Fairness and Bias in Online Selection

José Correa[1] , **Andrés Cristi[1]**, Paul Dütting[2] and Ashkan Norouzi-Fard[2]

[1]Universidad de Chile, Chile
[2]Google Research, Switzerland

Online Selection

Sequence of candidates

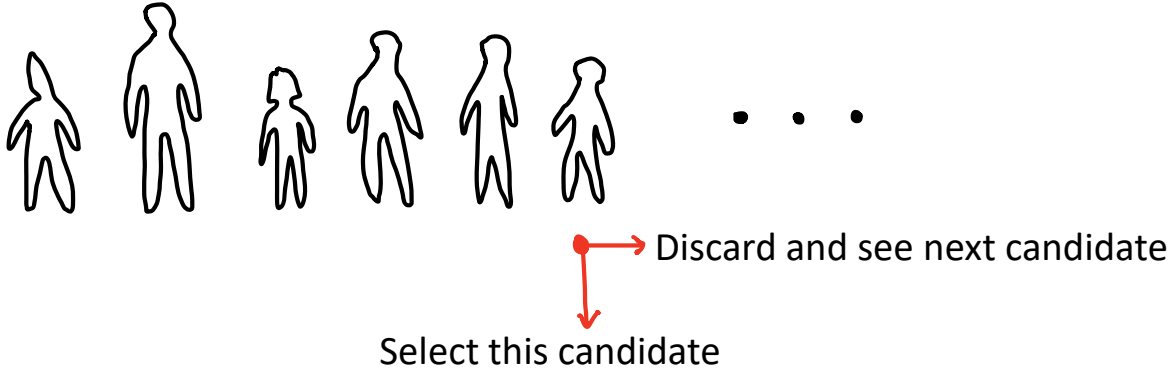

Two classic settings: **Secretary Problem** and **Prophet Inequality**

Online Selection

Sequence of candidates

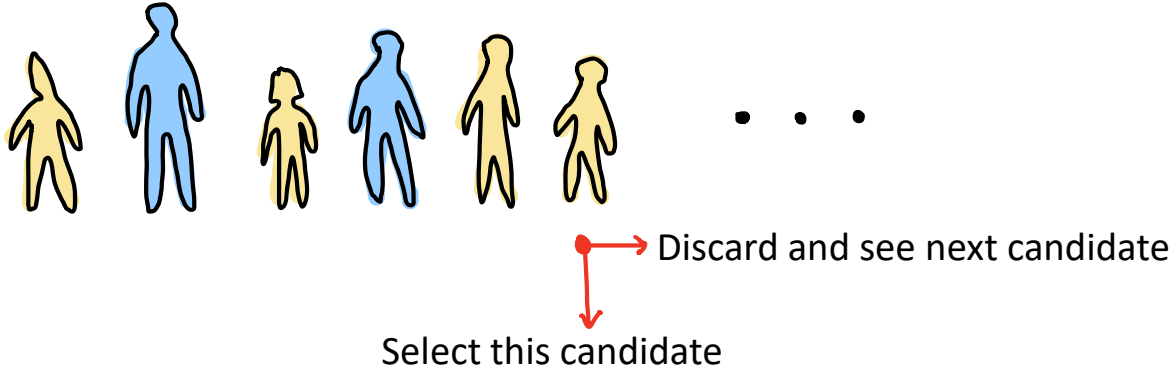

Two classic settings: **Secretary Problem** and **Prophet Inequality**

Multi-Color Secretary Problem

Sequence of candidates in **random order.**
Objective: Select the best

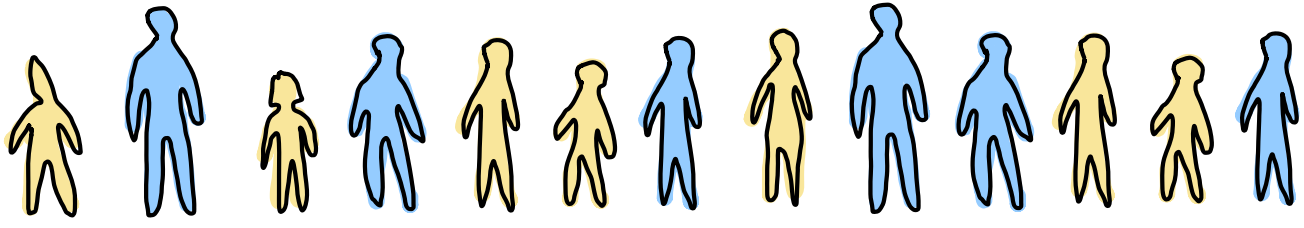

Men
Women

Prob. of having the best
p_m
p_w

We cannot compare across colors

Multi-Color Secretary Problem

Optimal strategy: set time-thresholds t_m and t_w

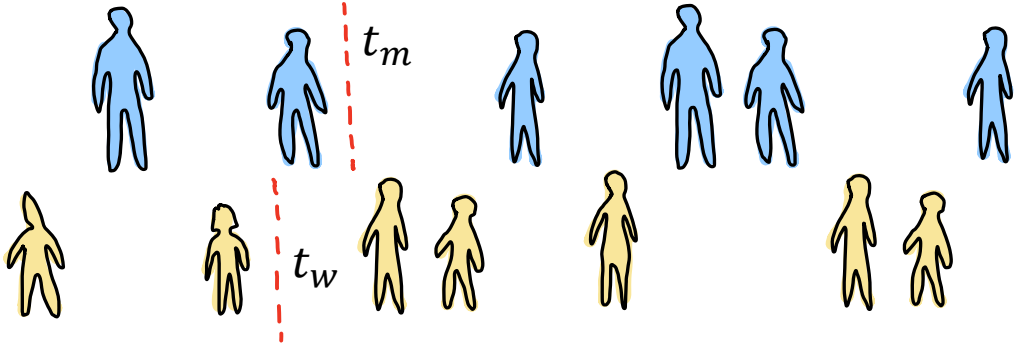

