROM}$ are the FOM and ROM purchase prices of an arborescence in \vec{K}_n , then*

$$\lim_{n \rightarrow \infty} \alpha_n^{FOM} = \lim_{n \rightarrow \infty} \alpha_n^{ROM} = 2.$$

Furthermore, in the AOM model we can w.h.p. construct an arborescence of cost $O(1)$.

Note that the w.h.p result does not necessarily imply an $O(1)$ purchase price for the arborescence in the AOM model. An interesting open question is whether $\lim_{n \rightarrow \infty} \alpha_n^{AOM} = 2$ as well.

As a final result, we prove w.h.p results for the Traveling Salesperson Problem (TSP) on the complete digraph \vec{K}_n . Again, it is not clear that this implies results for the purchase price (expected value).

Theorem 1.11. *In the directed case of the TSP we have that w.h.p.*

$$4 \lesssim H_n^{FOM} \leq H_n^{ROM} \leq H_n^{AOM} \lesssim 4c_2,$$

where H^{model} is the cost of the tour we can find w.h.p in the model.

2. PRELIMINARIES

For reference we use the following: Let $B(n, p)$ denote the binomial random variable with parameters n, p . Then for $0 \leq \varepsilon \leq 1$ and $\alpha > 0$, we have the following Chernoff bounds:

$$(2) \quad \Pr(B(n, p) \leq (1 - \varepsilon)np) \leq e^{-\varepsilon^2 np/2}.$$

$$(3) \quad \Pr(B(n, p) \geq (1 + \varepsilon)np) \leq e^{-\varepsilon^2 np/3}.$$

We will also make use of a simple application of the Azuma-Hoeffding martingale tale inequality, often referred to as McDiarmid's inequality [25]. Let $Z = Z(Y_1, Y_2, \dots, Y_N)$ where Y_1, Y_2, \dots, Y_N are independent random variables. Suppose that if $\mathbf{Y} = (Y_1, Y_2, \dots, Y_N)$ and $\mathbf{Y}' = (Y'_1, Y'_2, \dots, Y'_N)$ differ in only one coordinate that $|Z(\mathbf{Y}) - Z(\mathbf{Y}')| \leq c$. Then for any $u > 0$ we have

$$(4) \quad \Pr(|Z - \mathbf{E} Z| \geq u) \leq 2 \exp \left\{ -\frac{u^2}{2Nc^2} \right\}.$$

3. PURCHASING ANY k ITEMS

The symmetry of the k -uniform hypergraph means that the order in which we inspect the variables x_i is irrelevant. Thus, we suppose they are simply given in some fixed order x_N, x_{N-1}, \dots, x_1 . In other words, in this context, there is essentially no difference between any of the three models under consideration. At step i (beginning with $i = N$), we examine x_i , and either *accept* i and pay x_i , or *reject* i and continue to step $i - 1$. The process ends

after k acceptances. Note that since we are required to purchase k items, we will have to pay all of $x_\ell, x_{\ell-1}, \dots, x_1$ to accept $\ell, \dots, 1$ if only $k - \ell$ items have been accepted before index ℓ is reached.

Our goal is to minimize the expected total cost; we denote this expected value by $\rho_{N,k}$.

3.1. The case $k = 1$. We write $\rho_N := \rho_{N,1}$ for the expected cost in this case. When we inspect x_N , an optimum strategy to minimize our total expected cost is to accept x_N if and only if $x_N < \rho_{N-1}$, since ρ_{N-1} is the expected cost when we reject x_N . This simple dynamic will allow us to prove the case $k = 1$ of Theorem 1.1 by analyzing the process recursively.

Proof. Let Z denote the cost of the element selected by the optimal strategy so that $\rho_N = \mathbf{E}(Z)$. Then we have the following recurrence:

$$\begin{aligned} (5) \quad \rho_N &= \mathbf{E}(Z \mid x_N \leq \rho_{N-1}) \Pr(x_N \leq \rho_{N-1}) + \mathbf{E}(Z \mid x_N > \rho_{N-1}) \Pr(x_N > \rho_{N-1}) \\ &= \frac{\rho_{N-1}^2}{2} + \rho_{N-1}(1 - \rho_{N-1}) \\ (6) \quad &= \rho_{N-1} \left(1 - \frac{\rho_{N-1}}{2}\right). \end{aligned}$$

We can now easily get an upper bound on ρ_N by induction. Clearly, $\rho_1 = 1/2$ and if $\rho_{N-1} \leq 2/N$ for some $N > 1$ then we see from (6) that,

$$\rho_N \leq \frac{2}{N} \left(1 - \frac{1}{N}\right) \leq \frac{2}{N+1}.$$

For a corresponding lower bound let $\varepsilon_N = A/N$ where $A = 10$ suffices. We have checked numerically that $\rho_N \geq 2(1 - \varepsilon_{N-1})/(N - 1)$ for $N \leq 3(A + 1) = 33$. Now assume that $N \geq 3(A + 1)$ and $\rho_{N-1} \geq 2(1 - \varepsilon_{N-1})/(N - 1)$. Then, because $x(1 - x/2)$ increases monotonically in $[0, 1]$,

$$\begin{aligned} \rho_N &\geq \frac{2(1 - \varepsilon_{N-1})}{N - 1} - \frac{(1 - \varepsilon_{N-1})^2}{(N - 1)^2} \\ &= \frac{2(1 - \varepsilon_N)}{N} + \frac{2(\varepsilon_N - \varepsilon_{N-1})}{N - 1} + \frac{2(1 - \varepsilon_N)}{N(N - 1)} - \frac{(1 - \varepsilon_{N-1})^2}{(N - 1)^2} \\ &= \frac{2(1 - \varepsilon_N)}{N} - \frac{2A}{N(N - 1)^2} + \frac{2(N - A)}{N^2(N - 1)} - \frac{(N - 1 - A)^2}{(N - 1)^4} \\ &\geq \frac{2(1 - \varepsilon_N)}{N} - \frac{2A}{N(N - 1)^2} + \frac{2(N - A)}{N^2(N - 1)} - \frac{(N - A)^2}{N^2(N - 1)^2} \\ &= \frac{2(1 - \varepsilon_N)}{N} + \frac{N^2 - 2(A + 1)N - (A^2 - 2A)}{N^2(N - 1)^2} \\ &\geq \frac{2(1 - \varepsilon_N)}{N}, \end{aligned}$$

since $N \geq 3(A + 1)$. □

3.2. **The case $k \geq 2$.** We now complete the proof of Theorem 1.1.

Proof. We replace (5) by

$$\begin{aligned}
 \rho_{k,N} &= \mathbf{E}(\min \{x_N + \rho_{k-1,N-1}, \rho_{k,N-1}\}) \\
 &= \mathbf{E}(x_N + \rho_{k-1,N-1} \mid x_N < \rho_{k,N-1} - \rho_{k-1,N-1}) \mathbf{Pr}(x_N < \rho_{k,N-1} - \rho_{k-1,N-1}) \\
 &\quad + \rho_{k,N-1} \mathbf{Pr}(x_N \geq \rho_{k,N-1} - \rho_{k-1,N-1}) \\
 &= \left(\frac{\rho_{k,N-1} - \rho_{k-1,N-1}}{2} + \rho_{k-1,N-1} \right) (\rho_{k,N-1} - \rho_{k-1,N-1}) + \rho_{k,N-1} (1 - (\rho_{k,N-1} - \rho_{k-1,N-1})) \\
 (7) \quad &= \rho_{k,N-1} - \frac{(\rho_{k,N-1} - \rho_{k-1,N-1})^2}{2}.
 \end{aligned}$$

Suppose inductively that $\rho_{k,N-1} \leq \frac{c_k}{N}$, $\rho_{k-1,N-1} \leq \frac{c_{k-1}}{N}$. Then, (7) and $\rho_{k,N-1} \geq \rho_{k-1,N-1}$ implies that

$$\begin{aligned}
 (8) \quad \rho_{k,N} &= \rho_{k,N-1} \left(1 - \frac{\rho_{k,N-1}}{2} + \rho_{k-1,N-1} \right) - \frac{\rho_{k-1,N-1}^2}{2} \\
 &\leq \frac{c_k}{N} - \frac{(c_k - c_{k-1})^2}{2N^2} \\
 &\leq \frac{c_k}{N+1} + \frac{2c_k - (c_k - c_{k-1})^2}{2N^2} \\
 &= \frac{c_k}{N+1}.
 \end{aligned}$$

Here the form of (8) implies that we should put $\rho_{k,N-1}, \rho_{k-1,N-1}$ at their upper bounds to maximize the RHS.

For a lower bound, let $\varepsilon_N = 1/N^{1/2}$ and assume inductively that

$$\rho_{k,N-1} \geq (1 - \varepsilon_{N-1}) c_k / (N - 1), \rho_{k-1,N-1} \geq (1 - \varepsilon_{N-1}) c_{k-1} / (N - 1).$$

We will first prove that

$$(9) \quad \rho_{k,k} = \frac{k}{2} \geq \left(1 - \frac{1}{k^{1/2}} \right) \frac{c_k}{k}$$

for $k \geq 2$, which will allow us a base for our induction.

