

# Threshold rules for online sample selection

Eric Bach\*

Shuchi Chawla†

Seeun Umboh‡

## Abstract

We consider the following sample selection problem. We observe in an online fashion a sequence of samples, each endowed by a quality. Our goal is to either select or reject each sample, so as to maximize the aggregate quality of the subsample selected so far. There is a natural trade-off here between the rate of selection and the aggregate quality of the subsample. We show that for a number of such problems extremely simple and oblivious “threshold rules” for selection achieve optimal tradeoffs between rate of selection and aggregate quality in a probabilistic sense. In some cases we show that the same threshold rule is optimal for a large class of quality distributions and is thus oblivious in a strong sense.

---

\*Computer Sciences Dept., University of Wisconsin - Madison, [bach@cs.wisc.edu](mailto:bach@cs.wisc.edu).

†Computer Sciences Dept., University of Wisconsin - Madison, [shuchi@cs.wisc.edu](mailto:shuchi@cs.wisc.edu).

‡Computer Sciences Dept., University of Wisconsin - Madison, [seeun@cs.wisc.edu](mailto:seeun@cs.wisc.edu).

# 1 Introduction

Imagine a heterogeneous sequence of samples from an array of sensors, where samples have different utilities reflecting their accuracy, quality, or applicability to the task at hand. Then we may wish to discard all but the most relevant or useful samples. Further suppose that this sample selection is performed online — every time we receive a new sample we must make an irrevocable decision to keep it or discard it. What rules can we use for sample selection? There is a tradeoff here: while we want to retain only the most useful samples, we may not want to be overly selective and discard a large fraction. So we could either fix a rate of selection (the number of examples we want to retain as a function of the number we see) and ask for the best quality subsample, or fix a desirable level of quality as a function of the size of the subsample and ask to achieve this with the fewest samples rejected.

An example of online sample selection is the following “hiring” process that has been studied previously. Imagine that a company wishing to grow interviews candidates to observe their qualifications, work ethic, compatibility with the existing workforce, etc. How should the company make hiring decisions so as to obtain the highest quality workforce possible? As for the sensor problem, there is no single correct answer here. Rather a good hiring strategy depends on the rate at which the company plans to grow—again there is a trade-off between being overly selective and growing fast. Broder et al. [7] studied this hiring problem in a simple setting where each candidate’s quality is a one-dimensional random variable and the company wants to maximize the average or median quality of its workforce.

In general performing such selection tasks may require complicated rules that depend on the samples seen so far. Our main contribution is to show that in a number of settings an extremely simple class of rules that we call “threshold rules” is close to optimal (within constant factors) for this problem in a probabilistic sense.

Specifically, suppose that each sample is endowed with a “quality”, which is a random variable drawn from a known distribution. We are interested in maximizing the aggregate quality of a set of samples, which is a numerical function of the individual qualities. In the settings we consider the quality of a sample can be multi-dimensional and the aggregate quality of a set can be the average quality in that set or some other more complicated aggregation function. Suppose that we want to select a subset of  $n$  samples out of a total of  $T$  seen. Let  $Q_{T,n}^{\text{OPT}}$  denote the aggregate quality that the best offline algorithm can achieve having seen all  $T$  samples. Our goal is to design an online selection rule that approximates  $Q_{T,n}^{\text{OPT}}$  in expectation over the  $T$  samples. We use two measures of approximation — the ratio of the expected quality achieved by the offline optimum to that achieved by the online selection rule,  $E[Q_{T,n}^{\text{OPT}}]/E[Q_{T,n}^{\text{online}}]$ , and the expectation of the ratio of the qualities of the two rules,  $E[Q_{T,n}^{\text{OPT}}/Q_{T,n}^{\text{online}}]$ . Here the expectation is taken over the distribution from which the sample is drawn. In either case, we say that the online rule is optimal if the resulting approximation is bounded from above by a constant independent of  $n$ . Note that it is always at least 1.

To put this in context, consider the setting studied by Broder et al. [7]. Each sample is associated with a quality in the range  $[0, 1]$ , and the goal is to maximize the average quality of the subsample we pick. Broder et al. show that if the quality is distributed uniformly in  $[0, 1]$  a natural “select above the mean” rule is optimal to within constant factors with respect to the optimal offline algorithm that has the same selection rate as the rule. In fact they show that the same observation holds also for the “select above the median” rule. Both of these rules are adaptive in the sense that the next selection decision depends on the quality of the samples seen so far. In more general settings, adaptive rules of this kind can require unbounded space to store information about samples seen previously. For example, consider the following 2-dimensional skyline problem: each sample is a point in a unit square; the quality of a single point  $(x, y)$  is the area of its “shadow”  $[0, x] \times [0, y]$ , and the quality of a set of points is the area of the collective shadows of all the

points; the goal is to pick a subsample with the largest shadow. In this case, a natural selection rule may be to select a sample if it falls out of the shadow of the previously seen points. However, by [4], we can see that implementing this rule requires remembering on average  $O(\log n)$  samples out of  $n$  samples seen. We therefore study non-adaptive selection rules.

We focus in particular on so-called “threshold rules” for selection. A threshold rule specifies a certain criterion that a candidate must satisfy to get selected. Most crucially, the threshold is determined *a priori* (based on a desirable selection rate); it depends only on the number of samples picked so far and is otherwise independent of the samples. Threshold rules are extremely simple oblivious rules and can, in particular, be “hard-wired” into the selection process. This suggests the following natural questions. When are threshold rules optimal for online selection problems? Does the answer depend on the desired rate of selection? We answer these questions in three different settings in this paper.

The first setting we study is a single-dimensional-quality setting which is similar to Broder et al.’s model. In this setting, we study threshold rules of the form “*Pick the next sample whose quality exceeds  $f(i)$* ” where  $i$  is the number of samples picked so far. We show that for a large class of functions  $f$  these rules give constant factor approximations.

In the second setting, samples are nodes in a rooted infinite-depth tree. Each node is said to cover all the nodes on the unique path from the root to itself. The quality of a collection of nodes is the total number of distinct nodes that they collectively cover. This is different from the first setting in that the quality defines only a partial order over the samples. Once again, we study threshold rules of the form “*Pick the next sample whose quality exceeds  $f(i)$* ” and show that they are constant factor optimal.

Our third setting is a generalization of the skyline problem described previously. Specifically, consider a domain  $X$  with a probability measure  $\mu$  and a partial ordering  $\prec$  over it. For an element  $x \in X$ , the “shadow” of  $x$  is the set of all the points that it dominates in this partial ordering,  $\text{sh}(x) = \{y : y \prec x\}$ ; likewise the shadow of a sample  $S \subseteq X$  is  $\text{sh}(S) = \cup_{x \in S} \text{sh}(x)$ . Once again, as in the second setting, we can define the coverage of a single sample to be the measure of all the points in its shadow. However, unlike the tree setting, here it is usually easy to obtain a constant factor approximation to coverage—the maximum coverage achievable is 1 (i.e. the measure of the entire universe), whereas in many cases (e.g. for the uniform distribution over the unit square) a single random sample can in expectation obtain constant coverage. We therefore measure the quality of a subsample  $S \subset X$  by its “gap”:  $\text{Gap}(S) = 1 - \mu(\text{sh}(S))$ . In this setting, rules that place a threshold on the quality of the next sample to be selected are not constant-factor optimal. Instead, we study threshold rules of the form “*Pick the next sample  $x$  for which  $\text{sh}'(x)$  exceeds  $f(i)$* ”, where  $\text{sh}'(x) = \{y : x \prec y\}$  is the measure of elements that dominate  $x$ , and show that these rules obtain constant factor approximations.

