Columbia B9136 March 25, 2025

Lecture 7 — Constant Regret in Online Resource Allocation

Lecturer: Will Ma *Scribe: Jason Chatzitheodorou & Vlad Pyltsov*

1 Introduction

1.1 Overview of Lecture

In this lecture, we design and analyze algorithms that aim to minimize regret for the following problems:

1. Multi-Secretary under discrete valuations
   - Constant regret algorithm
2. Network Revenue Management
   - Constant regret algorithm
3. Multi-Secretary under continuous valuations
   - Logarithmic regret algorithm

1.2 Multi-Secretary

The instances $\mathcal{I}$ of the Multi-Secretary problem can be defined as follows:

1. **Inventory**: k units available
2. **Agents**: $T > k$ agents
3. **Valuations:** Valuations $V_t \overset{\text{iid}}{\sim} F \in \Delta(\mathbb{R}_{\geq 0})$ for all $t \in [T]$

On each iteration $t \in [T]$ we see the realisation of V_t and must accept or reject agent t, as long as we have accepted less than k agents in iterations $\{1, 2, \ldots, t-1\}$.

1.3 Regret vs Ratio

In lecture 3 we showed that for the Multi-Secretary problem there is an algorithm with ratio:

$$\frac{P(F^T)}{\mathsf{OFF}(F^T)} \geq \frac{P(F^T)}{\mathsf{FLU}(F^T)} \geq 1 - e^{-k} \cdot \frac{k^k}{k!} \simeq 1 - \frac{1}{\sqrt{2\pi k}}.$$

If we assume that $F \in \Delta([0, R])$, then we have that:

$$\underbrace{\mathsf{OFF}(F^T) - P(F^T)}_{\text{Regret}} \leq \frac{\mathsf{OFF}(F^T)}{\sqrt{2\pi k}} = O(\sqrt{k}) = O(\sqrt{T})$$

where the first equality is due to $\mathsf{OFF}(F^T)$ increasing linearly with k in the worst case and the asymptotic notation hides terms that depend on R.

Remark 1. We will develop algorithms with improved regret: $O(1)$ for discrete valuations (does not grow with k, T) and $O(\log T)$ for continuous valuations.

2 Multi-Secretary under Discrete Valuations

2.1 Model

The valuation distributions we are interested in for this section have the following form:

1. **Values**: $R \geq r_1 > \cdots > r_n \geq 0$
2. **Probabilities**: $\lambda_j := \Pr[V_t = r_j]$ for all $t \in [T]$ and $\sum_j \lambda_j = 1$

Assumption 2. We assume that $\lambda_{\min} = \min_j \lambda_j > 0$.

We use the following notation in the analysis:

1. $\theta_t \in [n]$: The realized type of agent t for all $t \in [T]$.
2. $X_t \in \{0, 1\}$: Indicator that algorithm accepts t.
3. B_t: Remaining inventory at the start of time t, with $B_1 = k$.
4. $D_j(t, T] := \sum_{t'>t} \mathbb{1}\{\theta_{t'} = j\}$: The number of agents after t with type j.
5. $J_t^{\mathsf{OFF}}(b)$: The maximum value accrued by OFF, starting from agent t with b inventory left.

Specifically, we define $J_t^{\mathsf{OFF}}(b)$ using an LP over variables $z_j, \bar{z}_j$, which represent the number of type j agents accepted/rejected respectively:

$$\begin{aligned} J_t^{\mathsf{OFF}}(b) := \max \sum_j r_j z_j & \\ \text{s.t.} \ \sum_j z_j \leq b & \\ z_j + \bar{z}_j = \sum_{t'=t}^{T} \mathbb{1}\{\theta_{t'} = j\} & \quad \forall j \in [n] \\ z_j, \bar{z}_j \geq 0 & \quad \forall j \in [n] \end{aligned}$$

Remark 3. The previous LP's optimal solutions are integral, since its matrix is totally unimodular and the right-hand sides are integral.

2.2 Algorithm

The algorithm can only observe the valuation V_t on iteration t, so it cannot use $J_t^{\mathsf{OFF}}(b)$, which is defined based on the realizations of the future valuations. It uses the following Fluid LP to decide whether to accept or reject on each iteration t:

$$\begin{aligned} \mathsf{FLU}(b) := \max \sum_j r_j y_j & \\ \text{s.t.} \ \sum_j y_j \leq b & \\ y_j + \bar{y}_j = \mathbb{1}\{\theta_t = j\} + (T - t)\lambda_j & \quad \forall j \in [n] \\ y_j, \bar{y}_j \geq 0 & \quad \forall j \in [n] \end{aligned}$$

Remark 4. Note that θ_t is known to the algorithm, as it only depends on V_t.

