Simulation Is All You Need

A Unifying algorithm for Network Revenue Management (NRM) and Dynamic Spatial Matching (DSM)

by Yash Kanoria (Columbia)
on

» Collaborators


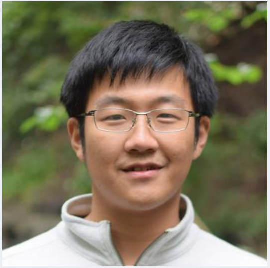

Omar Besbes Columbia DRO	Yilun Chen CUHK Shenzhen	Akshit Kumar Columbia DRO	Wenxin Zhang Columbia DRO

Network Revenue Management: Online Allocation with Resource Constraints

Network Revenue Management: Online Allocation with Resource Constraints

Yield Management

American Airlines

DELTA

UNITED

KLM Royal Dutch Airlines

Network Revenue Management: Online Allocation with Resource Constraints

Yield Management

American Airlines

DELTA

UNITED

KLM Royal Dutch Airlines

Budget Management in Ad Auctions

Google AdWords

Bing ads

Order Fulfillment

Flipkart

amazon.com

Walmart

AliExpress

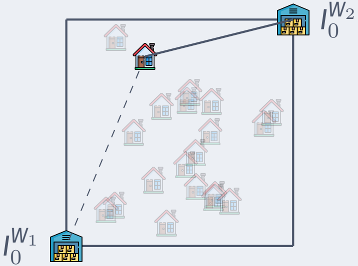

Landscape of NRM Problems

» Landscape of NRM Problems

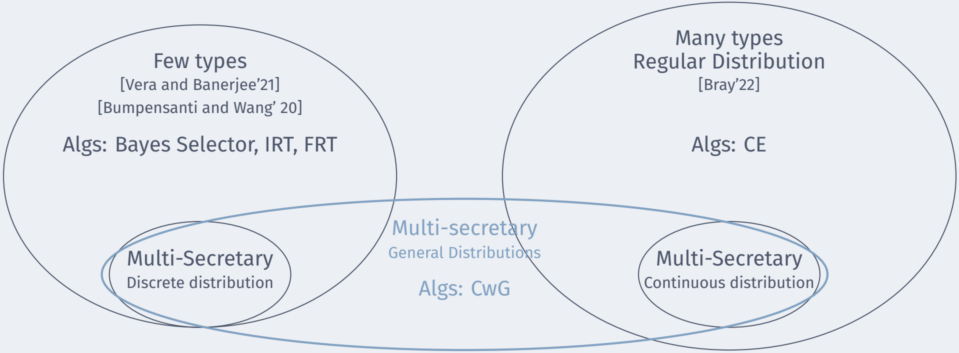

Landscape of NRM Problems

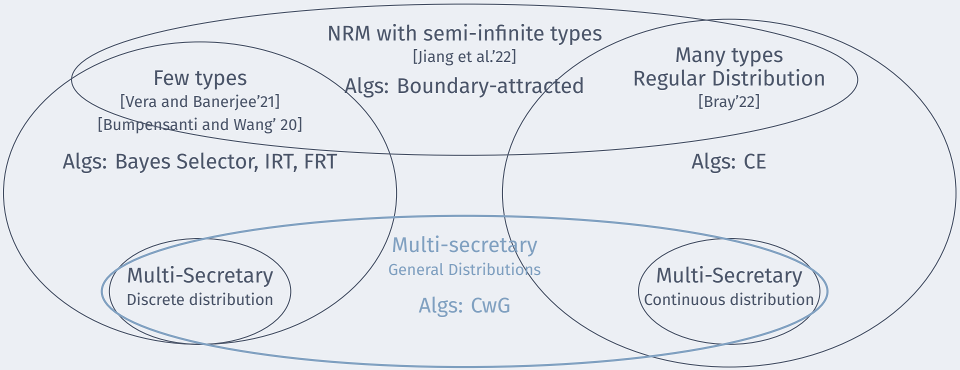

Landscape of NRM Problems

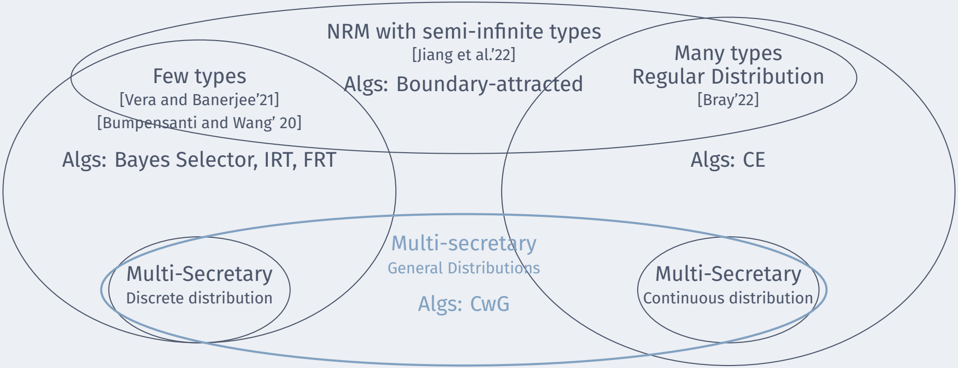

Research Question: One policy to solve them all?

» One Policy to solve them all and with a simulator we bind them

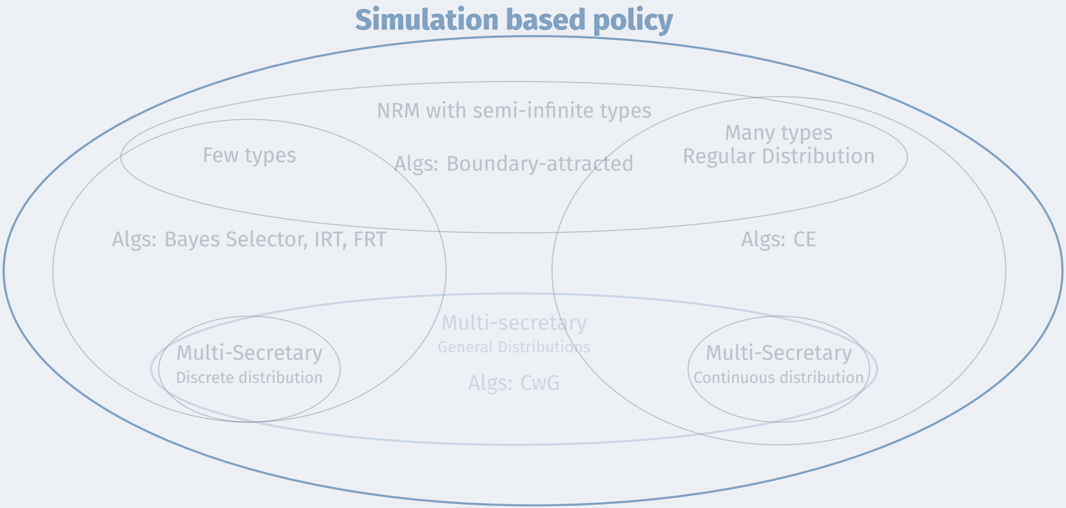

One Policy to solve them all and with a simulator we bind them

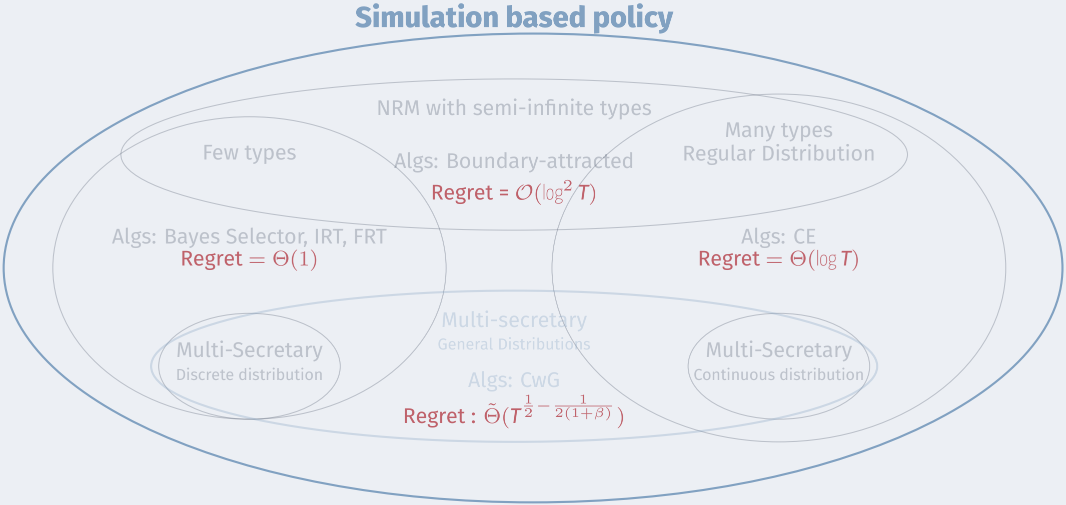

» One Policy to solve them all and with a simulator we bind them

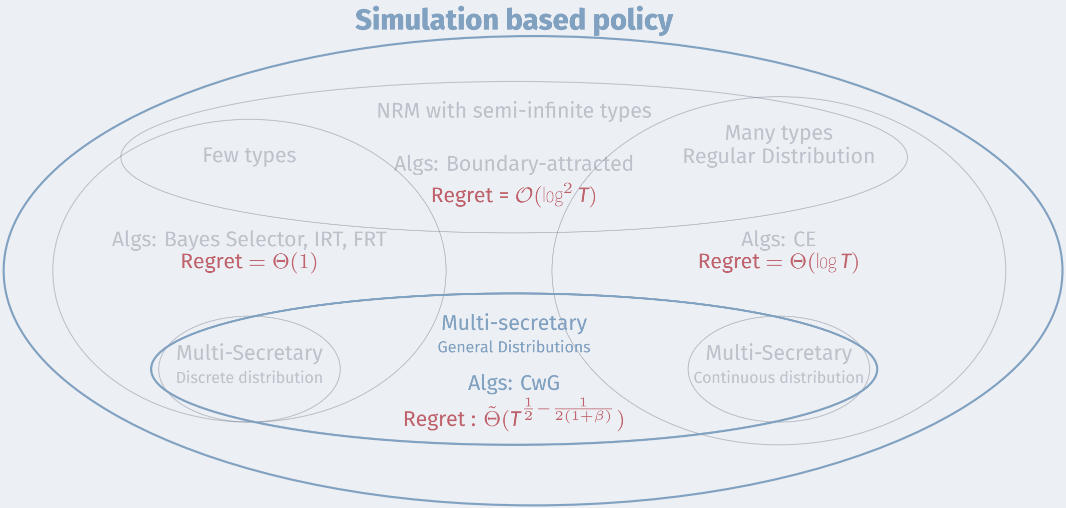

» Dynamic Spatial Matching (DSM): Motivation

- Ridehailing: spatial matching in two dimensions
- Matching platforms
  - Lodging e.g. Airbnb: supply and demand live in a multi-dimensional space (location, size, amenities, price, etc.)
  - Labor e.g. Upwork: (expertise dimensions, price, duration, etc.)
- Network revenue management with a large number of demand types

» DSM: Setting and Research Questions

* Supply and demand which live in d dimensional space.
* Cost of match distance between the matched pair.
* T supply units are present beforehand.
* Demand arrives sequentially. Needs to be matched immediately with a supply unit.

DSM: Setting and Research Questions

- Supply and demand which live in d dimensional space.
- Cost of match distance between the matched pair.
- T supply units are present beforehand.
- Demand arrives sequentially. Needs to be matched immediately with a supply unit.

- How to match demand and supply to minimize spatial costs of matching under dynamic arrivals?
- How large are the costs arising from spatial heterogeneity and uncertainty about the future in dynamic matching

Summary of findings and talk outline

- DSM with identical supply and demand distributions [K.]
  - Greedy matching suffices
  - Match distance $\sim$ Nearest-neighbor-distance achievable, except one case
- DSM with different supply and demand distributions [Chen, Akshit Kumar, K., Zhang]
  - Greedy fails
  - Simulate-Optimize-Assign-Repeat (SOAR) is near optimal
- Multisecretary problem with lumpy value distribution (a 1d DSM problem) [Besbes, Akshit Kumar, K.]
  - The Certainty Equivalent policy and SOAR with one sample path fail
  - RAMS with multiple sample paths achieves optimal regret scaling
  - RAMS works also for $d \geq 2$, and across NRM settings.

» Talk outline

- DSM with identical supply and demand distributions
  - Greedy matching suffices
  - Match distance $\sim$ Nearest-neighbor-distance achievable, except one case
- DSM with different supply and demand distributions
  - Greedy fails
  - Simulate-Optimize-Assign-Repeat (SOAR) is near optimal
- **Multisecretary problem with lumpy value distribution (a 1d DSM problem)** [O. Besbes, Akshit Kumar & K. '22]
  - SOAR with one sample path fails
  - RAMS with multiple sample paths achieves optimal regret scaling
  - Works also for $d \geq 2$, and across NRM settings.

» Multi-secretary Problem

Problem Statement

Given a sequence of T secretaries and a hiring budget B, a decision maker (DM) wants to hire the top B secretaries in terms of their ability.

Multi-secretary Problem

Problem Statement

Given a sequence of T secretaries and a hiring budget B, a decision maker (DM) wants to hire the top B secretaries in terms of their ability.

Details and Assumptions

* The secretaries arrive in an online fashion.

» Multi-secretary Problem

Problem Statement

Given a sequence of T secretaries and a hiring budget B, a decision maker (DM) wants to hire the top B secretaries in terms of their ability.

Details and Assumptions

* The secretaries arrive in an online fashion.
* The DM makes **irrevocable** *hire* or *reject* decisions.

» Multi-secretary Problem

Problem Statement

Given a sequence of T secretaries and a hiring budget B, a decision maker (DM) wants to hire the top B secretaries in terms of their ability.

Details and Assumptions

* The secretaries arrive in an online fashion.
* The DM makes **irrevocable** *hire* or *reject* decisions.
* The abilities (types) of the secretaries are drawn **independently** from a **common** and **known** distribution F over $[0, 1]$.

» Multi-secretary Problem

Problem Statement

Given a sequence of T secretaries and a hiring budget B, a decision maker (DM) wants to hire the top B secretaries in terms of their ability.

Details and Assumptions

* The secretaries arrive in an online fashion.
* The DM makes **irrevocable** *hire* or *reject* decisions.
* The abilities (types) of the secretaries are drawn **independently** from a **common** and **known** distribution F over $[0, 1]$.

Note: This is a 1d DSM problem, with an atomic "supply" distribution with B units at 1 and $T - B$ units at 0.

