Columbia B9136	February 11, 2025

Lecture 3 — Multiple Units and Online Contention Resolution

Lecturer: Will Ma *Scribe: Shouqiao Wang and Xintong Yu*

Here is a summary of results.

Prophet (known independent distribution)

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

$$\inf_{\mathbf{F}\in\Delta(\mathbb{R}_{\geq 0})^T} \sup_{P\in\Delta(\Pi)} \inf_{\sigma} \frac{\mathbb{E}_{\mathbf{V}\sim\mathbf{F}}[P(\sigma(\mathbf{V}))]}{\mathrm{OFF}(\mathbf{F})} \geq 1 - O\left(\sqrt{\frac{\log k}{k}}\right)$$

The $\frac{1}{2}$ result is proven in lecture 1, while the $1 - O\left(\sqrt{\frac{\log k}{k}}\right)$ result is proven in the class.

Prophet Secretary

$$\inf_{\mathbf{F}\in\Delta(\mathbb{R}_{\geq 0})^T} \sup_{\pi\in\prod} \frac{\mathbb{E}_{\sigma}[\pi(\sigma(\mathbf{F}))]}{\mathrm{OFF}(\mathbf{F})} \geqslant 1 - 1/e$$

The result is proven in lecture 2. With fluid relaxation discussed in this lecture, the denominator can be updated to FLU($\mathbf{F}$).

IID (is strictly easier than Prophet Secretary case)

$$\inf_{F\in\Delta(\mathbb{R}_{\geq 0})} \sup_{\pi\in\prod} \frac{\pi(F^T)}{\mathrm{OFF}(F^T)} \geq 1 - O\left(\frac{1}{\sqrt{k}}\right)$$

The result is proven in this lecture.

Secretary (random order):

$$\sup_{P\in\Delta(\Pi)} \inf_{\mathbf{V}\in\mathbb{R}^T_{\geq 0}} \frac{\mathbb{E}_{\sigma}[P(\sigma(\mathbf{V}))]}{\mathrm{OFF}(\mathbf{V})} = 1/e \quad (k=1)$$

The result is proven in lecture 2.

Unknown/Correlated Distribution: 0, guarantees do not improve with k. This result not proven in class.

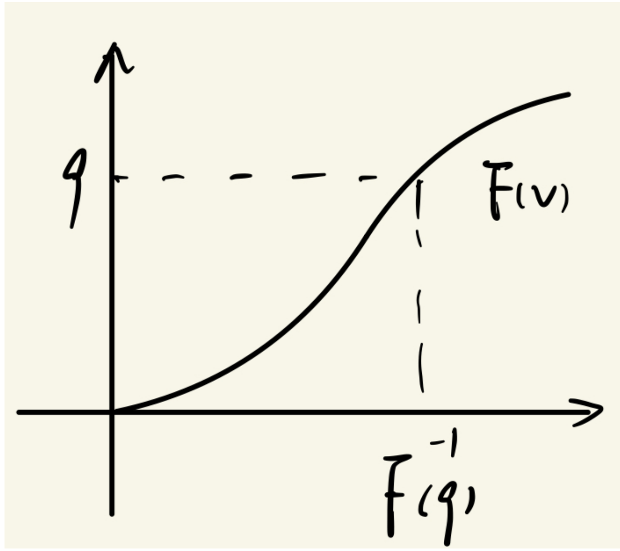

1 Fluid Relaxation

In this lecture, we consider the following *fluid relaxation* for a prophet-type problem with k selections. The fluid relaxation is defined as

$$\mathrm{FLU}(\mathbf{F}) = \max_{\{x_t \in [0,1] : \sum_{t=1}^{T} x_t \le k\}} \quad \sum_{t=1}^{T} \int_{1-x_t}^{1} F_t^{-1}(q)\, dq \ge \mathrm{OFF}(\mathbf{F}),$$

where $\mathrm{FLU}(\mathbf{F})$ is the expected value of the offline optimum, and $F_t^{-1}(q)$ denotes the q-quantile of the distribution F_t. In general, $\mathrm{FLU}(\mathbf{F})$ is more powerful than $\mathrm{OFF}(\mathbf{F})$. One intuition is that fluid can avoid "unluckiness".