If we let $1 + 2c_k = d_k^2$ then (1) implies that

$$(10) \quad d_k^2 = 1 + 2c_k = 1 + 2(c_{k-1} + 1 + d_{k-1}) = d_{k-1}^2 + 2d_{k-1} + 2 = (d_{k-1} + 1)^2 + 1.$$

Since $d_1 > 1$, this implies that

$$d_k > k.$$

On the other hand, (10) and $(1+x)^{1/2} \leq x^{1/2} + \frac{1}{2x}$ implies that

$$d_k \leq d_{k-1} + 1 + \frac{1}{2(d_{k-1} + 1)^2} \leq d_{k-1} + 1 + \frac{1}{2k^2}.$$

It follows from this that

$$d_k \leq k + \sum_{r=1}^{\infty} \frac{1}{2r^2} < k + 1 \text{ which implies that } c_k \leq \frac{k^2}{2} \left(1 + \frac{2}{k}\right).$$

We can see therefore that (9) holds if $(1 + \frac{2}{k})(1 - \frac{1}{k^{1/2}}) \leq 1$. This holds for $k \geq 2$ and the proof of (9) is complete.

We now use induction on N with $N = k$ as the base case. Referring to (7), the function $f(x, y) = x - (x - y)^2/2$ is monotone increasing in x for $x \leq 1 + y$ and in y for $y \leq x$. Thus we can substitute our lower bounds for $\rho_{k, N-1}$ and $\rho_{k-1, N-1}$ and use $2c_k = (c_k - c_{k-1})^2$ to get

$$\begin{aligned} \rho_{k, N} &\geq (1 - \varepsilon_{N-1}) \frac{c_k}{N-1} - \frac{(1 - \varepsilon_{N-1})^2}{2(N-1)^2} (c_k - c_{k-1})^2 \\ &= (1 - \varepsilon_{N-1}) \frac{c_k}{N-1} - \frac{(1 - \varepsilon_{N-1})^2 c_k}{(N-1)^2} \\ &= \frac{(1 - \varepsilon_N) c_k}{N} + \frac{(\varepsilon_N - \varepsilon_{N-1}) c_k}{N} + \frac{(1 - \varepsilon_{N-1}) c_k}{N(N-1)} - \frac{(1 - \varepsilon_{N-1})^2 c_k}{(N-1)^2} \\ &= \frac{(1 - \varepsilon_N) c_k}{N} + c_k g(N), \end{aligned}$$

where

$$\begin{aligned} g(x) &= \frac{\sqrt{1/x} - \sqrt{1/(x-1)}}{x} + \frac{1 - \sqrt{1/(x-1)}}{x(x-1)} - \frac{(1 - \sqrt{1/(x-1)})^2}{(x-1)^2} \\ &= \frac{1 - \sqrt{1 + \frac{1}{x-1}}}{x^{3/2}} + \left(\frac{1 - \frac{1}{\sqrt{x-1}}}{x-1} \right) \left(\frac{1}{x} - \frac{1 - \frac{1}{\sqrt{x-1}}}{x-1} \right) \\ &\geq -\frac{1}{2x^{3/2}(x-1)} + \frac{(\sqrt{x-1} - 1)(x - \sqrt{x-1})}{x(x-1)^{5/2}} \\ &\geq \frac{x(2\sqrt{x-1} - 5) + 2 + \sqrt{x-1}}{2x^{5/2}(x-1)} \\ &\geq 0 \text{ for } x \geq 7. \end{aligned}$$

By direct evaluation $\rho_{k, N} \geq (1 - \varepsilon_N) c_k / N$ for $2 \leq k < N \leq 6$. □

We will need the following estimates for c_k as k grows:

Lemma 3.1. *If $k \geq 1$ then*

$$\frac{k^2}{2} \leq c_k \leq 2k^2.$$

Proof. $c_1 = 2$ and then we inductively have for $k > 1$,

$$\begin{aligned} c_k &\leq 1 + 2(k-1)^2 + \sqrt{1 + 4(k-1)^2} \\ &= 1 + 2(k-1)^2 + 2(k-1) \left(1 + \frac{1}{4(k-1)^2}\right)^{1/2} \\ &\leq 1 + 2(k-1)^2 + 2(k-1) \left(1 + \frac{1}{8(k-1)^2}\right) \\ &\leq 2k^2 - 2k + \frac{5}{4} \\ &\leq 2k^2. \end{aligned}$$

This confirms the upper bound.

Similarly,

$$\begin{aligned} c_k &\geq 1 + \frac{(k-1)^2}{2} + (1 + (k-1)^2)^{1/2} \\ &\geq 1 + \frac{k^2}{2} - k + \frac{1}{2} + k - 1 \\ &\geq \frac{k^2}{2}. \end{aligned}$$

□

3.3. Purchasing at least one and two on average. In this section, we consider a peculiar variant of the purchasing problem. In this variant, we consider uniform costs x_1, \dots, x_N as before, and still must decide on the spot to accept or reject an index. Now, however, the total number of accepted indices is not required to be fixed. Instead, as a random variable (depending on the costs x_1, \dots, x_N), the number ν_N of accepted indices is required to satisfy:

- (1) $\nu_N \geq 1$
- (2) $\mathbf{E}(\nu_N) = 2$.

We call this the *average-two-purchase problem* and we consider this problem merely for technical reasons; it will arise in our analysis of the purchase price of spanning trees in Section 7.1.2. Let ϕ_N denote the minimum expected total cost of items purchased under these constraints. (The minimum is taken over all strategies which ensure that ν_N satisfies the two conditions above.)

Theorem 3.2.

$$\phi_N \gtrsim \frac{2.5}{N}.$$

Proof. The first item accepted must cost $\gtrsim 2/N$ in expectation, else we can improve the upper bound for $k = 1$ in Theorem 1.1.

Let $1 \leq T \leq N$ be the random variable equal to the step on which the first item is accepted. We are required to purchase, on average, one item from among $x_{T+1}, x_{T+2}, \dots, x_N$; let C denote the optimum expected cost to do this. (C depends on the distribution of T .)

Evidently C is at least the optimum expected cost to purchase an average of one item from x_1, \dots, x_N , and an optimum strategy for this task is to accept any item seen which has cost $\leq \frac{1}{N}$. Indeed, this strategy accepts an average of exactly one item, and each rejection of an item of cost $\leq \frac{1}{N}$ requires that the strategy is modified to sometimes accept some items of cost greater than $\frac{1}{N}$ to compensate; so that the expected total cost $\frac{1}{2N}$ can only be increased. Thus $C \geq \frac{1}{2N}$, and consequently,

$$\phi_N \gtrsim \frac{2}{N} + \frac{1}{2N}.$$

□

3.4. Proof of Corollary 1.2. Using $\widehat{\rho}_{k,N}$ to denote the expected minimum cost in this context, we write

$$(11) \quad \widehat{\rho}_{k,N} = \left(\frac{\widehat{\rho}_{k,N-1} - \widehat{\rho}_{k-1,N-1}}{2 + O\left(\frac{k^2}{N}\right)} + \widehat{\rho}_{k-1,N-1} \right) \left(\frac{\widehat{\rho}_{k,N-1} - \widehat{\rho}_{k-1,N-1}}{D} + O\left(\frac{k^2}{N}\right) \right) \\ + \widehat{\rho}_{k,N-1} \left(1 - \frac{\widehat{\rho}_{k,N-1} - \widehat{\rho}_{k-1,N-1}}{D} + O\left(\frac{k^2}{N}\right) \right).$$

Explanation: The k^2 in $O(k^2/N)$ is derived from $c_k = O(k^2)$ which will imply that $\widehat{\rho}_{k,N}$, being close to $\rho_{k,N}$ is of order k^2/N . Our assumptions on the density of the distribution of costs can be used to justify (11). We then carry out the analysis of Section 3.2 with $\rho_{N,k}$ replaced by $D\widehat{\rho}_{N,k}$, making adjustment for the error terms.

4. SHORTEST PATHS

In this section, we prove Theorem 1.3. We again contrast this with the fact that Janson [21] proved that offline, the shortest distance is $\approx \frac{\log n}{n}$ w.h.p.

