Columbia B9136 February 4, 2025

Lecture 2 — Unknown Distributions and/or Random Order

Lecturer: Will Ma *Scribes: Tao Luo and Alon Rieger*

1 Introduction

Preliminary Formalism. We discuss additional settings for the problem of stopping on a single number. In order to do so, we begin by laying out some formal definitions.

- assume T fixed and known.
- $\mathbf{V} := (V_t)_{t\in[T]}$: (realized) valuations of agents.
- π: (deterministic) policy that maps history to accept/reject decision.
- $\pi(\mathbf{V})$ reward collected by π on sequence $\mathbf{V}$. For example, if

  $$\pi \quad = \quad \text{Threshold}(5)$$

  is the threshold policy, with an associated threshold of 5, and

  $$\mathbf{V} \quad = \quad (3, 6, 4, 9)\ ,$$

  then $\pi(\mathbf{V}) = 6$.
- Π as space of all detereminstic policies.
- $\Delta(\cdot) :=$ distribution over $\cdot$.
- $\mathbf{F} := (F_t)_{t\in[T]} \in \Delta(\mathbb{R}_{\geq 0})^T$ is the sequence of independent distributions of nonnegative variables (we use $\Delta(\mathbb{R}^T_{\geq 0})$ when the distributions are not independent).
- $P \in \Delta(\pi)$: randomized policy.
- $\pi(\mathbf{F}) \quad = \quad \mathbb{E}_{\mathbf{V}\sim\mathbf{F}}[\pi(\mathbf{V})]$ is the expected value of a deterministic policy (where the expectation is taken with respect to the distribution $\mathbf{F}$).
- $P(\mathbf{V}) := \mathbb{E}_{\pi\sim P}[\pi(\mathbf{F})]$ is the expected value of a randomized policy (where $\mathbf{V}$ remains deterministic and the expectation is taken with respect to the distribution of P).
- $P(\mathbf{F}) := \mathbb{E}_{\pi\sim P, \mathbf{V}\sim\mathbf{F}}[\pi(\mathbf{V})]$ is the expected value of a random policy with random valuations.
- $\text{OFF}(\mathbf{V})$: optimal reward when the valuations are $\mathbf{V}$, that is

  $$\text{OFF}(\mathbf{V}) = \max_{t\in[T]} V_t\ .$$
- $\text{OFF}(\mathbf{F})$: optimal reward when valuation distributions are $\mathbf{F}$, that is

  $$\text{OFF}(\mathbf{F}) \quad = \quad \mathbb{E}_{\mathbf{V}\sim\mathbf{F}}\text{OFF}(\mathbf{V}) \quad = \quad \mathbb{E}_{\mathbf{V}\sim\mathbf{F}}[\max_{t\in[T]} V_t]\ .$$

The "Prophet" Setting Last class we showed that

$$\inf_{T\in\mathbb{N}} \inf_{\mathbf{F}\in\Delta(\mathbb{R}_{\geq 0})^T} \sup_{\pi\in\Pi} \frac{\pi(\mathbf{F})}{\mathrm{OFF}(\mathbf{F})} \quad = \quad \frac{1}{2} \ . \tag{1}$$

Indeed, we showed for any T and given independent distributions $\mathbf{F} \in \Delta(\mathbb{R}_{\geq 0})^T$ that there is a policy π for which $\pi(\mathbf{F}) \geq \mathrm{OFF}(\mathbf{F})/2$. On the other hand, we showed for $T = 2$ a sequence of distributions $\mathbf{F} \in \Delta(\mathbb{R}_{\geq 0})^2$ (parametrized by ε) for which $\pi(\mathbf{F})/\mathrm{OFF}(\mathbf{F}) \to 1/2$ as $\varepsilon \to 0$.

In this class we consider other settings for the problem of stopping on a single number, which can be obtained by:

1. Interchanging the order of the $\inf_{\mathbf{F}\in\Delta(\mathbb{R}_{\geq 0})^T}$ and $\sup_{\pi\in\Pi}$ in (1), which affects whether distributional knowledge about the valuations can be used to determine π;
2. Allowing for random arrival orders, to be discussed later.

2 First Setting: Unknown Distributions

In this setting, the adversary chooses distributions, but they are unknown to the policy. Therefore the order of the inf and the sup arguments changes. Moreover, the policy may need to be random, not deterministic:

$$\inf_{T\in\mathbb{N}} \sup_{P\in\Delta(\Pi)} \inf_{\mathbf{F}\in\Delta(\mathbb{R}_{\geq 0})^T} \frac{P(\mathbf{F})}{\mathrm{OFF}(\mathbf{F})} \ .$$

We claim that

$$\underbrace{\inf_{T\in\mathbb{N}} \sup_{P\in\Delta(\Pi)} \inf_{\mathbf{F}\in\Delta(\mathbb{R}_{\geq 0})^T} \frac{P(\mathbf{F})}{\mathrm{OFF}(\mathbf{F})}}_{\spadesuit} \quad \equiv \quad \underbrace{\inf_{T\in\mathbb{N}} \sup_{P\in\Delta(\Pi)} \inf_{\mathbf{V}\in\mathbb{R}^T_{\geq 0}} \frac{P(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})}}_{\blacklozenge} \ ,$$

meaning that the same bounds can be achieved when the valuations are deterministic and unknown to the policy. To see why this is true, we use a sandwich argument. First, we claim that

$$\inf_{\mathbf{V}\in\mathbb{R}^T_{\geq 0}} \frac{P(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})} \quad \geq \quad \inf_{\mathbf{F}\in\Delta(\mathbb{R}_{\geq 0})^T} \frac{P(\mathbf{F})}{\mathrm{OFF}(\mathbf{F})} \ ,$$

since $\Delta(\mathbb{R}_{\geq 0})^T$ is a larger feasible space than $\mathbb{R}^T_{\geq 0}$. To prove the other inequality, suppose that

$$\inf_{\mathbf{V}\in\mathbb{R}^T_{\geq 0}} \frac{P(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})} \quad = \quad \rho \ .$$