Algorithm for Discrete Valuations

For agents $t = 1, \ldots, T$:

1. Solve $\mathsf{FLU}(B_t)$, hereafter let $(y_j)_j, (\bar{y}_j)_j$ refer to the optimal solution.
2. Accept t iff $y_{\theta_t} \geq \bar{y}_{\theta_t} \Leftrightarrow y_{\theta_t} \geq \frac{1+(T-t)\lambda_{\theta_t}}{2}$.

2.3 Analysis

Theorem 5 (Constant Regret)**.** $\mathsf{OFF}(F^T) - \pi(F^T) \leq R \cdot \frac{2}{\lambda_{\min}}$.

Proof.

$$
\begin{aligned}
\mathsf{OFF}(F^T) - \pi(F^T) &= \mathbb{E}\left[J_1^{\mathsf{OFF}}(k)\right] - \mathbb{E}\left[\sum_{t=1}^{T} r_{\theta_t} X_t\right] && (1)\\
&= \mathbb{E}\left[\sum_{t=1}^{t}\left(J_t^{\mathsf{OFF}}(B_t) - J_{t+1}^{\mathsf{OFF}}(B_{t+1}) - r_{\theta_t} X_t\right)\right] && (2)\\
&\leq \sum_{t=1}^{T} R \cdot \Pr\left[J_t^{\mathsf{OFF}}(B_t) > J_{t+1}^{\mathsf{OFF}}(B_{t+1}) + r_{\theta_t} X_t\right] && (3)\\
&\leq R \cdot \sum_{s=1}^{T-1} e^{-\frac{s}{2}\lambda_{\min}^2} && (4)\\
&= R \cdot \sum_{s=1}^{T-1} \left(e^{-\frac{\lambda_{\min}^2}{2}}\right)^s && (5)\\
&\leq R \cdot \frac{e^{-\frac{\lambda_{\min}^2}{2}}}{1 - e^{-\frac{\lambda_{\min}^2}{2}}} && (6)\\
&\leq R \cdot \frac{2}{\lambda_{\min}^2} && (7)
\end{aligned}
$$

where (1) is by definition, (2) is by telescoping sum and linearity of expectation, (3) is by Homework 7, (4) is by $s = T - t$ and because the event $\{J_t^{\mathsf{OFF}}(B_t) > J_{t+1}^{\mathsf{OFF}}(B_{t+1}) + r_{\theta_t} X_t\}$ only happens when the algorithm makes a mistake and by theorems 6 and 7 we can upper bound the probability of this event and the fact that $\lambda_{\theta_t} + \frac{1}{s} \geq \lambda_{\min}$, (5) is by refactoring, (6) is by the geometric sum if we sum up to infinity and (7) is by the inequality $\frac{e^{-x}}{1-e^{-x}} = \frac{1}{e^x - 1} \leq \frac{1}{1+x-1} = \frac{1}{x}$. □

Lemma 6 (Mistaken acceptance)**.** *The algorithm accepts agent t mistakenly, i.e. the algorithm accepts t and $J_t^{\mathsf{OFF}}(B_t)$ rejects it, with probability at most* $\exp\left(-\frac{s}{2} \cdot (\lambda_{\theta_t} + \frac{1}{s})^2\right)$ *for $s = T - t$.*

Proof. Since the algorithm accepts, it must be true that for the variables of $\mathsf{FLU}(B_t)$ $y_{\theta_t} \geq \frac{1+(T-t)\lambda_{\theta_t}}{2}$, using which we show that:

$$B_t - \sum_{j<\theta_t} (T-t)\lambda_j \geq \frac{1 + (T-t)\lambda_{\theta_t}}{2} \qquad (1)$$

To see that consider the first constraint of $\mathsf{FLU}(B_t)$, which can be written $B_t \geq \sum_{j \neq \theta_t} y_j + y_{\theta_t} \geq \sum_{j<\theta_t} y_j + y_{\theta_t}$. Also, the second constraint gives $y_j \leq (T-t)\lambda_j$ for any $j < \theta_t$. Assume towards a contradiction that there exists $j < \theta_t$ such that $y_j < (T-t)\lambda_j$. Then we can define a new feasible solution by increasing y_j by $\min\{\frac{1+(T-t)\lambda_{\theta_t}}{2}, (T-t)\lambda_j - y_j\}$ and decreasing y_{θ_t} by the same amount and still have a feasible solution. But the new solution has a higher objective value due to $r_j > r_{\theta_t}$ which is a contradiction to optimality, therefore $y_j = (T-t)\lambda_j$ for all $j < \theta_t$. Overall, we have that $B_t \geq \sum_{j<\theta_t} (T-t)\lambda_j + \frac{1+(T-t)\lambda_{\theta_t}}{2}$.