The following fact will be used several times. If $\eta_1, \eta_2, \dots, \eta_M$ are independent $[0, 1]$ random variables and $\zeta \leq 1$ then

$$(12) \quad \Pr(\eta_1 + \eta_2 + \dots + \eta_M \leq \zeta) = \frac{\zeta^M}{M!}.$$

4.1. Upper Bound. We compute an upper bound of $n^{-2/3+o(1)}$ in the ROM model. Let

$$\alpha = \frac{1}{3} \text{ and } k = \log \log n.$$

We partition the edge set of K_n into X, Y, Z where X is the first $N/3$ edges and Y is the next $N/3$ edges and Z is the final $N/3$ edges in the given order. Then partition X into X_1, X_2, \dots, X_k where each set is of size $N/(3k)$ and where X_{i+1} follows X_i in the random order. Similarly partition Y into Y_1, Y_2, \dots, Y_k . Next let $p = n^{-1+\alpha/k}$. Now let $G(X_i)$ denote the graph induced by edges in X_i that are of cost at most p , $i = 1, 2, \dots, k$. Define $G(Y_i), i = 1, 2, \dots, k$ similarly, as well as $G(Z)$.

We grow Breadth First Search Trees T, T^* of a bounded depth from 1, n respectively. We only use low cost edges and then we find a low cost edge joining the leaves of these two trees. Our choice of α balances the cost of this edge with the cost of two root to leaf paths.

Let $S_0 = \{1\}$ and S'_1 be the set of neighbors of vertex 1 in $G(X_1)$. We can determine S'_1 in the ROM model. Let $\varepsilon = n^{-\alpha/3k}$ and S_1 be the first $(1 - \varepsilon)nq$ members of S'_1 , where $q = p/3k$, assuming that there are at least this many neighbors. Observe that $|S'_1|$ is distributed as $\text{Bin}(n - 1, q)$ and so from the Chernoff bound (2) we see that for $\ell = 1$ we have

$$(13) \quad \Pr(|S'_\ell| \leq ((1 - \varepsilon)nq)^\ell) \leq \exp\left\{-\frac{nq}{3n^{2\alpha/3k}}\right\} = o(n^{-4}).$$

We can now w.h.p. inductively define sets S_ℓ such that

$$(14) \quad |S_\ell| = ((1 - \varepsilon)nq)^\ell \text{ for } \ell = 2, \dots, k.$$

Given S_ℓ , we let $S'_{\ell+1}$ be the set of vertices not in $S_{\leq \ell} = \bigcup_{i=1}^{\ell} S_i$ which are neighbors of S_ℓ in $G(X_{\ell+1})$. Now $|S'_{\ell+1}|$ is distributed as the binomial $\text{Bin}(n - m_1, (1 - (1 - q)^{m_2}))$, where $m_1 = |S_{\leq \ell}|$ and $m_2 = |S_\ell|$. Then since $(1 - q)^{m_2} \leq 1 - m_2q + (m_2q)^2$ we see that $|S'_{\ell+1}|$ has expectation at least $(n - m_1)m_2q(1 - m_2q)/3k$. Putting $1 - \delta = \frac{(1-\varepsilon)nqm_2}{(n-m_1)m_2q(1-m_2q)} \leq 1 - \varepsilon/2$ (using $\ell < k$) we again see from (2) that (13) holds with ℓ replaced by $\ell + 1$. In which case we, can let $S_{\ell+1}$ be the first $((1 - \varepsilon)nq)^{\ell+1}$ members of $S'_{\ell+1}$. This completes the induction and we see that the construction of S_1, S_2, \dots, S_k succeeds with probability $1 - o(n^{-3})$.

Let E_ℓ be a set of $|S_{\ell+1}|$ edges joining S_ℓ to $S_{\ell+1}$ in $G(X_{\ell+1})$. Let T be the tree with vertex set $\bigcup_{\ell=1}^k S_\ell$ and edge set $\bigcup_{\ell=1}^{k-1} E_\ell$.

T can be constructed in our model. We simply check the edges between S_ℓ and $[n] \setminus S_{\leq \ell}$ as we see them and accept edges if they (i) have cost at most p , (ii) connect S_ℓ to a new vertex, and (iii) we have accepted fewer than $((1 - \varepsilon)nq)^{\ell+1}$ edges of this sort. Given (14) and the fact that each edge of T has cost at most p , we see that the cost of the edges of T will be

$$O((nq + (nq)^2 + \dots + (nq)^k)p) = O(n^{\alpha-1+\alpha/k}).$$

If $n \in T$ then we are done. Otherwise we build a corresponding tree rooted T^* rooted at n . We let $S_\ell^*, \ell \geq 1$ be the corresponding sets at graph distance ℓ from n .

Because $|S_k|, |S_k^*| \geq \left(\frac{1-\varepsilon}{3k}\right)^k n^\alpha = n^{\alpha-o(1)}$ w.h.p., we see that w.h.p. there will be at least $|S_k||S_k^*|/4$ edges joining S_k, S_k^* in $G(Z)$. Then, with failure probability at most $\left(1 - \frac{(\log n)^2}{|S_k||S_k^*|}\right)^{|S_k||S_k^*|/4} = o(n^{-3})$ we can find an edge in Z between S_k and S_k^* of cost at most $(\log n)^2/|S_k||S_k^*| = O(n^{-2\alpha+o(1)})$. Putting things together, we get a total cost of

$$O(n^{\alpha+\alpha/k-1}) + O(n^{-2\alpha}) = O(n^{-2/3+O(1/k)}),$$

with probability $1 - o(n^{-3})$.

We have explained the algorithm in terms of events that happen w.h.p. To get a bound on expectation, we have to also deal with the unlikely cases. So, we always check to see that the next edge to be considered is essential for connecting vertex 1 to vertex n . if it is, we take this edge and an arbitrary path P joining these two vertices. The cost of P is at most n and

we need to construct it only with probability $o(n^{-3})$. This adds $o(n^{-2})$ to the expectation and completes our proof of the upper bound.

4.2. Lower Bound. We now compute a lower bound of $n^{-2/3+o(1)}$ in the FOM model. Consider an algorithm for finding a short path from vertex 1 to vertex n . We let $C_1(t)$ and $C_2(t)$ be the components of the purchased graph at step t which contain the vertices 1 and n , respectively. (Thus, for each t , $C_1(t)$ and $C_2(t)$ are random variables depending on $x_1, \dots, x_{\binom{n}{2}}$.)

Observe that any algorithm which finds a path from 1 to n finishes its task by choosing an edge e^* from $C_1(t)$ to $C_2(t)$ for some step t . We will thus analyze the cost of building components $C_1(t)$ and $C_2(t)$ in terms of their size, and also the cost of the final edge e^* . We begin by considering the cost of e^* .

Cost of e^* : To bound the cost of e^* from below, we will even give the algorithm extra information, revealing the costs of all edges between $C_1(t)$ and $C_2(t)$ as soon as their endpoints belong to the respective components; in particular, we will assert that even the minimum cost edge between these two sets is not too small. As such, from the standpoint of choosing an inexpensive e^* , we view the algorithm's task as simply producing by time t_0 components $C_1(t_0), C_2(t_0)$ between which there exists an edge e^* of small cost.

We thus let $e_1, e_2, \dots, e_\ell = e^*$ be a sequence of edges consisting of all edges between $C_1(t_0)$ and $C_2(t_0)$, ordered such that if e_i appears before e_j (in the sense that e_i but not e_j joins the sets $C_1(t)$ and $C_2(t)$ for some t) then $i < j$. This is a sequence of independent (and unconditioned) uniform $[0, 1]$ random variables. We will simply argue that no term e_j in this sequence can be much less than $\frac{1}{j}$.

Indeed, for any sequence ξ_1, ξ_2, \dots , of independent uniform $[0, 1]$ random variables, we have for any n that

$$(15) \quad \Pr\left(\exists k : k\xi_k \leq \frac{1}{(\log \max\{k, n\})^2}\right) \leq \sum_{k=1}^{\infty} \frac{1}{k(\log \max\{k, n\})^2} = O\left(\frac{1}{\log n}\right).$$

Thus w.h.p. we have that the minimum $m_\ell = \min_{j \leq \ell} \{e_j\}$ satisfies

$$(16) \quad m_\ell \geq \frac{1}{\ell \log^2 n} \geq \frac{1}{|C_1(t_0)||C_2(t_0)|(\log n)^2}.$$

Cost of C_1, C_2 : For this we argue that if \mathcal{E}_m denotes the event that there is a sub-tree of K_n with at least $m \gg \log n$ vertices and total cost $\leq m/3n$ then

$$(17) \quad \Pr(\mathcal{E}_m) \leq \sum_{k \geq m} \binom{n}{k} k^{k-2} \frac{(m/3n)^{k-1}}{k!} \leq \frac{3n}{m} \sum_{k \geq m} \frac{1}{k^2} \left(\frac{me}{3k}\right)^k = O\left(\frac{ne^m}{3^m m^3}\right) = o(1).$$

Explanation: We choose our tree of K_n with k vertices in $\binom{n}{k} k^{k-2}$ ways. Then we use (12).

It follows from (16) and (17) that w.h.p. the algorithm must pay at least

$$\min_{ab \geq n^{1/2}} \left\{ \frac{1}{ab(\log n)^2} + \frac{a+b}{3n} \right\} = \Omega \left(\frac{1}{n^{2/3}(\log n)^2} \right).$$

Here $a = |C|, b = |C'|$ and we minimize over $ab \geq n^{1/2}$, say, because if $ab \leq n^{1/2}$ then we already pay $n^{-1/2-o(1)}$ in expectation from (16).