For instance, consider the following example:

$$V_1 = V_2 = \begin{cases} 2 & \text{w.p. } \frac{1}{2}, \\ 1 & \text{w.p. } \frac{1}{2}, \end{cases}$$
$$\mathrm{OFF}(\mathbf{F}) = 1.75, \quad \text{and} \quad \mathrm{FLU}(\mathbf{F}) = 2.$$

Claim 1. *For any threshold $\tau \ge 0$, the following inequality holds:*

$$\sum_{t=1}^{T} \mathbb{E}\left[(V_t - \tau)^+\right] + k\tau \ge \mathrm{FLU}(\mathbf{F}).$$

Proof. Let $\{x_t\}_{t=1}^{T}$ be an optimal solution for the fluid relaxation, so that

$$\mathrm{FLU}(\mathbf{F}) = \sum_{t=1}^{T} \int_{1-x_t}^{1} F_t^{-1}(q)\, dq \quad \text{with} \quad \sum_{t=1}^{T} x_t \le k.$$

Observe that for any real number a and any $\tau \ge 0$, we have

$$(a-\tau)^+ + \tau \ge a.$$

Indeed, if $a \geq \tau$, then $(a-\tau)^+ = a-\tau$, and the sum is exactly a. Otherwise, if $a < \tau$, then $(a-\tau)^+ = 0$ and $\tau > a$. Applying this pointwise with $a = F_t^{-1}(q)$ for $q \in [1-x_t, 1]$, we obtain

$$\left[F_t^{-1}(q) - \tau\right]^+ + \tau \geq F_t^{-1}(q) \quad \text{for all } q \in [1-x_t, 1].$$

Integrate over $q \in [1-x_t, 1]$:

$$\int_{1-x_t}^{1} \left[F_t^{-1}(q) - \tau\right]^+ dq + \tau\, x_t \geq \int_{1-x_t}^{1} F_t^{-1}(q)\, dq.$$

Summing over all $t = 1, \ldots, T$, we have

$$\sum_{t=1}^{T} \int_{1-x_t}^{1} \left[F_t^{-1}(q) - \tau\right]^+ dq + \tau \sum_{t=1}^{T} x_t \geq \sum_{t=1}^{T} \int_{1-x_t}^{1} F_t^{-1}(q)\, dq.$$

Since $\sum_{t=1}^{T} x_t \leq k$, it follows that

$$\sum_{t=1}^{T} \int_{1-x_t}^{1} \left[F_t^{-1}(q) - \tau\right]^+ dq + k\tau \geq \mathrm{FLU}(\mathbf{F}).$$

Moreover, note that for each t,

$$\int_{1-x_t}^{1} \left[F_t^{-1}(q) - \tau\right]^+ dq \leq \int_{0}^{1} \left[F_t^{-1}(q) - \tau\right]^+ dq = \mathbb{E}\left[(V_t - \tau)^+\right],$$

which implies

$$\sum_{t=1}^{T} \mathbb{E}\left[(V_t - \tau)^+\right] + k\tau \geq \mathrm{FLU}(\mathbf{F}).$$

This completes the proof. $\square$

Claim 2. *If $F_1 = F_2 = \cdots = F_T = F$ (i.e., the distributions are identical) and $T \geq k$, then*

$$\mathrm{FLU}(F^T) = T \int_{1-\frac{k}{T}}^{1} F^{-1}(q)\, dq.$$

Proof. By concavity, it is optimal to allocate the same fraction of acceptance to every time period. In other words, choose

$$x_t = \frac{k}{T} \quad \text{for every } t = 1, 2, \ldots, T.$$

Substituting this into the definition of FLU gives

$$\mathrm{FLU}(F^T) = \sum_{t=1}^{T} \int_{1-\frac{k}{T}}^{1} F^{-1}(q)\, dq = T \int_{1-\frac{k}{T}}^{1} F^{-1}(q)\, dq.$$

This completes the proof. $\square$

2 IID Setting Analysis

In the IID setting, we assume that each agent's valuation is drawn independently from the same distribution F. The policy is designed by tuning a threshold τ and a tie-breaking probability p so that each agent t "clears" the threshold with probability

$$\mathbb{P}(D_t = 1) = \frac{k}{T},$$

where $D_t \in \{0,1\}$ is the indicator for crossing the threshold (i.e., $D_t \sim \mathrm{Ber}(k/T)$). The policy is designed such way to construct a clear comparision with the fluid benchmark.

The performance of the policy is given by

$$P(F^T) = \sum_{t=1}^{T} \mathbb{E}[V_t \mid X_t = 1]\,\mathbb{E}[X_t],$$

where X_t is the indicator that agent t is accepted. By designing the policy so that an agent that clears the threshold has conditional expected value

$$\mathbb{E}[V_t \mid X_t = 1] = \frac{1}{\frac{k}{T}} \int_{1-\frac{k}{T}}^{1} F^{-1}(q)\,dq,$$

and noting that the total expected number of accepted agents is $\mathbb{E}\left[\sum_{t=1}^{T} X_t\right]$, we can write

$$P(F^T) = \left(\frac{1}{\frac{k}{T}} \int_{1-\frac{k}{T}}^{1} F^{-1}(q)\,dq\right) \mathbb{E}\left[\sum_{t=1}^{T} X_t\right].$$

This expression can be rearranged to yield

$$P(F^T) = \frac{\mathrm{FLU}(F^T)}{k}\,\mathbb{E}\Big[\min\Big(\mathrm{Bin}\Big(T, \frac{k}{T}\Big), k\Big)\Big],$$

where $\mathrm{FLU}(F^T)$ denotes the fluid relaxation value for T agents with distribution F.

