Individual Fairness in Prophet Inequalities

Makis Arsenis[1*] Robert Kleinberg[1]

*marsenis@cs.cornell.edu [1]Cornell University

Individual Fairness in Resource Allocation

Shopper: needs a ride 10% of the time (e.g. only when shopping and miss the bus).
Frequent traveler: needs a ride 90% of the time.
Assumptions: Rural area, limited rides available. *At most one* person receives a ride at a time.
Frequent traveler's perspective: Gets a ride almost every time they need one.
Shopper's perspective: Gets a ride $\sim 50\%$ of the time.

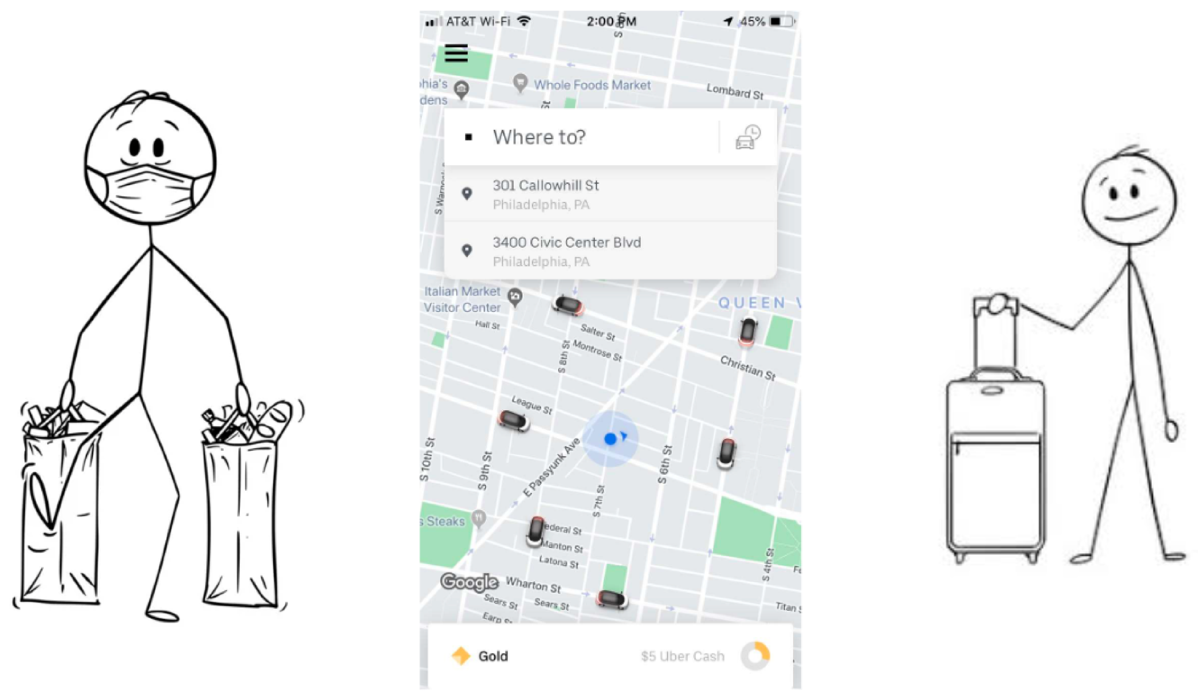

Grant applications with **no deadlines**: Two, otherwise *equal* applicants, have **different probability** of getting a grant depending on when they apply.

Online Decision Making / Optimal Stopping

- Sequence of r.v.s $X_1, X_2, \dots, X_n$, with known distr. $X_i \sim \mathcal{F}_i$
- E.g. X_i is the fitness of person i for a job they're being interviewed.
- Inspect values (interview) in order and pick (hire) at most one.
- Decisions are **immediate** and **irrevocable**.
- **Goal**: Maximize $\mathbb{E}$[value of variable picked] (denoted by $\mathbb{E}[\text{ALG}]$).
- Benchmark: *Prophet* — can see all X_i in an offline manner and choose $\max_i X_i$.

$$X_1 = \begin{cases} 1, & \text{w.p. } 9/10 \\ 0, & \text{w.p. } 1/10 \end{cases}, \quad X_2 = \begin{cases} 1, & \text{w.p. } 1/10 \\ 0, & \text{w.p. } 9/10 \end{cases}$$

Theorem (Prophet Inequality)

For any instance of the optimal stopping problem, there is an online algorithm ALG *s.t.:*

$$\mathbb{E}[\text{ALG}] \geq \underbrace{\frac{1}{2}}_{\textit{competitive ratio}} \cdot \underbrace{\mathbb{E}[\max_{i=1}^{n} X_i]}_{\textit{"Prophet"}}$$

Identity-Independent Fairness (IIF)

Instance $\mathcal{I} = (\mathcal{F}_i)_{i=1}^n$ of n r.v.s supported on $\mathcal{S}$.
Algorithm (*online* or *offline*) ALG for $\mathcal{I}$ satisfies **IIF** if

$$\Pr[\text{ALG hires } i \mid X_i = x] = \underbrace{p(x)}_{\text{Independent of } i}, \quad \forall i \in [n], x \in \mathcal{S}$$

Time-Independent Fairness (TIF)

Instance $\mathcal{I} = (\mathcal{F}_i)_{i=1}^n$ of n r.v.s supported on $\mathcal{S}$.
A family of algorithms $\{\text{ALG}^\pi\}_{\pi \in S_n}$ for $\mathcal{I}$ is **TIF** if

$$\Pr[\text{ALG}^\pi \text{ hires } i \mid X_i = x] = \underbrace{p(i, x)}_{\text{Independent of } \pi}, \quad \forall i \in [n], x \in \mathcal{S}, \pi \in S_n$$