4.3. Lower Bound in AOM. We now compute a lower bound of $\Omega(n^{-1/2})$ in the AOM model. Our strategy is to place the edges X incident with vertices $1, n$ last in the sequence. Suppose now that after we have seen all edges except those in X we have purchased components C_1, C_2, \dots, C_m in the subgraph induced by the vertices $[2, n-1]$. The expected cost of purchasing an edge between 1 and C_j is at least $(2 - o(1))/|C_j|$, by Theorem 1.1. We can assume therefore that at least one component is of size at least $n^{1/2}$. But then we have that w.h.p.

$$\text{every tree of size } n^{1/2} \text{ has cost at least } 1/10n^{1/2}.$$

Indeed, let $k = n^{1/2}$ and $C_n = 1/10n^{1/2}$. Then, if $p(T)$ denotes the price of tree T ,

$$\Pr(\exists T : p(T) \leq C_n) \leq \binom{n}{k} k^{k-2} \frac{C_n^{k-1}}{(k-1)!} \leq \frac{1}{C_n k} \left(\frac{ne^2 C_n}{k} \right)^k = o(1).$$

Explanation: The factor $\frac{C_n^{k-1}}{(k-1)!}$ comes from applying (12).

It follows that it costs at least $(1 - o(1))/10n^{1/2}$ in expectation to purchase a path from 1 to n in the AOM model.

This completes the proof of Theorem 1.3.

5. TRIANGLES AND PATHS OF LENGTH TWO

In this section, we prove Theorem 1.4. We first observe that the expected number of triangles of total length λ is bounded by $\binom{n}{3} \frac{\lambda^3}{6} \approx \frac{(\lambda n)^3}{6}$. And then another calculation shows via the Chebyshev inequality that if $\lambda n \rightarrow \infty$ then there will be a triangle of cost λ w.h.p. As we will see, it will not be possible to find such a triangle w.h.p., instead we will show that if $\tau_n^{\text{FOM}}, \tau_n^{\text{ROM}}$ are the expected minimum cost of a triangle constructible in the FOM and ROM settings, then

$$\frac{1}{n^{4/7} \log n} \leq \tau_n^{\text{FOM}} \leq \tau_n^{\text{ROM}} \leq \frac{10}{n^{4/7}}.$$

Our proof of this inequality will follow from an analysis of the problem of purchasing a large collection of paths of length 2; we will see that this is sufficient, and in some sense also necessary, to purchase a low-cost triangle.

5.1. Paths of length Two. We will prove the following:

Theorem 5.1. *Let $\kappa_{\ell,n}^{\text{FOM}}, \kappa_{\ell,n}^{\text{ROM}}$ denote the expected cost of purchasing $\ell = o(n)$ paths of length two in the FOM and ROM settings, respectively. Then*

$$\left(\frac{\ell}{16n(\log n)^4}\right)^{4/3} \leq \kappa_{\ell,n}^{\text{FOM}} \leq \kappa_{\ell,n}^{\text{ROM}} \leq 6 \left(\frac{\ell}{n}\right)^{4/3}.$$

5.1.1. Obtaining $\ell = o(n)$ paths of length two: Upper Bound. Our goal in this section is to show how to find ℓ distinct paths of length two. Our paths will be distinct, but not necessarily edge or node disjoint. We will subsequently use the case $k = 1$ of Theorem 1.1 to close one of these to a triangle at expected cost order $1/\ell$.

The first $N/3$ edges will be considered to be colored red, the next $N/3$ will be considered to be colored blue and the final $N/3$ edges will be considered to be colored green. We will use a parameter k and its value will be revealed shortly. We choose $k = o(n)$ disjoint red edges by examining the red edges and accepting the first k of cost at most $10k/n^2$. Then we choose ℓ blue edges incident with the selected red edges by examining the relevant blue edges and accepting the first ℓ edges of cost at most $2\ell/kn$. The Chernoff bounds imply that we will succeed with superpolynomial probability if $k, \ell = n^{\Omega(1)}$. Success will mean the creation of ℓ paths of length two.

The expected cost of creating ℓ paths of length two with this strategy is at most

$$(18) \quad \frac{5k^2}{n^2} + \frac{2\ell^2}{kn}.$$

We choose k to minimize (18). This gives $k = (\ell^2 n/5)^{1/3} \gg \ell$ and a total cost of at most $6(\ell/n)^{4/3}$, as claimed in Theorem 5.1.

If we ever get to the stage where it is impossible to obtain either a triangle or ℓ paths of length two without the inclusion of the next edge then we abort the above procedure and choose any triangle still available. This will cost at most 3 and the effect on the expectation is negligible.

We will assume from now on in our approach to proving an upper bound, that we will always check as to whether or not the next edge must be chosen in order to be able to construct the required object. If this is the case then we take this edge and build an object regardless of cost. It will be apparent that this happens with such a low probability that it makes a negligible contribution to the expectation.

5.1.2. Obtaining $(\log n)^5 \leq \ell = o(n)$ paths of length two: Lower Bound. We consider the cost of producing ℓ paths of length two. One natural way of producing many paths of length two is to create a collection of edge disjoint stars (i.e., trees with only one vertex of degree greater than 1). Most of our analysis in this section is aimed at constraining (from below) the cost of purchasing such a collection of stars. We begin by arguing that this will suffice to give a lower bound in general; i.e., that any strategy to produce many paths of length 2 is in some sense not too far from a strategy to produce many disjoint stars.

To make this reduction, let us consider a modified game: in particular, let us suppose that whenever we choose to purchase an edge, we must also assign it an orientation. Our goal is to purchase (and orient) a set of edges to ensure that

$$\mathcal{S} = \sum_v \binom{\text{outdeg}(v)}{2}$$

is large. Observe that any strategy in the modified game which achieves that $\mathcal{S} = \ell$ for some ℓ can be translated into a strategy in the standard game which purchases ℓ paths of length 2 at the same expected cost, while purchasing only edge disjoint stars. In particular, our analysis later will thus give a lower bound on the cost in the modified game of achieving that $\mathcal{S} = \ell$.

On the other hand, we argue that for any algorithm \mathcal{A} for the standard game to purchase ℓ edges at small expected cost, there is a strategy in the modified game to ensure that $\mathcal{S} \geq \ell/16$. Indeed, to see this, we consider two cases.

Case 1: \mathcal{A} produces at least $\ell/2$ paths from large stars.

Here a large star is a star with more than $L = (\log n)^2$ edges. Let there be d_i stars with i edges. We have that

$$\sum_{i=L+1}^n d_i \binom{i}{2} \geq \frac{\ell}{2}.$$

In conjunction with Lemma 5.2 below, this implies that the cost is at least

$$\sum_{i=L+1}^n \frac{d_i i^2}{3n} \geq \frac{\ell}{4n} \gg \left(\frac{\ell}{n}\right)^{4/3},$$

and the expected cost is already above our lower bound.

Case 2: \mathcal{A} produces at least $\ell/2$ paths from small stars.

We translate \mathcal{A} into a strategy for the modified game by orienting each purchased edge randomly. A path of length two is said to be good if both edges are oriented away from the center vertex, so that \mathcal{S} is simply the number of good paths. The expected number of good paths is at least $\ell/8$. Also, switching the orientation of an edge can only change the number of good paths by at most $2L$. It follows from McDiarmid's inequality (4) that

$$\Pr(\mathcal{S} \leq \ell/16) \leq 2 \exp \left\{ -\frac{\ell^2}{128\ell L^2} \right\}.$$

This reduction allows us to focus on strategies in the standard game which restrict themselves to purchasing edge disjoint stars. As a first step, we rule out the relevance of large stars, by showing that all large stars are quite expensive (even offline).

Let

$$L = (\log n)^2.$$

Lemma 5.2. *W.h.p. every star with $k > L$ edges has cost at least $\xi_k = k^2/3n$.*

Proof. The probability that there is a star of smaller cost than claimed is at most

$$n \sum_{k=L+1}^{n-1} \binom{n-1}{k} \frac{\xi_k^k}{k!} \leq n \sum_{k=L+1}^{n-1} \left(\frac{ne\xi_k}{k^2} \right)^k = o(1).$$

Here we have used (12). □

Suppose next that there are $k_i, 1 \leq i \leq L \leq \ell^{1/2}$ edges such that when the algorithm purchases them they are added to a star with $i-1$ edges already for $i = 1, 2, \dots, L$. Call these *type i edges*. If $i = 1$ then these edges are the first of their stars. When $i \geq 2$ and such an edge is added, $i-1$ new paths of length two are created.

Thus

$$\frac{\ell}{16} \leq \sum_{i=2}^L (i-1)k_i \leq \ell + L \leq 2\ell,$$

where the upper bound follows from the fact that we can stop purchasing edges once we have at least ℓ .