Taking the limit as $T \to \infty$, we obtain

$$P(F^T) \geq \mathrm{FLU}(F^T)\,\frac{1}{k} \lim_{T\to\infty} \mathbb{E}\left[\min\Big(\mathrm{Bin}\Big(T, \frac{k}{T}\Big), k\Big)\right].$$

Intuition: A splitting argument (details omitted) shows that the performance of the threshold policy is closely related to the fluid relaxation value.

Since a $\mathrm{Bin}(T, k/T)$ random variable converges in distribution to a $\mathrm{Pois}(k)$ random variable as $T \to \infty$, we have

$$\begin{aligned} P(F^T) &= \mathrm{FLU}(F^T)\,\frac{1}{k}\,\mathbb{E}\Big[\min\big(\mathrm{Pois}(k), k\big)\Big] \\ &= \mathrm{FLU}(F^T)\left(1 - e^{-k}\frac{k^k}{k!}\right) \\ &\approx \mathrm{FLU}(F^T)\left(1 - \frac{1}{\sqrt{2\pi k}}\right), \end{aligned}$$

where the final approximation uses Stirling's formula.

Thus, the performance of the IID policy is tightly linked to the fluid relaxation value, with the multiplicative factor $\left(1 - e^{-k}\frac{k^k}{k!}\right)$ (or its asymptotic equivalent) quantifying the effect of the stochastic acceptance process.

Note that as $k \to \infty$, can get converges 100%.

3 Concentration Inequality Analysis

Prophet Setting Recall: For a threshold policy, we have

$$P(\mathbf{F}) \geq \min\left\{ \Pr\Big[D \leq k\Big],\ \mathbb{E}\Big[\min\Big\{\frac{D}{k}, 1\Big\}\Big] \right\} \left(\sum_t \mathbb{E}\Big[(V_t - \tau)^+\Big] + \tau k \right), \tag{1}$$

where

- $D = \sum_t D_t$, the total number of people clearing the threshold, and
- $(x)^+ = \max\{x, 0\}$

Recap from Last Session: Last time we want to set $\Pr\Big[D \leq k\Big]$ equals to $\mathbb{E}\Big[\min\Big\{\frac{D}{k}, 1\Big\}\Big]$. Recall that we established

$$\sum_t \mathbb{E}\Big[(V_t - \tau)^+\Big] + \tau k \geq \text{FLU}(\mathbf{F}).$$

Focus on the Regime of Large k: In today's analysis we assume that k is large. In particular, we set

$$\mathbb{E}[D] = k\,(1 - \epsilon)$$

for some $\epsilon \in \left(0, \frac{1}{2}\right)$ and aim to show that the random variable D is tightly concentrated around its mean.

Concentration Bound for D: Fix $\epsilon \in (0,1)$ and suppose that D is the sum of independent Bernoulli random variables (i.e., a *Poisson binomial* random variable). Then, the multiplicative Chernoff bound states that

$$\Pr\Big[\,|D - \mathbb{E}[D]| \geq \epsilon\,\mathbb{E}[D]\,\Big] \leq 2\,\exp\Big(-\frac{1}{3}\,\epsilon^2\,\mathbb{E}[D]\Big).$$

Lower Bound on $\Pr\big[D < k\big]$**:** Substituting $\mathbb{E}[D] = k(1-\epsilon)$ yields

$$\Pr\Big[\,|D - \mathbb{E}[D]| \geq \frac{\epsilon}{1-\epsilon}\,k(1-\epsilon)\,\Big] \leq 2\,\exp\Big(-\frac{1}{3}\,\frac{\epsilon^2\,k}{1-\epsilon}\Big) \leq 2\,\exp\Big(-\frac{1}{3}\,\epsilon^2 k\Big).$$

Choosing

$$\epsilon^2 = \frac{3\log k}{k},$$

we deduce that

$$2\,\exp\Big(-\frac{1}{3}\,\epsilon^2 k\Big) = 2\,\exp(-\log k) = \frac{2}{k}.$$