Optimal, IIF/TIF Algorithm via LPs

(Online IIF LP)

$$\begin{cases} \max & \sum_{i=1}^{n} \sum_{x \in \mathcal{S}} x \cdot f_i(x) \cdot p_x \\ \text{s.t.} & p_x + \sum_{i=1}^{n-1} \sum_{y \in \mathcal{S}} f_i(y) \cdot p_y \leq 1, \ \forall x \\ & p_x \in [0,1] \end{cases}$$

(Online TIF LP)

$$\begin{cases} \max & \sum_{i=1}^{n} \sum_{x \in \mathcal{S}} x \cdot f_i(x) \cdot p_{i,x} \\ \text{s.t.} & p_{i,x} + \sum_{j \neq i} \sum_{y \in \mathcal{S}} f_j(y) \cdot p_{j,y} \leq 1, \ \forall i, x \\ & p_{i,x} \in [0,1] \end{cases}$$

Any solution to above LPs can be turned into an online algorithm:

ALGORITHM 1: IIF/TIF Algorithm for arrival ordering $X_{\pi(1)}, X_{\pi(2)}, \dots, X_{\pi(n)}$

for $t = 1, \dots, n$ **do**
  Inspect $X_{\pi(t)}$
  $Q_t \leftarrow 1 - \sum_{s=1}^{t-1} \sum_{y \in \mathcal{S}} f_{\pi(s)}(y) \cdot p(\pi(s), y)$; `// Pr[Algorithm arrives at` $\pi(t)$`]`
  Flip a biased coin with Heads probability equal to $q_t = \frac{p\left(\pi(t), X_{\pi(t)}\right)}{Q_t}$.
  if *coin comes up Heads* **then**
    Hire $X_{\pi(t)}$ and halt.
  else
    Reject and proceed.
  end
end

Summary of competitive ratio results

Off: *Offline* setting (prophet)
Off,IIF: *Offline* setting + *IIF* property
On: *Online* setting — but no fairness constraints
On,IFF/On,TIF: *Online* setting + *IIF/TIF* property

$$\mathcal{X} \to \mathcal{Y} : \rho = \inf_{\mathcal{F}_i} \frac{\sup_{\text{ALG}_\mathcal{Y} \in \mathcal{Y}} \mathbb{E}[\text{ALG}_\mathcal{Y}]}{\sup_{\text{ALG}_\mathcal{X} \in \mathcal{X}} \mathbb{E}[\text{ALG}_\mathcal{X}]}$$

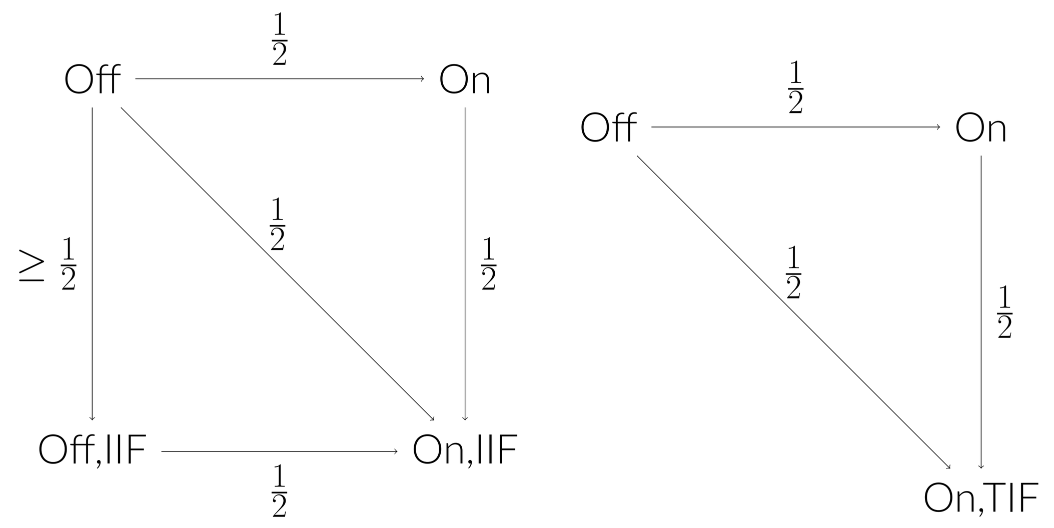

An impossibility result

Our algorithms are **randomized** and have a non-zero probability of selecting **no element**.
This is **inevitable**:

For any $\varepsilon > 0$, there is an instance of the optimal stopping problem such for any IIF or TIF algorithm that selects an element with probability 1:

$$\mathbb{E}[\text{ALG}] < \varepsilon \cdot \mathbb{E}\left[\max_i X_i\right]$$

Access to distributions via samples only

- Sometimes, the assumption that we have *exact* knowledge of distr. is **unrealistic**.
- Instead, can assume we have access to a constant number of **samples** from each distribution.
- The following is a **single-sample**, **offline**, **IIF**, $1/2$-**competitive** algorithm:

ALGORITHM 2: Single-sample **offline** algorithm

Data: $Y_i \sim F_i$ (samples)
Let $i^* \in [n]$ be such that $X_{i^*} = \max\{X_1, \dots, X_n\}$.
if $X_{i^*} > Y_{i^*}$ **then**
  Hire X_{i^*}.
else
  Hire no candidate.
end

- There exists a (similar) 2-**sample**, **online**, **IIF+TIF**, $1/9$-**competitive** algorithm.