The expected cost of constructing these stars is at least

$$(19) \quad a_0 \sum_{i=1}^L \frac{k_i^2}{k_{i-1}n},$$

where $k_0 = n$ and a_0 is an absolute constant.

Explanation of (19): We can assume that the k_i edges of type i are chosen before any edges of type $i+1$. This provides the largest choice for each edge and thus gives a lower bound. That said, the expected cost of adding k_i edges to k_{i-1} vertices, by choosing from $\Omega(k_{i-1}n)$ edges is as claimed in (19).

We are therefore left with considering the following optimization problem, where $x_i = k_i/n$ and $\lambda = \ell/16n$:

$$(20) \quad \text{Minimize } \sum_{i=1}^L \frac{x_i^2}{x_{i-1}} \text{ subject to } \sum_{i=2}^L (i-1)x_i \geq \lambda \text{ and } 1 = x_0 \geq x_1 \geq x_2 \geq \dots \geq x_L.$$

Observe now that

$$\sum_{i=2}^L (i-1)x_i \leq L^2 x_2 \text{ which implies that } x_2 \geq \frac{\lambda}{L^2}.$$

Going back to (20), we have a lower bound of

$$\text{Minimum: } x_1^2 + \frac{x_2^2}{x_1} \text{ subject to } x_2 \geq \frac{\lambda}{L^2}.$$

So optimizing with respect to x_1 for a given value of x_2 we get a lower bound of

$$x_2^{4/3} (2^{-2/3} + 2^{1/3}) \geq \left(\frac{\lambda}{L^2}\right)^{4/3} = \left(\frac{\ell}{16L^2n}\right)^{4/3},$$

as required in Theorem 5.1.

5.1.3. Creating a Triangle: Upper Bound. In this section we prove an upper bound of $O(n^{-4/7})$ in the ROM model. We use the green edges and the case $k = 1$ of Theorem 1.1 to find a triangle by selecting a low cost edge joining the paths of length two. Notice that by the construction in Section 5.1.1, no two paths of length two have the same endpoints. Thus the expected cost of creating a triangle is at most

$$(21) \quad 6 \left(\frac{\ell}{n}\right)^{4/3} + \frac{6}{\ell}.$$

We have a cost of $6/\ell$ because only $1/3$ of the paths of length two can be completed by a green edge. Optimizing our choice of $\ell = (3/4)^{3/7}n^{4/7}$ in (21) gives us the required upper bound in Theorem 1.4.

5.1.4. Creating a Triangle: Lower Bound. In this section we prove a lower bound of $\Omega(n^{-4/7+o(1)})$ in the FOM model. For the lower bound we suppose that our algorithm creates ℓ paths of length two before adding an edge that closes a triangle. Equation (15) implies that w.h.p. the cost of the edge closing a triangle is at least $\frac{1}{\ell(\log n)^2}$. Therefore, the total cost of the set of paths of length two and this closing edge is at least

$$(22) \quad \left(\frac{\ell}{16L^2n}\right)^{4/3} + \frac{1}{\ell(\log n)^2} \geq \frac{1}{n^{4/7}(\log n)^2}.$$

Inequality (22) holds for any value of ℓ .

5.2. Creating a Triangle: Lower Bound in AOM. In this section we prove a lower bound of $\Omega(n^{-1/2})$ in the AOM model.

The adversary's strategy is to first present all of the edges incident with vertex 1. Suppose that we accept the edges $\{1, v\}, v \in A$ where $|A| = k$. The expected cost of these edges is $\approx c_k/n$. The adversary will now present all edges within A . We have two choices now. We can accept one of the edges within A and create a triangle at expected cost

$$\Omega\left(\frac{c_k}{n} + \frac{2c_1}{k^2}\right) = \Omega\left(\frac{k^2}{n} + \frac{1}{k^2}\right) = \Omega\left(\frac{1}{n^{1/2}}\right).$$

Failing this we will have to build our triangle without using vertex 1. The adversary now presents the edges incident with vertex 2 that have not been presented, and so on. We can assume inductively that this costs $\Omega\left(\frac{1}{(n-1)^{1/2}}\right)$ in expectation and we are done.

This completes the proof of Theorem 1.4.

6. COMPLETE GRAPHS UPPER BOUND IN ROM

Recall that $A_{r,n}$ is the expected cost of purchasing a copy of K_r , $r \geq 3$ and we will prove that $A_{r,n} = O(n^{-d_r+o(1)})$. We have already proved that we can take $d_3 = 4/7$ and we start an induction from here.

We consider the first $N/2$ edges to be colored red and the remaining $N/2$ edges to be colored blue. We use Theorem 1.1 and the red edges to construct a star $K_{1,\ell}$ centered at vertex 1, at a cost of $\frac{(2+o(1))c_\ell}{n} \leq \frac{(4+o(1))\ell^2}{n}$. Here ℓ is to be determined. We then use the blue edges to find a low cost copy of K_r in the red neighborhood of 1 at a cost of $A_{r,\ell}$. We need to be a little careful here as the graph induced by the blue edges is disturbed as $G_{\ell,1/2}$ i.e. is not K_ℓ . (Note that this would not be an issue in the FOM mode, as we could let the red edges be those incident with vertex 1.) Our inductive assumption is that we can find w.h.p. a copy of K_r in $G_{\ell,p}$ of cost $\ell^{-d_r+o(1)}$, provided p is a constant independent of ℓ, n . The base case will be triangles. It is straightforward to amend the arguments of Sections 5.1.1 and 5.1.3 for this case. For example in Section 5.1.1 we accept the first k red edges of weight at most $10k/(n^2p)$ and then accepting the first ℓ blue edges of weight $2\ell/(knp)$. In so doing we increase the expected cost of our paths by a factor of $1/p$. In Section 5.1.3 we replace $6/\ell$ by $6/(\ell p)$ to account for missing edges. Thus the expected cost of our triangle is inflated by $1/p$ overall. Having dealt with the base case we proceed to the inductive step. Thus,

$$(23) \quad A_{r+1,n} \leq \min_{\ell} \left\{ \frac{(4+o(1))\ell^2}{n} + \frac{1}{\ell^{d_r+o(1)}} \right\}.$$

We optimize (23) by choosing $\ell = n^{1/(d_r+2+o(1))}$ which gives

$$A_{r+1,n} \leq \frac{1}{n^{d_r+o(1)/(d_r+2)}}.$$

(Here we absorbed a constant into the $o(1)$ term.)

This gives us Theorem 1.5 with

$$d_{r+1} = \frac{d_r}{d_r + 2}.$$

Because $d_3 = 4/7$, we get $d_4 = 2/9$ and $d_5 = 1/10$. Putting $d_r = \frac{1}{11 \times 2^{r-5} - 1}$ we get

$$\frac{1}{d_{r+1}} = 1 + 22 \times 2^{r-5} - 2 = 11 \times 2^{r-4} - 1.$$

This completes the proof of Theorem 1.5.

7. SPANNING TREES AND ARBORESCENCES

7.1. Spanning Trees. In this section we will prove Theorem 1.6.

7.1.1. *Spanning Tree Upper Bound.* We describe an algorithm that finds a tree of expected cost strictly less than $2\zeta(3)$. We realise that it can be improved by a more complicated analysis, but not to anything tight. As a result, we aim for a decent upper bound that can be obtained without too much computation. The upper bound that we compute here is in the ROM model.

To begin we choose $0 < \alpha < 1$ and β where $\alpha\beta > 1$. Their precise values will be given below at the end of this section and they were found by searching for values that appear to be close to minimising the sum of (26) and (30).

Our strategy is to build a forest using the first αN edges while only accepting cheap edges. This forest will w.h.p. contain a giant component C_1 and then we use Theorem 1.1 to cheaply join the remaining trees to C_1 .

Algorithm BUYTREE:

Step 1 Let $E_1 = \{e_1, e_2, \dots, e_{\alpha N}\}$. We go through E_1 in order and accept an edge if its length $X_e \leq \beta/n$ and it does not make a cycle with already accepted edges. Let F_1 be the forest induced by the accepted edges.

Step 2 E_1 induces a graph Γ_1 that is distributed as $G_{n,m}$ where $m \approx \gamma n/2, \gamma = \alpha\beta$. So w.h.p. Γ_1 and hence F_1 has a giant (tree) component

$$(24) \quad \text{W.h.p. } |C_1| \approx \left(1 - \frac{x}{\gamma}\right) n \text{ where } xe^{-x} = \gamma e^{-\gamma} \text{ and } 0 < x < 1 < \gamma.$$

This follows from Erdős and Rényi [10]. See also Chapter 2 of [14].

Outside of C_1 there will w.h.p. be ν_k components of size k where

$$(25) \quad |\nu_k - n \frac{k^{k-2}}{k!} \gamma^{k-1} e^{-\gamma k}| \leq n^{2/3} \text{ for } k = 1, 2, \dots, O(\log n).$$

This also follows from [10]

Complete the building of the tree by using the case $k = 1$ of Theorem 1.1 and the edges of $E(K_n) \setminus E_1$ to select a single cheap edge from each small component to the giant C_1 . We can do this in the ROM model, attacking each 1-purchasing problem separately.

