Multi-secretary problem with many types

Omar Besbes (ob2105@columbia.edu), Yash Kanoria (ykanoria@gmail.com) and Akshit Kumar (ak4599@columbia.edu)

Columbia Business School

Multi-secretary Problem

Given a hiring budget B and horizon T, choose the top B secretaries based on their realized abilities.
Offline Problem: Can see the entire future.

Online Problem: Non-anticipating.

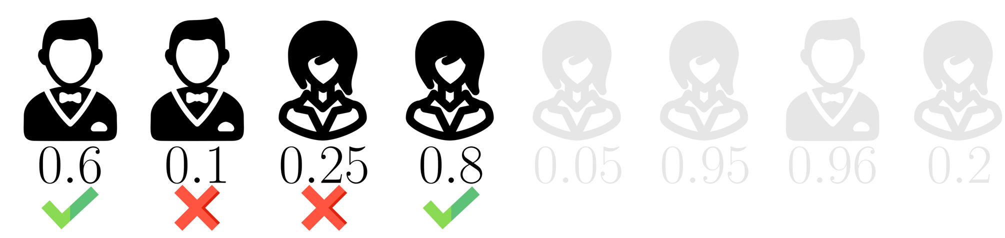

Realized abilities $\theta_t \overset{iid}{\sim} F$, F is known.

$$\text{Regret}(B,T;\texttt{ALG}) \triangleq \mathbb{E}[\texttt{OPT}] - \mathbb{E}[\texttt{ALG}]$$

What is known?

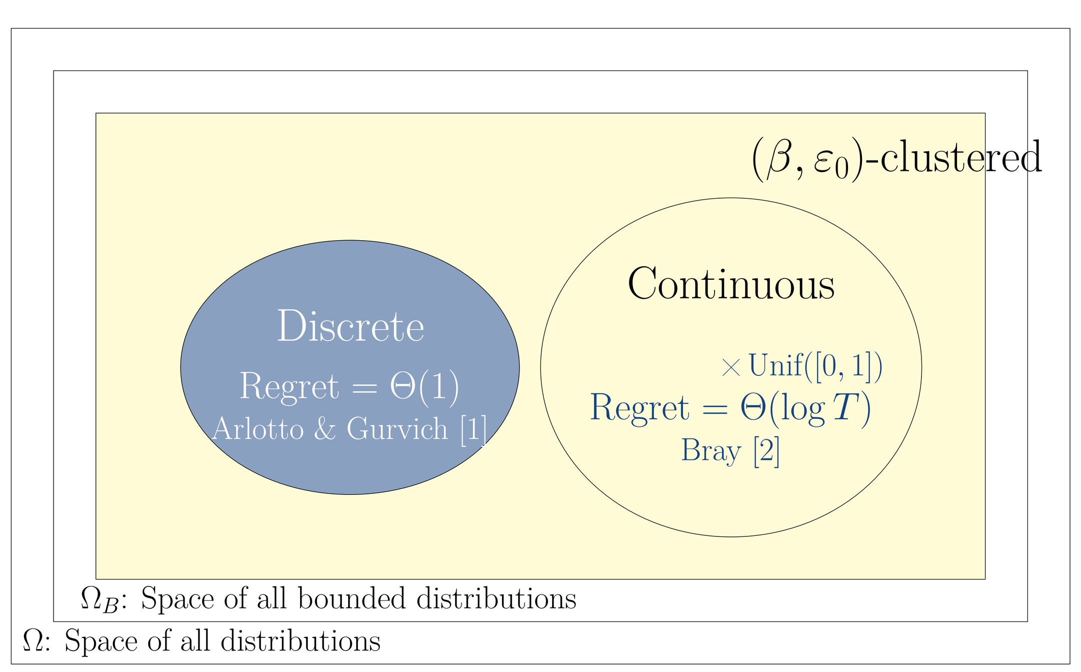

What is not known?

Low types and high types of secretaries (well separated) with uniform distribution over the types.