Since $J_t^{\mathsf{OFF}}(B_t)$ rejects t and its optimal solutions are integral we have that $z_{\theta_t} = 0$, using which we show that:

$$\sum_{j<\theta_t} D_j(t,T] \geq B_t \tag{2}$$

We first show that $z_j = 0$ for all $j > \theta_t$. Assume towards a contradiction that there is such j for which $z_j > 0$, then we can create a new feasible solution with $z'_j = 0$ and $z'_{\theta_t} = z_j$, which has a higher objective value as $r_{\theta_t} > r_j$ and contradicts optimality. Now notice that due to the second constraint we have that $z_j \leq \sum_{t'=t}^{T} \mathbb{1}\{\theta_t = j\} = \sum_{t'=t+1}^{T} \mathbb{1}\{\theta_t = j\} = D_j(t,T]$ for any $j < \theta_t$. Assume towards a contradiction that $B_t > \sum_{j<\theta_t} D_j(t,T]$. But then we could create a new feasible solution $z'_j = D_j(t,T]$ for all $j < \theta_t$, $z'_{\theta_t} = B_t - \sum_{j<\theta_t} D_j(t,T]$ and $z'_j = z_j = 0$ for all $j > \theta_t$, which has a higher objective value since it only increased variables that correspond to higher values and contradicts optimality.

Combining (1) and (2) we have that mistaken acceptance implies:

$$\sum_{j<\theta_t} D_j(t,T] - \sum_{j<\theta_t} (T-t)\lambda_j \geq \frac{1+(T-t)\lambda_{\theta_t}}{2}$$

But $\sum_{j<\theta_t} D_j(t,T] = \sum_{t'=t+1}^{T} \sum_{j<\theta_t} \mathbb{1}\{\theta_{t'} = j\} \sim \text{Bin}(T-t, \sum_{j<\theta_t} \lambda_j)$ since independently for each $t' \in (t,T]$, at most one of the indicators $\mathbb{1}\{\theta_{t'} = j\}$ is equal to 1 for $j < \theta_t$, which happens with probability $\sum_{j<\theta_t} \lambda_j$. Overall we have that:

$$\begin{aligned}
\Pr[\text{mistaken acceptance}] &\leq \Pr[\text{Bin}(T-t, \sum_{j<\theta_t} \lambda_j) - \mathbb{E}[\text{Bin}(T-t, \sum_{j<\theta_t} \lambda_j)] \geq \frac{1+(T-t)\lambda_{\theta_t}}{2}] \\
&\leq \exp\left(\frac{-2\left(\frac{1+(T-t)\lambda_{\theta_t}}{2}\right)^2}{T-t}\right) \\
&\leq \exp\left(-\frac{s}{2}\left(\lambda_{\theta_t} + \frac{1}{s}\right)^2\right)
\end{aligned}$$

where the second inequality is due to Hoeffding's Inequality and the third due to setting $s = T-t$. ∎

Lemma 7 (Mistaken rejection)**.** *The algorithm rejects agent t mistakenly with probability at most $\exp\left(-\frac{s}{2} \cdot (\lambda_{\theta_t} + \frac{1}{s})^2\right)$ for $s = T-t$.*

Proof. The proof is very similar to theorem 6 so only the main points are presented. Since the algorithm rejects, it must be true that for the variables of $\mathsf{FLU}(B_t)$ $y_{\theta_t} < \frac{1+(T-t)\lambda_{\theta_t}}{2}$, which implies:

$$B_t - \sum_{j \leq \theta_t} (T-t)\lambda_j < \frac{1+(T-t)\lambda_{\theta_t}}{2} \tag{1}$$

Since $J_t^{\mathsf{OFF}}(B_t)$ rejects t and its optimal solutions are integral we have that $z_{\theta_t} = 0$, which implies:

$$\sum_{j \leq \theta_t} D_j(t, T] < B_t \tag{2}$$

Combining (1) and (2) we have that mistaken rejection implies:

$$\sum_{j \leq \theta_t} D_j(t, T] - \sum_{j \leq \theta_t} (T-t)\lambda_j < \frac{1 + (T-t)\lambda_{\theta_t}}{2}$$

But $\sum_{j \leq \theta_t} D_j(t, T] = \sum_{t'=t+1}^{T} \sum_{j \leq \theta_t} \mathbb{1}\{\theta_{t'} = j\} \sim \text{Bin}(T-t, \sum_{j \leq \theta_t} \lambda_j)$ since independently for each $t' \in (t, T]$, at most one of the indicators $\mathbb{1}\{\theta_{t'} = j\}$ is equal to 1 for $j \leq \theta_t$, which happens with probability $\sum_{j \leq \theta_t} \lambda_j$. Overall we have that:

$$\begin{aligned}
\Pr[\text{mistaken rejection}] &\leq \Pr[\text{Bin}(T-t, \sum_{j \leq \theta_t} \lambda_j) - \mathbb{E}[\text{Bin}(T-t, \sum_{j \leq \theta_t} \lambda_j)] \leq \frac{1 + (T-t)\lambda_{\theta_t}}{2}] \\
&\leq \exp\left(\frac{-2\left(\frac{1+(T-t)\lambda_{\theta_t}}{2}\right)^2}{T-t} \right) \\
&\leq \exp\left(-\frac{s}{2}\left(\lambda_{\theta_t} + \frac{1}{s}\right)^2 \right)
\end{aligned}$$

where the second inequality is due to Hoeffding's Inequality and the third due to setting $s = T - t$. □