## 1.1 Related work

As mentioned earlier, our work is inspired by and extends the work of Broder et al. [7]. Broder et al. consider a special case of the one-dimensional selection problem described above. They assume that the quality of a sample is distributed uniformly over the interval  $(0, 1)$ ; this assumption is not without loss of generality. They analyze two adaptive selection rules—*select above the mean*, and *select above the median*—and show that both are constant-factor optimal, although they lead to different growth rates. These rules are adaptive in the sense that the next selection decision depends on the quality of the samples accepted so far. Note that the *select above the median* rule requires the algorithm to remember all of the samples accepted so far, and is therefore a computationally intensive rule. Even the relatively simpler *select above the mean* rule requires remembering the current mean and number so far accepted. In contrast we show (Section 3) that there exists a class of simple non-adaptive selection strategies that also achieves optimality and includes rules

with selection rates equal to those of the ones studied by Broder et al. These strategies make decisions based only on the number hired so far. Furthermore we extend these results to more general coverage problems.

Our third setting is closely related to the skyline problem that has been studied extensively in online settings by the database community (see, for example, [2] and references therein). Kung, et al. [12] gave an offline divide-and-conquer algorithm that finds the skyline of a given set of vectors in  $d$ -dimensional space. Their algorithm uses  $O(n \log_2 n)$  comparisons of vector components when  $d = 2, 3$  and  $O(n(\log_2 n)^{d-2})$  when  $d \geq 4$ . The implementation of a *Skyline* query for database systems was recently introduced by [6]. The closest in spirit to our work is [15]. They considered a stream of uncertain objects to model uncertainty in measurement. Each object has an associated set of possible instances and they are interested in the objects whose probability of being dominated by another object is at most some  $q$  supplied by the database user.

Online sample selection is closely related to secretary problems, however there are some key differences. In secretary problems (see, e.g., [8, 9, 13]) there is typically a fixed bound on the desired number of hires. In our setting on the other hand, the selection process is ongoing and we must pick more and more samples as we see more of them. This makes the tradeoff between the rate of hiring and the rate of improvement of quality interesting.

Finally, while our goal is to analyze a class of online algorithms in comparison to the optimal offline algorithms, our approach is different from the competitive analysis of online algorithms [5]. In competitive analysis the goal is to perform nearly as well as the optimal offline algorithm for *any arbitrary* sequence of input. In contrast, we bound the *expected* competitive ratio of the rules we study. Furthermore, a crucial aspect of the strategies that we study is that not only are they online, but they are also non-adaptive or oblivious. That is, the current acceptance threshold does not depend on the samples seen by the algorithm so far. In this sense, our model is closer in spirit to work on oblivious algorithms (see, e.g., [11, 3, 10]). Oblivious algorithms are highly desirable in practical settings because the rules can be hard-wired into the selection process, making them very easy to implement. The caveat is, of course, that for many optimization problems oblivious algorithms do not provide good approximations. Surprisingly, we show that in many scenarios related to sample selection, obliviousness has only a small cost.

## 2 Models and results

Let  $X$  be a domain with probability measure  $\mu$  over it. A *threshold rule*  $\mathcal{X}$  is specified by a sequence of subsets of  $X$  indexed by  $X = \mathcal{X}_0 \supseteq \mathcal{X}_1 \supseteq \mathcal{X}_2 \supseteq \dots \supseteq \mathcal{X}_n \supseteq \dots$ . A sample is selected if it belongs to  $\mathcal{X}_i$  where  $i$  is the number of samples previously selected.

Let  $\mathcal{T}$  be a random infinite sequence of samples drawn according to  $\mu$ . Let  $\mathcal{T}^{\mathcal{X}}(n)$  denote a prefix of  $\mathcal{T}$  of length  $T_n^{\mathcal{X}}$  such that the last sample on this prefix is the  $n$ th sample chosen by the threshold rule  $\mathcal{X}$ . Note that the length of this prefix, that we denote by  $T_n^{\mathcal{X}}$ , is a random variable. We drop the superscript and the subscript when they are clear from the context. The “selection overhead” of a threshold rule as a function of  $n$  is simply defined to be the expected waiting time to select  $n$  samples, or  $E[T_n^{\mathcal{X}}]$ , where the expectation is over the samples drawn according to  $\mu$ .

Let  $Q$  be a function denoting “quality”. Thus  $Q(x)$  denotes the quality of a sample  $x$  and  $Q(S)$  the aggregate quality of a set  $S \subset X$  of samples. Note that  $Q(x)$  is a random variable and we assume that it is drawn from a known distribution. Let  $Q_{\mathcal{T},n}^{\text{OPT}}$  denote the quality of an optimal subset of  $n$  out of a set  $\mathcal{T}$  of samples with respect to measure  $Q$ . We use  $Q_n^{\text{OPT}}$  as shorthand for  $Q_{\mathcal{T}^{\mathcal{X}}(n),n}^{\text{OPT}}$  where  $\mathcal{X}$  is clear from the context. Let  $Q_{\mathcal{T}^{\mathcal{X}}(n),n}^{\mathcal{X}}$  or  $Q_n^{\text{online}}$  denote the quality of a sample of size  $n$  selected by threshold rule  $\mathcal{X}$  with respect to measure  $Q$ . Note that these quantities are all random variables that depend on the stream of samples drawn according to  $\mu$ .

We look at both maximization and minimization problems. When our goal is to maximize the quality of a set of samples, we say that a threshold rule  $\mathcal{X}$  achieves a *competitive ratio of  $\alpha$  in expectation* with respect to  $Q$  if for all  $n$ ,

$$E_{\mathcal{T} \sim \mu} \left[ \frac{Q_n^{\text{OPT}}}{Q_n^{\text{online}}} \right] \leq \alpha$$

On the other hand, a threshold rule  $\mathcal{X}$   *$\alpha$ -approximates expected coverage* with respect to  $Q$  if for all  $n$ ,

$$\frac{E_{\mathcal{T} \sim \mu}[Q_n^{\text{OPT}}]}{E_{\mathcal{T} \sim \mu}[Q_n^{\text{online}}]} \leq \alpha$$

Likewise, for minimization problems, the competitive ratio and approximation with respect to  $Q$  are defined similarly:

$$E_{\mathcal{T} \sim \mu} \left[ \frac{Q_n^{\text{online}}}{Q_n^{\text{OPT}}} \right] \leq \alpha \qquad \frac{E_{\mathcal{T} \sim \mu}[Q_n^{\text{online}}]}{E_{\mathcal{T} \sim \mu}[Q_n^{\text{OPT}}]} \leq \alpha$$

We now describe the specific settings we study and the results we obtain.