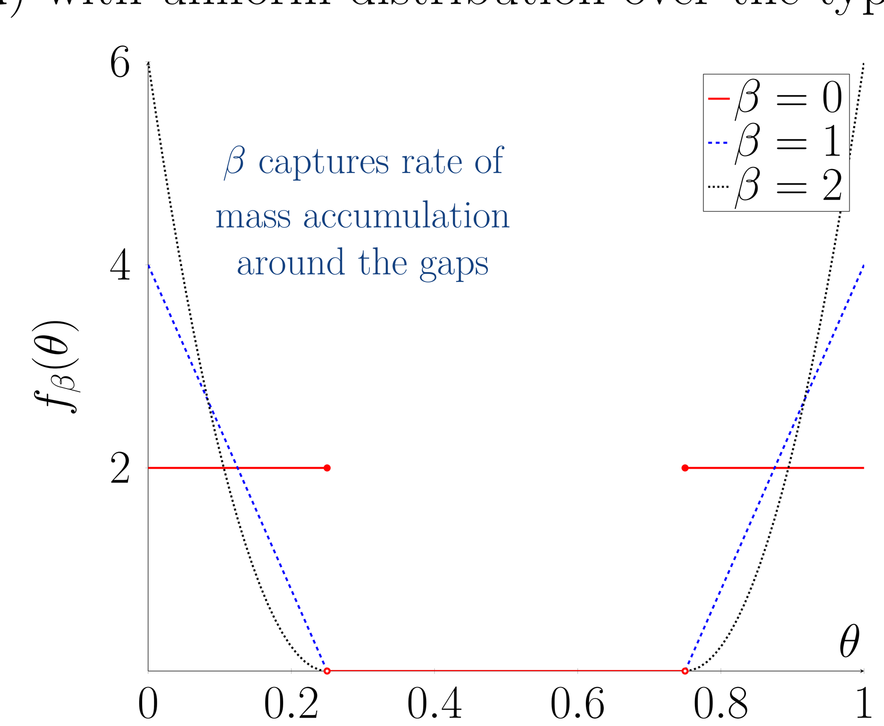

Common Heuristic

$\texttt{OPT}$: $\max_{\boldsymbol{x}} \sum_{t=1}^{T} \theta_t \boldsymbol{x_t}$ s.t $\sum_{t=1}^{T} \boldsymbol{x_t} \le B, \boldsymbol{x_t} \in \{0,1\}$

Difficulty: Online algorithm does not know the future i.e does not know all the θ_t in advance.

Certainty Equivalent Principle

Replace the stochastic quantities by their expectations and the constraints by their realized values; solve the opt. problem and use the solution.

For uniform distribution, CE is a ***threshold*** policy.

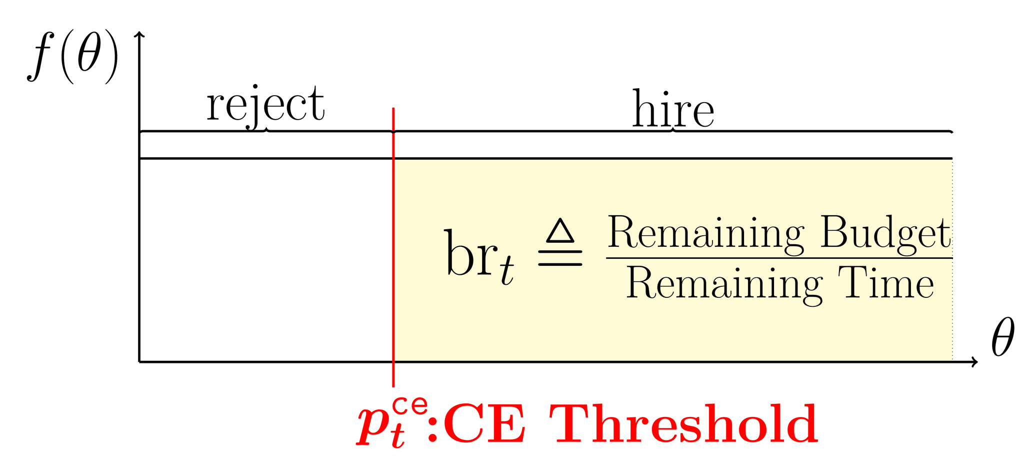

$$\text{CE Policy}(t) = \begin{cases} \text{hire}, & \text{if } \theta_t \ge p_t^{\text{ce}} \\ \text{reject}, & \text{if } \theta_t < p_t^{\text{ce}} \end{cases}$$

Conservatism wrt Gaps

CE fails in the case of distr. with gaps.

Conservatism Principle

If the CE threshold is *close* to a gap, use the gap as a threshold.