2.4 Remarks

Some remarks related to Theorem 5:

- Adaptivity, i.e. deciding whether to accept/reject based on the realization of the remaining inventory B_t, is necessary for achieving $O(1)$ regret (Arlotto and Gurvich, 2019, Thm. 1).
- The regret must scale with $\frac{1}{\lambda_{\min}}$ (see Arlotto and Gurvich, 2019, Lem. 1).

3 Network Revenue Management

3.1 Model

The instances $\mathcal{I}$ of this problem are:

1. **Resources**: m resources.
2. **Products**: n products.
3. **Inventory**: Resource i starts with k_i units, for $i \in [m]$.
4. **Revenue**: Product j has revenue r_j for $j \in [n]$, where $R \geq r_1 > \cdots > r_n \geq 0$.
5. **Probabilities**: $\lambda_j := \Pr[V_t = r_j]$ for all $t \in [T]$ and $\sum_j \lambda_j = 1$.

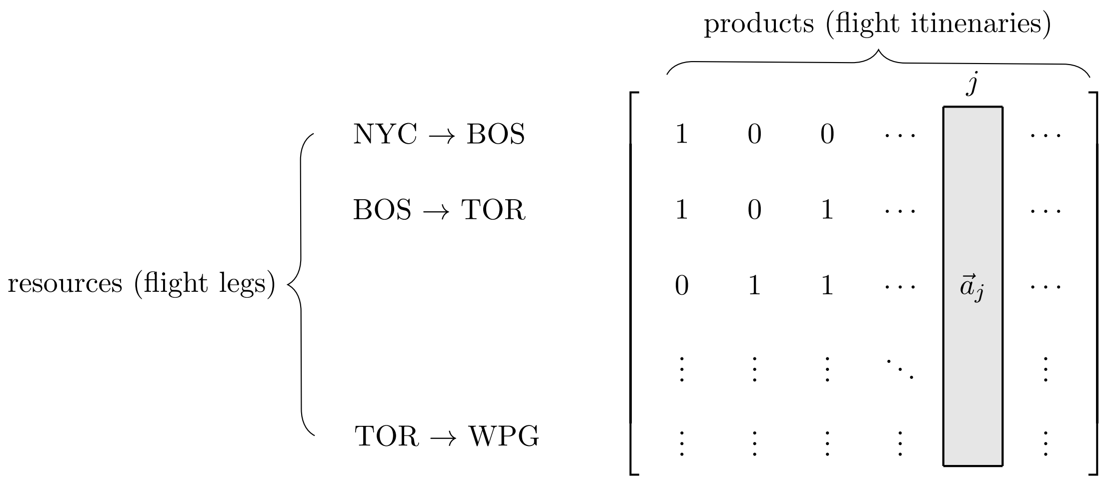

Figure 1: Example of an NRM Setting

6. **Requirements**: $a_{ij} \in \{0, 1\}$ indicate whether product j requires resource i.

On each step t we accept/reject in order to maximize revenue.

Remark 8. This problem is a generalization of Multi-Secretary for discrete valuations.

Assumption 9. We assume that $\lambda_{\min} = \min_j \lambda_j > 0$.

We use the following notation in the analysis:

1. $\theta_t \in [n]$: The realized type of agent t for all $t \in [T]$.
2. $X_t \in \{0, 1\}$: Indicator that algorithm accepts t.
3. $\vec{B}_t$: Remaining inventory for each resource at the start of time t, with $\vec{B}_1 = [k_1, \dots, k_m]$ and $\vec{B}_{t+1} = \vec{B}_t - X_t \cdot \vec{a}_{\theta_t}$.
4. $D_j(t, T] \coloneqq \sum_{t' > t} \mathbb{1}\{\theta_{t'} = j\}$: The number of agents after t with type j.
5. $J_t^{\mathsf{OFF}}(b)$: The maximum fractional value accrued by OFF, starting from agent t with b inventory left.

Specifically, we define $J_t^{\mathsf{OFF}}(\vec{b})$ using an LP over variables $z_j, \bar{z}_j$, which represent the number of type j agents accepted/rejected respectively:

$$
\begin{aligned}
J_t^{\mathsf{OFF}}(\vec{b}) \coloneqq \max \sum_j r_j z_j & \\
\text{s.t.} \ \sum_j z_j a_{ij} \le b_i & \quad \forall i \in [m] \\
z_j + \bar{z}_j = \sum_{t'=t}^{T} \mathbb{1}\{\theta_{t'} = j\} & \quad \forall j \in [n] \\
z_j, \bar{z}_j \ge 0 & \quad \forall j \in [n]
\end{aligned}
$$

Remark 10. The previous LP's optimal solutions are not necessarily integral.