Then

$$\begin{aligned}
P(\mathbf{F}) \quad &= \quad \mathbb{E}_{\mathbf{V}\sim\mathbf{F}} P(\mathbf{V}) \\
&= \quad \int_{\mathbf{F}} P(\mathbf{V}) \cdot \mathrm{f}_{\mathrm{F}}(\mathbf{V}) d\mathbf{V} \\
&\geq \quad \int_{\mathbf{F}} \rho \cdot \mathrm{OFF}(\mathbf{V}) \cdot \mathrm{f}_{\mathrm{F}}(\mathbf{V}) d\mathbf{V} \\
&= \quad \rho \cdot \int_{\mathbf{F}} \mathrm{OFF}(\mathbf{V}) \cdot \mathrm{f}_{\mathrm{F}}(\mathbf{V}) d\mathbf{V} \\
&= \quad \rho \cdot \mathbb{E}_{\mathbf{V}\sim\mathbf{F}} \mathrm{OFF}(\mathbf{V}) \quad = \quad \rho \cdot \mathrm{OFF}(\mathbf{F})
\end{aligned}$$

Therefore for any realization of $\mathbf{F} \in \Delta(\mathbb{R}_{\geq 0})^T$ we have

$$\frac{P(\mathbf{F})}{\mathrm{OFF}(\mathbf{F})} \quad \geq \quad \rho \quad = \quad \inf_{\mathbf{V} \in \mathbb{R}^T_{\geq 0}} \frac{P(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})} \ ,$$

and in particular

$$\inf_{\mathbf{F} \in \Delta(\mathbb{R}_{\geq 0})^T} \frac{P(\mathbf{F})}{\mathrm{OFF}(\mathbf{F})} \quad \geq \quad \inf_{\mathbf{V} \in \mathbb{R}^T_{\geq 0}} \frac{P(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})} \ .$$

The result of this equivalence is that if we want to construct a policy that works well on ♠, we can show that it works well on ⧫.

2.1 Upper Bounding the Ratio

Since the analysis of all randomized policies is hard to perform, it is hard to upper bound the value of ⧫. We therefore use Yao's minimax Lemma (Yao, 1977), which provides further equivalent ratios, that are easier to upper bound. Yao's minimax Lemma enables to change the order of the inf and the sup (which player goes first), and resort to a deterministic policy, and a (correlated) random valuation distribution. The Lemma is stated below:

Lemma 1 (Yao's Minimax Lemma)**.**

$$\underbrace{\sup_{P \in \Delta(\Pi)} \inf_{\mathbf{V} \in \mathbb{R}^T} \mathbb{E}_{\pi \sim P}\left[\frac{\pi(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})}\right]}_{\blacklozenge} \quad = \quad \inf_{\mathbf{F} \in \Delta(\mathbb{R}^T)} \sup_{\pi \in \Pi} \mathbb{E}_{\mathbf{V} \sim \mathbf{F}}\left[\frac{\pi(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})}\right] \quad = \quad \underbrace{\inf_{\mathbf{F} \in \Delta(\mathbb{R}^T)} \frac{\sup_{\pi \in \Pi} \mathbb{E}_{\mathbf{V} \sim \mathbf{F}}\left[\pi(\mathbf{V})\right]}{\mathbb{E}_{\mathbf{V} \sim \mathbf{F}}\left[\mathrm{OFF}(\mathbf{V})\right]}}_{\diamondsuit}$$

A very important thing to note about ♢ is that the underlying distribution is correlated and not independent. Notice that if the setting was independent, then the ratio will be equivalent to the ratio in the prophet problem, and the upper bound would be $\frac{1}{2}$. We are now ready to present a correlated setting that achieves the upper bound.

"Weight-Dependent" Setting : Let $V_t \in \mathcal{R} = \{r_1, \cdots, r_n, 0\}$ where $r_1 > \cdots > r_n > 0$. We use the following construction of $\mathbf{F} \in \Delta(\mathcal{R}^T)$ to upper-bound $\frac{\sup_{\pi \in \Pi} \mathbb{E}_{\mathbf{V} \sim \mathbf{F}}[\pi(\mathbf{V})]}{\mathbb{E}_{\mathbf{V} \sim \mathbf{F}}[\mathrm{OFF}(\mathbf{V})]}$:

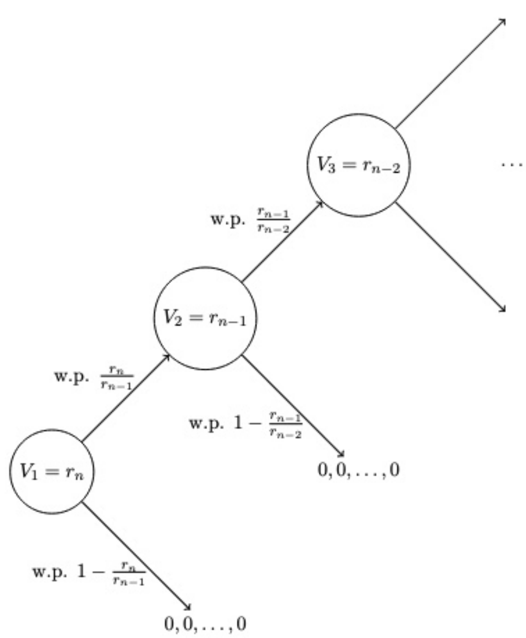

Explained in words, at the first node the valuation of agent 1 is r_n, then with probability $\frac{r_n}{r_{n-1}}$ the valuation of agent 2 is r_{n-1}, and with probability $1 - \frac{r_n}{r_{n-1}}$ all the agents after agent 1 have a valuation of zero. Then, if agent 2 has a valuation of r_{n-1}, then with probability $\frac{r_{n-1}}{r_{n-2}}$ the valuation of agent 3 is r_{n-2}, and with probability $1 - \frac{r_{n-1}}{r_{n-2}}$ all the agents after agent 2 have a valuation of zero, and so on and so forth, up to the last agent. Notice that this setting is indeed correlated, since if $V_j = 0$, then all the subsequent valuations are also zero.