**Model 1: Unit interval (Section 3).** Our first setting is the one-dimensional setting studied by Broder et al. [7]. Specifically, each sample is associated with a quality drawn from a distribution over the unit line. Our measure of success is the mean quality of the subsample we select. Note that in the context of approximately optimal selection rules this is a weak notion of success. For example when  $\mu$  is the uniform distribution over  $[0, 1]$ , even in the absence of any selection rule we can achieve a mean quality of  $1/2$ , while the maximum achievable is 1. So instead of focussing on the mean quality, we focus on the *mean quality gap*— $\text{Gap}(S) = 1 - (\sum_{x \in S} x)/|S|$ —of the subsample.

We focus in particular on *power-law* distributions on the unit line, i.e. distributions with c.d.f.  $\mu(1 - x) = 1 - x^k$  for some constant  $k$ , and study threshold rules of the form  $\mathcal{X}_i = \{x \mid x \geq 1 - c_i\}$  where  $c_i = \Omega(1/\text{poly}(i))$ . We show that these threshold rules are constant factor optimal simultaneously for *any* power-law distribution. This is remarkable since it gives an oblivious selection algorithm that is optimal. In other words, we can pick a threshold rule from the class with the preferred tradeoff between selection speed and increase in mean quality between accepted samples without knowing which power-law gap distribution the interviewees are coming from, and we would still be choosing nearly optimally. In contrast, Broder et al.'s optimality result works only for the uniform distribution<sup>1</sup> We obtain the following result (see definitions in Section 3).

**Theorem 1** *For the unit line equipped with a power-law distribution, any threshold rule  $\mathcal{X}_i = \{x \mid x \geq 1 - c_i\}$ , where  $c_i = 1/i^\alpha$  with  $0 \leq \alpha < 1$  for all  $i$ , achieves an  $O(1)$  approximation to the expected gap.*

**Dominance and shadow.** For the next two settings, we need some additional definitions. Let  $\prec$  be a partial order over the universe  $X$ . As defined earlier, the shadow of an element  $x \in X$  is the set of all the points that it dominates,  $\text{sh}(x) = \{y : y \prec x\}$ ; likewise the shadow of a sample  $S \subseteq X$  is  $\text{sh}(S) = \cup_{x \in S} \text{sh}(x)$ . Let  $\text{sh}'(x) = \{y : x \prec y\}$  be the set of points that shadow  $x$ ;  $\text{sh}'(S)$  for a set  $S$  is defined similarly. Note that  $\text{sh}'(x)$  is a subset of  $X \setminus \text{sh}(x)$  and  $\mu(\text{sh}'(x))$  is the probability that a random sample covers  $x$ .

<sup>1</sup>While the result can be extended to any arbitrary distribution via a standard transformation from one space to another, the resulting selection rule becomes distribution dependent.

**Model 2: Random tree setting (Section 4).** While the previous setting was in a continuous domain, next we consider a discrete setting, where the goal is to maximize the cardinality of the shadow set. Specifically, our universe  $X$  is the set of all nodes in an rooted infinite-depth binary tree. The following random process generates samples. Let  $0 < p < 1$ . We start at the root and move left or right at every step with equal probability. At every step, with probability  $p$  we terminate the process and output the current node. A node  $x$  in the tree dominates another node  $y$  if and only if  $y$  lies on the unique path from the root to  $x$ . For a set  $S$  of nodes, we define coverage as  $\text{Cover}(S) = |\text{sh}(S)|$ . Note, that unlike in the previous setting, there is no notion of a gap in this setting.

Once again the threshold rules we consider here are based on sequences of integers  $\{c_i\}$ . For any such sequence, we define  $\mathcal{X}_i = \{x \mid |\text{sh}(x)| \geq c_i\}$ . We show that constant-factor optimality can be achieved with subexponential as well as exponential selection overheads.

**Theorem 2** *For the binary tree model described above, any threshold rule based on a sequence  $\{c_i\}$  with  $c_i = O(\text{poly}(i))$  achieves an  $O(1)$  competitive ratio in expectation with respect to coverage, as well as an  $O(1)$  approximation to the expected coverage.*

**Model 3: Skyline problem (Section 5).** Finally, we consider another continuous domain that is a generalization of the skyline problem mentioned previously. We are interested in selecting a set of samples with a large shadow. Specifically, we define the “gap” of  $S$  to be  $\text{Gap}(S) = 1 - \mu(\text{sh}(S))$ . Our goal is to minimize the gap.

We show that a natural class of threshold rules obtains near-optimal gaps in this setting. Recall that  $\text{sh}'(x)$  for an element  $x \in X$  denotes the set of elements that dominate  $x$ . We consider threshold rules of the form  $\mathcal{X}_i = \{x \in X \mid \mu(\text{sh}'(x)) \leq c_i\}$  for some sequence of numbers  $\{c_i\}$ . We require the following continuity assumption on the measure  $\mu$ .

**Definition 1 (Measure continuity)** *For all  $x \in X$  and  $c \in [0, \mu(\text{sh}'(x))]$ , there exists an element  $y \in \text{sh}'(x)$  such that  $\mu(\text{sh}'(y)) = c$ . Furthermore, there exist elements  $\underline{x}, \bar{x} \in X$  with  $\text{sh}'(\underline{x}) = X$  and  $\text{sh}'(\bar{x}) = \emptyset$ .*

Measure continuity ensures that the sets  $\mathcal{X}_i$  are all non-empty and proper subsets of each other.

**Theorem 3** *For the skyline setting with an arbitrary measure satisfying measure continuity, any threshold rule based on a sequence  $\{c_i\}$  with  $c_i = i^{-(1/2 - \Omega(1))}$  achieves a  $1 + o(1)$  competitive ratio in expectation with respect to the gap.*

We note that the class of functions  $c_i$  specified in the above theorem includes all functions for which  $1/c_i$  grows subpolynomially. In particular, this includes threshold rules with selection overheads that are slightly superlinear.

For the special case of the skyline setting over a two-dimensional unit square  $[0, 1]^2$  bestowed with a product distribution and the usual precedence ordering— $(x_1, y_1) \prec (x_2, y_2)$  if and only if  $x_1 \leq x_2$  and  $y_1 \leq y_2$ —we are able to obtain a stronger result that guarantees constant-factor optimality for any polynomial selection overhead:

**Theorem 4** *For the skyline setting on the unit square with any product distribution, any threshold rule based on a sequence  $\{c_i\}$  with  $c_i = \Omega(1/\text{poly}(i))$  achieves a  $1 + o(1)$  competitive ratio in expectation with respect to the gap.*

### 3 Sample selection in one dimension

In this section we study one-dimensional sample selection problems. Specifically, each sample is associated with a quality drawn from a distribution over the unit line. Recall that our goal is to minimize the *mean quality gap*,  $\text{Gap}(S) = 1 - (\sum_{x \in S} x)/|S|$ , of the subsample that we pick. We focus in particular on *power-law* distributions on the unit line, i.e. distributions with c.d.f.  $\mu(1 - x) = 1 - x^k$  for some constant  $k$ , and study threshold rules of the form  $\mathcal{X}_i = \{x \mid x \geq 1 - c_i\}$  where  $c_i = \Omega(1/\text{poly}(i))$ .