3.2 Algorithm

The algorithm can only observe the valuation V_t on iteration t, so it cannot use $J_t^{\mathsf{OFF}}(\vec{b})$, which is defined based on the realizations of the future valuations. It uses the following Fluid LP to decide whether to accept or reject on each iteration t:

$$\begin{aligned}\mathsf{FLU}(\vec{b}) := \max & \sum_j r_j y_j \\ \text{s.t.} & \sum_j y_j a_{ij} \le b_i \quad \forall i \in [m] \\ & y_j + \bar{y}_j = \mathbb{1}\{\theta_t = j\} + (T-t)\lambda_j \quad \forall j \in [n] \\ & y_j, \bar{y}_j \ge 0 \quad \forall j \in [n]\end{aligned}$$

Remark 11. Note that θ_t is known to the algorithm, as it only depends on V_t.

Algorithm for Network Revenue Management

For agents $t = 1, \dots, T$:

1. Solve $\mathsf{FLU}(\vec{B}_t)$, hereafter let $(y_j)_j, (\bar{y}_j)_j$ refer to the optimal solution.

2. Accept t iff $y_{\theta_t} \ge \bar{y}_{\theta_t} \Leftrightarrow y_{\theta_t} \ge \frac{1+(T-t)\lambda_{\theta_t}}{2}$.

3.3 Analysis

Theorem 12. *(Constant Regret)* $\mathsf{OFF}(F^T) - \pi(F^T) \le LR \cdot \left(\frac{2}{\lambda_{\min}} + n \cdot \frac{16K(A)^2}{\lambda_{\min}^2}\right)$.

Proof.

$$\begin{aligned}\mathsf{OFF}(F^T) - \pi(F^T) &= \mathbb{E}\left[J_1^{\mathsf{OFF}}(\vec{k})\right] - \mathbb{E}\left[\sum_{t=1}^{T} r_{\theta_t} X_t\right] && (1)\\ &= \mathbb{E}\left[\sum_{t=1}^{T}\left(J_t^{\mathsf{OFF}}(\vec{B}_t) - J_{t+1}^{\mathsf{OFF}}(\vec{B}_{t+1}) - r_{\theta_t} X_t\right)\right] && (2)\\ &\le \sum_{t=1}^{T} LR \cdot \Pr\left[J_t^{\mathsf{OFF}}(\vec{B}_t) > J_{t+1}^{\mathsf{OFF}}(\vec{B}_t - X_t \vec{a}_{\theta_t}) + r_{\theta_t} X_t\right] && (3)\\ &\le \sum_{s=1}^{2/\lambda\min} LR + \sum_{s=2/\lambda\min+1}^{T-1} LR \cdot 2n \cdot \exp\left(-\frac{2}{s}\left(\frac{s\lambda_{\min} - 1}{2K(A)}\right)^2\right) && (4)\\ &\le LR \cdot \frac{2}{\lambda\min} + LR \cdot 2n \cdot \sum_{s=2/\lambda\min+1}^{T-1} \exp\left(-\frac{2}{s}\left(\frac{\frac{s\lambda_{\min}}{2}}{2K(A)}\right)^2\right) && (5)\\ &\le LR \cdot \left(\frac{2}{\lambda_{\min}} + 2n \cdot \sum_{s=1}^{\infty} \exp\left(-\frac{s}{8}\frac{\lambda_{\min}^2}{K(A)^2}\right)\right) && (6)\\ &\le LR \cdot \left(\frac{2}{\lambda_{\min}} + 2n \cdot \frac{8K(A)^2}{\lambda_{\min}^2}\right) && (7)\end{aligned}$$

where (1) is by definition, (2) is by telescopic sum and linearity of expectation, (3) is for $L = \max_j \sum_{i=1}^m a_{ij}$ and upper bounding each term of the sum by LR whenever it is positive by Homework 7, (4) is by splitting the sum at $s = 2/\lambda_{\min}$, upper bounding the terms of the left sum by setting the probability to 1 and using union bound and theorem 13 on the terms of the right sum, (5) is due to $s\lambda_{\min} - 1 \geq \frac{s\lambda_{\min}}{2}$ for $s > \frac{2}{\lambda_{\min}}$, (6) is by extending the right sum from $s = 1$ to ∞ and (7) is by the geometric sum. □

Lemma 13 (Mistaken acceptance/rejection)**.** *Let t be such that $(T-t)\lambda_{\min} \geq 1$. The algorithm makes a mistake at t with probability at most* $2n \cdot \exp\left(-\frac{2}{s}\left(\frac{s\lambda_{\min}-1}{2K(A)}\right)^2\right)$.