Thus

$$\Pr\big[D < k\big] \geq 1 - \Pr\Big[\,|D - \mathbb{E}[D]| \geq \epsilon\,\mathbb{E}[D]\,\Big] \geq 1 - \frac{2}{k}$$

Lower Bound on $\mathbb{E}\Big[\min\{\frac{D}{k}, 1\}\Big]$**:** Using the above concentration result, we can lower bound the expected acceptance ratio. In particular, observe that

$$\mathbb{E}\Big[\min\Big\{\frac{D}{k}, 1\Big\}\Big] \geq \frac{\mathbb{E}[D](1-\epsilon)}{k}\;\Pr\Big[D \geq \mathbb{E}[D](1-\epsilon)\Big].$$

Since $\mathbb{E}[D] = k(1-\epsilon)$ and with our choice of ϵ the deviation probability is at most $2/k$, we obtain

$$\mathbb{E}\Big[\min\Big\{\frac{D}{k}, 1\Big\}\Big] \geq \Big(1 - \sqrt{\frac{3\log k}{k}}\Big)^2\Big(1 - \frac{2}{k}\Big).$$

Expanding for large k gives

$$\mathbb{E}\Big[\min\Big\{\frac{D}{k}, 1\Big\}\Big] \geq \Big(1 - 2\sqrt{\frac{3\log k}{k}}\Big)\Big(1 - \frac{2}{k}\Big) \geq 1 - 2\sqrt{\frac{3\log k}{k}} - \frac{2}{k} = 1 - O\Big(\sqrt{\frac{\log k}{k}}\Big).$$

This analysis shows that for large k the term $\mathbb{E}\big[\min\{D/k, 1\}\big]$ is very close to 1, up to an error term of order $O\Big(\sqrt{\frac{\log k}{k}}\Big)$. The $O\Big(\sqrt{\frac{\log k}{k}}\Big)$ error term is associated with the "threshold" policy, and requires an "adaptive policy" to eliminate the error term. This concentration result is crucial for relating the performance of the threshold policy to the fluid relaxation benchmark.

4 Online contention Resolution Schemes

4.1 Online Contention Resolution Schemes (OCRS)

We consider a setting with T agents, indexed by $t \in [T]$, each of which becomes active independently with probability $x_t \in [0, 1]$. Define

$$D_t = \begin{cases} 1, & \text{if agent } t \text{ is active,} \\ 0, & \text{otherwise.} \end{cases}$$

The algorithm can accept only active agents and is subject to a capacity constraint of k; that is, at most k agents can be accepted. We also assume the budget condition

$$\sum_{t=1}^{T} x_t \leq k.$$

Objective: c-Selectable

The goal is to design an online algorithm that, for each agent t, ensures

$$\Pr[\text{agent } t \text{ is accepted}] \geq c \cdot x_t,$$

where c is as large as possible. When this property is achieved, we say that the algorithm is *c-selectable.*

Special Case ($k = 1$): It is known that for $k = 1$, one can achieve $c = \frac{1}{2}$.

Algorithm: A common algorithm works as follows: process agents in order, and if agent t arrives active (i.e., $D_t = 1$) and no previous agent has been accepted (i.e., $D_{t' \prec t} = 0$), accept agent t with probability

$$\frac{1/2}{1 - \frac{1}{2}\sum_{t' < t} x_{t'}},$$

which is always at most 1.
Inductively Prove:

$$\mathbb{E}[X_t] = x_t(1 - \sum_{t' \prec t} \frac{1}{2} x_{t'}) \frac{1/2}{1 - \frac{1}{2}\sum_{t' < t} x_{t'}} = \frac{1}{2} x_t$$

4.1.1 Example with $k = 2$

We illustrate the tension between accepting agents early versus preserving capacity for later arrivals with a small example:

$$k = 2, \quad T = 3, \quad x_1 = 1, x_2 = x_3 = \frac{1}{2}.$$

Thus, each of the three agents arrives *with probability 1.* We have two available slots in total. We want to show $c = \frac{3}{4}$ is possible

Naïve Algorithm Consider the following simple algorithm as outlined below. The decision tree is sketched.

- **Agent 1:** Accepted with probability 3/4.
- **Agent 2:**
  - If Agent 1 was accepted and $x_2 = 1$, we accept Agent 2 with probability 2/3.
  - If Agent 1 was rejected and $x_2 = 1$, we again accept Agent 2 with probability 1.

  Overall, Agent 2 also has acceptance probability $\frac{3}{8} = x_2(\frac{3}{4} \cdot \frac{2}{3} + \frac{1}{4}) = \frac{3}{4} x_2$.
- **Agent 3:**
  - If there are any slot remains, and $x_3 = 1$, we accept Agent 3.