Gotta catch'em ~~all~~ some

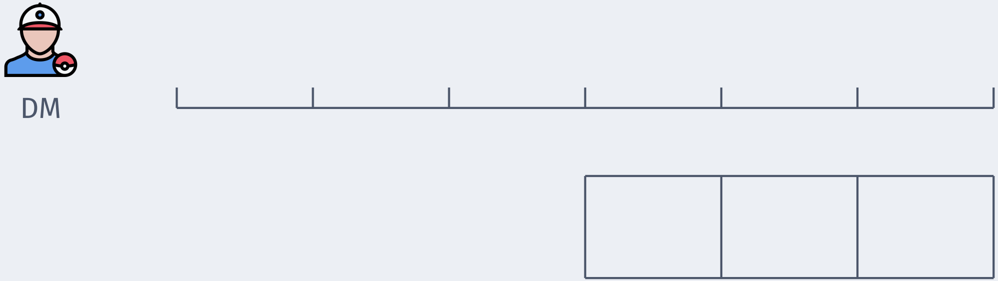

» Gotta catch'em ~~all~~ some

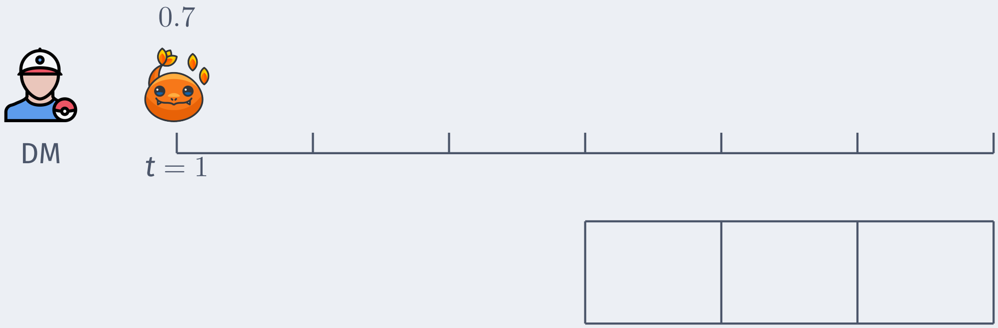

» Gotta catch'em ~~all~~ some

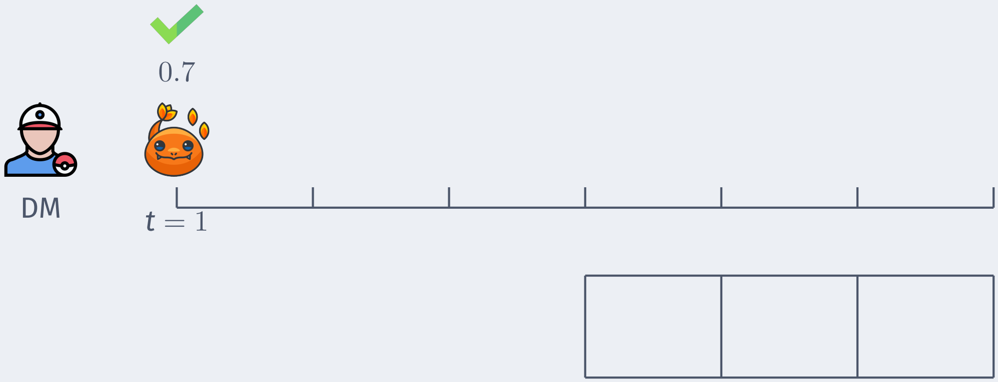

» Gotta catch'em ~~all~~ some

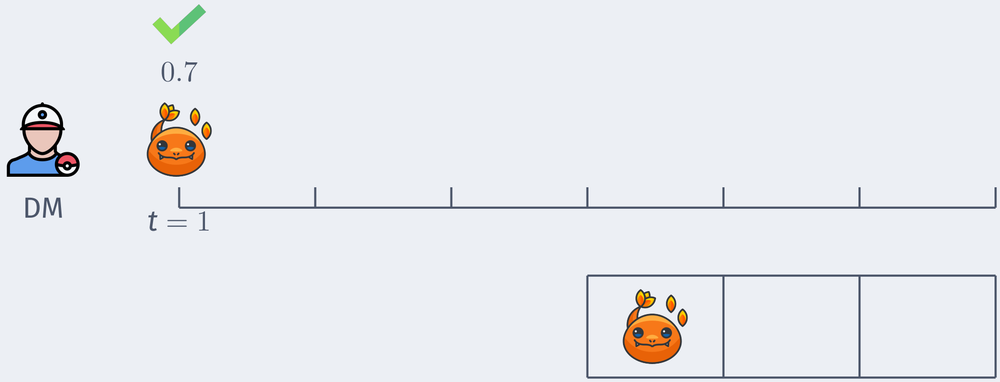

Gotta catch'em ~~all~~ some

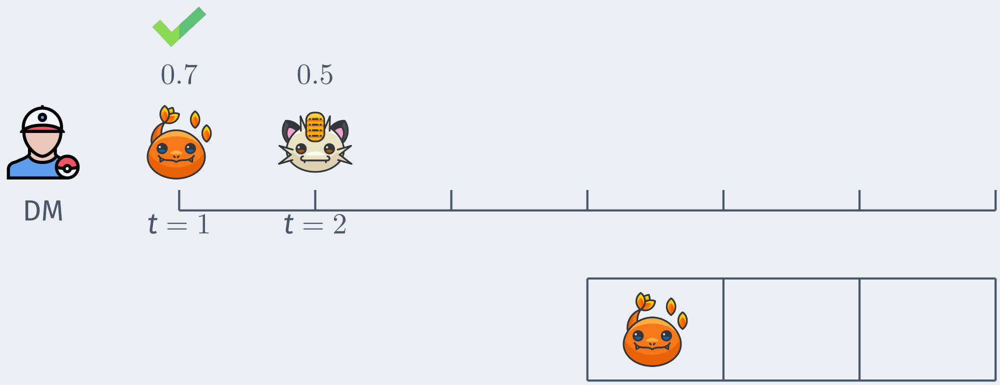

Gotta catch'em ~~all~~ some

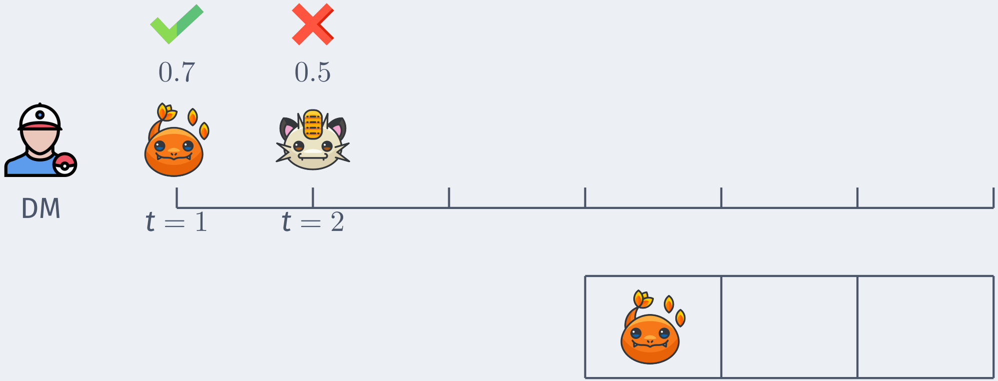

Gotta catch'em ~~all~~ some

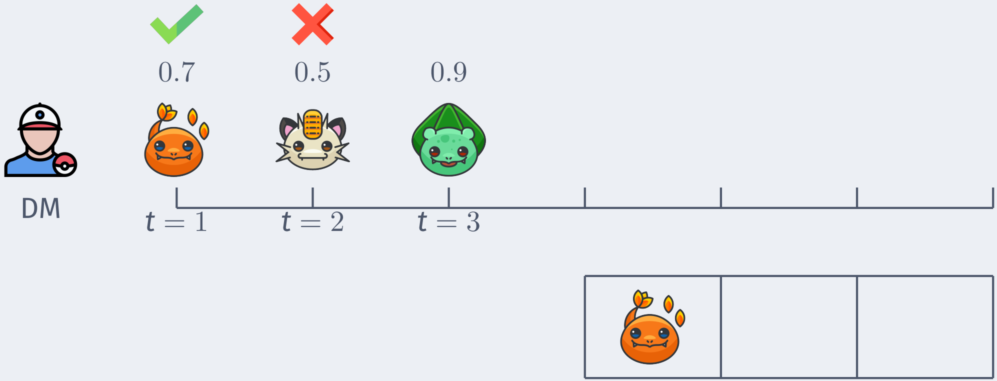

Gotta catch'em ~~all~~ some

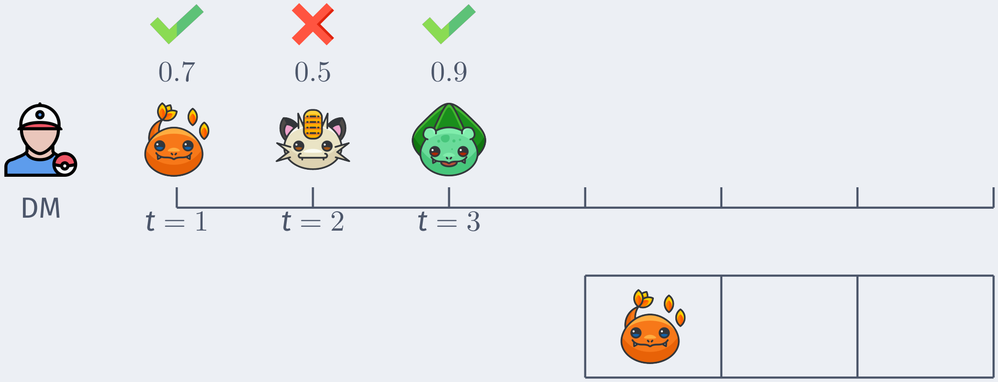

» Gotta catch'em ~~all~~ some

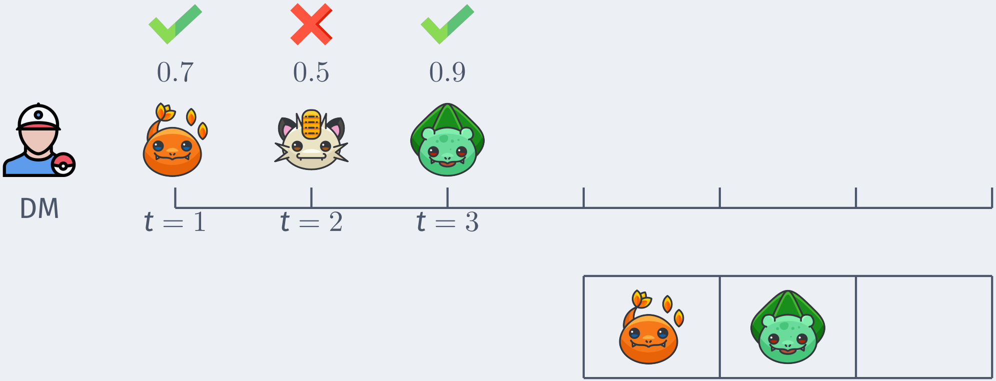

» Gotta catch'em ~~all~~ some

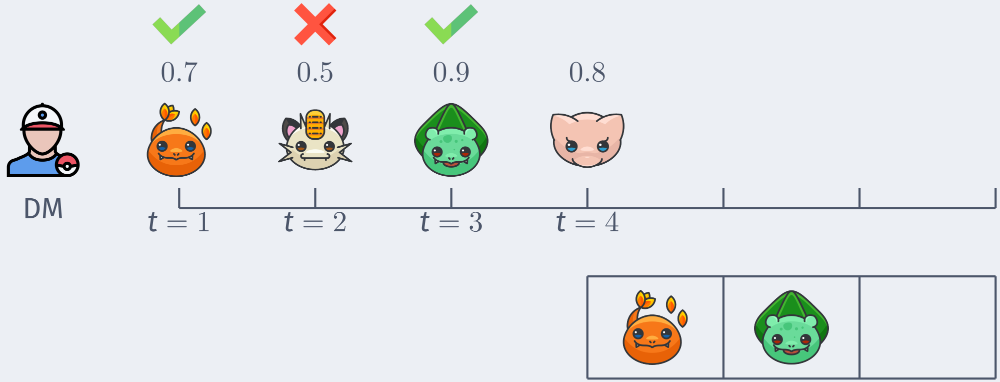

Gotta catch'em ~~all~~ some

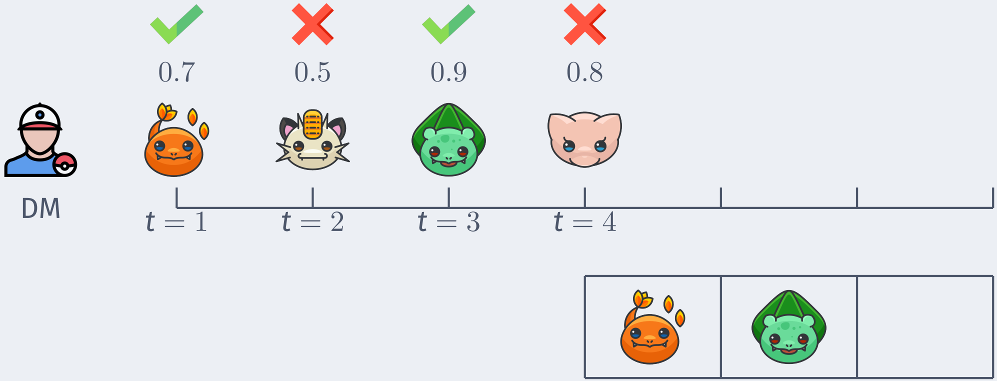

Gotta catch'em ~~all~~ some

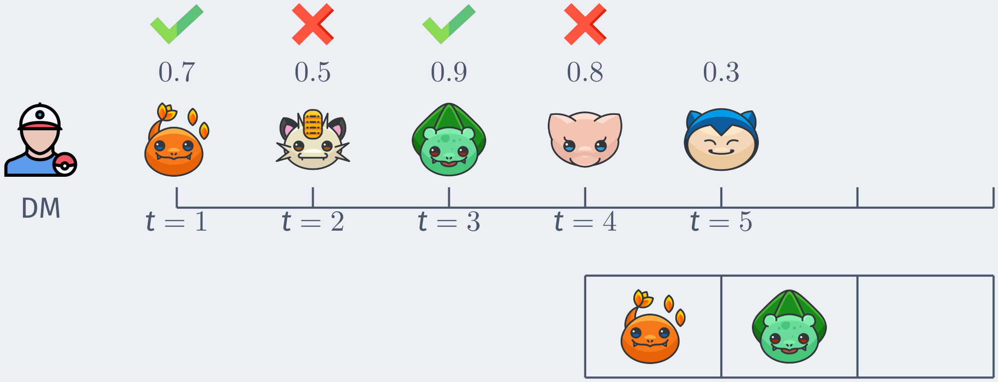

» Gotta catch'em ~~all~~ some

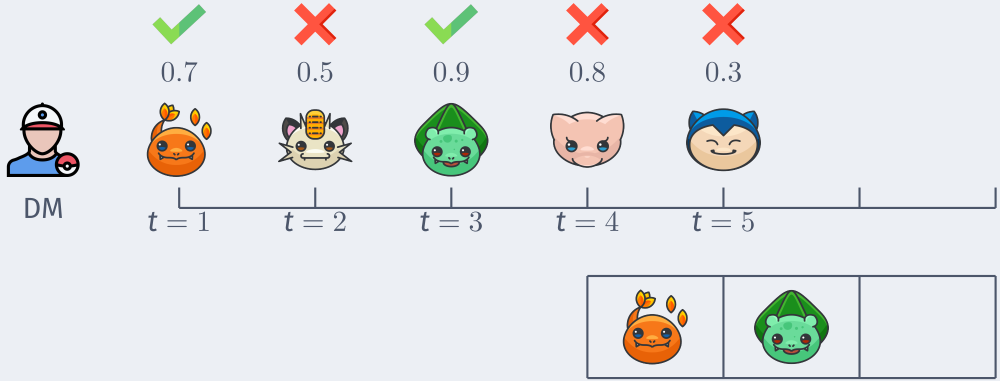

» Gotta catch'em ~~all~~ some

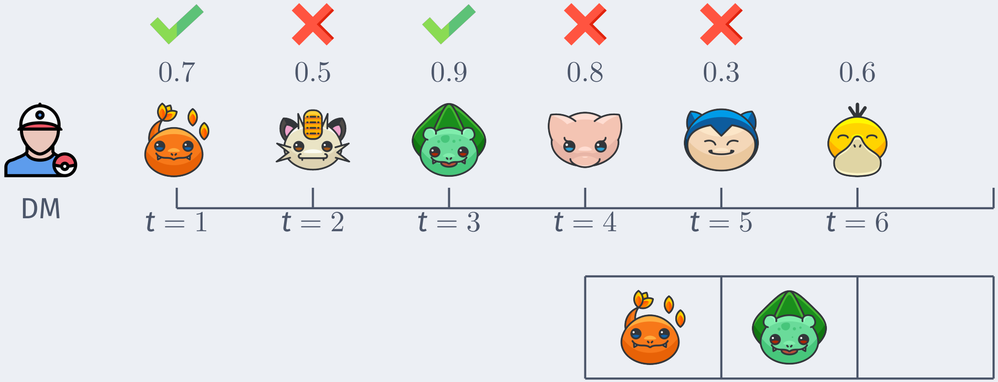

Gotta catch'em ~~all~~ some

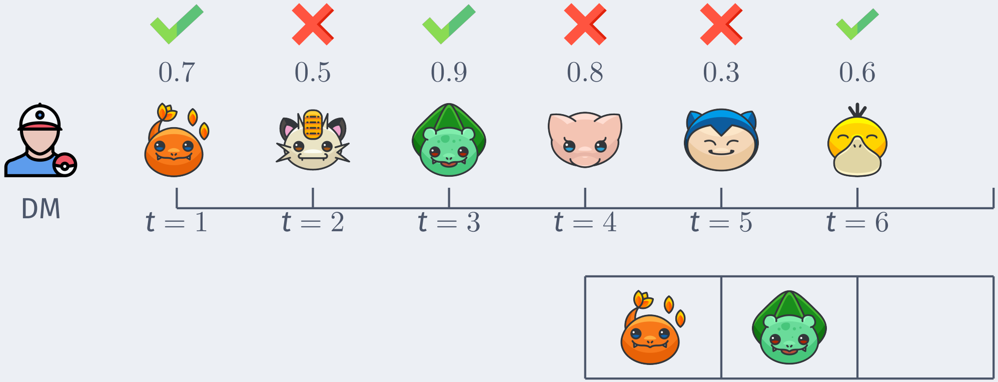

Gotta catch'em ~~all~~ some

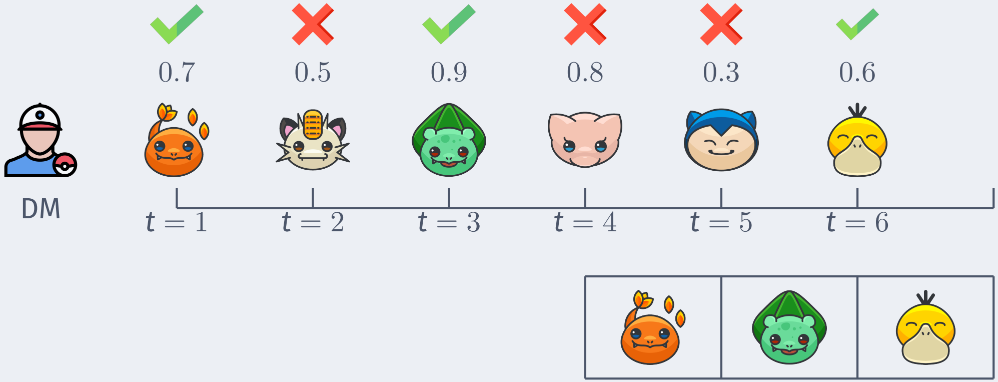

» Gotta catch'em ~~all~~ some

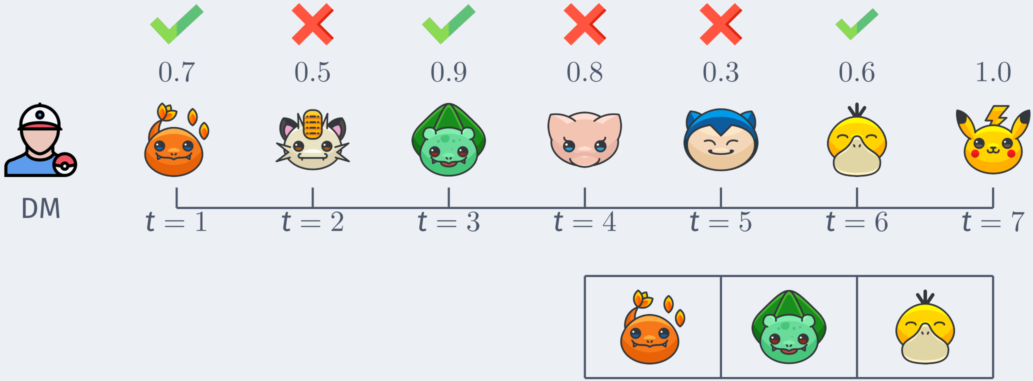

» Gotta catch'em ~~all~~ some

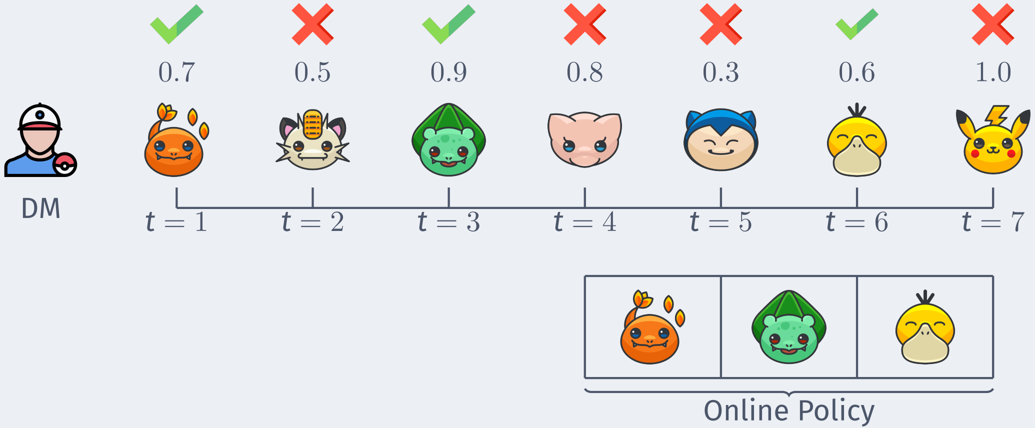

» Gotta catch'em ~~all~~ some

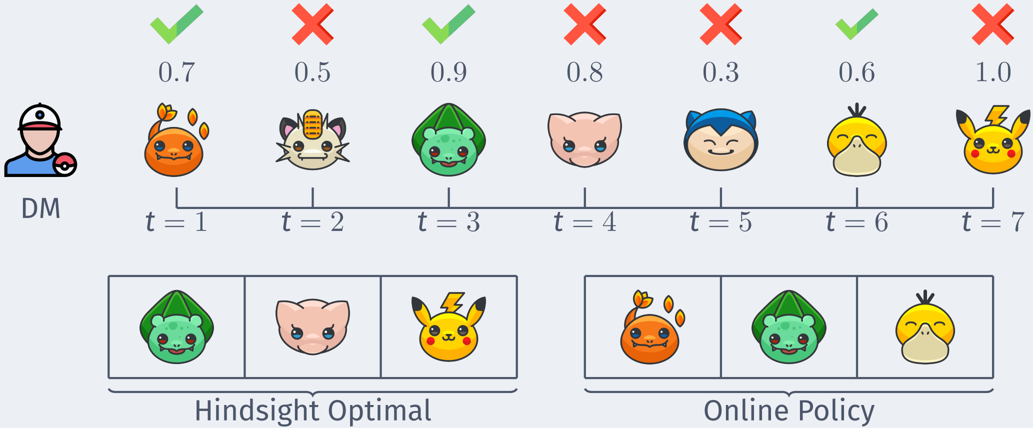

» Gotta catch'em ~~all~~ some

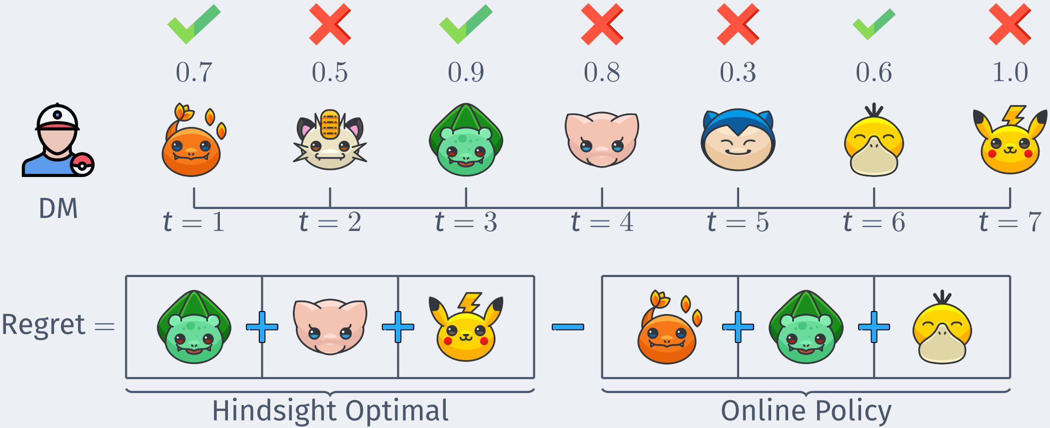

Gotta catch'em ~~all~~ some

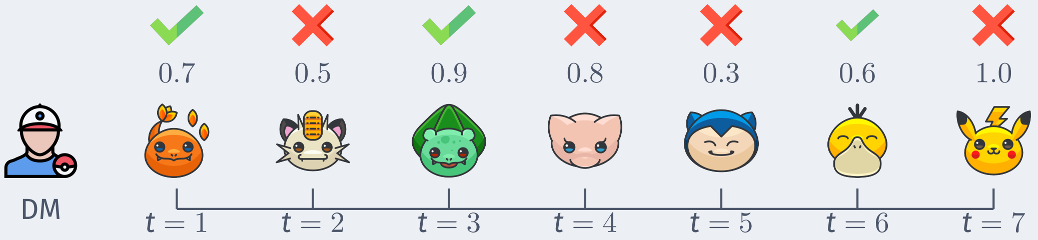

$$\text{Regret} = \underbrace{\boxed{0.9 + 0.8 + 1.0}}_{\text{Hindsight Optimal}} - \underbrace{\boxed{0.7 + 0.9 + 0.6}}_{\text{Online Policy}} = 0.5$$

» Gotta catch'em ~~all~~ some

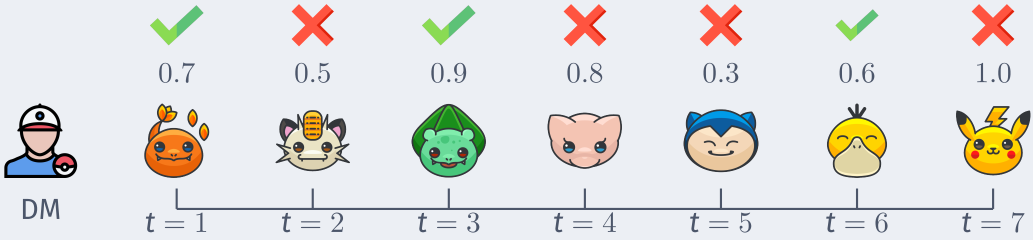

$$\text{Regret} = \underbrace{\boxed{0.9 + 0.8 + 1.0}}_{\text{Hindsight Optimal}} - \underbrace{\boxed{0.7 + 0.9 + 0.6}}_{\text{Online Policy}} = 0.5$$

We will focus on $\mathbb{E}\left[\text{Regret}\right]$

» Hindsight-based regret