For a (random) variable  $x \in [0, 1]$ , let  $\bar{x}$  denote its complement  $1 - x$ . For a cumulative distribution  $\mu$  with domain  $[0, 1]$ , we use  $\bar{\mu}$  to denote the cumulative distribution for the complementary random variable:  $\bar{\mu}(x) = 1 - \mu(1 - x)$ .

Let  $X_n$  denote the random variable representing the quality of the  $n$ th sample selected by the online rule, and let  $\text{Gap}_n^{\text{online}}$  denote the mean quality gap of the first  $n$  selected samples. That is,  $\text{Gap}_n^{\text{online}} = \frac{1}{n} \sum_{i=1}^n \bar{X}_i$ .

We will now derive an expression for  $\text{Gap}_n^{\text{online}}$ . Let  $X$  denote a random variable drawn from the distribution  $\mu$ . Then recall that for power-law distributions  $\mu$  we have that  $\bar{X}_i$  is statistically identical to  $\bar{X} \cdot c_i$ . This implies the following series of equations:

$$\text{Gap}_n^{\text{online}} = \sum_{i=1}^n \frac{\bar{X}_i}{n} = \sum_{i=1}^n \frac{\bar{X} \cdot c_i}{n}.$$

So, by linearity of expectation, we have  $E[\text{Gap}_n^{\text{online}}] = \frac{E[\bar{X}]}{n} \sum_{i=1}^n c_i$ .

We start by determining the optimal mean quality gap achievable when we pick a subsample of size  $n$  from a sample of size  $E[T_n]$ . See Appendix A for a proof of this lemma.

**Lemma 5** *The expected mean gap of the largest  $n$  out of  $T$  samples drawn from a distribution with  $\bar{\mu}(x) = x^k$  is asymptotically equal to  $\frac{1}{1+1/k} \left(\frac{n}{T+1}\right)^{1/k}$ .*

We also note the following bound on the selection rate.

**Fact 6**  $E[T_n] = \sum_{i=1}^n 1/\bar{\mu}(c_i) = \sum_{i=1}^n 1/c_i^k$ .

The next lemma bounds the gap of a large class of threshold rules.

**Lemma 7** *For selection thresholds  $c_i = 1/i^\alpha$  with  $0 \leq \alpha < 1$ , we have that*

$$E[\text{Gap}_n^{\text{online}}] \leq \frac{1}{1 - \alpha} \left(\frac{n}{E[T_n]}\right)^{1/k}.$$

*Proof:* Using Fact 6 we have

$$\begin{aligned} E[\text{Gap}_n^{\text{online}}] \cdot (E[T_n])^{1/k} &= \left(\frac{E[\bar{X}]}{n} \sum_{i=1}^n i^{-\alpha}\right) \left(\sum_{i=1}^n i^{\alpha k}\right)^{1/k} \\ &\sim \frac{E[\bar{X}]}{n} \cdot \frac{n^{1-\alpha}}{1 - \alpha} \cdot \left(\frac{n^{\alpha k + 1}}{\alpha k + 1}\right)^{1/k} \leq \frac{n^{1/k}}{1 - \alpha}, \end{aligned}$$

where the second line follows from the Euler-Maclaurin formula and the third follows from noting that  $E[\bar{X}] \leq 1$ . This gives the inequality that we want.  $\blacksquare$

Lemmas 5 and 7 together show that the expected mean gap of these threshold rules is only a small constant factor bigger than the optimal offline selection rule that picks the best  $n$  out of  $E[T_n]$  samples. Next we show that these threshold rules are in fact constant factor optimal in the following stricter sense: if an adversary were allowed to choose any  $n$  out of  $T_{n+1} - 1$ , its expected mean gap is only a constant factor smaller than that of the online algorithm. We denote this optimal offline gap by  $\text{Gap}_{n+1}^{\text{OPT}}$ . The reason we are interested in  $T_{n+1} - 1$  is because the adversary should be able to use the samples we rejected while waiting for the  $(n + 1)$ th selection.

**Lemma 8** For  $c_i$  satisfying  $c_i = 1/i^\alpha$  with  $0 \leq \alpha < 1$  for all  $i$ , we have

$$E[\text{Gap}_{n+1}^{\text{OPT}}] \geq \frac{1}{16} \left( \frac{n}{E[T_n]} \right)^{1/k}$$

*Proof:* By Markov's inequality and Lemma 5, we have that

$$E[\text{OPT}_n] \geq \frac{1}{2} E[\text{OPT}_n | T_{n+1} \leq 2E[T_{n+1}]] \geq \frac{1}{2} \frac{1}{(1 + 1/k)} \cdot \left( \frac{n}{2E[T_{n+1}]} \right)^{1/k}$$

using Fact 6,

$$\approx \left( \frac{1}{2} \right)^{1+1/k} \frac{1}{1 + 1/k} \left( \frac{n}{E[T_n]} \right)^{1/k} \left( 1 + \frac{1}{n} \right)^{-\alpha k} \geq \frac{1}{16} \left( \frac{n}{E[T_n]} \right)^{1/k}$$

■

Together with Lemma 7, this proves Theorem 1 with the constant in the  $O(1)$  equal to  $16/(1 - \alpha)$ .

## 4 Sample selection in binary trees

We begin with some notation and observations. For a sample (node)  $x$  in the tree let  $\text{sh}(x)$  denote both the unique path from the root to  $x$  as well as the set of nodes covered by  $x$  (the shadow of  $x$ ). Let  $\text{sh}_k(x)$  be the  $k$ th node on  $\text{sh}(x)$ ,  $\text{sh}_{\leq k}$  be the first  $k$  nodes of  $\text{sh}(x)$ , and  $\text{sh}_{\geq k} = \text{sh}(x) \setminus \text{sh}_{< k}$ .

We say that a set of  $n$  paths associated with nodes  $x_1, \dots, x_n$  is *independent* at level  $k$  if  $|\cup_{i=1}^n \text{sh}_k(x_i)| = n$ . In other words, no two paths share the same vertex at level  $k$ . Thus, the paths are disjoint after level  $k$ . This implies the following fact.

**Fact 9** If a set of  $n$  paths  $\{\text{sh}(x_i)\}$  of length at least  $k'$  is independent at level  $k < k'$ , then

$$|\text{sh}(\{x_1, \dots, x_n\})| = \left| \bigcup_{i=1}^n \text{sh}(x_i) \right| \geq \left| \bigcup_{i=1}^n \text{sh}_{\geq k}(x_i) \right| \geq n(k' - k).$$

Our analysis depends on whether  $c_i$  is a slow-growing or fast-growing function. We first consider the case of  $c_i = O(\text{poly } n)$  but with  $c_i \geq \log_2 i$  for all  $i$ .

**Theorem 10** For the binary tree model described above, any threshold rule based on a sequence  $\{c_i\}$ , with  $c_i = O(\text{poly}(i))$  and  $c_i \geq \log i$  for all  $i$ , achieves an  $O(1)$  competitive ratio in expectation with respect to coverage, as well as an  $O(1)$  approximation to the expected coverage.