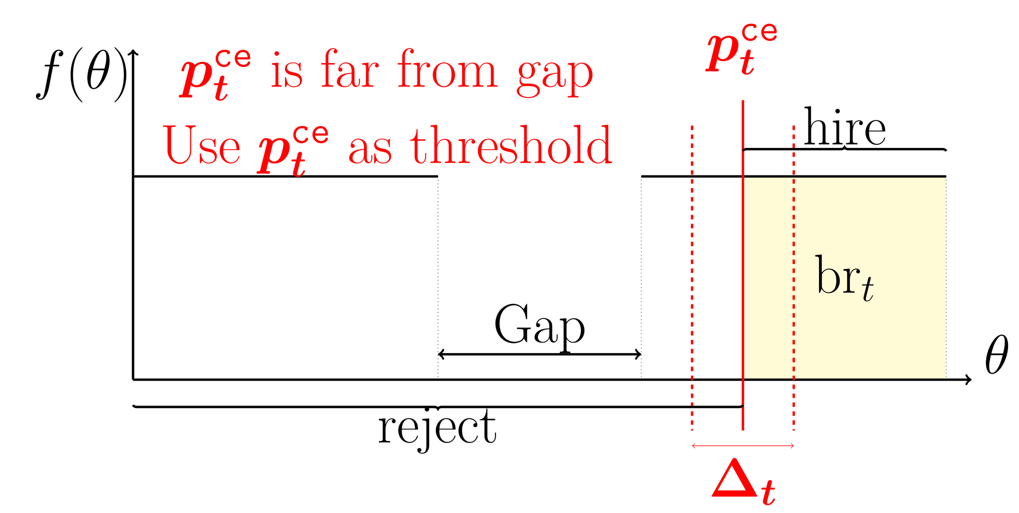

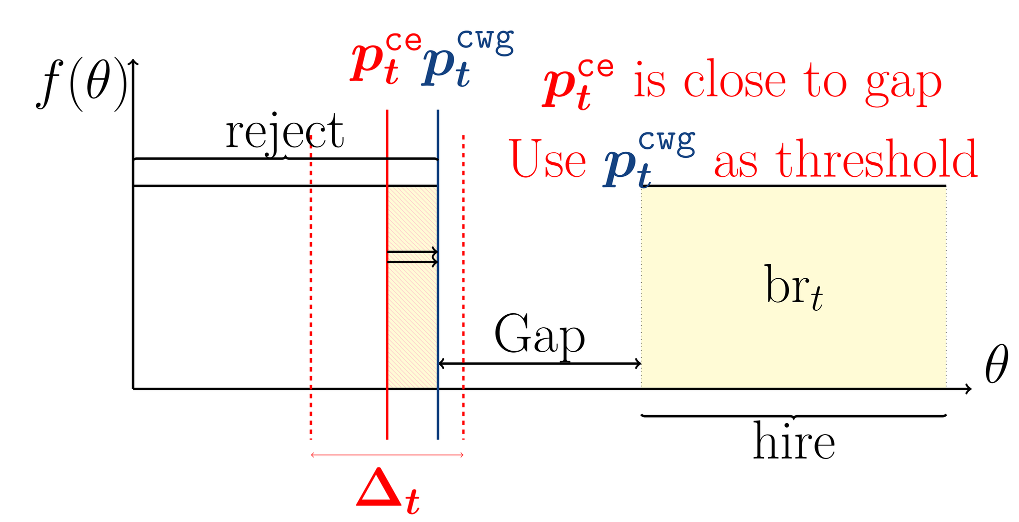

Failure of CE Policy For Many Types w/ Gaps i.e CE incurs large regret

For the CE policy, there exists a distribution F such that $\textbf{Regret}(\boldsymbol{B},\boldsymbol{T};\textbf{CE}) = \boldsymbol{\Omega}(\sqrt{\boldsymbol{T}})$.

Universal Lower Bound i.e the best any online policy can do

Consider any $\beta \in [0,\infty)$ and $\varepsilon_0 \le 1/2$. Then there exists a distribution F_{β,ε_0} and a budget B such that

$$\textbf{Regret}(\boldsymbol{B},\boldsymbol{T};\boldsymbol{\pi}) = \boldsymbol{\Omega}\left(\boldsymbol{T}^{1/2-1/(2(1+\beta))} \cdot \mathbb{I}\{\beta > 0\} + \log \boldsymbol{T} \cdot \mathbb{I}\{\beta = 0\}\right)$$