Multisecretary is a 1d DSM problem, with an atomic "supply" distribution with B units at 1 and $T - B$ units at 0. $\Theta(\sqrt{T})$ optimal regret wrt fluid benchmark, which can be achieved by a trivial static policy.

Gap between fluid and hindsight benchmarks is already $\Omega(\sqrt{T})$.

As in the recent NRM literature, we adopt the tighter hindsight benchmark.

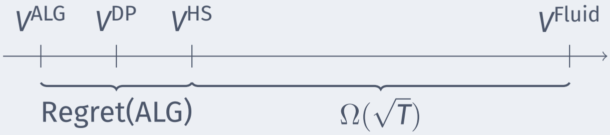

Multi-secretary Problem: What's known and research question?

Multi-secretary Problem: What's known and research question?

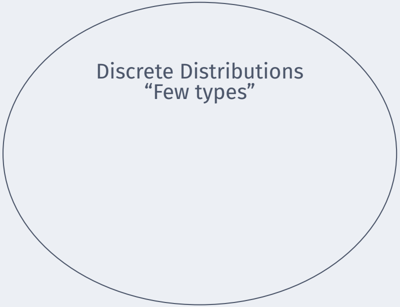

Multi-secretary Problem: What's known and research question?

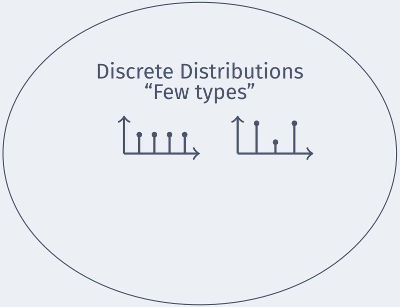

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$

[Arlotto and Gurvich'19]

Regret can scale with # of types

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]
Regret can scale with # of types

Continuous Distributions
"Many Types"

Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$

[Arlotto and Gurvich'19]

Regret can scale with # of types

Continuous Distributions
"Many Types"

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]
Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$
[Leuker'98, Bray'22]

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]
Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$
[Leuker'98, Bray'22]
Need density to be bounded below

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$

[Arlotto and Gurvich'19]

Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$

[Leuker'98, Bray'22]

Need density to be bounded below

Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]
Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$
[Leuker'98, Bray'22]
Need density to be bounded below
Do not admit distributions with gaps

» Multi-secretary Problem: What's known and research question?

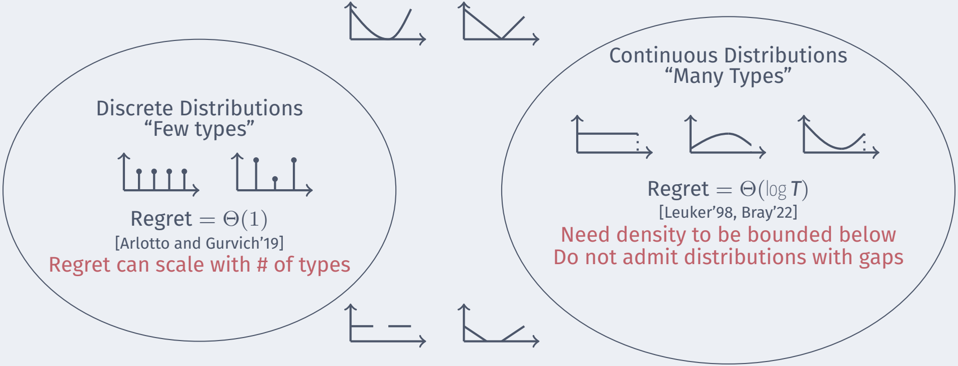

» Multi-secretary Problem: What's known and research question?

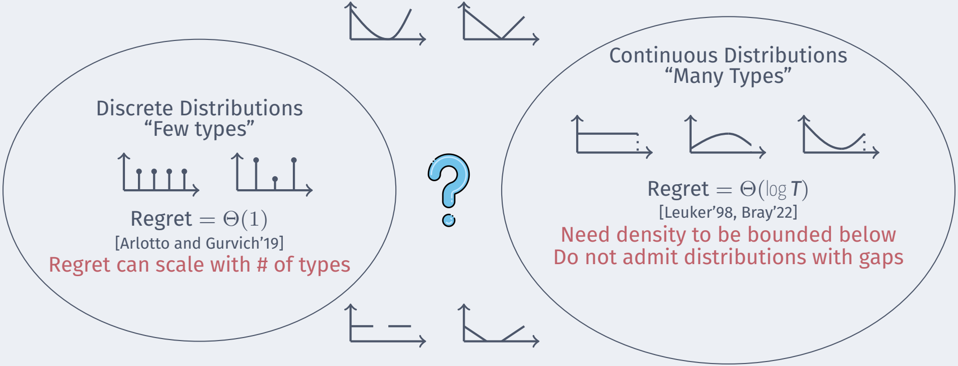

How to interpolate between these distribution classes?
How to deal with gaps? What are the possible regret scalings?