Proof. Since the algorithm makes a mistake, it is either a mistaken acceptance, i.e. it accepts t but $J_t^{\mathsf{OFF}}(\vec{B}_t)$ only partially accepts it, or a mistaken rejection, i.e. it rejects it but $J_t^{\mathsf{OFF}}(\vec{B}_t)$ only partially rejects it. We start with the first case. Since the algorithm accepts t, then we must have $y_{\theta_t} \geq \frac{1+(T-t)\lambda_{\theta_t}}{2}$. Also, since $J_t^{\mathsf{OFF}}(\vec{B}_t)$ rejects t we know that $z_{\theta_t} < 1$. Summing the two inequalities:

$$y_{\theta_t} - z_{\theta_t} \geq \frac{(T-t)\lambda_{\theta_t} - 1}{2} \Rightarrow ||y - z||_\infty \geq \frac{(T-t)\lambda_{\min} - 1}{2}.$$

We proceed to the second case. Since the algorithm rejects t, then we must have $y_{\theta_t} < \frac{1+(T-t)\lambda_{\theta_t}}{2}$. Also, since $J_t^{\mathsf{OFF}}(\vec{B}_t)$ only partially rejects t we know that $\bar{z}_{\theta_t} < 1$. Summing the two inequalities and substituting $\bar{z_{\theta_t}}$ from the second constraint of $J_T^{\mathsf{OFF}}(\vec{B}_t)$:

$$y_{\theta_t} - z_{\theta_t} < -\frac{(T-t)\lambda_{\theta_t} - 1}{2} \Rightarrow ||y - z||_\infty \geq \frac{(T-t)\lambda_{\min} - 1}{2}.$$

In either case we can apply the result of Mangasarian & Shiau to $\mathsf{FLU}(\vec{B}_t)$ and $J_t^{\mathsf{OFF}}(\vec{B}_t)$ since the lower bound of $||y - z||_\infty$ is the same and we get:

$$||d - d'||_\infty \geq \frac{(T-t)\lambda_{\min} - 1}{2K(A)} \quad \Rightarrow \quad \bigcup_{j=1}^{n} \{|D_j(t,T] - s\lambda_j| \geq \frac{s\lambda_{\min} - 1}{2K(A)}\}$$

by substituting the expressions of d, d' from the LP's, where $s = T - t$. Notice that $D_j(t,T] \sim \text{Bin}(s, s\lambda_j)$, then:

$$\begin{aligned}
\Pr[\text{mistake}] &\leq \Pr\left[\bigcup_{j=1}^{n} \{|D_j(t,T] - s\lambda_j| \geq \frac{s\lambda_{\min} - 1}{2K(A)}\}\right] \\
&\leq n \cdot \Pr\left[|\text{Bin}(s, s\lambda_{\min}) - s\lambda_{\min}| \geq \frac{s\lambda_{\min} - 1}{2K(A)}\right] \\
&\leq n \cdot 2\exp\left(-\frac{2}{s}\left(\frac{s\lambda_{\min} - 1}{2K(A)}\right)^2\right)
\end{aligned}$$

where the second inequality is by union bound and the third by Hoeffding's Inequality. □

Mangasarian & Shiau '87 (Lipschitz Continuity of Solutions of Linear Inequalities, Programs and Complementarity Problems)

For all $A \in \{0,1\}^{m\times n}$, $\exists$ constant $\kappa(A)_{>0}$ dependent only on A such that given any $\vec{\mathbf{r}} \in \mathbb{R}^n_{\geq 0}$, $\vec{\mathbf{b}} \in \mathbb{R}^m_{\geq 0}$, and $\vec{\mathbf{d}}$, $\vec{\mathbf{d}}' \in \mathbb{R}^n_{\geq 0}$.

(LP1)

$$\begin{aligned} \max \quad & \vec{\mathbf{r}} \cdot \vec{\mathbf{z}} \\ \text{s. t.} \quad & A\mathbf{z} \leq \mathbf{b} \\ & 0 \leq \mathbf{z} \leq \mathbf{d} \end{aligned}$$

(LP2)

$$\begin{aligned} \max \quad & \vec{\mathbf{r}} \cdot \vec{\mathbf{y}} \\ \text{s. t.} \quad & A\mathbf{y} \leq \mathbf{b} \\ & 0 \leq \mathbf{y} \leq \mathbf{d}' \end{aligned}$$

Given an optimal solution y to (LP2), $\exists$ an optimal solution z to (LP1), such that:

$$||\mathbf{z} - \mathbf{y}||_\infty \leq \kappa(A)||\mathbf{d} - \mathbf{d}'||_\infty$$

4 Multi-Secretary under Continuous Valuations

4.1 Model

The model has the following definitions:

1. **Valuations:** Valuations $V_t \sim \text{Unif}[0,1]$ for all $t \in [T]$ (or simply saying the valuations follow the uniform distribution).
2. **Remaining inventory:** B_t at the start of time t, with $B_1 = k$.
3. **The maximum value:** $J_t^{\mathsf{OFF}}(b) = V^{(1)} + \dots V^{(b)}$, where $V^{(i)}$ is the i-th highest order statistic among $V_1 \dots, V_T$.
4. **Indicator:** $X_t \in \{0,1\}$ that algorithm accepts t.