This setting is designed to ensure that the expected value of each agent's valuation is equal to the expected value of the valuation of the agent that precedes him, that is

$$\mathbb{E}[V_j] \quad = \quad \mathbb{E}[V_{j-1}], \quad \forall j \in [1, n] \ .$$

Now since $V_1 = r_n$ with probability 1, the result is that

$$\mathbb{E}[V_j] \quad = \quad \mathbb{E}[V_1] \quad = \quad r_n, \quad \forall j \in [n] \ .$$

This argument establishes the following claim

Claim 2. Let $\mathbf{F}$ be the distribution defined by the above diagram. Then

$$\sup_{\pi \in \Pi} \pi(\mathbf{F}) \quad = \quad r_n \ .$$

We also claim that an offline algorithm will achieve an objective value of

$$\mathbb{E}_{\mathbf{V} \sim \mathbf{F}}[\text{OFF}(\mathbf{V})] \quad = \quad \sum_{j=1}^{n} \underbrace{(r_j - r_{j+1})}_{\text{additive value from reaching } r_n} \underbrace{\frac{r_n}{r_j}}_{\text{probability of reaching } r_n} \quad = \quad \underbrace{r_n}_{\sup_{\pi \in \Pi} \pi(\mathbf{F})} \cdot \sum_{j=1}^{n} \left(1 - \frac{r_{j+1}}{r_j}\right) ,$$

where r_{n+1} is set to zero by assumption.

Consequently the competitive ratio is

$$\mathbf{CR} \quad = \quad \frac{1}{\sum_{j=1}^{n} \left(1 - \frac{r_{j+1}}{r_j}\right)}$$

2.2 A Policy that Achieves the Upper Bound

For the best possible lower bound, our analysis goes back to the ♦ ratio. Since all these ratios are equivalent, the result is that the policy achieves the upper bound. We construct a randomized threshold policy. The idea is to assign weights W_j for each reward r_j, and pick $\tilde{\tau} = r_j$ as the threshold, with probability that is proportional to W_j, that is $\frac{W_j}{\sum_{j'=1}^{n} W_{j'}}$. We now move on to the analysis, in order to understand how to choose the weights properly. Let $\tilde{\tau}$ be the randomized threshold and suppose that

$$\text{OFF}(\mathbf{V}) = r_{j^*} \quad \text{for some } j^* \in [n] \ .$$

Lower bounds for realized thresholds:

- If $\tilde{\tau} > r_{j^*}$, then $\pi(\mathbf{V}) = 0$.
- If $\tilde{\tau} = r_{j^*}$, then $\pi(\mathbf{V}) = r_{j^*}$.
- If $\tilde{\tau} < r_{j^*}$, then $\pi(\mathbf{V}) \geq \tilde{\tau}$.

Consequently we have

$$P(\mathbf{V}) \quad \geq \quad \sum_{j=j^*}^{n} r_j \cdot \frac{W_j}{\sum_{j'=1}^{n} W_{j'}} \ .$$

Here, the inequality holds since the RHS is the weighted probability of each threshold, and if the threshold is r_j (and $r_j \leq r_{j^*}$), then the reward is at least r_j. Consequently

$$\frac{P(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})} \quad \geq \quad \frac{1}{r_{j^*}} \sum_{j=j^*}^{n} r_j \frac{W_j}{\sum_{j'=1}^{n} W_{j'}} \ .$$

We now wish to find the appropriate weights that maximize the above ratio. The intuition is to have an identical ratio for all $j^* \in [n]$. First, we set $W_n = 1$ and $\Delta = \sum_{j=1}^{n} W_j$. We look at the different realizations of j^*:

- *Case 1:* $j^* = n$: We have
$$\frac{P(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})} = \frac{1}{\Delta}$$
- *Case 2:* $j^* = n-1$: We want to achieve
$$\begin{aligned} \frac{P(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})} \quad &= \quad \frac{1}{r_{n-1}} \sum_{j=n-1}^{n} r_j \frac{W_j}{\sum_{j'=1}^{n} W_{j'}} \\ &\underset{\sum_{j'=1}^{n} W_{j'} = \Delta}{=} \quad \frac{1}{r_{n-1}} \cdot \frac{r_n \cdot W_n + r_{n-1} \cdot W_{n-1}}{\Delta} \\ &\underset{W_n = 1}{=} \quad \frac{r_n + r_{n-1} \cdot W_{n-1}}{r_{n-1} \cdot \Delta} \quad = \quad \frac{1}{\Delta} \end{aligned}$$
In order to achieve this, we must have
$$\frac{r_n + r_{n-1} \cdot W_{n-1}}{r_{n-1}} \quad = \quad 1 \ ,$$
which means that $W_{n-1} = 1 - \frac{r_n}{r_{n-1}}$.
- *Case 3:* $j^* = n-2$: Following the same line of argument, we have
$$\frac{P(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})} \quad = \quad \frac{r_n + r_{n-1} \cdot W_{n-1} + +r_{n-2} \cdot W_{n-2}}{r_{n-2} \cdot \Delta} \quad = \quad \frac{1}{\Delta} \ ,$$
which results in $W_{n-2} = 1 - \frac{r_{n-1}}{r_{n-2}}$.
- *General Case*: The argument remains the same. Overall we get that $W_j = 1 - \frac{r_{j+1}}{r_j}$.

Overall we established that

$$\inf_{T \in \mathbb{N}} \sup_{P \in \Delta(\Pi)} \inf_{\mathbf{V} \in \mathbb{R}^T} \frac{P(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})} \quad = \quad f(\mathcal{R}) \quad = \quad \frac{1}{\sum_{j=1}^{n} \left(1 - \frac{r_{j+1}}{r_j}\right)} \ ,$$

if $\mathcal{R} = \{r_1, r_2, \ldots, r_n, r_{n+1}\}$ and $r_1 > r_2 > \cdots > r_n > r_{n+1} = 0$.