*Proof:* Let  $f(i) = c_i - \log i$ . We will first obtain an upper bound on  $\text{Cover}_n^{\text{OPT}}$ . Let  $S_n$  be the  $n$  selected nodes,  $O_n$  the optimal set of  $n$  paths, and  $R_n$  the paths that are rejected and are not covered by  $S_n$ .

$$\begin{aligned} \text{Cover}_n^{\text{OPT}} &= |(\text{sh}(O_n) \cap \text{sh}(R_n)) \cup (\text{sh}(O_n) \cap \text{sh}(S_n))| \\ &\leq |\text{sh}(O_n) \cap \text{sh}(R_n)| + |\text{sh}(S_n)| \leq (2n + nf(n)) + \text{Cover}_n^{\text{online}}, \end{aligned}$$

Here the last inequality follows by noting that  $\text{sh}(O_n) \cap \text{sh}(R_n)$  forms a binary tree with at most  $2n$  vertices in the first  $\log n$  levels, and at most  $nf(n)$  other vertices since it is the union of  $n$  paths of length at most  $\log n + f(n)$ .

Next we obtain a lower bound on  $\text{Cover}_n^{\text{online}}$ . Consider the last  $n/2$  selected nodes  $s_{n/2+1}, \dots, s_n$ . By definition,  $|\text{sh}(s_{n/2+i})| \geq c_{n/2} = \log n/2 + f(n/2)$ . Let  $N = |\cup_{i=1}^{n/2} \text{sh}_{\log n/2}(s_{n/2+i})|$  be the number of paths  $\text{sh}(s_{n/2+i})$  that are independent at level  $\log n/2$ . Since  $\text{sh}_{\log n/2}(s_{n/2+i})$  chooses from each of the  $n/2$  nodes at level  $\log n/2$  equiprobably,  $N$  has the same distribution as the number of occupied bins when  $n/2$  balls are thrown into  $n/2$  bins uniformly at random. The expected number of unoccupied bins is  $n/2e$ . By Markov's inequality, with probability at least  $1/2$ , the number of empty bins is at most  $2\frac{n}{2e}$ . So, we have that  $\Pr[N \geq \frac{n}{2}(1 - 2/e)] \geq 1/2$ .

Thus, we have that

$$\begin{aligned} E \left[ \frac{\text{Cover}_n^{\text{OPT}}}{\text{Cover}_n^{\text{online}}} \right] &\leq E \left[ \frac{\text{Cover}_n^{\text{OPT}}}{\text{Cover}_n^{\text{online}}} \mid N \geq \frac{n}{2}(1 - 2/e) \right] + \Pr[N \leq \frac{n}{2}(1 - 2/e)] \\ &\leq E \left[ \frac{n(2 + f(n)) + \text{Cover}_n^{\text{online}}}{\text{Cover}_n^{\text{online}}} \mid N \geq \frac{n}{2}(1 - 2/e) \right] + \frac{1}{2} \end{aligned}$$

and by Fact 9,

$$\leq \frac{3}{2} + \frac{2(2 + f(n))}{(1 - 2/e)f(n/2)} = O(1),$$

where the constant depends on  $f(n)$ . Likewise we can obtain a bound on the approximation factor by noting that

$$E [\text{Cover}_n^{\text{online}}] \geq E \left[ \text{Cover}_n^{\text{online}} \mid N \geq \frac{n}{2}(1 - 2/e) \right] \cdot \Pr[N \geq \frac{n}{2}(1 - 2/e)] \geq \frac{1}{2} \frac{n}{2} (1 - 2/e) f(n/2)$$

and therefore,

$$\frac{E [\text{Cover}_n^{\text{OPT}}]}{E [\text{Cover}_n^{\text{online}}]} \leq 1 + \frac{2n + nf(n)}{\frac{1}{2} \frac{n}{2} (1 - 2/e) f(n/2)} = O(1)$$

where once again the constant depends on  $f(n)$ . ■

We defer the case when  $c_i$  is a slow-growing function to Appendix A.2. In particular we assume that  $c_i \leq c_{i/2} + O(1)$  for all  $i$ .

**Theorem 11** *For the binary tree model described above, any threshold rule based on a sequence  $\{c_i\}$ , with  $c_i \leq c_{i/2} + O(1)$  for all  $i$ , achieves an  $O(1)$  competitive ratio in expectation with respect to coverage, as well as an  $O(1)$  approximation to the expected coverage.*

## 5 Sample selection in the skyline model

In this section we focus on the following “skyline” model. We first consider the case where the universe  $X$  is the unit square  $[0, 1]^2$  and  $(x_1, y_1) \prec (x_2, y_2)$  for  $(x_1, y_1), (x_2, y_2) \in X$  if and only if  $x_1 \leq x_2$  and  $y_1 \leq y_2$ . In Section 5.2 we discuss general spaces.

### 5.1 Uniform and product distributions over 2 dimensions

As mentioned earlier, we consider threshold rules of the form  $\mathcal{X}_i = \{(x, y) \in X \mid \mu(\text{sh}'(x, y)) \leq c_i\}$  for some sequence of numbers  $\{c_i\}$ , where  $\text{sh}'(x, y)$  is the set of points that dominate  $(x, y)$ .

For simplicity, we first prove the following version of Theorem 4 for the uniform distribution over  $X$ , and then describe how it extends to general product distributions.

**Theorem 12** *For the skyline setting on the unit square with the uniform distribution, any threshold rule based on a sequence  $\{c_i\}$  with  $c_i = \Omega(1/\text{poly}(i))$  achieves a  $1 + o(1)$  competitive ratio in expectation with respect to the gap.*

Let  $S_n$  denote the set of samples selected by the (implicit) threshold rule out of the set  $T_n$  of samples seen. Let  $R_n = T_n \setminus S_n$  denote the samples rejected by the threshold rule, and  $O_n$  denote an optimal subset of  $T_n$  of size  $n$ . Recall that our goal is to maximize the shadow of the selected subsample, and so all points in  $O_n$  must be undominated by other points. Let  $\mathcal{E}_n$  denote the event that  $O_n$  contains a point in  $R_n$ , that is, there is a point in  $R_n \setminus \text{sh}(S_n)$ . It is immediate that  $\text{Gap}_n^{\text{online}} \neq \text{Gap}_n^{\text{OPT}}$  if and only if the event  $\mathcal{E}_n$  happens. We will show that the event  $\mathcal{E}_n$  occurs with very low probability and use this fact to prove Theorem 12.

We first give a lemma that shows how the approximation factor and expected competitive ratio with respect to the gap of a threshold rule relates to the probability of the event  $\mathcal{E}_n$ . We defer its proof to Appendix A.3

**Lemma 13** *For the skyline model with an arbitrary distribution the gap of a threshold rule based on the sequence  $\{c_i\}$  satisfies the following, where  $\mathcal{E}_n$  is the event that  $\text{Gap}_n^{\text{online}} \neq \text{Gap}_n^{\text{OPT}}$ .*

$$E_\mu \left[ \frac{\text{Gap}_n^{\text{online}}}{\text{Gap}_n^{\text{OPT}}} \right] \leq 1 + \frac{1}{c_n} \Pr[\mathcal{E}_n]$$

To complete the proof of Theorem 12 we give an upper bound on the probability of the event  $\mathcal{E}_n$ . Our goal is to show that with high probability, every sample in  $R_n$  is dominated by some sample in  $S_n$ . We start with a simple observation about the number of rejected samples.