4.2 Algorithm

At time t, accept V_t if and only if:

$$V_t \geq \mathbb{E}_{V_{t+1},\dots,V_T}\left[J_{t+1}^{\mathsf{OFF}}(B_t) - J_{t+1}^{\mathsf{OFF}}(B_t - 1)\right]$$

4.3 Analysis

Theorem 14 (Regret)**.** $\mathsf{OFF}(F^T) - \pi(F^T) \leq \frac{1}{8}\ln(T+1)$.

Proof. Note that any time t, if B_t

$$
\begin{aligned}
\mathsf{OFF}(F^T) - \pi(F^T) &= \sum_{t=1}^{T} \mathbb{E}\left[J_t^{\mathsf{OFF}}(B_t) - J_{t+1}^{\mathsf{OFF}}(B_t - X_t) - V_t X_t\right] \\
&= \sum_{t=1}^{T} \mathbb{E}\left[J_t^{\mathsf{OFF}}(B_t) - J_{t+1}^{\mathsf{OFF}}(B_t) + J_{t+1}^{\mathsf{OFF}}(B_t) - J_{t+1}^{\mathsf{OFF}}(B_t - X_t) - V_t X_t\right] \\
&\le \sum_{s=1}^{T-1} \max_{b\in\{1,\dots,s\}} \left(\sum_{l=1}^{b}\left(\frac{s+2-l}{s+2} - \frac{s+1-l}{s+1}\right) - \frac{1}{2}\left(\frac{b}{s+1}\right)^2\right) \\
&= \sum_{s=1}^{T-1} \max_{b\in\{1,\dots,s\}} \left(\frac{1}{(s+1)(s+2)}\sum_{l=1}^{b} l - \frac{1}{2}\left(\frac{b}{s+1}\right)^2\right) \\
&= \sum_{s=1}^{T-1} \max_{b\in\{1,\dots,s\}} \left(\frac{b(b+1)}{2(s+1)(s+2)} - \frac{b^2}{2(s+1)^2}\right) \\
&= \sum_{s=1}^{T-1} \max_{b\in\{1,\dots,s\}} \frac{b}{2(s+1)}\left(\frac{(b+1)(s+1) - b(s+2)}{(s+2)(s+1)}\right) \\
&\le \sum_{s=1}^{T-1} \frac{1}{2(s+1)^2(s+2)} \sup_{b\in\mathbb{R}} b(bs + b + s + 1 - bs - 2b) \\
&= \sum_{s=1}^{T-1} \frac{1}{2(s+1)^2(s+2)} \cdot \frac{(s+1)^2}{4} \\
&= \frac{1}{8}\sum_{s=1}^{T-1} \frac{1}{s+2} \\
&\le \frac{1}{8}\sum_{s=2}^{T} \frac{1}{s} \le \frac{1}{8}\int_1^T \frac{1}{s} ds = \frac{1}{8}\ln T
\end{aligned}
$$

□

The main work comes in the first inequality. Here, we are switching the sum to let $s = T - t$ denote the number of time steps in the future. We are also allowing an adversary to select a worst-case value b for B_t at each time t, noting that if $b = 0$ or if $b > s$ then there is no regret. Using properties of the Uniform Distribution (see •), we can express the terms $\mathbb{E}[J_t^{\mathsf{OFF}}(B_t)]$ and $\mathbb{E}[J_{t+1}^{\mathsf{OFF}}(B_t)]$ as a function of b and s (where the expectations are conditioned on $B_t = b$). For the expression $\mathbb{E}[J_{t+1}^{\mathsf{OFF}}(B_t) - J_{t+1}^{\mathsf{OFF}}(B_t - X_t) - V_t X_t]$, the algorithm sets $X_t = 1$ exactly when doing so would make this expression negative. Therefore, this expression is equal to

$$\mathbb{E}_{B_t,V_t}[\min\{\mathbb{E}_{V_{t+1},\dots,V_T}[J_{t+1}^{\mathsf{OFF}}(B_t) - J_{t+1}^{\mathsf{OFF}}(B_t - 1)] - V_t, 0\}]$$

which equals $\mathbb{E}[\min\{\frac{s+1-B_t}{s+1} - V_t, 0\}]$, again using •. Because V_t is drawn from Unif[0,1], it is easy to see that this expectation (conditional on $B_t = b$) equals $-(b/(s+1))^2/2$.