» Punchline for the Multi-secretary Problem

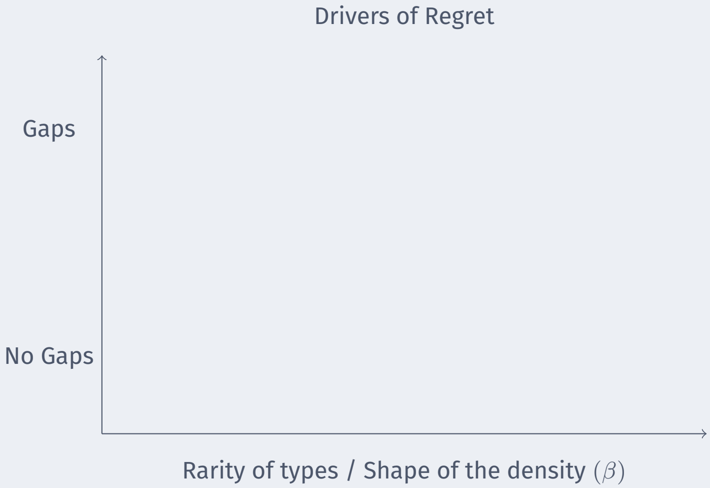

Punchline for the Multi-secretary Problem

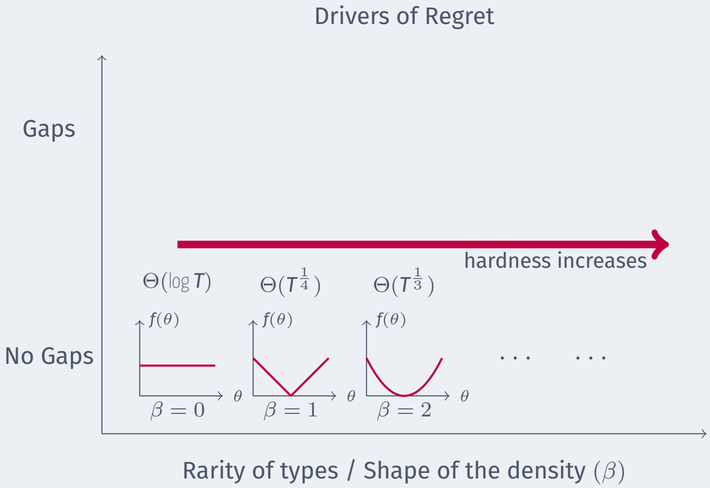

Punchline for the Multi-secretary Problem

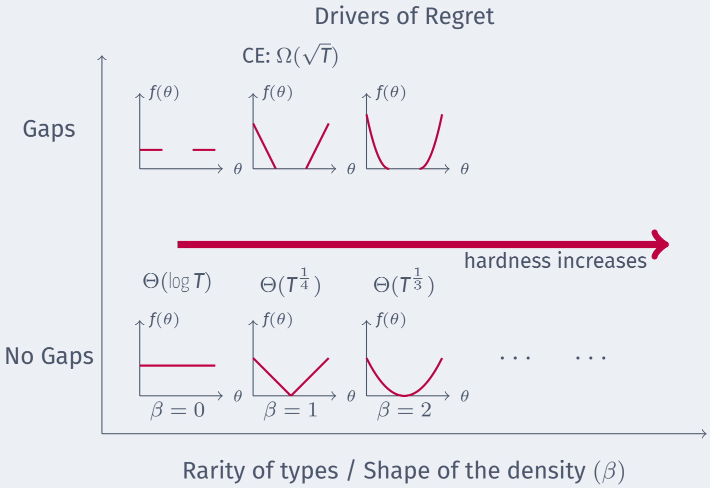

» Punchline for the Multi-secretary Problem

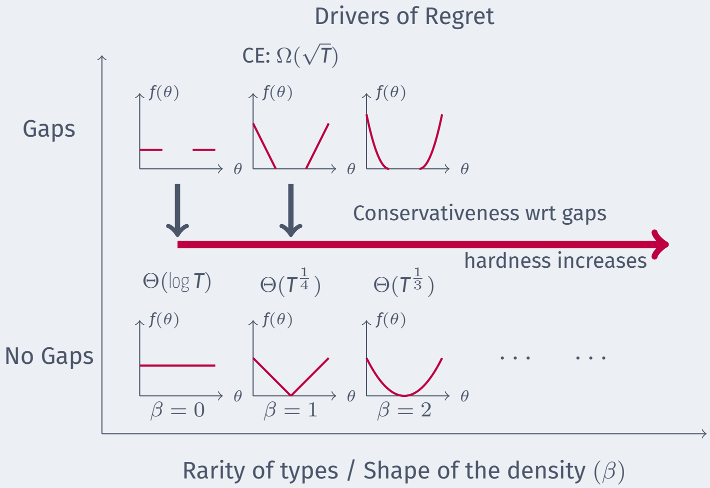

» Punchline for the Multi-secretary Problem

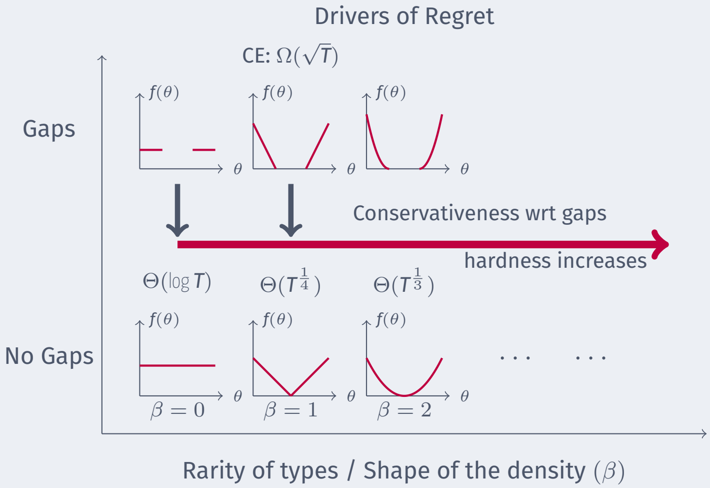

- Distribution shape is a **fundamental** driver of regret.
- Dealing with gaps is an algorithmic challenge.
- Novel Principle: Conservativeness wrt gaps (CwG)
- Simulation-based approach automatically pursues CwG

» Towards unification of the multi-secretary problem

Discrete Distributions
"Few types"

Regret $= \Theta(1)$

[Arlotto and Gurvich'19]

Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$

[Bray'22, Leuker'98]

Need density to be bounded below

Do not admit distributions with gaps

» Towards unification of the multi-secretary problem

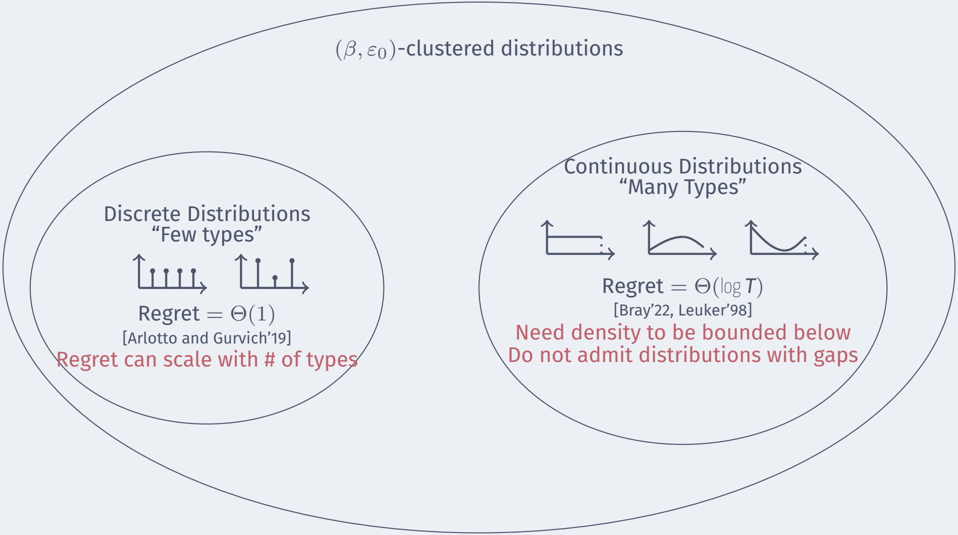

» Fundamental Limits

Universal Lower Bound

For every $\beta \in [0, \infty)$, there exists a distribution F_β such that

$$\sup_{B \in [T]} \mathbb{E}_{F_\beta} [\text{Regret(DP)}] = \begin{cases} \Omega(\log T), & \beta = 0, \\ \Omega\left(T^{\frac{1}{2} - \frac{1}{2(1+\beta)}}\right), & \beta > 0. \end{cases}$$

Fundamental Limits

Universal Lower Bound

For every $\beta \in [0, \infty)$, there exists a distribution F_β such that

$$\sup_{B \in [T]} \mathbb{E}_{F_\beta}\left[\text{Regret(DP)}\right] = \begin{cases} \Omega\left(\log T\right), & \beta = 0, \\ \Omega\left(T^{\frac{1}{2} - \frac{1}{2(1+\beta)}}\right), & \beta > 0. \end{cases}$$

$\Omega(\log T)$ $\Omega(T^{\frac{1}{4}})$ $\Omega(T^{\frac{1}{3}})$

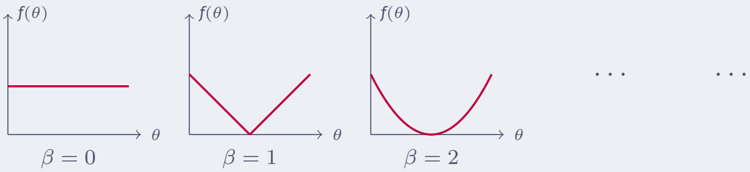

» Certainty Equivalent Control

- Let B_t be the remaining budget at time t
- Compute the budget ratio
  $br_t = \frac{\text{Remaining Budget}}{\text{Remaining Time}} = \frac{B_t}{T-t}$
- Define a quantile threshold $p_t^{ce} = 1 - br_t$
- Define a ability threshold $\gamma_t^{ce} = F^{-1}(p_t^{ce})$
- hire $\iff \theta_t \geq \gamma_t^{ce}$

For $(\beta, 1)$-clustered distributions

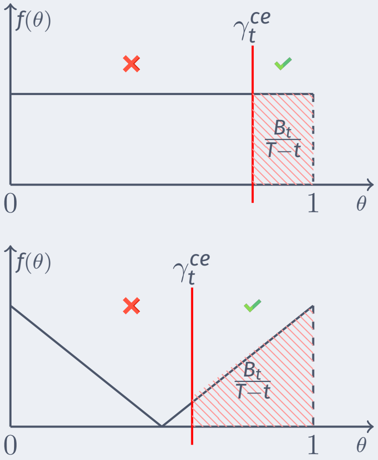

» Certainty Equivalent Control

For Bi-modal Uniform Distribution

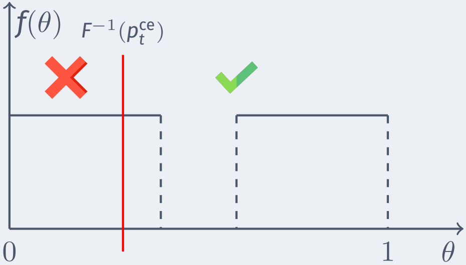

Let B_t be the remaining budget at time t

$$\text{Budget Ratio} = \frac{\text{Remaining Budget}}{\text{Remaining Time}} = \frac{B_t}{T-t}$$

$$\text{CE Quantile Threshold} = 1 - \frac{B_t}{T-t} \triangleq p_t^{ce}$$

$$\text{Decision: hire} \iff \theta_t \geq F^{-1}(p_t^{ce})$$

» Failure of Certainty Equivalent Control

Regret Lower Bound

Insufficiency of Certainty Equivalent Control

Assume that $F = \text{Unif}([0, \frac{1}{4}] \cup [\frac{3}{4}, 1])$, for $B = T/2$, we have

$$\mathbb{E}\left[\text{Regret}(CE)\right] = \Omega\left(\sqrt{T}\right)$$

» Failure of Certainty Equivalent Control

Regret Lower Bound

Insufficiency of Certainty Equivalent Control

Assume that $F = \text{Unif}([0, \frac{1}{4}] \cup [\frac{3}{4}, 1])$, for $B = T/2$, we have

$$\mathbb{E}\left[\text{Regret}(CE)\right] = \Omega\left(\sqrt{T}\right)$$

Remark

* Same scaling is achievable under a static threshold policy.

Why does CE fail?

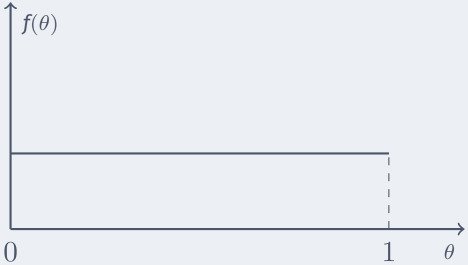

» Why does CE fail?

$$B/\tau = \frac{1}{2} + \epsilon$$

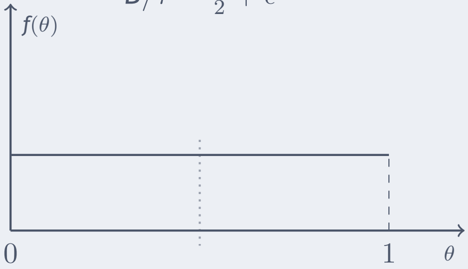

Why does CE fail?

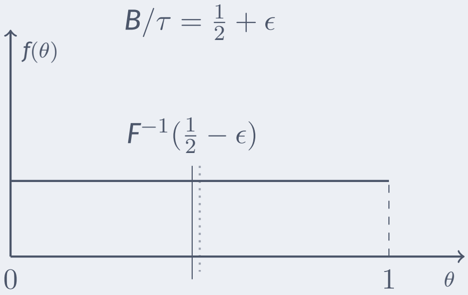

Why does CE fail?

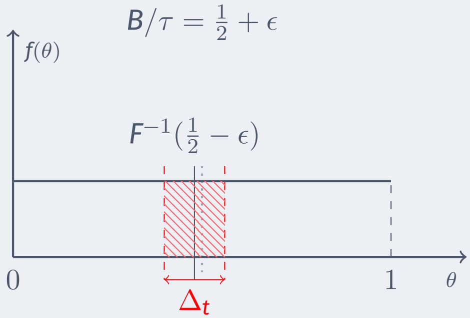

Why does CE fail?

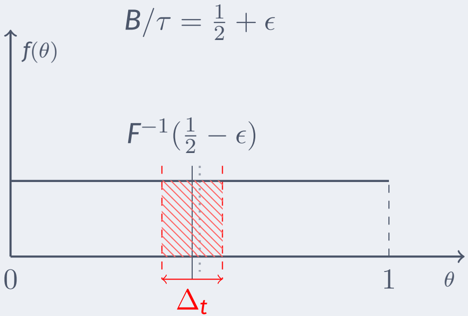

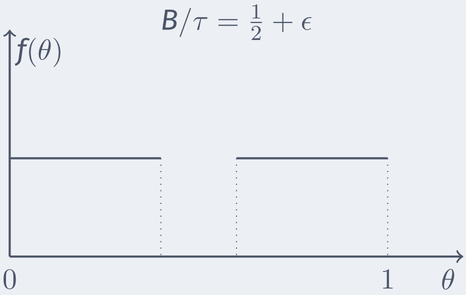

» Why does CE fail?

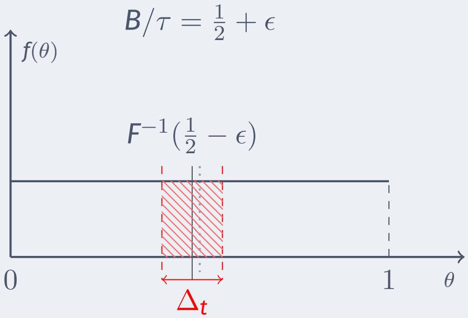

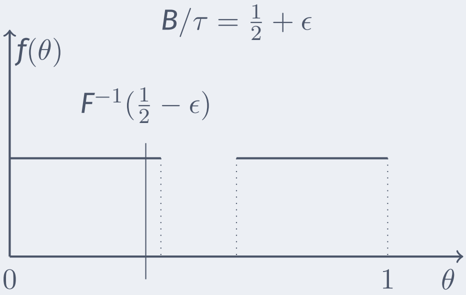

Why does CE fail?