  A quick calculation shows that the *overall* acceptance probability for Agent 3 is $\frac{3}{4} x_3$.

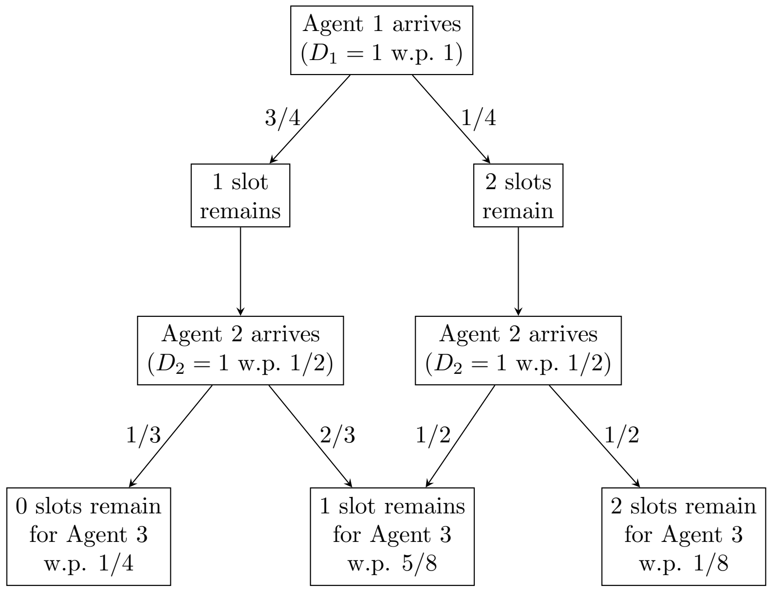

Observation: Coordination Between "Parallel Worlds" The **key insight** is that an OCRS algorithm must *coordinate* its acceptance probabilities across the different possible "worlds" (branches of the decision tree) so that no agent is at a disadvantage. If we want each agent t to have acceptance probability $\geq c\,x_t$ for some c, we must carefully tune the acceptance probabilities to "save room" for later agents and avoid saturating capacity too soon.

4.2 k-Unit OCRS via Linear Programming

To analyze the general k-unit OCRS problem, we introduce an LP formulation that captures the acceptance probabilities and the evolution of available capacity.

Notation

- β_t^l: the probability that l units remain when agent t arrives.
- α_t^l: the probability of accepting agent t when there are l units remaining.

We impose the convention $\alpha_t^{k+1} := 0$ and require $\alpha_t^l \geq 0$ for all t, l.

Constraints

(1) **OCRS Guarantee:**

$$c\,x_t \leq \sum_{l=1}^{k} \alpha_t^l, \quad \forall\, t.$$

This constraint ensures that the overall acceptance probability for agent t is at least $c\,x_t$.

(2) **OCRS Rules:**

$$\alpha_t^l \leq \beta_t^l\, x_t, \quad \forall\, t,\, l.$$

Since an agent can only be accepted if it is active and there are l units available, this constraint bounds α_t^l.

Evolution of Capacity

The available capacity evolves according to the rule:

$$\beta_t^l = \begin{cases} 1, & \text{if } t = 1 \text{ and } l = k, \\ 0, & \text{if } t = 1 \text{ and } l < k, \\ \beta_{t-1}^l - \alpha_{t-1}^l + \alpha_{t-1}^{l+1}, & \text{if } t > 1,\ l \in [k]. \end{cases}$$

Here, β_t^l is the probability of having l units remaining when agent t arrives, and α_t^l represents the probability of accepting agent t with l units remaining.

Optimal Structure

Theorem 1 (Optimal Structure, Proof Omitted)**.** *In an optimal solution to the LP, the variables α_t^l satisfy*

$$\alpha_t^l = \min\Big\{\beta_t^l\, x_t,\ c\, x_t - \sum_{l'>l} \alpha_t^{l'}\Big\} \quad \forall\, t,\ l.$$

This structure can be interpreted as a "water-filling" procedure: for each agent t, we allocate acceptance probability starting from the highest available capacity, ensuring that we do not exceed $\beta_t^l\, x_t$, and we fill until the total acceptance probability reaches $c\, x_t$.