The reader will see immediately that there is room for improvement. We can consider using the edges between the small trees. This however leads to a more complicated analysis that we are trying to avoid.

Let us now carefully examine the expected cost of the tree found by the above algorithm. We break this into $P_1 + P_2$ where P_i is the expected cost of the edges added in Step $i=1,2$. We observe first that

$$(26) \quad P_1 = \frac{\beta(|F_1|)}{2n} \approx \frac{\beta}{2} \left(1 - \frac{x}{\gamma} + \frac{x^2}{2\gamma}\right).$$

Explanation of (26): We claim that in Step 1, Algorithm BUYTREE purchases the following number of edges asymptotically:

$$(27) \quad n \left(1 - \frac{x}{\gamma} + \frac{x^2}{2\gamma} \right).$$

The expected cost of each edge is $\frac{\beta}{2n}$ and so verifying (27) will verify (26). For this we rely on the following identities:

$$(28) \quad \frac{1}{\gamma} \sum_{k=1}^{\infty} \frac{k^{k-1}}{k!} (\gamma e^{-\gamma})^k = \begin{cases} 1 & \gamma < 1. \\ \frac{x}{\gamma} & \gamma > 1. \end{cases}$$

$$(29) \quad \frac{1}{\gamma} \sum_{k=1}^{\infty} \frac{k^{k-2}}{k!} (k-1) (\gamma e^{-\gamma})^k = \begin{cases} \frac{\gamma}{2} & \gamma < 1. \\ \frac{x^2}{2\gamma} & \gamma > 1. \end{cases}$$

When $\gamma < 1$, equation (28) is derived from the fact that the expected number of vertices on tree components is asymptotically n . When $\gamma > 1$ we replace $(\gamma e^{-\gamma})^k$ by $(x e^{-x})^k$ to get the value x/γ .

When $\gamma < 1$, equation (29) is derived from the fact that the expected number of edges on tree components is asymptotically $\gamma n/2$. When $\gamma > 1$ we replace $(\gamma e^{-\gamma})^k$ by $(x e^{-x})^k$ to get the value $x^2/2\gamma$.

We see from (28) that the number of edges in the giant component C_1 is w.h.p. asymptotically equal to $n \left(1 - \frac{x}{\gamma} \right)$. We then see from (29) that w.h.p. the number of edges selected that are in small trees is asymptotically equal to $n x^2/2\gamma$. There are w.h.p. $o(n)$ edges in F_1 that are obtained as spanning trees of small components with s vertices and $\geq s$ edges.

As for the cost of Step 2, we have that

$$(30) \quad P_2 \approx \frac{2}{(1-\alpha) \left(1 - \frac{x}{\gamma} \right)} \sum_{k=1}^{\infty} \frac{k^{k-3}}{k!} \gamma^{k-1} e^{-\gamma k}.$$

Explanation of (30): The algorithm seeks a low cost edge between each small tree T in F_1 to the giant component C_1 . Observe first that there are w.h.p.

$$m_1 \approx (1-\alpha) |C_1| (n - |C_1|)$$

edges between C_1 and $\bar{C}_1 = [n] \setminus C_1$. Indeed, for a given small constant $\varepsilon > 0$ and set $S \subseteq [n]$ with $\varepsilon n \leq |S| \leq (1-\varepsilon)n$ let $\bar{e}(S)$ be the number of edges in $G_{n,(1-\alpha)}$ that belong to $S : \bar{S}$, the edges of K_n between S and \bar{S} . Then the Chernoff bounds (2), (3) imply that

$$(31) \quad \Pr(|\bar{e}(S) - (1-\alpha)|S|(n - |S|)| \geq n^{3/2} \log n) \leq 2 \exp \left\{ -\frac{n^3 (\log n)^2}{3|S|(n - |S|)(1-\alpha)} \right\} \leq e^{-n(\log n)^2}.$$

We take a union bound and multiply the RHS of (31) by n^n to account for all possible choices of S and small components and so w.h.p. m_1 is as claimed.

Note next that these m_1 edges are distributed uniformly over the set of $|C_1|(n - |C_1|)$ edges in K_n with one endpoint in $|C_1|$. This is because given the fact that there are no Γ_1 edges between C_1 and \bar{C}_1 , we can interchange E_1 edges with non- E_1 edges within $S : \bar{S}$ without

changing F_1 . So, each such set of m_1 edges is equally likely. Under these circumstances, the number of edges between a small tree of size k and C_1 is distributed as a hypergeometric with mean $(1 - \alpha)k|C_1| \pm O(n^{1/2} \log n)$, given (31). As such, this will be concentrated around its mean (see Section 6 of Hoeffding [18]). Thus, using the case $k = 1$ of Theorem 1.1, we see that the expected cost of connecting a tree with k vertices to C_1 is at most $\approx \frac{2}{(1-\alpha)k|C_1|}$. Equation (30) now follows from (24) and (25).

Putting $\alpha = 0.69$ and $\beta = 3.5$ gives a total cost of less than 2.31 which is less than $2\zeta(3)$. This verifies the upper bound in Theorem 1.6.

Our upper bound can be improved by using edges between small components in Step 2.

Finally, in the AOM model, we can appeal to Theorem 1.9 and find a Hamilton path at expected cost $O(1)$.

7.1.2. Spanning Tree Lower Bound. In this section, we use the result of Theorem 3.2 to obtain a lower bound strictly greater than $\zeta(3)$ on the expected cost of purchasing a spanning tree. Suppose we have an algorithm A . We apply it first to the edges contained in $S_1 = [n + 1] \setminus \{x\}$ where x is chosen uniformly at random from $[n + 1]$. Suppose that A produces a spanning tree T_1 of S_1 . Now use the algorithm of Section 3.2 to find two edges e_1, e_2 from x to S_1 to create a set of edges T_2 . Now consider a fixed vertex v . Its degree in the random unicyclic graph T_2 is at least one and averages two. Applying Theorem 3.2 we see that the expected cost of T_2 is at least 1.25, since the cost of each edge is counted exactly twice here. It follows that the expected cost of T_1 is at least $1.25 - O(1/N)$ and the lower bound in Theorem 1.6 follows.

7.1.3. Improved Spanning Tree Lower Bound. The proof we outline here is from Aldous, Angel and Berestycki [1]. It sharpens the lower bound in Theorem 3.2. Consider an online algorithm. As in the proof of Theorem 3.2, let T be the index of the first selected item and let $q_k = \mathbf{Pr}(T > k)$ be the probability that the first k edges are all rejected. This sequence is decreasing from $q_0 = 1$ to $q_n = 0$. The threshold θ_k for accepting the k 'th edge is given by $\theta_k = \frac{q_{k-1} - q_k}{q_{k-1}}$. Conditional on $\{T \geq k\}$, the expected cost of accepting item k equals $\theta_k^2/2$.

On the event $\{T = k\}$ let a_k be the conditional expected number of subsequent edges accepted. The conditional expected cost of the subsequent edges is at least $a_k \times \frac{a_k}{2(n-k)}$, since we must accept edges of cost $\leq \frac{a_k}{n-k}$ to get a_k edges on average. Thus

$$\mathbf{E}(\text{total cost}) \geq \sum_{k=1}^n (q_{k-1} - q_k) \left(\frac{q_{k-1} - q_k}{2q_{k-1}} + \frac{a_k^2}{2(n-k)} \right).$$

Any choice of

$$(32) \quad 1 = q_0 \geq q_1 \geq \dots \geq q_n = 0; \quad 0 \leq a_k \leq n - k, \quad 1 \leq k \leq n$$

is feasible, and the constraint $\mathbf{E}|S| = 2$ becomes the constraint

$$(33) \quad \sum_{k=1}^n (q_{k-1} - q_k) a_k = 1.$$

Thus if ϕ_n is as in Theorem 3.2 then

$$\phi_n \geq \min \sum_{k=1}^n (q_{k-1} - q_k) \left(\frac{q_{k-1} - q_k}{2q_{k-1}} + \frac{a_k^2}{2(n-k)} \right)$$

minimized over (q_k) and (a_k) satisfying (32), (33). A somewhat difficult analysis gives $\phi_n \geq 2.73747$.

This completes the proof of Theorem 1.6.

7.2. Spanning Arborescence. In this section, we prove Theorem 1.10 viz. upper and lower bounds of $2 \pm o(1)$ for the expected cost of building an arborescence.

Our strategy is to construct a random mapping via the cheapest edge leaving each vertex. We keep back some edges and use them to convert this mapping into an arborescence.