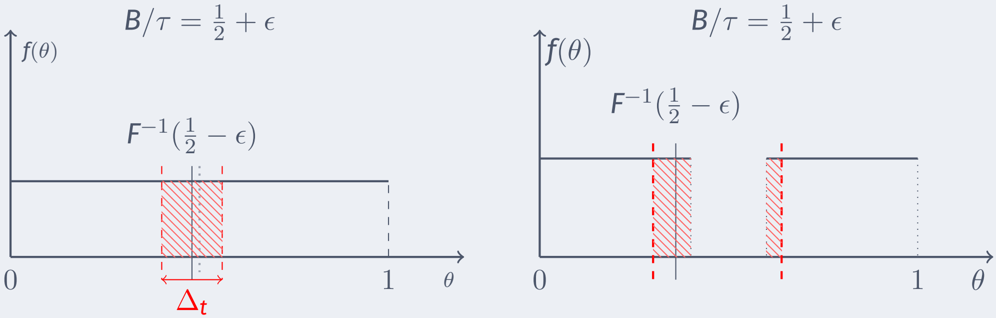

Why does CE fail?

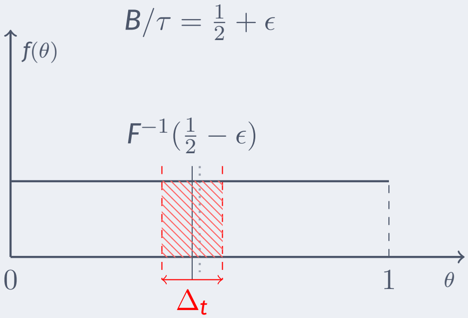

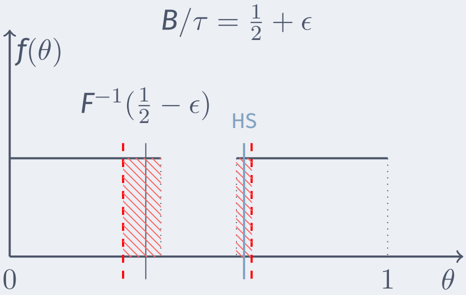

» Why does CE fail?

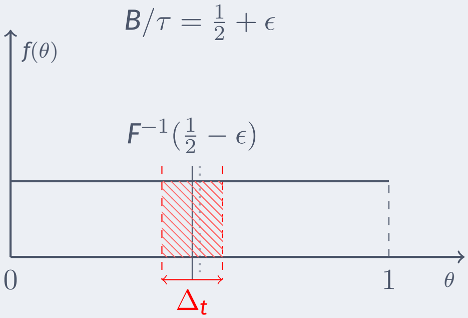

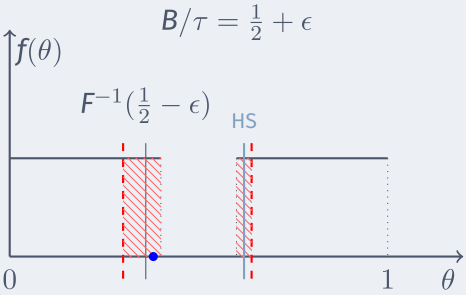

» Why does CE fail?

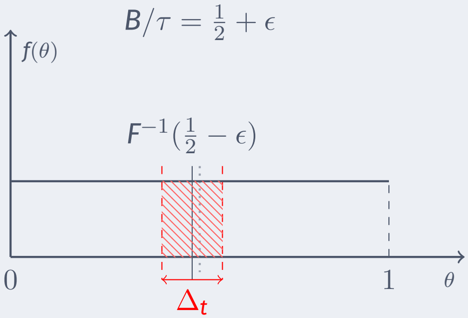

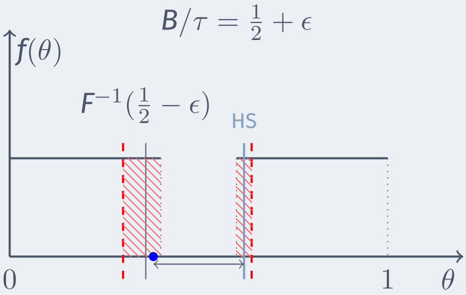

» Why does CE fail?

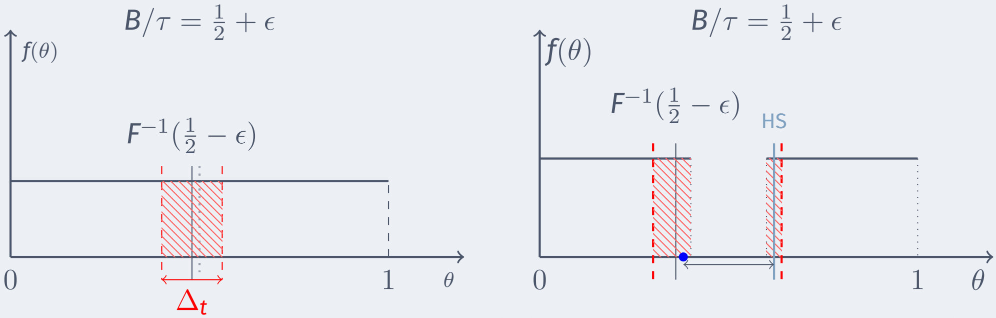

This is $\Omega(1/\sqrt{\tau})$ expected compensation

Why does CE fail?

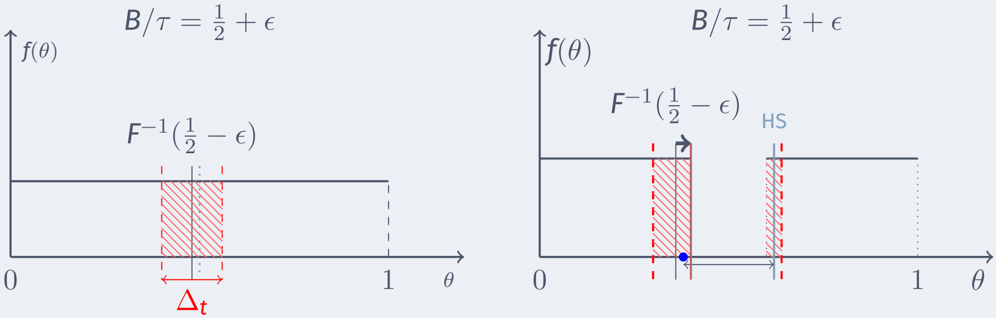

This is $\Omega(1/\sqrt{\tau})$ expected compensation

» Why does CE fail?

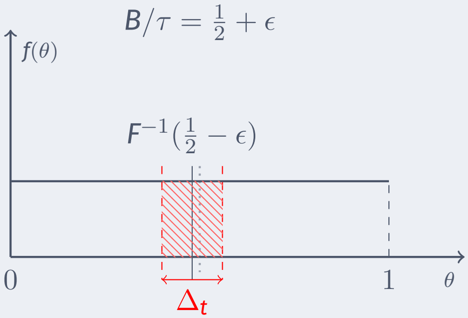

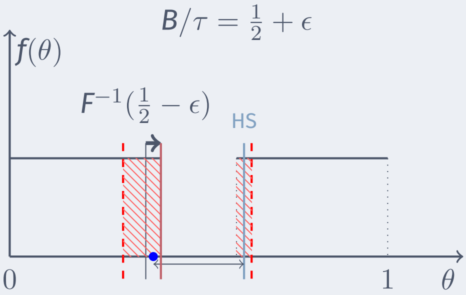

This is $\Omega(1/\sqrt{\tau})$ expected compensation

Conservativeness wrt gaps

» Why does CE fail?

What if?

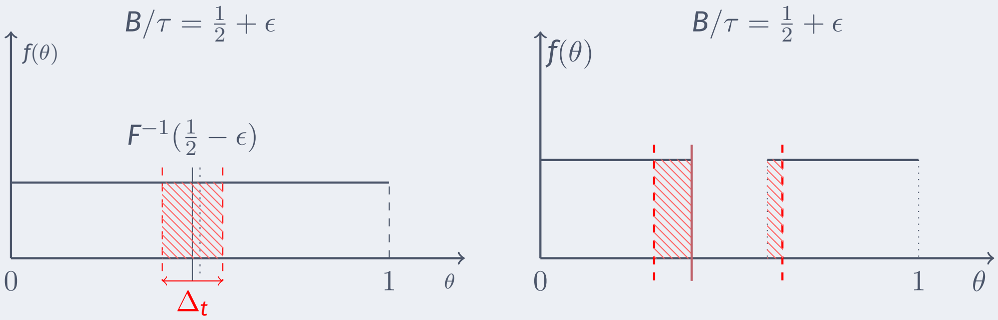

» Why does CE fail?

What if?

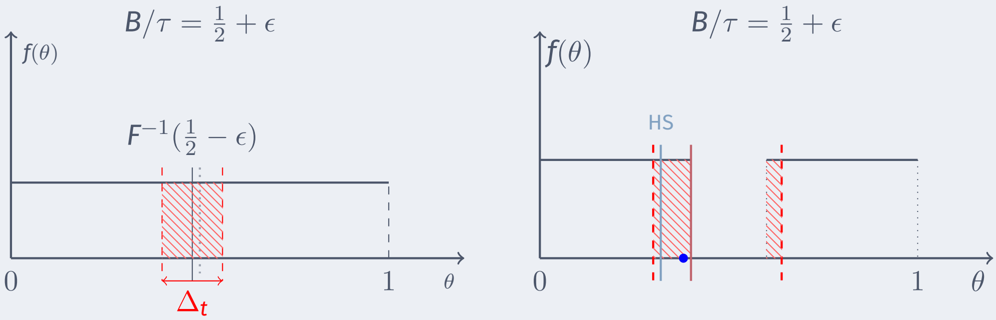

» Why does CE fail?

What if?

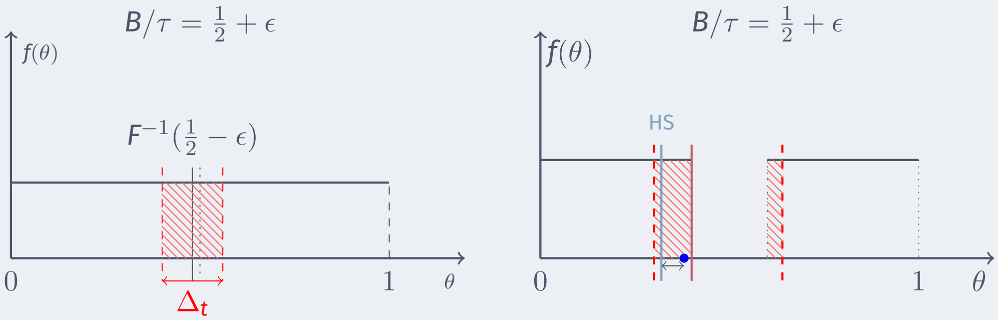

This is $\tilde{\mathcal{O}}(1/\tau)$ expected compensation

» Good in theory but practically infeasible

* What is the conservativeness parameter I should use?
* How to find where these gaps are? What happens if gaps shift?
* E.g., no chance of deploying for Amazon's fulfillment problem