Instance Optimality and Worst-Case Analysis: The LP gives an *instance-optimal* guarantee c. It turns out that the worst-case (smallest) c occurs when

$$x_t = \frac{k}{T} \quad \forall\, t, \quad \text{and } T \to \infty,$$

a regime that is often referred to as the "Poisson regime." In this case, the acceptance probabilities must be coordinated carefully to avoid exhausting capacity too early, hence the familiar intuition of "always saving a spot for the last person." Note, this is typical when comparing to fluid benchmark, which addresses a potential weakness of this analysis.
The worst case performance is

$$1 - O\big(\frac{1}{\sqrt{k}}\big)$$

4.3 1-Unit Random-Order CRS (RCRS)

In the random-order setting, agents arrive in a uniformly random order. We consider the case $k = 1$ and show that one can achieve a $(1 - \frac{1}{e})$-selectable CRS.

Setup

Assume that each agent t is assigned an arrival time

$$Y_t \sim \mathrm{Unif}[0, 1],$$

so that agents are observed in increasing order of Y_t. Since the agent arrives in random order, we no longer need to save the spot for the last one. The algorithm works as follows:

Algorithm for $k = 1$

When agent t arrives at time $Y_t = y$:

1. If the slot is still available and agent t is active (i.e., $D_t = 1$), then accept agent t with probability given by the outcome of a Bernoulli random variable with parameter

$$e^{-y\,x_t}.$$

Theorem

Under the above algorithm, assuming $\sum_{t=1}^{T} x_t \le 1$, we have

$$\mathbb{E}[X_t] \ge \Big(1 - \frac{1}{e}\Big) x_t.$$

That is, the scheme is $(1 - \frac{1}{e})$-selectable. (Recall that X_t is the indicator of t being selected.)

Proof

We condition on the arrival time $Y_t = y$:

$$\mathbb{E}[X_t] = \int_0^1 \mathbb{E}\big[\Pr(\text{accept } t \mid Y_t = y)\big]\, dy.$$

When agent t arrives at time y, it is active with probability x_t and is accepted with probability $e^{-y\,x_t}$ provided the slot is still available. The probability that the slot is free at time y can be expressed (after some calculations) as

$$\prod_{t' \ne t} \Big(1 - x_{t'} \int_0^y e^{-z\,x_{t'}}\, dz\Big) = \exp\Big(-y \sum_{t'=1}^{T} x_{t'}\Big).$$

Thus, we have

$$\mathbb{E}[X_t] = x_t \int_0^1 e^{-y\,x_t} \exp\Big(-y \sum_{t' \ne t} x_{t'}\Big) dy = x_t \int_0^1 \exp\Big(-y \sum_{t'=1}^{T} x_{t'}\Big) dy.$$

Since $\sum_{t'=1}^{T} x_{t'} \le 1$, it follows that

$$\int_0^1 \exp\Big(-y \sum_{t'=1}^{T} x_{t'}\Big) dy \ge \int_0^1 e^{-y} dy = 1 - \frac{1}{e}.$$

Hence,

$$\mathbb{E}[X_t] \ge \Big(1 - \frac{1}{e}\Big) x_t.$$

4.4 Equivalence with Prophet Inequalities Relative to Fluid

In this section, we explore how *Online Contention Resolution Schemes* (OCRS) connect to *Prophet Inequalities* through a fluid (fractional) benchmark. We consider a family of downward-closed feasible sets and show how the fluid-based LP bound (often called the "FLU" benchmark) upper-bounds the offline optimal solution and can be matched (up to a constant factor) by an online algorithm when a suitable OCRS is available.

General Problem Definition.

- We have T agents, each agent $t \in [T]$ arriving (or being "active") with probability $x_t \in [0, 1]$. Let D_t be the indicator for agent t being active.
- **Feasibility Constraints:** $\mathcal{F}$ is a *downward-closed* family of subsets of $[T]$. That is, if $S \in \mathcal{F}$ and $S' \subseteq S$, then $S' \in \mathcal{F}$ as well. e.g. $\mathcal{F} = \{S \subseteq [T] : |S| \leq k\}$.
- We can only accept active agents, and the set of accepted agents at the end must lie in $\mathcal{F}$. Because $\mathcal{F}$ is downward-closed, this imposes that the set of accepted agents at any point in time must belong to $\mathcal{F}$.

Examples In general, there are a wide variety of "packing" or "resources" constraints

1. **Knapsack.** Suppose we have a single knapsack of capacity 1, and each agent t has a size $s_t \in [0, 1]$. Then
$$\mathcal{F} \;=\; \Big\{ S \subseteq [T] \;\Big|\; \sum_{t \in S} s_t \;\leq\; 1 \Big\}.$$
2. **Network Revenue Management.** Suppose there are multiple resources i (each with capacity k_i). Agent t requires 1 unit of resource i if $a_{it} = 1$ (and requires none if $a_{it} = 0$). The feasible sets are those that do not exceed capacity on any resource:
$$\mathcal{F} \;=\; \Big\{ S \subseteq [T] \;\Big|\; \sum_{t \in S} a_{it} \;\leq\; k_i \;\; \forall\, i \Big\}.$$
3. **Matroid or Graph Matching.** In a matroid (or matching) setting, $\mathcal{F}$ is the collection of all independent sets. For graph matching, $\mathcal{F}$ is the family of matchable sets of edges. Can think agents are edges in a graph, and each edge uses exactly 2 resources.