CwG Policy is near-optimal

For any $\beta \in [0,\infty)$ and $\varepsilon_0 \in (0,1]$, suppose the distribution F with associated gaps is (β,ε_0)-clustered. Then for all $T \in \mathbb{N}$ and for all $B \in [T]$, the regret of our CwG policy scales as

$$\textbf{Regret}(\boldsymbol{B},\boldsymbol{T};\textbf{CwG}) = \tilde{\boldsymbol{O}}\left(\boldsymbol{T}^{1/2-1/(2(1+\beta))} \cdot \mathbb{I}\{\beta > 0\} + (\log \boldsymbol{T})^2 \cdot \mathbb{I}\{\beta = 0\}\right)$$

Corollary: If the distribution has a (small) discrete support, $\textbf{Regret}(\boldsymbol{B},\boldsymbol{T};\textbf{CwG}) \le \boldsymbol{C}\sqrt{\log(1/\varepsilon_0)}/\varepsilon_0$

Numerical Simulations

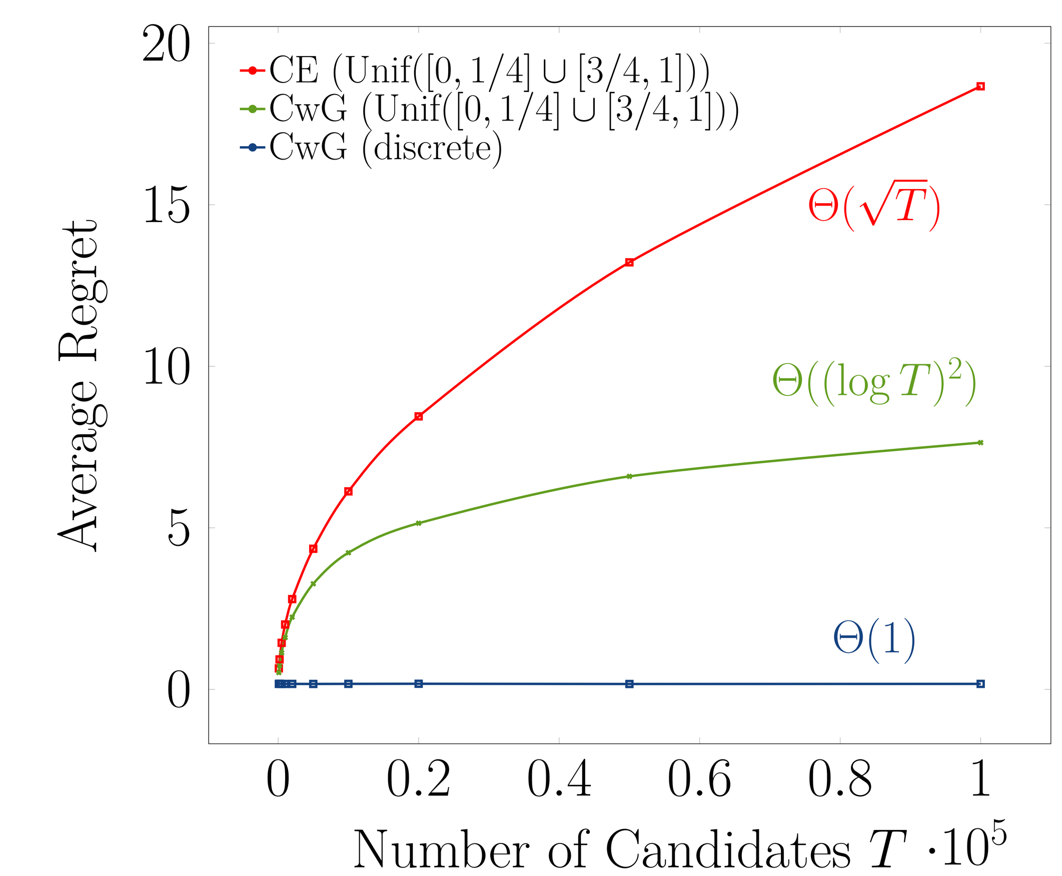

Contributions

- *Analytical*: We introduce the class of (β,ε_0)-clustered distributions which subsume previously considered distributions. Identify β as a key driver of the regret scaling. β also captures the *hardness* of the problem.
- *Algorithmic*: Devise a new algorithmic principle called *Conservatism wrt Gaps* to deal with distribution which have gaps and achieve near optimal performance.
- *Extensions*: Our results also extend to the setting with many small types which are relevant to other NRM problems like order fulfillment.

References

[1] Alessandro Arlotto and Itai Gurvich.
Uniformly bounded regret in the multisecretary problem.
Stochastic Systems, 9(3):231–260, 2019.

[2] Robert Bray.
Does the multisecretary problem always have bounded regret?
Available at SSRN 3497056, 2019.

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