Example: for $\mathcal{R} = \{2, 1, 0\}$ we get

$$f(\mathcal{R}) = \frac{1}{1 + (1 - \frac{1}{2})} = \frac{2}{3} \ .$$

This example shows that a randomized algorithm is necessary.

Now Suppose:

$$\mathcal{R} = \{M^k, M^{k-1}, \ldots, M, 0\}, \quad M > 0 \ .$$

Then:

$$f(\mathcal{R}) = \frac{1}{1 + \left(1 - \frac{1}{M}\right)(k-1)} \ .$$

For large M and k we get

$$M \to \infty, \quad k \to \infty \quad \Rightarrow \quad f(\mathcal{R}) \to 0 \ .$$

3 Second Setting: "Secretary" Problem

We consider the following (easier setting):

$$\inf_{T \in \mathbb{N}} \sup_{P \in \Delta(\Pi)} \inf_{\mathbf{V} \in \mathbb{R}^T_{\geq 0}} \mathbb{E}_\sigma \left[\frac{P(\sigma(\mathbf{V}))}{\mathrm{OFF}(\mathbf{V})} \right] \ ,$$

where

$$\sigma(\mathbf{V}) := (V_{\sigma(1)}, \ldots, V_{\sigma(T)})$$

is a random order of arrival drawn uniformly from all $T!$ possible orders. Essentially, the adversary picks valuations that are unknown to the policy, but then the order of their arrival is random. The random order eases the setting, in comparison to the unknown distributions setting that was presented in Section 2. The policy in this setting is random.

$$\sup_{P \in \Delta(\Pi)} \inf_{\mathbf{V} \in \mathbb{R}^T_{\geq 0}} \mathbb{E}_\sigma \left[\frac{P(\sigma(\mathbf{V}))}{\mathrm{OFF}(\mathbf{V})} \right] \quad \geq \quad \sup_{P \in \Delta(\Pi)} \inf_{\mathbf{V} \in \mathbb{R}^T_{\geq 0}} \mathbb{P}_\sigma \left[P(\sigma(\mathbf{V})) = \mathrm{OFF}(\mathbf{V}) \right] \ ,$$

since the RHS restricts the reward to the setting where the optimal agent is accepted. It could be shown that the inequality is tight, but the proof is very complex, and would not be discussed in the scope of this lecture.

3.1 The Arrival Times Technique

Imagine that each agent t arrives at time Y_t, where:

$$Y_t \underset{\text{i.i.d.}}{\sim} \mathrm{Unif}[0,1] \ .$$

Clearly if we order the agents by their arrival times, we get a random distribution of the $T!$ orders, which adheres to the setting of the problem.

Algorithm

- Pass on all candidates before time τ.
- After τ, stop on the first candidate that is the best so far.

Analysis We have

$$\begin{aligned}
\mathbb{P}_\sigma\left[P(\sigma(\mathbf{V})) = \mathrm{OFF}(\mathbf{V})\right] &= \int_0^1 \mathbb{P}\left[P(\sigma(\mathbf{V})) = \mathrm{OFF}(\mathbf{V}) \mid Y_{t^*} = y\right] dy \\
&= \int_\tau^1 \mathbb{P}\left[P(\sigma(\mathbf{V})) = \mathrm{OFF}(\mathbf{V}) \mid Y_{t^*} = y\right] dy \\
&\geq \int_\tau^1 \frac{\tau}{y} dy \quad = \quad \tau \log\left(\frac{1}{\tau}\right)
\end{aligned}$$

Here, the second equality follows from the fact that if $Y_{t^*} < \tau$, then the algorithm cannot accept t^*. The inequality that follows refers to the case in which the highest value so far (that is not t^*) was in the interval $[0, \tau]$ and not in $[\tau, t^*]$, which is a sufficient condition for $\pi(\sigma(\mathbf{V})) = \mathrm{OFF}(\mathbf{V})$. Maximizing over τ we get

$$\max_{\tau \in [0,1]} \tau \log\left(\frac{1}{\tau}\right) \quad = \quad \frac{1}{e} \log e \quad = \quad \frac{1}{e} \quad \approx \quad 0.37 \ .$$

Remark 3. T must be known (otherwise the competitive ratio is zero).

3.2 Proving that the Algorithm is Best Possible

Technique: Encapsulate all possible algorithms using Linear Programming (LP). For any algorithm, let α_t denote:

$$\mathbb{P}\Big[\text{Accept } t^{\text{th}} \text{ agent to arrive} \cap \text{they are the best so far}\Big] \ .$$

It is important to note that if the t^{th} agent is the best so far, then the probability that it is also the best out of all the T agents is $\frac{t}{T}$, because the order is random. Since we only stop on a best so far agent, if we stop at the t^{th} agent, then the probability that it is best o far is $\frac{t}{T}$. To upper bound the objective value, we solve the following LP:

$$\begin{aligned}
&\textbf{Maximize:} && \sum_{t=1}^{T} \alpha_t \frac{t}{T} \\
&\textbf{Subject to:} && \alpha_t \leq \frac{1}{t}\left(1 - \sum_{t'<t} \alpha_{t'}\right), \quad \forall t \in [T] \\
& && \alpha_t \geq 0, \quad \forall t \in [T]
\end{aligned}$$