**Fact 14**  $E[|R_n|] \leq E[T_n] = \sum_{i=1}^n 1/\mu(\mathcal{X}_i) \leq n/c_n$

**Fact 15** *Let  $\mu$  be the uniform measure over  $[0, 1]^2$ . Then for all  $n$ ,  $\mu(\mathcal{X}_n) = c_n(1 + \ln 1/c_n)$ .*

We first note that many of the samples in  $S_n$  are in fact in  $\mathcal{X}_n$ . We defer the proof of the following lemma to Appendix A.3.

**Lemma 16** *Let  $\alpha$  be a constant satisfying  $\frac{c_i}{c_{i/2}} \geq \alpha$  for large enough  $i$ . Then with probability  $1 - o(c_n)$ ,  $\alpha n/4$  of the samples in  $S_n$  belong to  $\mathcal{X}_n$ .*

The following lemma shows that given sufficient number of samples in  $S_n$  belonging to  $\mathcal{X}_n$ , with high probability  $R_n$  is dominated by these samples. Let  $\mathcal{E}'_n$  denote the event that at least one point in  $R_n$  is not dominated by  $S_n \cap \mathcal{X}_n$  and let  $z = |S_n \cap \mathcal{X}_n|$ .

**Lemma 17** *Conditioned on  $z$ ,*

$$\Pr[\mathcal{E}'_n] \leq \exp(-zc_n/\mu(\mathcal{X}_n)) \cdot E[|R_n|]$$

We defer the proof to Appendix A.3.

Finally we are ready to prove Theorem 12.

*Proof of Theorem 12:* Using Lemma 13, we have that

$$E_\mu[\text{Gap}_n^{\text{online}}/\text{Gap}_n^{\text{OPT}}] \leq 1 + \frac{1}{c_n} (\Pr[\mathcal{E}_n | z \geq \alpha(n/4)] + \Pr[z < \alpha(n/4)]).$$

where  $z = |S_n \cap \mathcal{X}_n|$ . By Lemma 16, the second term in the parentheses is  $o(c_n)$ . Event  $\mathcal{E}_n$  implies event  $\mathcal{E}'_n$ . So applying Lemma 17 and Fact 14 we get

$$\Pr[\mathcal{E}_n | z \geq \alpha(n/4)] \leq \exp(-zc_n/\mu(\mathcal{X}_n)) \cdot E[|R_n|] \leq \exp(-\alpha(n/4)) \cdot \frac{n}{c_n} = o(c_n),$$

since  $c_i = \Omega(1/\text{poly}(i))$ . ■

**General product distributions.** We now consider the skyline model with  $\mu$  being an arbitrary product distribution. In particular, for a point  $(a, b) \in X$ , let  $\mu(a, b) = \mu^x(a)\mu^y(b)$  for one-dimensional measures  $\mu^x$  and  $\mu^y$ . Our proof of Theorem 4 is nearly identical to our argument for the uniform case. We note first that as before

$$E_\mu \left[ \frac{\text{Gap}_n^{\text{online}}}{\text{Gap}_n^{\text{OPT}}} \right] \leq 1 + \frac{1}{c_n} \Pr[\mathcal{E}_n]$$

To bound the probability of  $\mathcal{E}_n$ , we give a reduction from the product measure setting to the uniform measure setting. In particular, consider mapping  $X$  into  $X' = [0, 1] \times [0, 1]$  by mapping a point  $(a, b) \in X$  to  $(\mu^x(a), \mu^y(b)) \in X'$ . Then it is easy to see that  $\mathcal{X}_i$  in  $X$  gets mapped to  $\mathcal{X}_i$  in  $X'$  for the same sequence  $\{c_i\}$ . Then, the probability of the event  $\mathcal{E}_n$  under the transformation remains the same as before, and is once again  $o(c_n)$ .

## 5.2 General spaces

In this section, we show how our results from the skyline model generalize and prove Theorem 3.

Once again we note that Lemmas 13 and 17 carry over to this setting. So our main approach is to bound the probability of the event  $\mathcal{E}'_n$ . The main difference from the previous analysis is that the best bound on the size of  $\mathcal{X}_i$  we can obtain is  $c_i \leq \mu(\mathcal{X}_i) \leq 1$ . This means that we can no longer claim that a constant fraction of the samples in  $S_n$  belong to  $\mathcal{X}_n$ . Instead we will show that under a stronger condition on the sequence  $\{c_i\}$ , namely  $c_i = i^{-(1/2-\Omega(1))}$ , the number of samples in  $\mathcal{X}_n$  is  $\Omega(nc_n)$  with a high probability. This will suffice to give us the bound we need. In particular, we will show that

$$\Pr[\mathcal{E}_n \mid |S_n \cap \mathcal{X}_n| \geq nc_n/2] = o(c_n) \tag{1}$$

$$\Pr[|S_n \cap \mathcal{X}_n| < nc_n/2] = o(c_n) \tag{2}$$

The following Lemma is a weaker version of Lemma 16 that implies Equation 2. We defer its proof to Appendix A.3

**Lemma 18** For  $c_i = \Omega(1/i^\epsilon)$  with  $\epsilon < 1$ , at least  $nc_n/2$  of the samples in  $S_n$  belong to  $\mathcal{X}_n$  with probability  $1 - o(c_n)$ .

Finally, using Fact 14 and Lemma 17, the following shows that Equation 1 holds for  $c_i = i^{-(1/2-\Omega(1))}$ .

$$\begin{aligned} \Pr[\mathcal{E}_n \mid |S_n \cap \mathcal{X}_n| \geq nc_n/2] &\leq \Pr[\mathcal{E}'_n \mid |S_n \cap \mathcal{X}_n| \geq nc_n/2] \\ &\leq \exp(-nc_n^2/2\mu(\mathcal{X}_n)) \cdot E[|R_n|] \leq \exp(-n \cdot n^{-2(1/2-\Omega(1))}) \cdot n/c_n = o(c_n) \end{aligned}$$

This proves Theorem 3.

## References

- [1] M. Abramowitz and I.A. Stegun. *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*. Courier Dover Publications, 1965.
- [2] Mikhail J. Atallah and Yinian Qi. Computing all skyline probabilities for uncertain data. In *PODS '09: Proceedings of the twenty-eighth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 279–287, New York, NY, USA, 2009. ACM.
- [3] Yossi Azar, Edith Cohen, Amos Fiat, Haim Kaplan, and Harald Räcke. Optimal oblivious routing in polynomial time. *J. Comput. Syst. Sci.*, 69(3):383–394, 2004.
- [4] J. L. Bentley, H. T. Kung, M. Schkolnick, and C. D. Thompson. On the average number of maxima in a set of vectors and applications. *J. ACM*, 25(4):536–543, 1978.
- [5] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.
- [6] Stephan Borzsonyi, Konrad Stocker, and Donald Kossmann. The skyline operator. *Data Engineering, International Conference on*, 0:0421, 2001.
- [7] Andrei Z. Broder, Adam Kirsch, Ravi Kumar, Michael Mitzenmacher, Eli Upfal, and Sergei Vassilvitskii. The hiring problem and lake wobegon strategies. In *SODA '08: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1184–1193, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- [8] T. Ferguson. Who solved the secretary problem. *Statist. Sci.*, 4(3):282–289, 1989.
- [9] P. R. Freeman. The secretary problem and its extensions: a review. *International Statistical Review*, 51(2):189–206, 1983.
- [10] Anupam Gupta, Mohammad T. Hajiaghayi, and Harald Räcke. Oblivious network design. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 970–979, 2006.
- [11] Lujun Jia, Guolong Lin, Guevara Noubir, Rajmohan Rajaraman, and Ravi Sundaram. Universal approximations for tsp, steiner tree, and set cover. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 386–395, 2005.