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

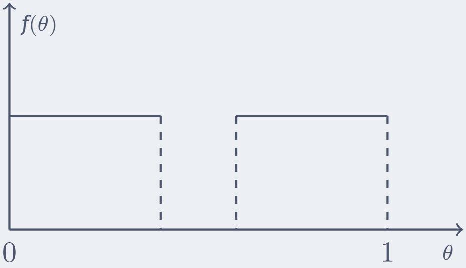

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

$B/\tau = 1/4$

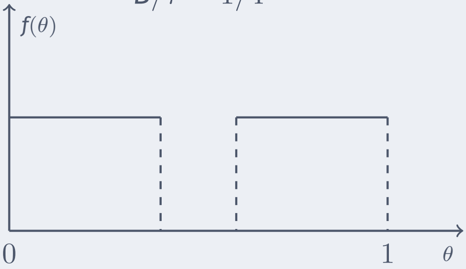

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

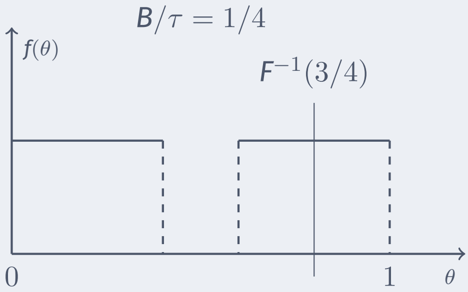

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

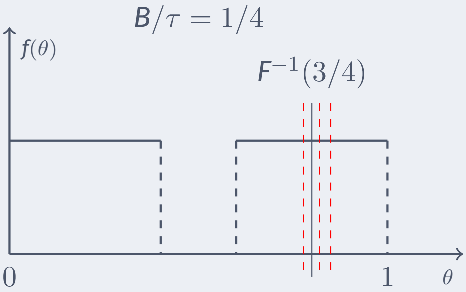

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

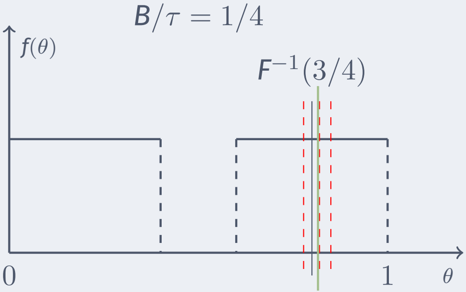

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

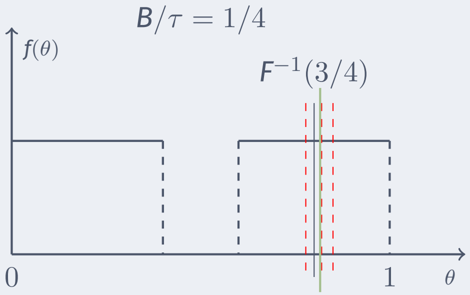

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

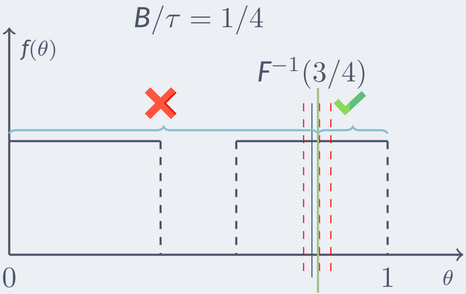

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

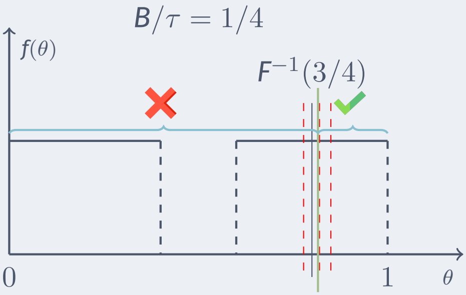

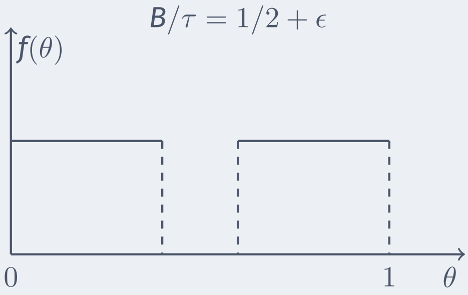

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

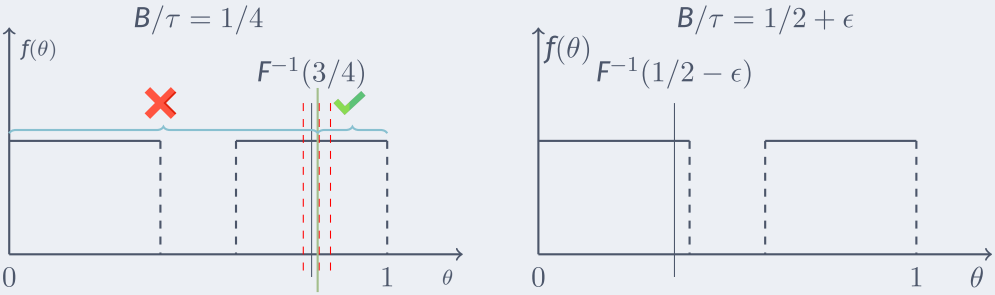

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

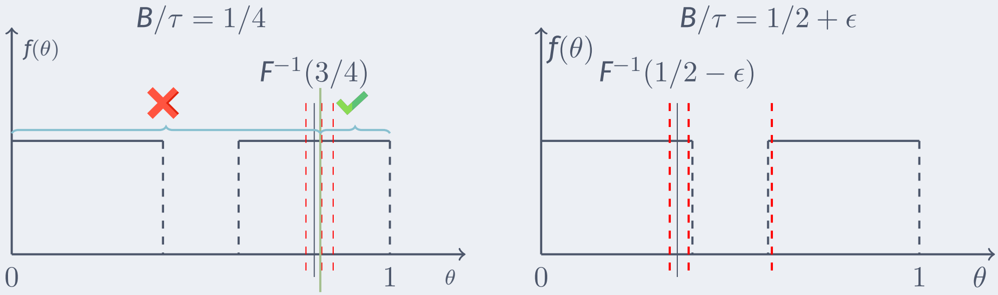

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

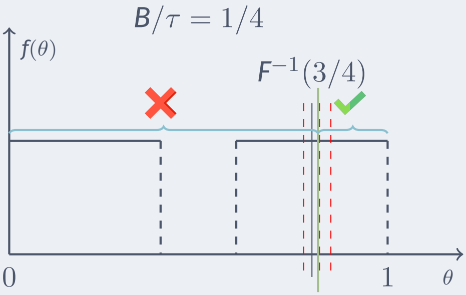

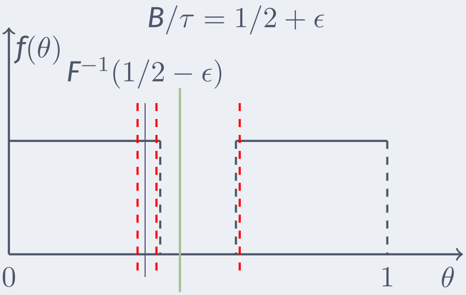

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

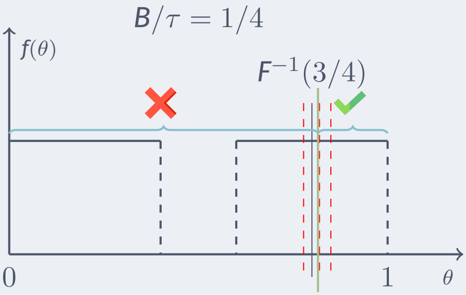

If far from a gap, use the CE threshold

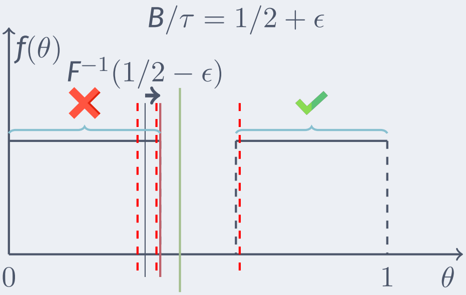

If close to gap, use the gap as threshold

» Conservativeness with respect to gaps

Punchline

Regret of RAMS Policy

If F is a (β, ε_0)-clustered distribution, then

$$\mathbb{E}\left[\text{Regret(RAMS)}\right] = \begin{cases} \mathcal{O}\left((\log T)^2\right), & \beta = 0, \\ \mathcal{O}\left(\text{poly}(\log T) T^{\frac{1}{2} - \frac{1}{2(1+\beta)}}\right), & \beta > 0 \end{cases}$$

If F is a discrete distribution, $\mathbb{E}\left[\text{Regret (RAMS)}\right] = \mathcal{O}(1/\varepsilon_0)$

» Conservativeness with respect to gaps

Punchline

Regret of RAMS Policy

If F is a (β, ε_0)-clustered distribution, then

$$\mathbb{E}\left[\text{Regret(RAMS)}\right] = \begin{cases} \mathcal{O}\left((\log T)^2\right), & \beta = 0, \\ \mathcal{O}\left(\text{poly}(\log T) T^{\frac{1}{2} - \frac{1}{2(1+\beta)}}\right), & \beta > 0 \end{cases}$$

If F is a discrete distribution, $\mathbb{E}\left[\text{Regret (RAMS)}\right] = \mathcal{O}(1/\varepsilon_0)$

Remark

- $F = \text{Unif}([0, \frac{1}{4}] \cup [\frac{3}{4}, 1])$, RAMS ($\mathcal{O}((\log T)^2)$) outperforms CE ($\Omega(\sqrt{T})$).
- Matches the universal lower bound upto polylog factors $\implies$ RAMS is near-optimal.

» Multi-secretary with general distributions

Brief Summary

(β, ε_0)-clustered distributions

Discrete Distributions
“Few types”

Regret $= \Theta(1)$
[?]
Regret can scale with # of types

Continuous Distributions
“Many Types”

Regret $= \Theta(\log T)$
[?], [?]
Need density to be bounded below
Do not admit distributions with gaps

Multi-secretary with general distributions

Brief Summary

(β, ε_0)-clustered distributions

$$\mathbb{E}\left[\text{Regret}(DP)\right] = \begin{cases} \Omega\left(\log T\right), & \beta = 0, \\ \Omega\left(T^{\frac{1}{2} - \frac{1}{2(1+\beta)}}\right), & \beta > 0 \end{cases}$$

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[?]
Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$
[?], [?]
Need density to be bounded below
Do not admit distributions with gaps

» Multi-secretary with general distributions

Brief Summary

(β, ε_0)-clustered distributions

$$\mathbb{E}\left[\text{Regret}(DP)\right] = \begin{cases} \Omega\left(\log T\right), & \beta = 0, \\ \Omega\left(T^{\frac{1}{2} - \frac{1}{2(1+\beta)}}\right), & \beta > 0 \end{cases}$$

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[?]
Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$
[?], [?]
Need density to be bounded below
Do not admit distributions with gaps

$$\mathbb{E}\left[\text{Regret(RAMS)}\right] = \begin{cases} \mathcal{O}\left((\log T)^2\right), & \beta = 0, \\ \mathcal{O}\left(\text{poly}(\log T) T^{\frac{1}{2} - \frac{1}{2(1+\beta)}}\right), & \beta > 0 \end{cases}$$

» One Policy to solve them all?

* The multi-secretary problem is special but RAMS is general: in each period, simulate several futures and choose the action which minimizes the expected "compensation" in hindsight. Compensation $\equiv$ How to much we need to pay an agent who knows the future to take a particular action, for a given future.
* Can be applied to NRM and stochastic online matching problems to recover almost all known guarantees in the literature.

» One Policy to solve them all?

Proposition (RAMS is as good as any algorithm)

Given an NRM setting P, consider *any* algorithm A for P, such that with τ periods remaining, uniformly over the state, the expected compensation under A is bounded above by $\delta_\tau(\mathsf{A})$. Then RAMS achieves an expected compensation bounded uniformly by $\delta_\tau + 1/\tau^{1.1}$. As a result the regret of RAMS is bounded above by a constant plus the regret guarantee for algorithm A,

$$\text{Regret(RAMS)} \leq \text{Constant} + \sum_{\tau=1}^{T} \delta_\tau(\mathsf{A}) \, .$$

» What to take away from this talk?

Simple and practical simulation-based policy SOAR is broadly applicable:

* RAMS (Repeatedly Act based on Multiple Sims) recovers the guarantees for almost all settings in the NRM literature (e.g., constant regret for finite types, $\log^2 T$ for semi-infinite types)
* Establishes novel guarantees for dynamic spatial matching problems

Thank you!

APPENDIX on (β, ε_0)-clustered distributions

» (β, ε_0)-clustered distribution

Examples

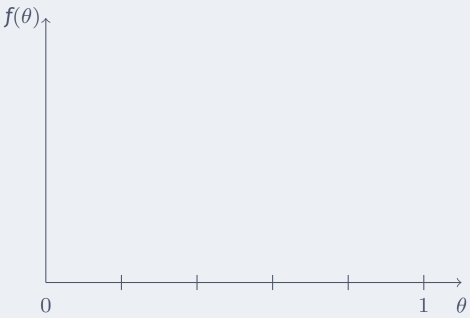

» (β, ε_0)-clustered distribution

Examples

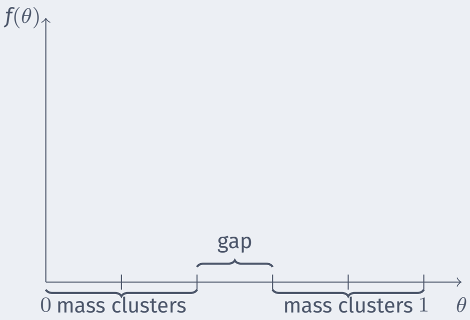

» (β, ε_0)-clustered distribution

Examples

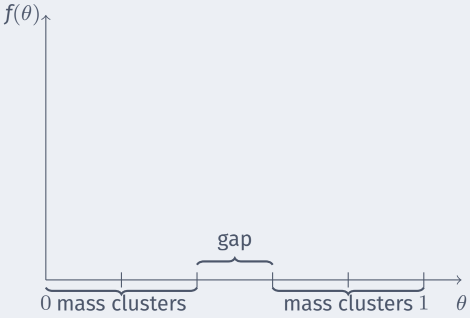

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

» (β, ε_0)-clustered distribution

Examples

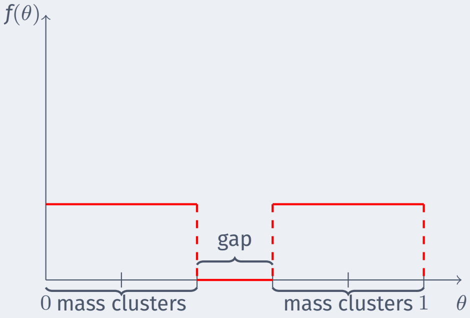

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 0$ (mass accumulation around gaps)

» (β, ε_0)-clustered distribution

Examples

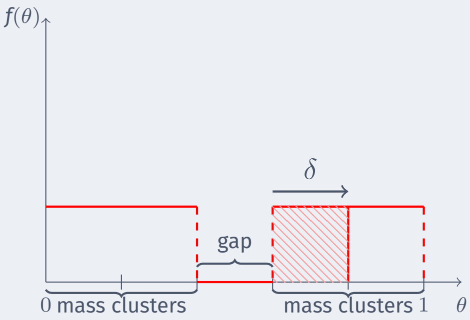

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 0$ (mass accumulation around gaps)

» (β, ε_0)-clustered distribution

Examples

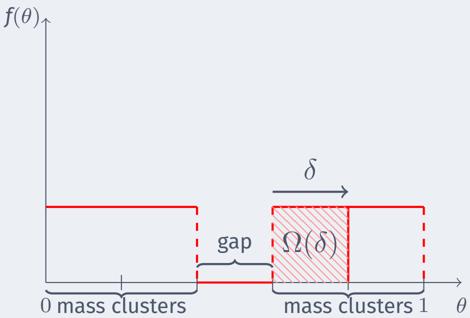

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 0$ (mass accumulation around gaps)

$|F(m+\delta) - F(m)| \geq \delta$ on the same mass cluster

» (β, ε_0)-clustered distribution

Examples

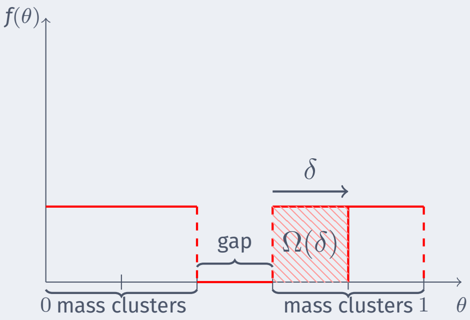

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 0$ (mass accumulation around gaps)

$|F(m+\delta) - F(m)| \geq \delta$ on the same mass cluster

$\mu(\text{mass clusters}) \geq \varepsilon_0$

» (β, ε_0)-clustered distribution

Examples

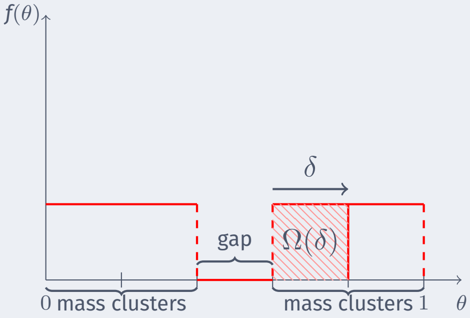

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 0$ (mass accumulation around gaps)

$|F(m+\delta) - F(m)| \geq \delta$ on the same mass cluster

μ(mass clusters) $\geq \varepsilon_0$

For discrete distrbutions, $\beta = 0, \varepsilon_0 = \min_j p_j$

» (β, ε_0)-clustered distribution

Examples

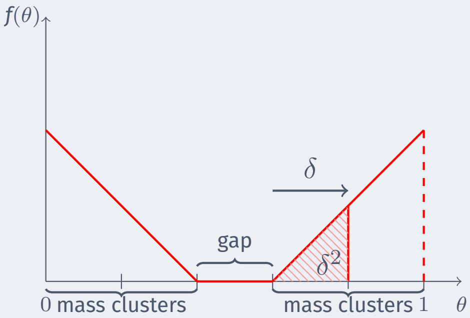

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 1$ (mass accumulation around gaps)

$|F(m+\delta) - F(m)| \geq \delta^2$ on the same mass cluster

μ(mass clusters) $\geq \varepsilon_0$

APPENDIX on PART II

Feature Based Dynamic Matching

Feature Based Dynamic Matching

- Platform has a pool of T service providers, who live in d dimensional feature space.
  - e.g., $Y_k = (\text{price}, \text{rating})$

Feature Based Dynamic Matching

- Platform has a pool of T service providers, who live in d dimensional feature space.
  - e.g., $Y_k = (\text{price}, \text{rating})$
- T customers arrive online and have i.i.d. preferences, i.e., weights over the features.
  - e.g., $X_i = -(\text{sensitivity to price}, \text{sensitivity to rating})$

Feature Based Dynamic Matching

- Platform has a pool of T service providers, who live in d dimensional feature space.
  - e.g., $Y_k = (\text{price, rating})$
- T customers arrive online and have i.i.d. preferences, i.e., weights over the features.
  - e.g., $X_i = -(\text{sensitivity to price, sensitivity to rating})$
- Match value is given as $\langle X_i, Y_k \rangle$

Feature Based Dynamic Matching

- Platform has a pool of T service providers, who live in d dimensional feature space.
  - e.g., $Y_k = (\text{price, rating})$
- T customers arrive online and have i.i.d. preferences, i.e., weights over the features.
  - e.g., $X_i = -(\text{sensitivity to price, sensitivity to rating})$
- Match value is given as $\langle X_i, Y_k \rangle$
- Both service provider and customer leave upon matching.

» Feature Based Dynamic Matching

- Platform has a pool of T service providers, who live in d dimensional feature space.
  - e.g., $Y_k = (\text{price, rating})$
- T customers arrive online and have i.i.d. preferences, i.e., weights over the features.
  - e.g., $X_i = -(\text{sensitivity to price, sensitivity to rating})$
- Match value is given as $\langle X_i, Y_k \rangle$
- Both service provider and customer leave upon matching.
- Supply and demand distributions are known and possibly *different*.

» Performance metric: regret with respect to fluid benchmark

* We aim to maximize the expected average match value $\frac{1}{T}\sum_{i=1}^{T}\langle X_i, Y_{\pi(i)}\rangle$
* Fluid benchmark is the value of the optimal transport between the demand distribution and the supply distribution
* We aim to minimize the additive regret wrt the fluid benchmark. We want $o(1)$ regret.

* Problem is equivalent to minimizing $\frac{1}{T}\sum_{i=1}^{T}\|X_i - Y_{\pi(i)}\|^2$

» 1-dimensional example of maximizing $\langle \boldsymbol{X_i}, \boldsymbol{Y_{\pi(i)}} \rangle$

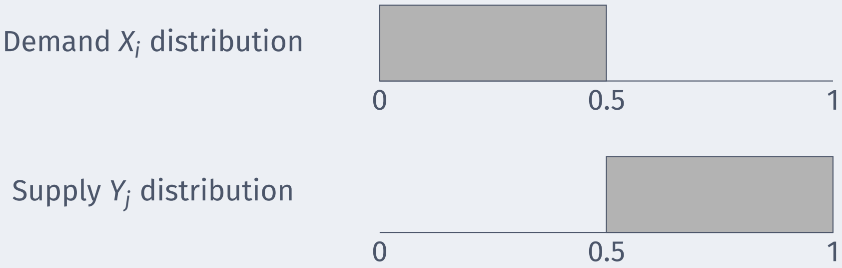

1-dimensional example of maximizing $\langle X_i, Y_{\pi(i)} \rangle$

Demand X_i distribution

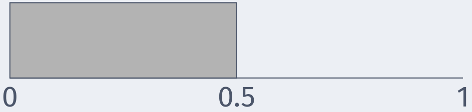

Optimal transport?