To understand why this formulation achieves an upper bound on the objective value, notice that the objective is exactly the probability of picking the best agent (the sum of probabilities that we stop at time t times the probability that agent t has the best valuation- law of total probability). The constraint ensures that agent t is the best so far (the $\frac{1}{t}$ term), and that no agent was chosen so far (the $1 - \sum_{t'<t} \alpha_{t'}$ term). It is convenient to change the constraint to the following (equivalent) formulation

$$t\alpha_t \quad \leq \quad 1 - \sum_{t'<t} \alpha_{t'} \ ,$$

for the purpose of taking the dual LP. The dual problem is

$$\textbf{Minimize:} \qquad \sum_{t=1}^{T} \beta_t$$

$$\textbf{Subject to:} \qquad t\beta_t + \sum_{t'>t} \beta_{t'} \geq \frac{t}{T}, \quad \forall t \in [T]$$

$$\beta_t \geq 0, \quad \forall t \in [T]$$

We now form a feasible dual solution, whose objective value is at least the objective value of the primal problem, whose objective value (the primal's) is at least the probability that the policy picks the best agent. It is important to note that this analysis focuses on T large enough. For a small value of T a better competitive ratio is achievable. The idea for deriving a dual feasible solution is to guess which α_t values are positive, and use complementary slackness to derive the dual variables. We set

$$\alpha_t \quad = \quad 0 \quad if \quad t < \frac{T}{e} \ ,$$

which means by complementary slackness that

$$\beta_t \quad = \quad 0 \quad if \quad t < \frac{T}{e} \ ,$$

since the inequalities

$$t\alpha_t \quad \leq \quad 1 - \sum_{t'<t} \alpha_{t'} \ ,$$

are not tight. In addition, we set

$$\alpha_t \quad \geq \quad 0 \quad if \quad t > \frac{T}{e} \ ,$$

which means that the constraint

$$t\beta_t + \sum_{t'>t} \beta_{t'} \quad \geq \quad \frac{t}{T} \ ,$$

must be tight, since α_t is the corresponding dual variable of that constraint. We now claim that β_t are set to satisfy the following equality:

$$\sum_{t'>t} \beta_{t'} \quad = \quad \frac{t}{T} \sum_{t'=t}^{T-1} \frac{1}{t'} \tag{2}$$

Therefore the dual objective value is

$$\sum_{t'>t} \beta_{t'} \quad = \quad \frac{t}{T} \sum_{t'=t}^{T-1} \frac{1}{t'} \quad \approx \quad \frac{t}{T} \cdot \log \frac{T}{t} \quad \underset{\frac{t}{T}=\frac{1}{e}}{\approx} \quad \frac{1}{e} \ .$$

Consequently, we get that $\frac{1}{e}$ is an upper bound on the dual objective, which means that it is also an upper bound on the primal objective, and that no algorithm can obtain a competitive ratio better than $\frac{1}{e}$. Combining this result with the fact that the algorithm that we presented achieves a competitive ratio of $\frac{1}{e}$, we get that the bound is tight. It remains to prove that equation (2) indeed holds. We do this by backward induction on T:

Induction Base: We set

$$\beta_T = \frac{1}{T} \ ,$$

and equation (2) holds trivially.

induction Hypothesis We assume that

$$\sum_{t'>t} \beta_{t'} \quad = \quad \frac{t}{T} \sum_{t'=t}^{T-1} \frac{1}{t'} \ .$$

Induction Step We know that the constraint must be tight (for $t > \frac{T}{e}$) and therefore

$$\beta_t \quad = \quad \frac{1}{T} - \frac{1}{t} \sum_{t'>t} \beta_{t'} \underset{i.h.}{=} \frac{1}{T} - \frac{1}{t} \cdot \frac{t}{T} \cdot \sum_{t'=t}^{T-1} \frac{1}{t'} \quad = \quad \frac{1}{T} - \frac{1}{T} \cdot \sum_{t'=t}^{T-1} \frac{1}{t'} \ .$$

Therefore we have

$$\begin{aligned}
\sum_{t'>t-1} \beta_{t'} \quad &= \quad \beta_t + \sum_{t'>t} \beta_{t'} \\
&= \quad \underbrace{\frac{1}{T} - \frac{1}{T} \sum_{t'=t}^{T-1} \frac{1}{t'}}_{\beta_t} + \underbrace{\frac{t}{T} \sum_{t'=t}^{T-1} \frac{1}{t'}}_{i.h.} \\
&= \quad \frac{1}{T} \cdot \left(1 + \sum_{t'=t}^{T-1} \frac{t-1}{t'} \right) \\
&= \quad \frac{1}{T} \cdot \sum_{t'=t-1}^{T-1} \frac{t-1}{t'} \quad = \quad \frac{t-1}{T} \cdot \sum_{t'=t-1}^{T-1} \frac{1}{t'}
\end{aligned}$$

This completes our proof.

4 Third Setting: "Prophet Secretary" Problem

We consider the following setting

$$\inf_{T \in \mathbb{N}} \inf_{\mathbf{F} \in \Delta(\mathbb{R})^T} \sup_{P \in \Delta(\Pi)} \mathbb{E}_\sigma \left[\frac{P(\sigma(\mathbf{F}))}{\mathrm{OFF}(\mathbf{F})} \right]$$

This is the easiest setting, since it has both random order and (random) valuations whose distributions are known in advance (in the "prophet" setting, discussed in Lecture 1, we had an infimum over the orders instead of an expectation). We generalize the setting for the case where we can accept k agents. Recalling that we use X_t to denote the indicator that agent t is accepted, we define

$$D_t \quad = \quad \mathbf{1}(V_t > \tau) + \mathbf{1}(V_t = \tau) \cdot \mathrm{Ber}(\rho) \ ,$$

as an indicator that agent t clears the threshold. We also define

$$D \quad = \quad \sum_{t=1}^{T} D_t \ ,$$

to be the total demand (the number of agents that clear the threshold. If $D \leq k$ then all of them get accepted, and otherwise only k get accepted). The reward of a policy P given a distribution $\mathbf{F} \in \Delta(\mathbb{R}_{\geq 0})^T$ is

$$\begin{aligned} P(\mathbf{F}) &= \mathbb{E}\left[\sum_t V_t X_t\right] \\ &= \mathbb{E}\left[\sum_t [V_t - \tau]^+ X_t\right] + \tau \sum_t \mathbb{E}[X_t] \\ &\stackrel{I}{=} \mathbb{E}\left[\sum_t \mathbb{E}[V_t - \tau]^+ X_t\right] + \tau \mathbb{E}[\min\{D, k\}] \\ &\stackrel{II}{=} \sum_t \mathbb{E}[X_t (V_t - \tau)^+ \mid D_t = 1] \cdot \mathbb{P}[D_t = 1] + \tau \mathbb{E}[\min\{D, k\}] \\ &= \sum_t \mathbb{E}[(V_t - \tau)^+ \mid D_t = 1] \cdot \mathbb{P}[D_t = 1] \mathbb{E}[X_t \mid D_t = 1] + \tau \mathbb{E}[\min\{D, k\}] \end{aligned} \tag{3}$$