7.2.1. Spanning Arborescence Upper Bound in ROM. We let $\varepsilon = 1/\log n$ and consider the first $(1 - \varepsilon)n(n - 1)$ edges of \vec{K}_n to be colored red and the remaining $\varepsilon n(n - 1)$ edges to be colored blue. Then for each vertex v we use the red edges and case $k = 1$ of Theorem 1.1 to construct a random mapping digraph D_f with vertex set $[n]$ and edges $\{(v, f(v)) : v \in [n]\}$ where $f(v)$ is v 's selection. The total expected cost of these edges is ≈ 2 .

Remark 7.1. It will be important to examine the correlation between the edges of the digraph D_f and the blue edges. What we claim is that one can see from the construction of $f(v)$ for $v \in [n]$ that if there are k_v blue edges directed out of v , then these form a uniform random choice from $\binom{[n] \setminus \{v\}}{k_v}$. Furthermore, B_v, B_w will be independent if $v \neq w$.

It is known, see Chapter 15 of [14] or Chapter 14 of [4] that if Z is the number of components of D_f then $\mathbf{E}(Z) \approx \mathbf{Var}(Z) \approx \frac{1}{2} \log n$. Thus, the Chebyshev inequality implies that Z is concentrated around its mean. We will however need an upper bound on Z that holds with probability $1 - o(1/n^3)$. It is known (see e.g. Theorem 15.1 of [14]) that the probability generating function (p.g.f.) G_n of Z is given by $G_n(x) = \mathbf{E}(x^Z) = \frac{x(x+1) \cdots (x+n-1)}{n!}$. So, for any positive integer u we have that

$$\mathbf{Pr}(Z \geq u) = \mathbf{Pr}(2^Z \geq 2^u) \leq \frac{\mathbf{E}(2^Z)}{2^u} = \frac{2 \cdot 3 \cdots (2+n-1)}{2^u n!} = \frac{(n+1)!}{2^u n!} = \frac{n+1}{2^u}.$$

We take $u = (\log n)^2$ and find that $\mathbf{Pr}(Z \geq (\log n)^2) = o(n^{-3})$ as required.

The digraph D_f consists of Z digraphs D_1, D_2, \dots, D_Z , each of which can be described as a directed cycle $C_i, i = 1, 2, \dots, Z$ with oriented trees attached to each vertex of each cycle. The edges of the trees are oriented towards the cycle. We first delete an edge (x_i, y_i) from each cycle. We are now left with Z trees rooted at $X = \{\rho_1, \rho_2, \dots, \rho_Z\}$. We will use the blue edges to merge the trees into a spanning arborescence. We do this by adding a blue edge from the root of one of the trees to a vertex in another tree. This has the effect of reducing the number of trees by one.

We go through the blue edges in the given random order. Suppose we have examined $t - 1$ blue edges. Let $X(t)$ denote the current set of roots of components. We stop when $|X(t)| = 1$.

For $\rho \in X(t)$ we let V_ρ denote the set of vertices of the tree that contains ρ . If the t th edge $e = (\rho, \sigma)$ is such that (i) $\rho \in X$, (ii) $\sigma \notin V_\rho$ and (iii) the cost $x_e \leq n^{-3/4}$ then we purchase e and reduce $X(t)$ by one. If we are successful in reducing $|X(t)|$ to one, in this way, then we will entail an additional cost of $(\log n)^2 \times n^{-3/4}$, which is negligible.

Now suppose that $|X(t)| = k$ and that the components are T_1, T_2, \dots, T_k and the tree sizes are m_1, m_2, \dots, m_k . Then the probability that the t th blue edge $e = (\rho, \sigma)$ is *good* i.e. joins a root to a vertex in a different tree is at least

$$\frac{n-2-m_1}{n^{11/4}} + \dots + \frac{n-2-m_k}{n^{11/4}} = \frac{k-1}{n^{7/4}} - \frac{2k}{n^{11/4}} \geq \frac{k-1}{2n^{7/4}}.$$

To see this, observe that there is a $1/n$ chance that ρ is the root of T_1 . There is then the probability that $\tau \notin T_1$ is at least $(n-2-m_1)/(n-2)$ (-2 as opposed to -1 from avoiding $f(\rho)$.) Finally, there is an $n^{-3/4}$ chance that the cost is at most $n^{-3/4}$.

So, with probability $1 - o(n^{-3})$, the number of blue edges needed is dominated by the sum Υ of $\ell = (\log n)^2$ independent geometric random variables with success probabilities $\lambda_k = \frac{k-1}{2n^{7/4}}$, $k = 2, 3, \dots, (\log n)^2$. The geometric random variable $Geo(\lambda)$ has p.g.f. $\frac{\lambda x}{1-(1-\lambda)x}$. The p.g.f. of Υ is therefore $\prod_{k=1}^{\ell} \frac{\lambda_k x}{1-(1-\lambda_k)x}$. So, for any $u > 0$,

$$\begin{aligned} \Pr(\Upsilon \geq u) &= \Pr\left(\left(1 + \frac{1}{4n^{7/4}}\right)^\Upsilon \geq \left(1 + \frac{1}{4n^{7/4}}\right)^u\right) \leq \\ &\left(1 + \frac{1}{4n^{7/4}}\right)^{-u} \prod_{k=2}^{\ell} \frac{\lambda_k \left(1 + \frac{1}{4n^{7/4}}\right)}{1 - (1-\lambda_k) \left(1 + \frac{1}{4n^{7/4}}\right)} \leq \exp\left\{-\frac{u}{5n^{7/4}}\right\} \prod_{k=2}^{\ell} \frac{1}{1 - \frac{1}{\lambda_k(4n^{7/4}+1)}} \\ &\leq \exp\left\{-\frac{u}{5n^{7/4}}\right\} \prod_{k=2}^{\ell} \left(1 + \frac{1}{k-1}\right) \leq \exp\left\{-\frac{u}{5n^{7/4}} + 1 + \log \ell\right\}. \end{aligned}$$

Putting $u = 100 \log n$ we see that with probability $1 - o(n^{-3})$, we only need $O(n^{7/4} \log n)$ blue edges. On the other hand we have $\Omega(n^2 / \log n)$ available, with this probability.

Finally, in the AOM model, we can appeal to the directed case of Theorem 1.9 and find a Hamilton path at cost $O(1)$, w.h.p.

7.2.2. Spanning Arborescence Lower Bound in RAM. For the lower bound we color the edges red and blue as in Section 7.2.1. Then if we want to solve the one-purchase problem of Section 3.1, we assume the edges out of vertex 1 have the costs x_1, x_2, \dots, x_N where N is the blue out-degree of vertex 1. An algorithm for finding an arborescence will provide a solution to the one-purchase problem, unless 1 is the root of the arborescence, which happens with probability $1/n$. In this case we just purchase a red edge of cost $\leq 2 \log n/n$. If the expected cost of the arborescence found was $c < 2$, then we would have a solution to the one-purchase problem with expected cost at most $c + 2 \log n/n^2$.

This completes the proof of Theorem 1.10.

8. PERFECT MATCHINGS

In this section and in Section 9 we follow the following strategy. We choose a cheap set of $O(n)$ edges that we know from prior theory contains our sought after object w.h.p. These will comprise the $O(1)$ cheapest edges incident with each vertex.

We deal with bipartite and non-bipartite separately.

8.1. Perfect Matchings in $K_{n,n}$. In this section, we prove Theorem 1.7.

8.1.1. Perfect Matching Lower Bound. For the lower bound observe that if $\mu_n \leq c < 2$ for some constant c then the average cost of each edge in an optimal algorithm is at most c/n . But then, just as in Section 7.2.2, we could use this to give an algorithm to improve the upper bound for $k = 1$ in Theorem 1.1. Given x_1, x_2, \dots, x_n we would simply make these values the costs of the edges incident to vertex $1 \in U$. The matching edge incident with vertex 1 would have expectation c/n .

8.1.2. Perfect Matching Upper Bound. For the upper bound we will use the following result of Walkup [29]: Suppose that we label the partition of the vertex set of $K_{n,n}$ as U, V . Let B_{k-out} denote the random bipartite graph where each vertex independently chooses k random neighbors from the opposite part of the bipartition giving a graph with $2kn - O(1)$ distinct edges in expectation. Note that B_{k-out} is chosen uniformly from some set of bipartite digraphs Ω_k . Then

$$\Pr(B_{3-out} \text{ does not have a perfect matching}) = o\left(\frac{1}{n}\right).$$

Walkup actually proved that B_{2-out} has a perfect matching w.h.p., but the failure probability is too high for our application.

To apply this result, we build something close to B_{3-out} as follows: As we see an edge, we color it red or blue with probability $1/2$. A vertex $v \in U$ uses the red edges incident with it and the algorithm of Theorem 1.1 to choose three edges incident with it. A vertex $w \in V$ uses the blue edges incident with it and the algorithm of Theorem 1.1 to choose three edges incident with it. Let H be the bipartite graph created.

The expected total cost of the edges in H is $\approx 2n \times c_3/(n/2) = 4c_3$ as claimed in Theorem 1.7. It will have a perfect matching with probability $1 - o(1/n)$.