Notation for Valuations and Benchmarks. Each agent t has a random valuation V_t drawn from some distribution F_t. Let $\mathbf{F} = (F_1, \ldots, F_T)$ be the product distribution. We consider:

- $\mathrm{OFF}(\mathbf{F}) \;=\; \mathbb{E}_{\mathbf{V} \sim \mathbf{F}}\Big[\max_{S \in \mathcal{F}} \sum_{t \in S} V_t\Big]$,
- A *fluid* (fractional) benchmark $\mathrm{FLU}(\mathbf{F})$ based on solving an LP with variables $\{x_t\}$ subject to $x \in P$, where $P \subseteq [0, 1]^T$ is a convex relaxation of $\mathcal{F}$.

Proposition 1. *(Fluid upper bound on the offline optimum.)* Suppose there is a convex set $P \subseteq [0, 1]^T$ such that for every $S \in \mathcal{F}$, the characteristic vector $\mathbb{1}_S \in P$. Then for any valuation distribution $\mathbf{F}$,

$$\mathrm{OFF}(\mathbf{F}) \;=\; \mathbb{E}_{\mathbf{V} \sim \mathbf{F}}\Big[\max_{S \in \mathcal{F}} \sum_{t \in S} V_t\Big] \;\leq\; \max_{\substack{\mathbf{x} \in P \\ x_t \in [0,1],\, t \in [T]}} \sum_{t=1}^{T} \int_{1 - x_t}^{1} F_t^{-1}(q)\, dq \;=\; \mathrm{FLU}(\mathbf{F}).$$

Proof. Let $X_t^* \in \{0,1\}$ be the indicator that agent t is chosen by the optimal offline solution on a particular valuation profile $\mathbf{V}$. Thus,

$$\mathrm{OFF}(\mathbf{F}) \;=\; \mathbb{E}_{\mathbf{V}\sim\mathbf{F}}\Big[\sum_{t=1}^{T} V_t\, X_t^*\Big].$$

Since the chosen set $S = \{t : X_t^* = 1\}$ must lie in $\mathcal{F}$, we have $\mathbb{1}_S \in P$. By convexity of P, it follows that the expectation of these indicators also lies in P, i.e. $\big(\mathbb{E}[X_t^*]\big)_{t\in[T]} \in P$. Define

$$x_t^* \;=\; \mathbb{E}[X_t^*].$$

Using an integral representation of expectations for $V_t \sim F_t$, we have

$$\sum_{t=1}^{T} \mathbb{E}[V_t X_t^*] \;\le\; \sum_{t=1}^{T} \int_{1-x_t^*}^{1} F_t^{-1}(q)\, dq.$$

Since $(x_t^*) \in P$ is a feasible solution for the fluid LP, we obtain

$$\mathrm{OFF}(\mathbf{F}) \;\le\; \max_{\mathbf{x}\in P} \sum_{t=1}^{T} \int_{1-x_t}^{1} F_t^{-1}(q)\, dq \;=\; \mathrm{FLU}(\mathbf{F}).$$

Proposition 2. *(From c-Selectable OCRS to Prophet Inequality.)* If there exists a c-selectable OCRS for a downward-closed family $\mathcal{F}$, then for *any* product distribution $\mathbf{F} \in \Delta(\mathbb{R}_{\ge 0})^T$, we can design an *online* accept/reject algorithm whose expected performance $P(\mathbf{F})$ is at least $c \cdot \mathrm{FLU}(\mathbf{F})$.

Proof.

1. **Solve the Fluid LP.** Compute the fractional solution $(x_t)_{t=1}^{T}$ that maximizes

   $$\sum_{t=1}^{T} \int_{1-x_t}^{1} F_t^{-1}(q)\, dq$$

   subject to $x \in P$.

2. **Construct "Active" Agents.** For each agent t, draw an independent $Q_t \sim \mathrm{Unif}[0,1]$ and define $D_t = 1$ ("active") if and only if $Q_t \ge 1 - x_t$. Thus, $\Pr[D_t = 1] = x_t$.

3. **Run the c-Selectable OCRS.** Given that each agent t is "active" with probability x_t, the c-selectable OCRS guarantees that the probability of agent t being accepted is at least $c\, x_t$.