Here, equality I follows from the fact that the sum of X_t is the smallest between k or D. Equality II holds since the agents whose valuation is more than τ (and therefore have $[V_t - \tau]^+ \neq 0$) are those that have $D_t = 1$. We now divide our analysis to two cases:

1. *General Case*: in the general case we do not take advantage of the random order. We have

$$\begin{aligned} P(\mathbf{F}) &\underset{\text{equation(3)}}{=} \sum_t \mathbb{E}[(V_t - \tau)^+ \mid D_t = 1] \cdot \mathbb{P}[D_t = 1] \mathbb{E}[(X_t \mid D_t = 1] + \tau \mathbb{E}[\min\{D, k\}] \\ &\geq \sum_t \mathbb{E}\left[(V_t - \tau)^+ \mathbb{P}\left[\sum_{t' \neq t} D_{t'} < k\right]\right] + \tau \mathbb{E}[\min\{D, k\}] \\ &\geq \mathbb{P}[D < k] \sum_t \mathbb{E}[(V_t - \tau)^+] + \tau k \frac{\mathbb{E}[\min\{D, k\}]}{k} \end{aligned}$$

   Here, the first inequality holds since having at most k candidates that clear the threshold is a sufficient condition to be accepted (an agent may be among the k accepted agents even if more than k agents clear the threshold). The final inequality follows from the fact that $D \geq \sum_{t' \neq t} D_{t'}$. We recall the following Lemma from last class:

   Lemma 4.

$$\forall \tau \geq 0, \quad \sum_t [V_t - \tau]^+ + \tau k \geq \mathrm{OFF}(\mathbf{V})$$

   We have

$$\begin{aligned} P(\mathbf{F}) &\geq \mathbb{P}[D < k] \sum_t \mathbb{E}[(V_t - \tau)^+] + \tau k \frac{\mathbb{E}[\min\{D, k\}]}{k} \\ &\geq \min\left\{\mathbb{P}[D < k], \mathbb{E}\left[\min\left\{\frac{D}{k}, 1\right\}\right]\right\} \cdot \mathbb{E}[\mathrm{OFF}(\mathbf{V})] \end{aligned}$$

   Here, the second inequality follows from Lemma 4. Our goal now is to maximize the lower bound on the competitive ratio

$$\min\left\{\mathbb{P}[D < k], \mathbb{E}\left[\min\left\{\frac{D}{k}, 1\right\}\right]\right\} .$$

We want to set the threshold τ in order to control D, and ensure that the ratio is maximized. We notice that the $\mathbb{P}[D < k]$ argument is monotonically increasing in τ (decreasing in D), and the expectation is monotonically decreasing in τ (increasing in D). So the minimum function looks like the following drawing:

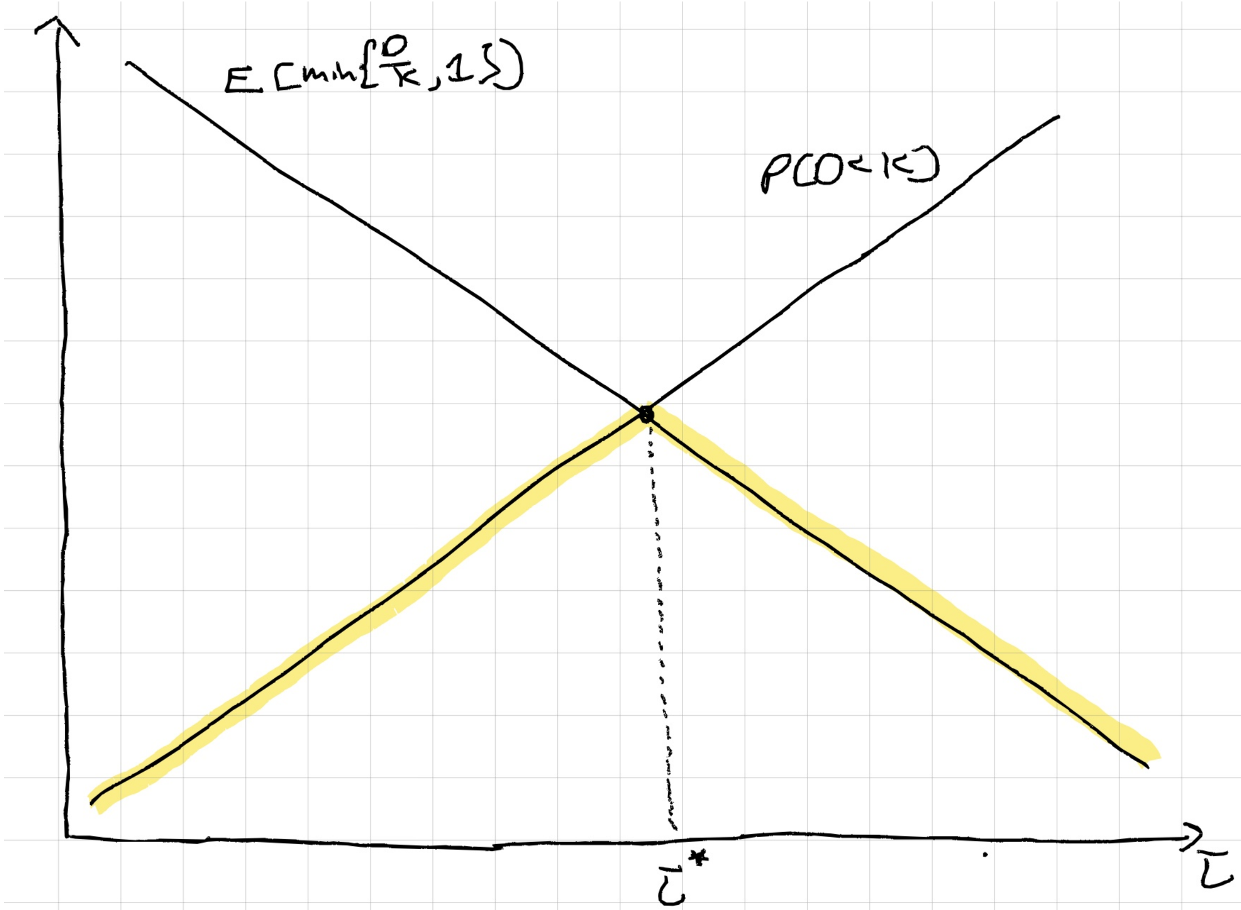

The highlighted part is the minimum function. It's maximum is attained at τ^*, which is the point where

$$\mathbb{P}[D < k] \quad = \quad \mathbb{E}\left[\min\left\{\frac{D}{k}, 1\right\}\right]$$

(Chawla et al., 2024). If $k = 1$, we can set τ according to the "median rule" such that

$$P[D = 0] \quad = \quad P[D \geq 1] \ ,$$

and get a $\frac{1}{2}$ competitive ratio.

2. *Random Order Case*: We go back to equation (3):

$$P(\mathbf{F}) = \sum_t \mathbb{E}[(V_t - \tau)^+ \mid D_t = 1] \cdot \mathbb{P}[D_t = 1]\mathbb{E}[X_t \mid D_t = 1] + \tau\mathbb{E}[\min\{D, k\}] \ .$$

We claim that when the order is random, we can increase $\mathbb{E}[(X_t \mid D_t = 1]$ and achieve a higher bound. When the order is random, we have

$$\mathbb{E}[(X_t \mid D_t = 1] \quad = \quad \mathbb{E}\left[\min\left\{\frac{k}{1 + \sum_{t' \neq t} D_{t'}}, 1\right\}\right] \ ,$$

since if agent t passes the threshold, its probability of acceptance is equal to the probability of any other threshold cleared agent. Consequently

$$\begin{aligned} P(\mathbf{F}) &\underset{\text{equation(3)}}{=} \sum_t \mathbb{E}[(V_t - \tau)^+ \mid D_t = 1] \cdot \mathbb{P}[D_t = 1]\mathbb{E}[(X_t \mid D_t = 1] + \tau\mathbb{E}[\min\{D, k\}] \\ &= \sum_t \mathbb{E}[(V_t - \tau)^+] \cdot \mathbb{E}\left[\min\left\{\frac{k}{1 + \sum_{t' \neq t} D_{t'}}, 1\right\}\right] + \tau\mathbb{E}[\min\{D, k\}] \\ &\underset{\text{Lemma4}}{\geq} \underbrace{\min\left\{\mathbb{E}\left[\min\left\{\frac{k}{1+D}, 1\right\}\right], \mathbb{E}\left[\min\left\{\frac{D}{k}, 1\right\}\right]\right\}}_{*} \mathbb{E}[\text{OFF}(\mathbf{V})] \end{aligned}$$

Now suppose that $k = 1$ and that the expected number of agents that clear the threshold ($\mathbb{E}[D]$) is $1 - 1/e$. We have

$$\begin{aligned} * &= \min\left\{\mathbb{E}\left[\frac{1}{1+D}\right], 1 - \frac{1}{e}\right\} \\ &\geq \min_{T, x \in [0,1]^T} \min\left\{\mathbb{E}\left[\frac{1}{1+D}\right], 1 - \frac{1}{e}\right\} \\ &\quad \textbf{Subject to:} \\ &\quad \mathbb{P}[D \geq 1] = 1 - \frac{1}{e} \\ &\quad D = \sum_{t=1}^{T} \text{Ber}(x_t) \end{aligned}$$

This is a Bernoulli optimization problem, where the Bernoulli variables are independent across t. We now present a theorem that we use in our analysis:

Theorem 5. *For a fixed T, $f, g : \{0, 1, \ldots, T\} \to \mathbb{R}$ and $c \in \mathbb{R}$, we present the following optimization problem:*

$$\begin{aligned} &\min_{x \in [0,1]^T} \mathbb{E}[f(D)] \\ &\quad s.t. \quad \mathbb{E}[g(D)] = c \\ &\quad\quad D = \sum_{t=1}^{T} Ber(x_t) \end{aligned}$$

The optimal solution always takes the following form:

$$\exists p \in (0, 1) \quad s.t. \forall t, x_t \in \{0, p, 1\}$$

Applying the theorem to our setting, we must have $x_t \in \{0, p\}$, since if by contradiction $x_t = 1$, then we would have $\mathbb{P}[D \geq 1] = 1 \geq 1 - \frac{1}{e}$, which violates the constraints of our optimization problem. Now let T denote the number of x_t's that are equal to p. The constraint is

$$1 - (1 - p)^T = 1 - \frac{1}{e} \Rightarrow p = 1 - e^{-1/T} .$$

It now remains to compute $\mathbb{E}[\frac{1}{1+D}]$:

$$\begin{aligned}
\mathbb{E}\left[\frac{1}{1+D}\right] &= \sum_{d=0}^{T}\binom{T}{d}p^d(1-p)^{T-d}\frac{1}{d+1} \\
&= \sum_{d=0}^{T}\frac{T!}{d!(d+1)(T-d)!}p^d(1-p)^{T-d} \\
&= \frac{1}{T+1}\frac{1}{p}\sum_{d=0}^{T}\binom{T+1}{d+1}p^{d+1}(1-p)^{T-d} \\
&= \frac{1}{(T+1)p}(1-(1-p)^{T+1}) \\
&\underset{p=1-e^{-1/T}}{=} \frac{\left(1-\frac{e^{-1/T}}{e}\right)}{(T+1)(1-e^{-1/T})}
\end{aligned}$$