Supply Y_j distribution

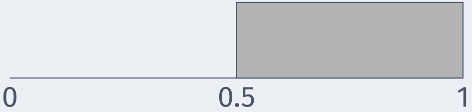

» 1-dimensional example of maximizing $\langle \boldsymbol{X_i}, \boldsymbol{Y_{\pi(i)}} \rangle$

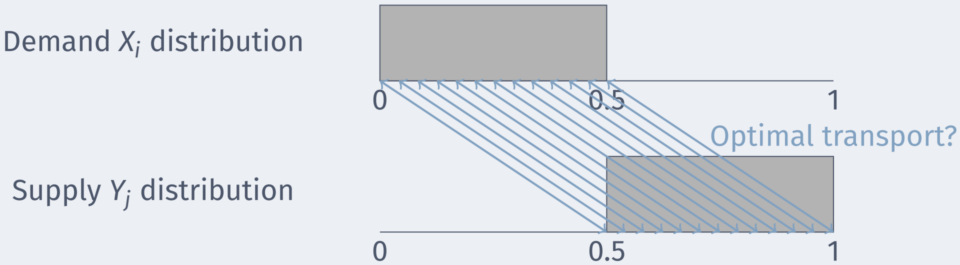

» 1-dimensional example of maximizing $\langle \boldsymbol{X_i}, \boldsymbol{Y_{\pi(i)}} \rangle$

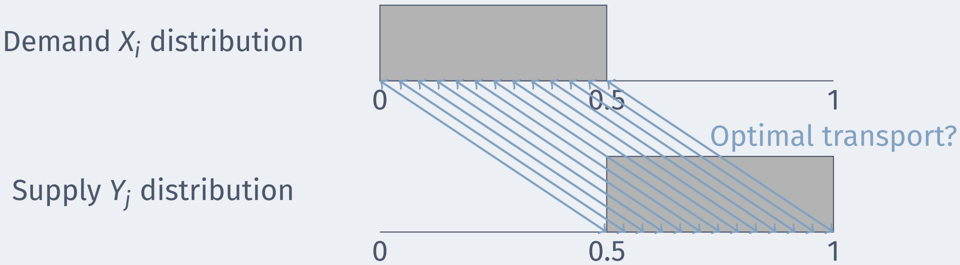

- Optimal transport has value per match 0.208
- Greedy fails: produces a random matching, expected value per match is only 0.188

» SOAR: a simple future-aware algorithm

We introduce a simple forward-looking algorithm dubbed SOAR

Simulate

Optimize

Assign

Repeat

SOAR calculates each matching decision based on a simulation of the future, and hindsight optimization on that future.

» SOAR: One simulation to optimize them, One simulation to assign them and with repetition, solve them

SOAR: One simulation to optimize them, One simulation to assign them and with repetition, solve them

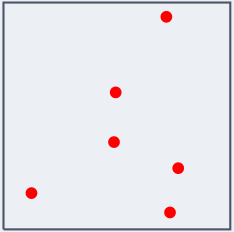

» SOAR: One simulation to optimize them, One simulation to assign them and with repetition, solve them

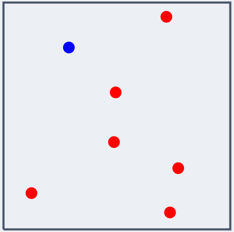

» SOAR: One simulation to optimize them, One simulation to assign them and with repetition, solve them

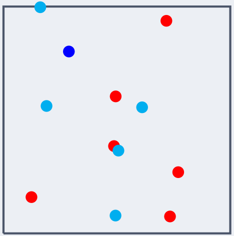

Simulate

» SOAR: One simulation to optimize them, One simulation to assign them and with repetition, solve them

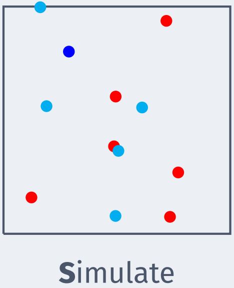

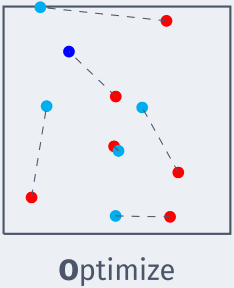

» SOAR: One simulation to optimize them, One simulation to assign them and with repetition, solve them

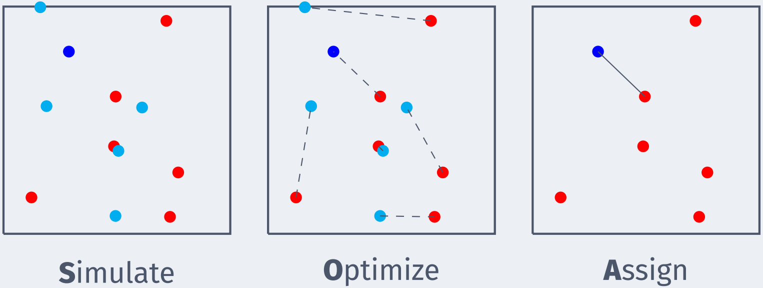

» SOAR: One simulation to optimize them, One simulation to assign them and with repetition, solve them

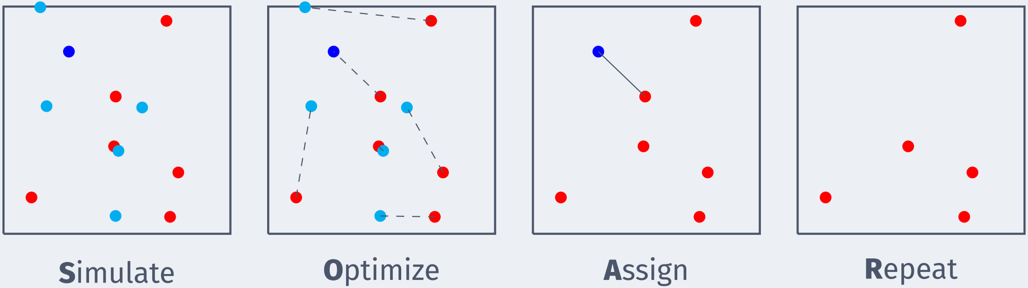

SOAR is provably near optimal

	P, Q Regular	P, Q Arbitrary
Lower Bound SOAR	$\tilde{\Omega}(T^{-(\frac{2}{d} \wedge 1)})$	$\tilde{\Omega}(T^{-(\frac{2}{d} \wedge \frac{1}{2})})$

* $\frac{1}{\mathrm{NND}^2}$ is a lower bound on the regret.
* For $d = 1$, the matching constraint leads to a tighter lower bound.
* For irregular distributions, a simple example tells us $1/\sqrt{T}$ is a lower bound. $1/\sqrt{T} \gg 1/\mathrm{NND}^2$ for $d \leq 3$.

» SOAR is provably near optimal

	P, Q Regular	P, Q Arbitrary
Lower Bound	$\tilde{\Omega}(T^{-(\frac{2}{d}\wedge 1)})$	$\tilde{\Omega}(T^{-(\frac{2}{d}\wedge \frac{1}{2})})$
SOAR	$\tilde{O}(T^{-(\frac{2}{d}\wedge 1)})$	$\tilde{O}(T^{-(\frac{2}{d}\wedge \frac{1}{2})})$

* $\frac{1}{\text{NND}^2}$ is a lower bound on the regret.
* For $d = 1$, the matching constraint leads to a tighter lower bound.
* For irregular distributions, a simple example tells us $1/\sqrt{T}$ is a lower bound. $1/\sqrt{T} \gg 1/\text{NND}^2$ for $d \leq 3$.

SOAR achieves the optimal regret scaling in all cases.

» SOAR is provably near optimal

	P, Q Regular	P, Q Arbitrary
Lower Bound	$\tilde{\Omega}(T^{-(\frac{2}{d} \wedge 1)})$	$\tilde{\Omega}(T^{-(\frac{2}{d} \wedge \frac{1}{2})})$
SOAR	$\tilde{O}(T^{-(\frac{2}{d} \wedge 1)})$	$\tilde{O}(T^{-(\frac{2}{d} \wedge \frac{1}{2})})$

* $\frac{1}{\text{NND}^2}$ is a lower bound on the regret.
* For $d = 1$, the matching constraint leads to a tighter lower bound.
* For irregular distributions, a simple example tells us $1/\sqrt{T}$ is a lower bound. $1/\sqrt{T} \gg 1/\text{NND}^2$ for $d \leq 3$.

SOAR achieves the optimal regret scaling in all cases.

Proof idea: Expected regret incurred by SOAR's match when t periods remain is the same as the regret for offline matching of t pairs (which is larger than $1/t$). Sum over t and divide by T. Result $\sim$ regret for offline matching of T pairs.

» Numerical evaluation of SOAR's performance

Figure: $d = 1$, demand $\sim$ Unif$(0, 1/2)$, supply $\sim$ Unif$(0, 1)$

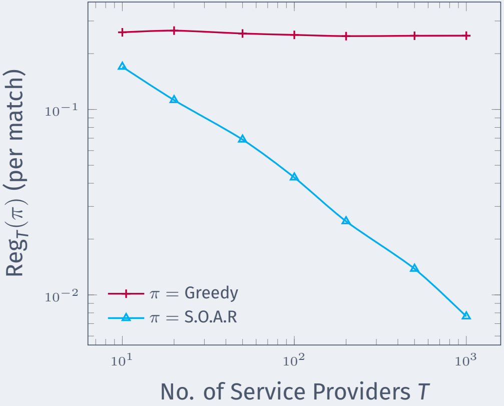

» Talk outline

- DSM with identical supply and demand distributions
  - Greedy matching suffices
  - Match distance $\sim$ Nearest-neighbor-distance achievable, except one case
- DSM with different supply and demand distributions
  - Greedy fails
  - Simulate-Optimize-Assign-Repeat (SOAR) is near optimal
- **Multisecretary problem with lumpy value distribution (a 1d DSM problem)** [O. Besbes, Akshit Kumar & K. '22]
  - SOAR with one sample path fails
  - RAMS with multiple sample paths achieves optimal regret scaling
  - Works also for $d \geq 2$, and across NRM settings.

Multi-secretary Problem

Problem Statement

Given a sequence of T secretaries and a hiring budget B, a decision maker (DM) wants to hire the top B secretaries in terms of their ability.

Multi-secretary Problem

Problem Statement

Given a sequence of T secretaries and a hiring budget B, a decision maker (DM) wants to hire the top B secretaries in terms of their ability.

Details and Assumptions

* The secretaries arrive in an online fashion.

» Multi-secretary Problem

Problem Statement

Given a sequence of T secretaries and a hiring budget B, a decision maker (DM) wants to hire the top B secretaries in terms of their ability.

Details and Assumptions

* The secretaries arrive in an online fashion.
* The DM makes **irrevocable** *hire* or *reject* decisions.

» Multi-secretary Problem

Problem Statement

Given a sequence of T secretaries and a hiring budget B, a decision maker (DM) wants to hire the top B secretaries in terms of their ability.

Details and Assumptions

- The secretaries arrive in an online fashion.
- The DM makes **irrevocable** *hire* or *reject* decisions.
- The abilities (types) of the secretaries are drawn **independently** from a **common** and **known** distribution F over $[0, 1]$.

» Multi-secretary Problem

Problem Statement

Given a sequence of T secretaries and a hiring budget B, a decision maker (DM) wants to hire the top B secretaries in terms of their ability.

Details and Assumptions

* The secretaries arrive in an online fashion.
* The DM makes **irrevocable** *hire* or *reject* decisions.
* The abilities (types) of the secretaries are drawn **independently** from a **common** and **known** distribution F over $[0,1]$.

Note: This is a 1d DSM problem, with an atomic "supply" distribution with B units at 1 and $T-B$ units at 0. $\Theta(\sqrt{T})$ optimal regret wrt fluid benchmark, which is trivial to achieve. We'll adopt a tighter benchmark to obtain algorithmic insights.