- [12] H. T. Kung, F. Luccio, and F. P. Preparata. On finding the maxima of a set of vectors. *J. ACM*, 22(4):469–476, 1975.
- [13] S. Samuels. Secretary problems. In *Handbook of Sequential Analysis*, pages 381–405. Marcel Dekker, Inc., 1991.
- [14] S. S. Wilks. *Mathematical Statistics*. Wiley, New York, NY, 1962.
- [15] Wenjie Zhang, Xuemin Lin, Ying Zhang, Wei Wang, and Jeffrey Xu Yu. Probabilistic skyline operator over sliding windows. In *ICDE '09: Proceedings of the 2009 IEEE International Conference on Data Engineering*, pages 1060–1071, Washington, DC, USA, 2009. IEEE Computer Society.

## A Missing proofs

We present the technical proofs we skipped over in the main article.

### A.1 Sample selection in one dimension

In section 3, we had some technical derivations that we prove in detail here.

*Proof of Fact 5:* Let  $G(x) = x^k$  be the distribution of the gap of any sample. The expectation of the  $m$ th order statistic for  $t_n$  samples from this distribution is as follows. See Claim 1 below for a proof.

$$\mathcal{E}[x_{(m)}] = \frac{\Gamma(t_n + 1)\Gamma(m + 1/k)}{\Gamma(t_n + 1 + 1/k)\Gamma(m)}.$$

So, we have that the expected mean gap of the  $n$  smallest-gap subsamples out of  $t_n$  samples is

$$\begin{aligned} \mathcal{E}\left[\frac{\sum_{i=1}^n x_{(i)}}{n}\right] &= (1/n) \sum_{i=1}^n \mathcal{E}[x_{(i)}] \\ &= (1/n) \frac{\Gamma(t_n + 1)}{\Gamma(t_n + 1 + 1/k)} \sum_{m=1}^n \frac{\Gamma(m + 1/k)}{\Gamma(m)} \end{aligned}$$

we postpone the proof of the following step to Claim 2 below,

$$\begin{aligned} &= (1/n) \frac{\Gamma(t_n + 1)}{\Gamma(t_n + 1 + 1/k)} \frac{n\Gamma(n + 1 + 1/k)}{(1 + 1/k)\Gamma(n + 1)} \\ &\sim \frac{1}{1 + 1/k} \left(\frac{n}{t_n + 1}\right)^{1/k}, \end{aligned}$$

■ where the last line follows by equation 6.1.46 of [1].

We now prove the two claims missing in the above proof.

**Claim 1** *The expectation of the  $m$ th order statistic with  $F(x) = x^k$  is*

$$\mathcal{E}[x_{(m)}] = \frac{\Gamma(t_n + 1)\Gamma(m + 1/k)}{\Gamma(t_n + 1 + 1/k)\Gamma(m)}.$$

*Proof:* From page 236 of [14], we have that the  $m$ th order statistic of a sample of size  $t_n$  from a population having continuous distribution function  $F(x)$  and probability distribution function  $f(x)$  has the probability distribution function:

$$\frac{\Gamma(t_n + 1)}{\Gamma(m)\Gamma(t_n - m + 1)} [F(x_{(m)})]^{m-1} [1 - F(x_{(m)})]^{t_n - m} f(x_{(m)}) dx_{(m)}$$

So, we have that the expectation of the  $m$ th order statistic with  $F(x) = x^k$  is

$$\begin{aligned} \mathcal{E}[x_{(m)}] &= \int_0^1 x_{(m)} \cdot \frac{\Gamma(t_n + 1)}{\Gamma(m)\Gamma(t_n - m + 1)} \cdot [F(x_{(m)})]^{m-1} \cdot [1 - F(x_{(m)})]^{t_n - m} \cdot f(x_{(m)}) dx_{(m)} \\ &= \frac{\Gamma(t_n + 1)}{\Gamma(m)\Gamma(t_n - m + 1)} \int_0^1 x_{(m)} [x_{(m)}^k]^{m-1} [1 - x_{(m)}^k]^{t_n - m} k x_{(m)}^{k-1} dx_{(m)} \\ &= \frac{\Gamma(t_n + 1)}{\Gamma(m)\Gamma(t_n - m + 1)} \int_0^1 x_{(m)}^{k(m-1)+1} [1 - x_{(m)}^k]^{t_n - m} k x_{(m)}^{k-1} dx_{(m)} \end{aligned}$$

using the substitution  $y = x^k$ , we have

$$\begin{aligned} &= \frac{\Gamma(t_n + 1)}{\Gamma(m)\Gamma(t_n - m + 1)} \int_0^1 y^{m-1+1/k} [1 - y]^{t_n - m} dy \\ &= \frac{\Gamma(t_n + 1)}{\Gamma(m)\Gamma(t_n - m + 1)} \cdot B(m + 1/k, t_n - m + 1) \end{aligned}$$

where  $B()$  is the Beta function

$$\begin{aligned} &= \frac{\Gamma(t_n + 1)}{\Gamma(m)\Gamma(t_n - m + 1)} \cdot \frac{\Gamma(m + 1/k)\Gamma(t_n - m + 1)}{\Gamma(t_n + 1 + 1/k)} \\ &= \frac{\Gamma(t_n + 1)\Gamma(m + 1/k)}{\Gamma(t_n + 1 + 1/k)\Gamma(m)} \end{aligned}$$

■

### Claim 2

$$\sum_{m=1}^n \frac{\Gamma(m + 1/k)}{\Gamma(m)} = \frac{n\Gamma(n + 1 + 1/k)}{(1/k + 1)\Gamma(n + 1)}.$$