There is a small point to clarify. The bipartite graph produced by the algorithm is B_{3-out} with the condition that no edge is chosen by both of its endpoints. The expected number of edges chosen twice in B_{3-out} is at most $n^2 \times (2/n)^2 = 6$. By computing higher moments we see that the probability no edge is chosen twice is $\approx e^{-6}$ and so the probability there is no perfect matching in the algorithm's graph is at most $e^6 + o(1)$ times the probability there is no perfect matching in B_{3-out} .

This completes the proof of Theorem 1.7.

8.2. Perfect Matchings in K_n . We can assume that $n = 2m$ is even. We can then look for a perfect matching between $[1, m]$ and $[m + 1, 2m]$. For this we can use the approach of Section 8.1.2 and this will give us an upper bound of $\approx 4c_3$. For a lower bound we can use $\approx n/2 \times 2/(n/2) = 2$, since we could use any algorithm for finding a perfect matching to find a solution to the one-purchase problem in Section 3.1 as we did for the lower bound in Section 1.7.

This completes the proof of Theorem 1.8.

9. THE TRAVELING SALESPERSON PROBLEM

9.1. TSP Upper Bound. For this we first replace Walkup’s result by the following: let G_{k-out} be the random graph constructed by allowing each vertex to independently choose k random neighbors. It is known, Bohman Frieze [3] that w.h.p. G_{k-out} is Hamiltonian w.h.p. if $k \geq 3$. We want this probability to be $1 - o(1/n)$ and to be sure of this we can use the result of Frieze [13] where we take $k = 10$.

We then write the uniform $[0, 1]$ random variables X_e as $X_e = \min \{Z_{e,j} : j = 1, 2, \dots, 10\}$ where the $Z_{e,j}$ are independently distributed as the random variable $Z \in [0, 1]$ where $\Pr(Z \geq x) = (1 - x)^{10}$. Assume first that when we examine an edge, we see these 10 values. Next, for $j = 1, 2, \dots, 10$ we use Corollary 1.2 to choose an edge $e_{v,j}$ for each $v \in [n]$. The edges chosen will create a graph distributed as G_{10-out} and the total cost will be at most $\approx 10 \times 10 \times 2 = 200$. The first 10 arises because we do this for 10 values of j . The next 10 arises because the density of Z near zero is $10x + o(x)$ and the 2 arises because $c_1 = 2$.

Of course when we examine an edge, we only see one value, ξ say. To get around this, we generate another 9 values of Z , viz. Z_2, Z_3, \dots, Z_{10} , but we condition here on $Z_j \geq \xi, j = 2, 3, \dots, 10$.

9.2. TSP Lower Bound. For a lower bound, we see that we could use an algorithm for finding a low cost Hamilton cycle to find a solution to the 2-purchase problem in Theorem 1.1.

This completes the proof of Theorem 1.9.

9.3. Directed case of TSP. In this case we replace the results of [3], [13] by the result of Cooper and Frieze [7]. It is shown in [7] that w.h.p. the random digraph $G_{2-in,2-out}$ is Hamiltonian. Here each vertex independently chooses 2 random out-neighbors and 2 random in-neighbors. To apply this, we replace each uniform edge-cost $X_{e.e} = (u, v)$ by $\min \{Z_{e,j} : j = 1, 2, \dots, 4\}$ where the $Z_{e,j}$ are independently distributed as the random variable $Z \in [0, 1]$ where $\Pr(Z \geq x) = (1 - x)^4$. We will then use $Z_{e,1}, Z_{e,2}$ to give 2 random out-neighbors to u and $Z_{e,3}, Z_{e,4}$ to give 2 random in-neighbors to v . The rest of the argument is as for the undirected case. For the lower bound we see that each vertex chooses an in-neighbor and an out-neighbor. For the upper bound we replace $\approx 10 \times 10 \times 2$ by $\approx 2 \times 2 \times c_2$. We cannot claim a bound in terms of expectation because the proofs in [7] and the related [6] do not give a small enough probability of failure.

This completes the proof of Theorem 1.11.

10. FINAL REMARKS

We have described several problems that can be analyzed within our framework. It would be of some interest to

- (1) Tighten the bounds, especially for minimum spanning trees.
- (2) Replace K_n by other graphs.
- (3) Extend the analysis to hypergraph structures.

Acknowledgement: We thank Colin Cooper for his comments and his contribution to the proof of Theorem 1.1.

REFERENCES

- [1] D. Aldous, O. Angel and N. Berestycki, Online Random Weight Minimal Spanning Trees and a Stochastic Coalescent, unpublished manuscript, 2008.
- [2] M. Babaioff, N. Immorlica and R. Kleinberg, Matroids, secretary problems and on-line mechanisms, *Proceedings of the eighteenth annual ACM-SIAM Symposium on Discrete Algorithms* (2007) 434-443.
- [3] T. Bohman and A.M. Frieze, Hamilton cycles in 3-out, *Random Structures and Algorithms* 35 (2009) 393-417.
- [4] B. Bollobás, *Random Graphs*, First Edition, Academic Press, London 1985, Second Edition, Cambridge University Press, 2001.
- [5] A. Cayley, Mathematical questions with their solutions, *The Educational Times* 23 (1875) 18-19. See *The Collected Mathematical Papers of Arthur Cayley* 10 (1896) 587-588, Cambridge University Press.
- [6] C. Cooper and A.M. Frieze, Hamilton cycles in a class of random directed graphs *Journal of Combinatorial Theory B* 62 (1994) 151-163
- [7] C. Cooper and A.M. Frieze, Hamilton cycles in random graphs and directed graphs, *Random Structures and Algorithms* 16 (2000) 369-401.
- [8] E.B. Dynkin, The optimum choice of the instant for stopping a Markov process, *Sov. Math. Dokl.* 4 (1963).
- [9] H. Esfandiari, M. Hajiaghayi, V. Liaghat and M. Monemizadeh, Prophet Secretary, arXiv:1507.01155.
- [10] P. Erdős and A. Rényi, On the evolution of random graphs, *Publ. Math. Inst. Hungar. Acad. Sci.* 5 (1960) 17-61.
- [11] T.S. Ferguson, Who solved the secretary problem?, *Statistical Science* 4 (1989) 282-296.
- [12] A.M. Frieze, On the value of a minimum spanning tree problem, *Discrete Applied Mathematics* 10 (1985) 47-56.
- [13] A.M. Frieze, Finding hamilton cycles in sparse random graphs, *Journal of Combinatorial Theory B* 44 (1988) 230-250.
- [14] A.M. Frieze and M. Karoński, *Introduction to Random Graphs*, Cambridge University Press, Cambridge, 2015.
- [15] A.M. Frieze, W. Pegden and G.B. Sorkin, On the distribution of the minimum weight clique, to appear.
- [16] T.P. Hill and R.P. Kertz, Comparisons of stop rule and supremum expectations of i.i.d. random variables, *The Annals of Probability* 10 (1982) 336-345.
- [17] T.P. Hill and R.P. Kertz, A Survey of Prophet Inequalities in Optimal Stopping Theory, *Contemporary Mathematics* 5 (1992) 191-207.
- [18] W. Hoeffding, Probability inequalities for sums of bounded random variables, *Journal of the American Statistical Association* 58 (1963) 13-30.

- [19] T. Huynh and P. Nelson, The matroid secretary problem for minor-closed classes and random matroids, arXiv:1603.06822.
- [20] G. Kamath and D. Karger, <http://www.gautamkamath.com/writings/matroidsec.pdf>.
- [21] S. Janson, One, two and three times $\log n/n$ for paths in a complete graph with random weights, *Combinatorics, Probability and Computing* 8 (1999) 347-361.
- [22] R.P. Kertz, Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality (unpublished).
- [23] U. Krengel and L. Sucheston, Semiamarts and finite values, *Bulletin of the American Mathematical Society* 83 (1977) 745-747.
- [24] R. Kleinberg and S. Weinberg, Matroid Prophet Inequalities, *Proceedings of the forty-fourth annual ACM symposium on Theory of computing* (2012) 123-136 .
- [25] C. McDiarmid, On the method of bounded differences, in *Surveys in Combinatorics, ed. J. Siemons, London Mathematical Society Lecture Notes Series* 141, Cambridge University Press, 1989.
- [26] L. Moser, On a problem of Cayley, *Scripta Mathematica* 22 (1956) 289-292.
- [27] E. Samuel-Cahn, Comparison of threshold stop rules and maximum for independent non-negative random variables, *The Annals of Probability* 12 (1984) 1213-1216.
- [28] A. Rubinstein, Beyond matroids: secretary problem and prophet inequality with general constraints, arXiv:1604.00357.
- [29] D.W. Walkup, Matchings in random regular bipartite graphs, *Discrete Mathematics* 31 (1980) 59-64.

DEPARTMENT OF MATHEMATICAL SCIENCES,, CARNEGIE MELLON UNIVERSITY,, PITTSBURGH PA 15213.