We can verify the the infimum over $T \in \mathbb{N}$ is attained when $T \to \infty$:

$$\lim_{T\to\infty}\frac{\left(1-\frac{e^{-1/T}}{e}\right)}{(T+1)(1-e^{-1/T})} = (1-1/e)\lim_{T\to\infty}\frac{\frac{1}{T+1}}{1-e^{-1/T}} = 1-\frac{1}{e} .$$

It is important to note that his result is not tight. The bound is an open problem (there exists an algorithm with competitive ratio > 0.669 (Correa et al., 2021)). If π is a static threshold policy, the $1-\frac{1}{e}$ bound is tight.

5 Discussion

We summarize the results that we showed:

"Prophet" setting (known distributions, deterministic policy, adversarial time horizon, adversarial order of arrival):

$$\inf_{T\in\mathbb{N}}\inf_{\mathbf{F}\in\Delta(\mathbb{R}_{\geq 0})^T}\sup_{\pi\in\Pi}\frac{\pi(\mathbf{F})}{\mathrm{OFF}(\mathbf{F})} = \frac{1}{2} . \qquad (\bigstar)$$

In fact most of our proofs showed the stronger fact that

$$\inf_{T\in\mathbb{N}}\inf_{\mathbf{F}\in\Delta(\mathbb{R}_{\geq 0})^T}\sup_{P\in\Delta(\Pi)}\inf_{\sigma}\frac{P(\sigma(\mathbf{F}))}{\mathrm{OFF}(\mathbf{F})} = \frac{1}{2} .$$

"Unknown Distributions" setting (unknown distributions, randomized policy, adversarial time horizon, adversarial order of arrival):

$$\inf_{T\in\mathbb{N}}\sup_{P\in\Delta(\Pi)}\inf_{\mathbf{V}\in\mathbb{R}^T_{\geq 0}}\frac{P(\mathbf{V})}{\mathrm{OFF}(\mathbf{V})} = f(\mathcal{R}) \to 0 . \qquad (\diamondsuit)$$

"Secretary" setting (unknown valuations, randomized policy, adversarial time horizon, random order of arrival):

$$\inf_{T\in\mathbb{N}} \sup_{P\in\Delta(\Pi)} \inf_{\mathbf{V}\in\mathbb{R}^T_{\geq 0}} \mathbb{E}_\sigma \left[\frac{P(\sigma(\mathbf{V}))}{\mathrm{OFF}(\mathbf{V})}\right] = \inf_{T\in\mathbb{N}} \sup_{P\in\Delta(\Pi)} \inf_{\mathbf{V}\in\mathbb{R}^T_{\geq 0}} \mathbb{P}_\sigma \left[P(\sigma(\mathbf{V})) = \mathrm{OFF}(\mathbf{V})\right] = \frac{1}{e} \approx 0.37 \,. \qquad (\|)$$

"Prophet Secretary" setting (known distributions, randomized policy, adversarial time horizon, adversarial order of arrival):

$$\inf_{T\in\mathbb{N}} \inf_{\mathbf{F}\in\Delta(\mathbb{R}_{\geq 0})^T} \sup_{\pi\in\Pi} \mathbb{E}_\sigma \left[\frac{\pi(\sigma(\mathbf{F}))}{\mathrm{OFF}(\mathbf{F})}\right] > \underbrace{0.669}_{\text{we showed } \approx\, 0.63} , \quad \text{upper bound unknown} \,. \qquad (\S)$$

Relationship :

We have

$$(\S) \geq (\|) \geq (\diamondsuit)$$

and

$$(\S) \geq (\bigstar).$$

And if we can take $(\bigstar)$ take infimum over σ, we still have

$$\inf_{T\in\mathbb{N}} \inf_{\mathbf{F}\in\Delta(\mathbb{R}_{\geq 0})^T} \sup_{P\in\Delta(\Pi)} \inf_{\sigma} \frac{P(\sigma(\mathbf{F}))}{\mathrm{OFF}(\mathbf{F})} = \frac{1}{2} \geq (\diamondsuit)$$

Bibliographical notes. The weight-dependent competitive ratios for unknown distributions are based on Ball and Queyranne (2009). The slick lower bound proof for secretary is due to Correa (2020), with its ideas originating from Bruss (1984). The upper bound proof for secretary using the LP comes from Buchbinder et al. (2014). The prophet secretary analysis presented here based on Bernoulli optimization is from Arnosti and Ma (2023).

References

N. Arnosti and W. Ma. Tight guarantees for static threshold policies in the prophet secretary problem. *Operations research*, 71(5):1777–1788, 2023.

M. O. Ball and M. Queyranne. Toward robust revenue management: Competitive analysis of online booking. *Operations Research*, 57(4):950–963, 2009.

F. T. Bruss. A unified approach to a class of best choice problems with an unknown number of options. *The Annals of Probability*, pages 882–889, 1984.

N. Buchbinder, K. Jain, and M. Singh. Secretary problems via linear programming. *Mathematics of Operations Research*, 39(1):190–206, 2014.

S. Chawla, N. Devanur, and T. Lykouris. Static pricing for multi-unit prophet inequalities. *Operations Research*, 72(4):1388–1399, 2024.

J. Correa. Private communication. 2020.

J. Correa, R. Saona, and B. Ziliotto. Prophet secretary through blind strategies. *Mathematical Programming*, 190(1):483–521, 2021.

A. C.-C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, pages 222–227. IEEE Computer Society, 1977.