*Proof:* To simplify a sum involving gamma functions, we can use the idea that

$$\Gamma(s) = \int_0^\infty t^{s-1} e^{-t} dt$$

Now, we apply the Laplace transform. Using  $s = m + \alpha$ , and then interchanging summation and integration,

we get

$$\begin{aligned}
\sum_{m=1}^n \Gamma(m + \alpha)/\Gamma(m) &= \int_0^\infty t^\alpha e^{-t} \sum_{m=0}^{n-1} \frac{t^m}{m!} dt \\
&= \int_0^\infty t^\alpha \sum_{m=0}^{n-1} e^{-t} \frac{t^m}{m!} dt \\
&= \int_0^\infty t^\alpha \Pr[\text{Poisson}(t) < n] dt \\
&= \int_0^\infty t^\alpha \Pr[G(n) \geq t] dt,
\end{aligned}$$

where  $G(n)$  is a Gamma( $n$ ) random variable. The last identity follows by the Poisson process. Plugging in

$$\Pr[G(n) \geq t] = \int_t^\infty u^{n-1} e^{-u} du$$

and then interchanging the order of integration gives an integral that evaluates to

$$\frac{n\Gamma(n + 1 + 1/k)}{(1/k + 1)\Gamma(n + 1)}.$$

■

## A.2 Sample selection in binary trees

*Proof of Theorem 11:* We follow the outline of the previous proof. Consider the last  $n/2$  selected nodes  $s_{n/2+1}, \dots, s_n$ . If  $c_{n/2} \geq \log(n/2)$ , we consider the number of independent paths at level  $\log(n/2)$  and the proof goes through exactly as before. So suppose that  $c_{n/2} < \log(n/2)$ . Let  $N = |\cup_{i=1}^{n/2} \text{sh}_{c_{n/2}}(s_{n/2+i})|$ . Then there are at most  $2^{c_{n/2}} < n/2$  nodes at level  $c_{n/2}$ , so we can use the same balls-and-bins argument as in the previous proof to obtain  $\Pr[N \geq 2^{c_{n/2}}(1 - 2/e)] \geq 1/2$ .

Since there are at most  $2^{c_n+1}$  nodes in the first  $c_n$  levels of a binary tree, and  $\text{Cover}_n^{\text{online}} \geq N$ , we have that

$$\begin{aligned}
E \left[ \frac{\text{Cover}_n^{\text{OPT}}}{\text{Cover}_n^{\text{online}}} \right] &\leq E \left[ \frac{2^{c_n+1} + \text{Cover}_n^{\text{online}}}{\text{Cover}_n^{\text{online}}} \mid N \geq 2^{c_{n/2}}(1 - 2/e) \right] + \frac{1}{2} \\
&\leq \frac{3}{2} + \frac{2^{c_n+1}}{2^{c_{n/2}}(1 - 2/e)} = O(1)
\end{aligned}$$

We can also use the same argument as in the previous proof to prove the claimed bound on the approximation factor.

$$\frac{E [\text{Cover}_n^{\text{OPT}}]}{E [\text{Cover}_n^{\text{online}}]} \leq 1 + \frac{2^{c_n+1}}{\frac{1}{2} \cdot 2^{c_{n/2}} \cdot (1 - 2/e)} = O(1)$$

■

### A.3 Sample selection in the skyline model

*Proof of Lemma 13:* We apply Bayes' rule to get

$$E_\mu \left[ \frac{\text{Gap}_n^{\text{online}}}{\text{Gap}_n^{\text{OPT}}} \right] \leq 1 + E_\mu \left[ \frac{\text{Gap}_n^{\text{online}}}{\text{Gap}_n^{\text{OPT}}} \mid \mathcal{E}_n \right] \Pr[\mathcal{E}_n]$$

If the event  $\mathcal{E}_n$  happens then by definition  $O_n$  contains a point in  $R_n$ , say  $x$ . Then,  $\text{Gap}_n^{\text{OPT}} = \text{Gap}(O_n) \geq \mu(\text{sh}'(x)) > c_n$ , where the last inequality follows from noting that  $x \notin \mathcal{X}_n$ . On the other hand,  $\text{Gap}_n^{\text{online}}$  is always less than 1. Therefore the claim follows. ■

*Proof of Lemma 16:* Consider samples belonging to  $S_n \cap \mathcal{X}_{n/2}$ ; these are at least  $n/2$  in number. We claim that a constant fraction of these are in  $\mathcal{X}_n$  with high probability. In particular,

$$\Pr_{x \sim \mu} [x \in \mathcal{X}_n \mid x \in \mathcal{X}_{n/2+i}] = \frac{\mu(\mathcal{X}_n)}{\mu(\mathcal{X}_{n/2+i})} \geq \frac{\mu(\mathcal{X}_n)}{\mu(\mathcal{X}_{n/2})} \geq \alpha,$$

Here we used the fact that  $\frac{1+\ln 1/c_n}{1+\ln 1/c_{n/2}} \geq 1$ . Therefore, the expected number of samples in  $S_n \cap \mathcal{X}_n$  is at least  $\alpha n/2$ . Since this is a binomial random variable, by Chernoff bounds,

$$\Pr \left[ |S_n \cap \mathcal{X}_n| < \frac{1}{2} \alpha (n/2) \right] < \exp \left\{ -\alpha (n/2) \left( \frac{1}{2} \right)^2 / 2 \right\} = o(c_n),$$

since  $c_i = \Omega(1/\text{poly}(i))$ . ■

For the subsequent proof, we first state the following consequence of measure continuity.

**Fact 19** *Let  $\mu$  satisfy measure continuity. Then for all  $k \in \mathbb{N}$ , and for all  $y \notin \mathcal{X}_k$ , we have  $\mu(\text{sh}'(y) \cap \mathcal{X}_k) \geq c_k$ .*

*Proof:* By measure continuity, there exists a  $z \in \text{sh}'(y)$  such that  $\mu(\text{sh}'(z)) = c_k \leq \mu(\text{sh}'(y))$ . Thus, we have that  $z \in \mathcal{X}_k$ ,  $\text{sh}'(z) \subseteq \text{sh}'(y) \cap \mathcal{X}_k$ , and so  $\mu(\text{sh}'(y) \cap \mathcal{X}_k) \geq \mu(\text{sh}'(z)) = c_k$ . ■

*Proof of Lemma 17:* For any sample  $y$  in  $R_n$ , the probability that it is dominated by a uniformly random sample in  $\mathcal{X}_n$  is

$$\Pr_{x \sim \mu} [y \in \text{sh}(x) \mid x \in \mathcal{X}_n] = \mu(\text{sh}'(y) \cap \mathcal{X}_n) / \mu(\mathcal{X}_n) \geq c_n / \mu(\mathcal{X}_n).$$

Then the probability that  $y$  is not dominated by any point in  $S_n \cap \mathcal{X}_n$  is  $(1 - c_n / \mu(\mathcal{X}_n))^z$ . Since this bound holds regardless of the specific value of  $y$ , applying Wald's identity,

$$\begin{aligned} \Pr[\mathcal{E}'_n] &\leq E[\text{number of samples in } R_n \text{ not dominated by } S_n \cap \mathcal{X}_n] \\ &= (1 - c_n / \mu(\mathcal{X}_n))^z \cdot E[|R_n|] \leq \exp(-z c_n / \mu(\mathcal{X}_n)) \cdot E[|R_n|] \end{aligned}$$

■

*Proof of Lemma 18:* For all  $i \leq n$ , we have that

$$\Pr_{x \sim \mu} [x \in \mathcal{X}_n \mid x \in \mathcal{X}_i] \geq \mu(\mathcal{X}_n) \geq c_n.$$

Thus, by Chernoff bounds,

$$\Pr \left[ |S_n \cap \mathcal{X}_n| < \frac{nc_n}{2} \right] \leq \exp(-\Omega(n^{1-\epsilon})) = o(c_n).$$

■