4. **Bound on Expected Performance.** By definition of the fluid objective,

   $$P(\mathbf{F}) \;\ge\; \sum_{t=1}^{T} \Big(c\, x_t\Big)\, \frac{1}{x_t} \int_{1-x_t}^{1} F_t^{-1}(q)\, dq \;=\; c \cdot \mathrm{FLU}(\mathbf{F}).$$

   - cx_t: the probability of accepting agent t
   - the expected fluid value given $Q_t \ge 1 - x_t$ is $\int_{1-x_t}^{1} F_t^{-1}(q)\, dq$

   Hence we obtain the desired prophet inequality guarantee.

Interpretation. Proposition 2 shows that *if* we can resolve contention online with a fraction c of the fluid benchmark for each agent, *then* we can translate this into a c-approximate prophet inequality for the original problem. Conversely, developing improved OCRS implies improved prophet inequalities relative to the fluid LP bound.

4.5 Polytope

Ideal Polytope: $\mathcal{P}^* = \text{conv}\{\mathbb{1}_S : S \in \mathcal{F}\}$

- $\mathcal{P}$ for k-unit: $\{\mathbf{x} \in [0,1]^T : \sum_t x_t \leq k\} = \mathcal{P}^*$
- $\mathcal{P}$ for knapsack: $\{\mathbf{x} \in [0,1]^T : \sum_t s_t x_t \leq k\} \supsetneq \mathcal{P}^*$
- $\mathcal{P}$ for NRM (Network Revenue Management): $\{\mathbf{x} \in [0,1]^T : \sum_t a_{i,t} x_t \leq k_i\}$ for all i. $\mathcal{P} \supsetneq \mathcal{P}^*$

The problem with loose $\mathcal{P}$ is loose guarantee.

Proposition 3. Suppose we have algorithm that can gurantee $\pi(\mathbf{F}) \geq c\,\text{FLU}(\mathbf{F})$ for all $\mathbf{F} \in \Delta(\mathbb{R}_{\geq 0})^T$. Then we have c-selectable OCRS.

Proof. (Assume known order)
Given $\mathbf{x} \in \mathcal{P}$. Consider instances where $F_t = r_t \cdot \text{Ber}(x_t)\ \forall t$. Further let $\sum_t x_t r_t = 1$.
With the $\text{Ber}(x_t)$ operation, we know $\text{FLU} = \sum_t x_t r_t$
Define p_t^π : probability that policy π accepts t.

$$\begin{aligned}
&\min J \\
\text{s.t.}\quad & J \geq \sum_t p_t^\pi r_t = \pi(\mathbf{F}) && \text{dual variable: } z_\pi,\ \pi \in \Pi\{0,1\} \\
& \sum_{t=1} x_t r_t = 1 && \text{dual variable: } c' \\
& r_t \geq 0\ \forall t
\end{aligned}$$

By assumption, optimal objective value is at least c. For now, we can assume the LP is of finite dimension, since π maps history to zero or nonzero. The dual program is

$$\begin{aligned}
&\min c' \\
\text{s.t.}\quad & c' x_t \leq \sum_\pi p_t^\pi z_\pi\ \forall t \\
& \sum_{\pi \in \Pi\{0,1\}} z_\pi = 1 \\
& z_\pi \geq 0\ \forall \pi \in \Pi
\end{aligned}$$

The constraint $c' x_t \leq \sum_\pi p_t^\pi z_\pi\ \forall t$ shows that if we take expectation over a distribution of 0-1 accept/reject policy, the probability of accepting each agent t is at least $c' x_t$. By strong duality, the solution to dual program is at least c. So there exists solution $c' = c$ and we know such policy exists.

Bibliographical notes. The analysis for the IID special case is copied from Ma et al. (2021), with the main technical lemma originating from Yan (2011). The concentration inequality analysis for multiple units comes from Hajiaghayi et al. (2007). The multi-unit OCRS results come from Jiang et al. (2022). The slick analysis for 1-unit RCRS and the equivalence with fluid-relative prophet inequalities come from Lee and Singla (2018).

References

M. T. Hajiaghayi, R. Kleinberg, and T. Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, volume 7, pages 58–65, 2007.

J. Jiang, W. Ma, and J. Zhang. Tight guarantees for multi-unit prophet inequalities and online stochastic knapsack. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1221–1246. SIAM, 2022.

E. Lee and S. Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *26th European Symposium on Algorithms, ESA 2018*, page . Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, 2018.

W. Ma, D. Simchi-Levi, and J. Zhao. Dynamic pricing (and assortment) under a static calendar. *Management Science*, 67(4):2292–2313, 2021.

Q. Yan. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 710–719. SIAM, 2011.