Theorem: If $p_m < p_w < (1+\varepsilon)p_m$, then $t_m > t_w > (1-\varepsilon)t_w$ and $0 < \mathbb{P}(\text{ALG selects a woman}) - \mathbb{P}(\text{ALG selects a man}) < \varepsilon$

Multi-Color Prophet Problem

Objective: maximize $\mathbb{E}[\mathrm{ALG}]$

Each candidate has a value drawn from a known distribution

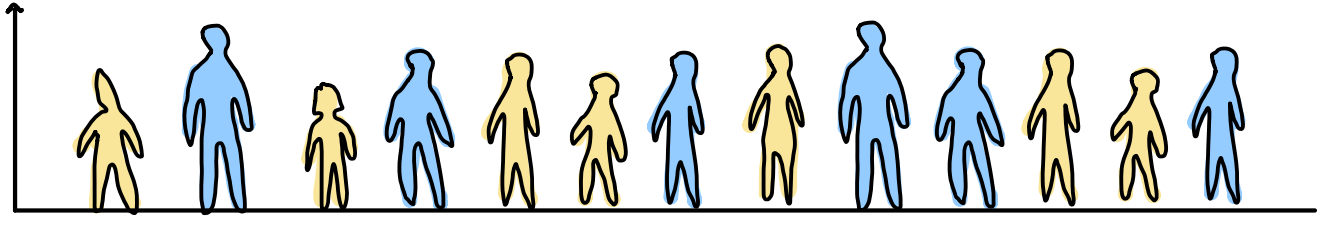

Men
Women

Marginals

p_m

p_w

← We want to respect these marginals

Multi-Color Prophet Problem

We design algorithms with good approximation guarantees

$$\mathbb{P}(ALG\ selects\ \text{  }\ ) = \alpha \cdot \mathbb{P}(FairOPT\ selects\ \text{  }\ )$$

Setting	Approximation factor
General	2
IID values, p_j proportional to group size	1.5
IID within group, p_j proportional, random order	1.707

Empirical Evaluation

Multi-Color Prophet, IID setting