- **Fact (from Probability Theory)**

  1. Let there be $s \overset{\text{iid}}{\sim}$ Unif[0,1] Random Variables
  2. Let $V^{(l)}$ denote the l'th largest of these, $l \in [s]$

Then the distribution of $V^{(l)}$ has the following PDF ("Beta distribution"):

$$f(\theta) = \frac{(1-\theta)^{l-1}\theta^{s-l}}{\int_0^1 (1-\theta')^{l-1}(\theta')^{s-l}d\theta'}$$

and has a mean of:

$$\frac{s-l+1}{s+1}$$

4.4 Concluding Remarks

The concluding remarks state a few facts:

1. This analysis works for any distribution (with the final bound depending on that distribution) as long as $J_{t+1}^{\mathsf{OFF}}(B_t)$ can be calculated (see Besbes et al., 2024).
2. The lower bound of regret for the derivation provided for the continuous case can be shown to be $\Omega(\log(T))$ for Unif[0,1] (Bray, 2024, Prop. 2).
3. Suppose $k = T/2$ and $F = \text{Ber}(1/2)$. Then $\mathsf{FLU}(F^T) = k$ and $\mathsf{OFF}(F^T) = \mathbb{E}[\min\{\text{Bin}(T, 1/2), k\}] = \mathbb{E}[\min\{\text{Bin}(2k, 1/2), k\}] \approx k - \sqrt{k}$. Subtracting, we can get a lower bound of $\Omega(\sqrt{T})$ on regret. This comparison is reflected in Figure 2.
4. For the discrete case, $\sup_{k,F\in\Delta([0,1])}(\mathsf{OFF}(F^T)) - \sup_{\pi}(\pi(F^T)) = \Omega(\sqrt{T})$ (Jiang et al., 2025, Prop. 2). That is, it was important for our results that *the distribution was held fixed while T and k increased.* If the distribution is allowed to change with T, then a regret better than $\sqrt{T}$ is impossible.

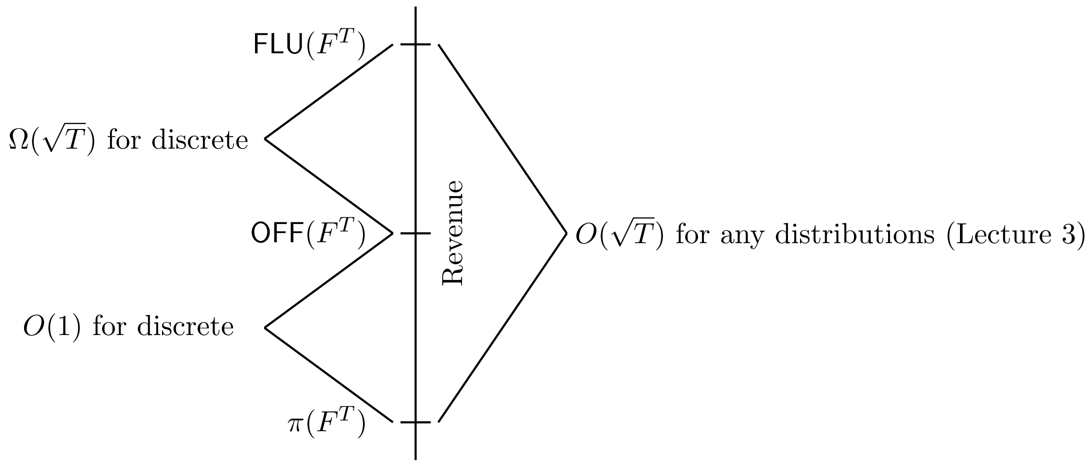

Figure 2: Summary of regret comparison

Bibliographical Notes. The analysis for the discrete cases presented here is based on Vera and Banerjee (2021), using the "delayed" version of the algorithm/analysis from Xie et al. (2025). The analysis for the Uniform[0,1] case is based on Bray (2024).

References

A. Arlotto and I. Gurvich. Uniformly bounded regret in the multisecretary problem. *Stochastic Systems*, 9(3):231–260, 2019.

O. Besbes, Y. Kanoria, and A. Kumar. Dynamic resource allocation: Algorithmic design principles and spectrum of achievable performances. *Operations Research*, 2024.

R. L. Bray. Logarithmic regret in multisecretary and online linear programs with continuous valuations. *Operations Research*, 2024.

J. Jiang, X. Li, and J. Zhang. Online stochastic optimization with wasserstein-based nonstationarity. *Management Science*, 2025.

A. Vera and S. Banerjee. The bayesian prophet: A low-regret framework for online decision making. *Management Science*, 67(3):1368–1391, 2021.

Y. Xie, W. Ma, and L. Xin. The benefits of delay to online decision-making. *Management Science*, 2025.