Gotta catch'em ~~all~~ some

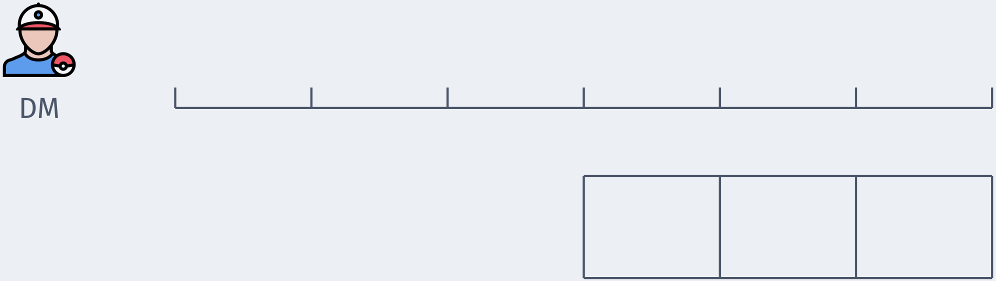

» Gotta catch'em ~~all~~ some

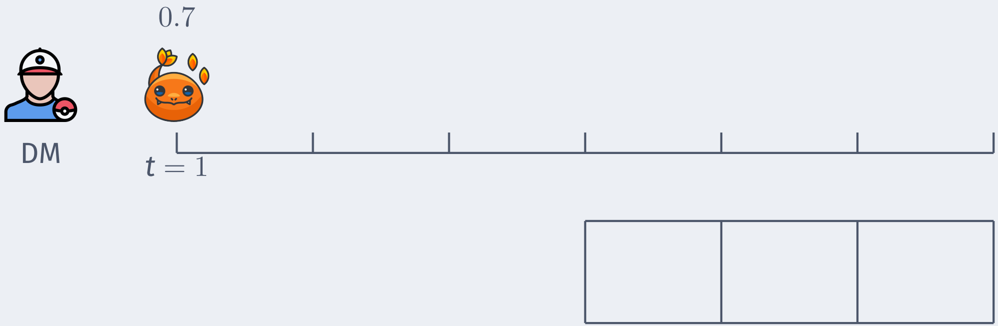

Gotta catch'em ~~all~~ some

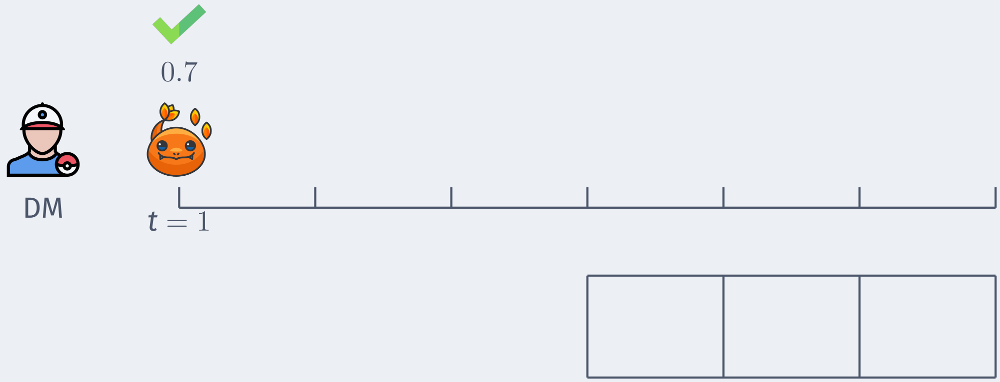

» Gotta catch'em ~~all~~ some

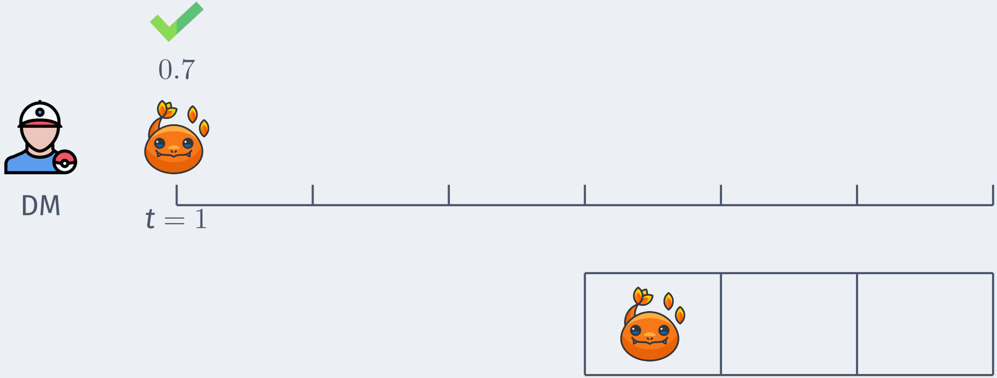

» Gotta catch'em ~~all~~ some

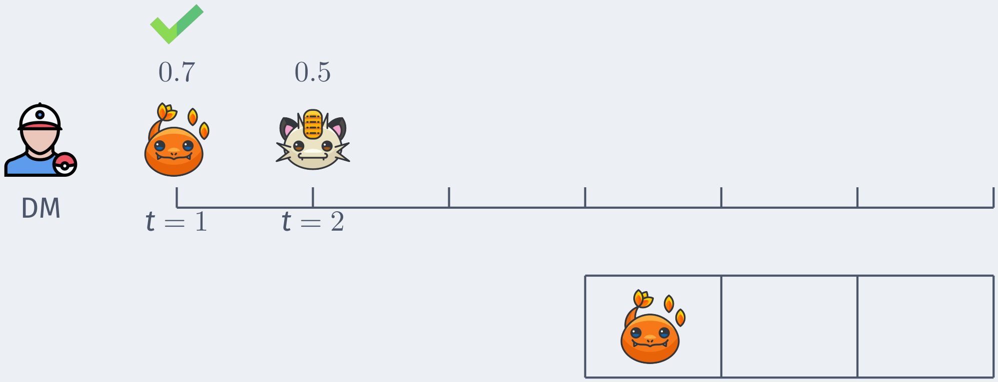

Gotta catch'em ~~all~~ some

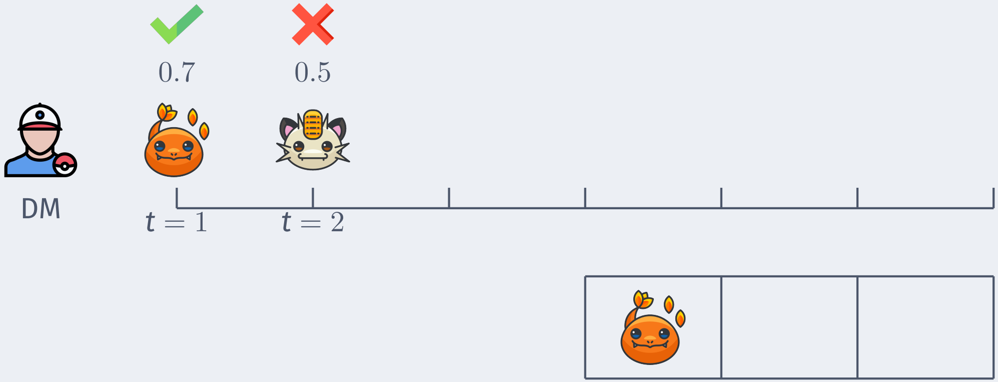

» Gotta catch'em ~~all~~ some

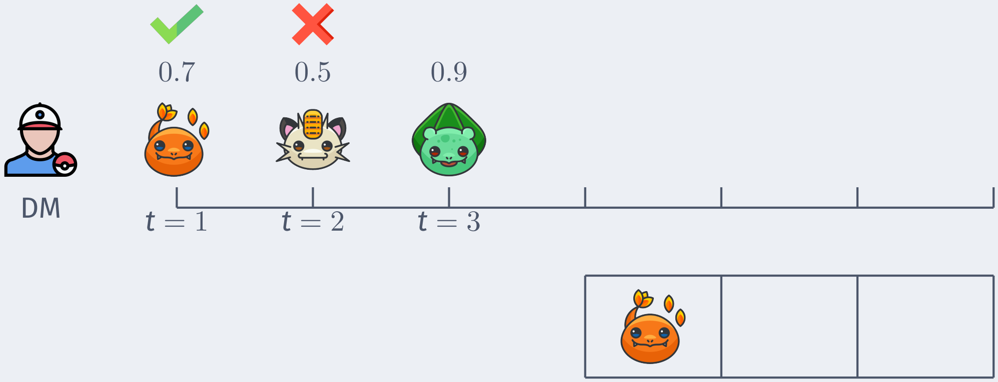

» Gotta catch'em ~~all~~ some

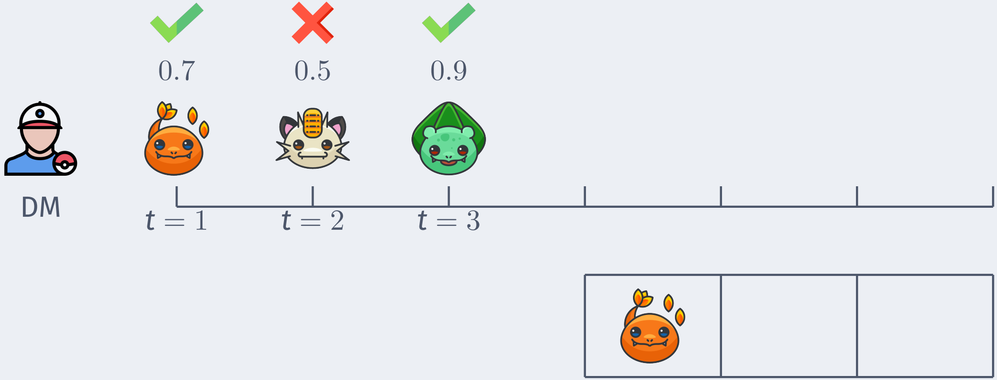

» Gotta catch'em ~~all~~ some

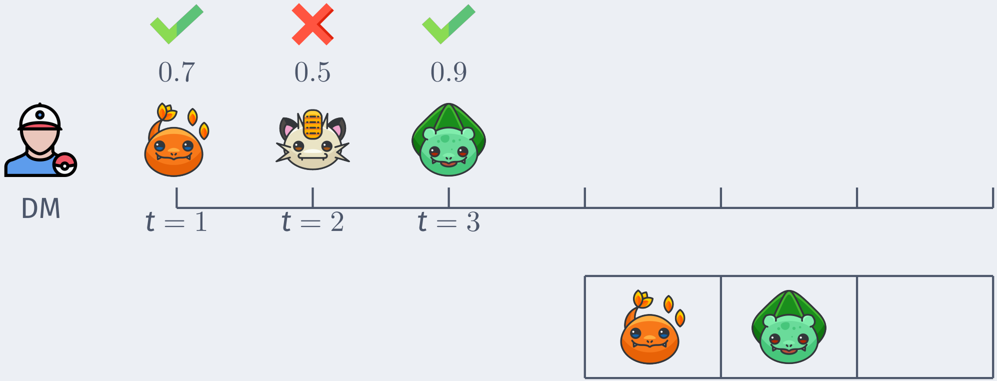

» Gotta catch'em ~~all~~ some

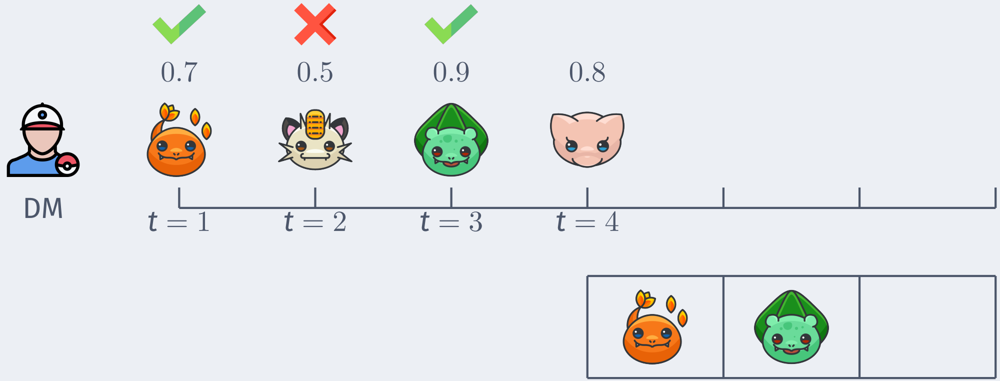

Gotta catch'em ~~all~~ some

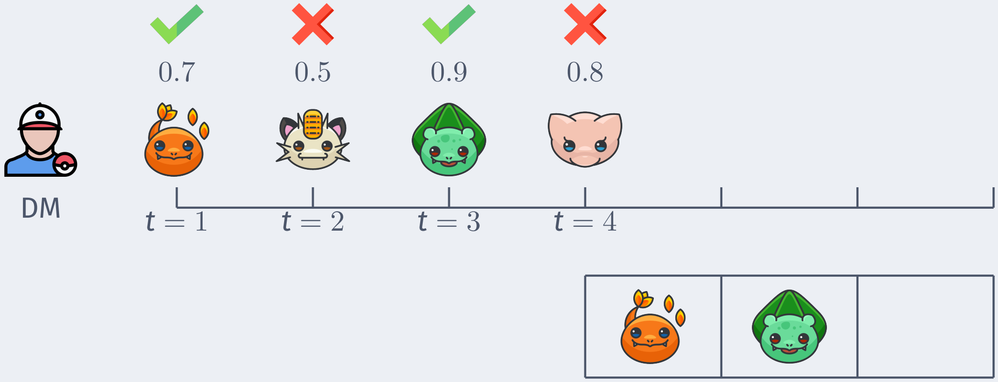

Gotta catch'em ~~all~~ some

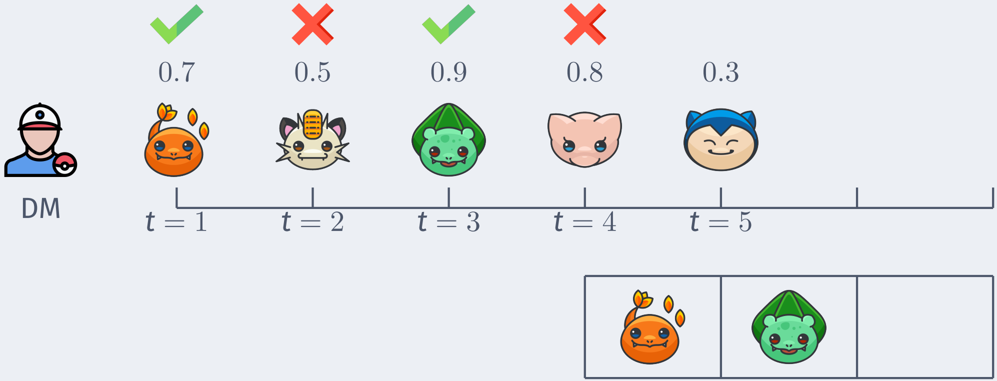

Gotta catch'em ~~all~~ some

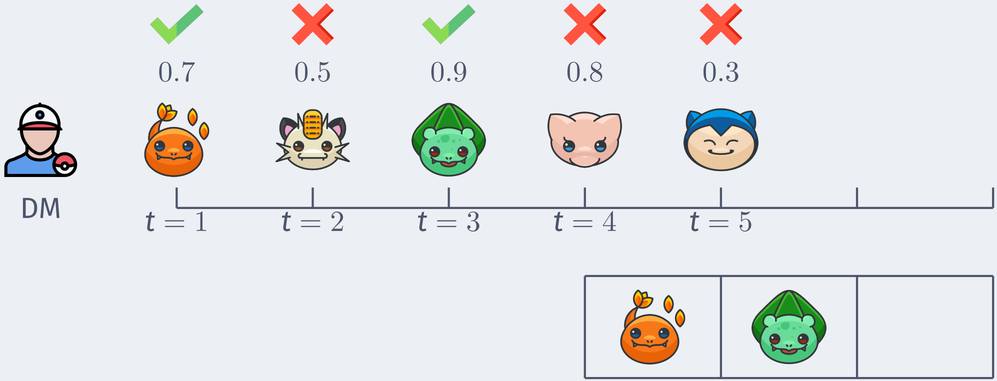

Gotta catch'em ~~all~~ some

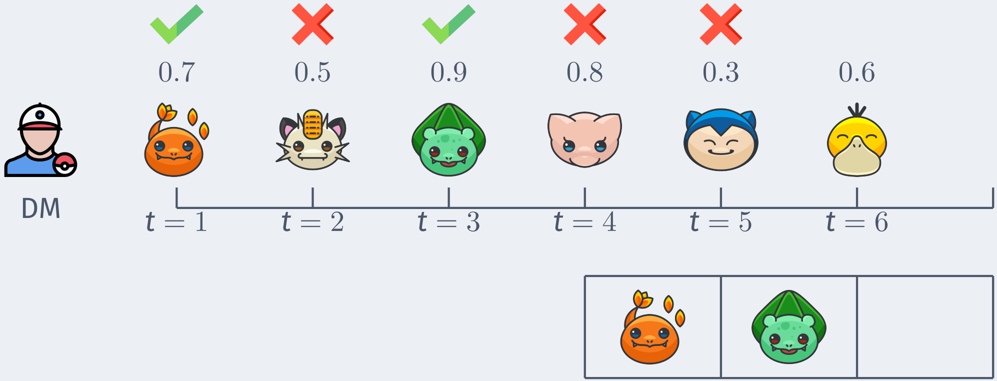

» Gotta catch'em ~~all~~ some

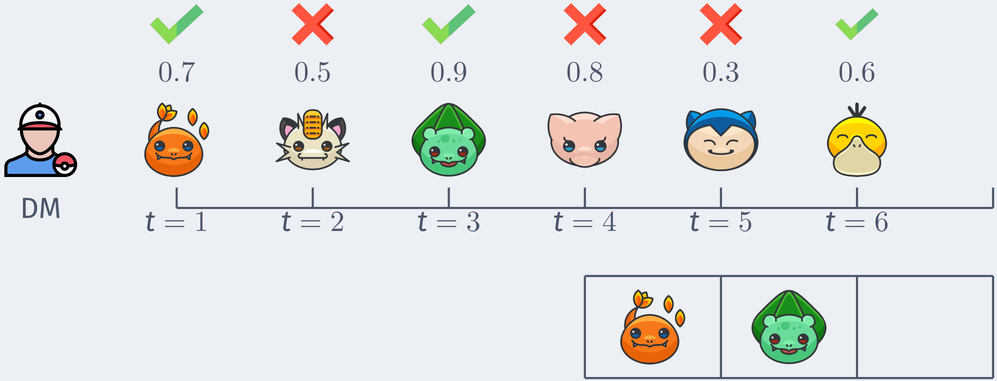

» Gotta catch'em ~~all~~ some

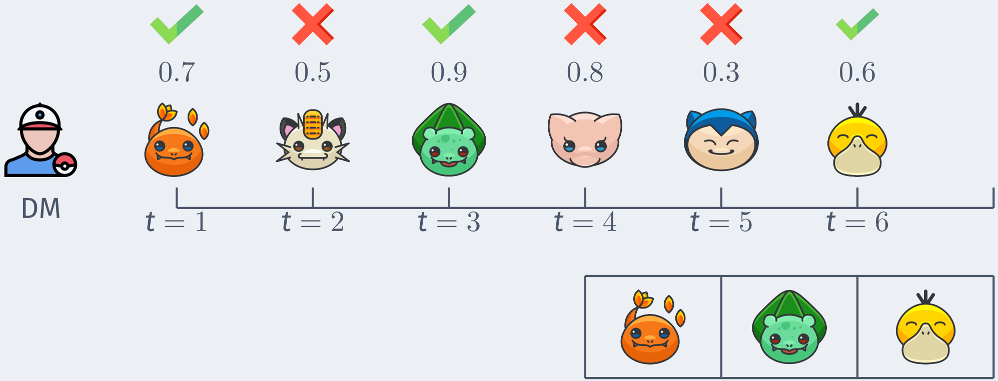

» Gotta catch'em ~~all~~ some

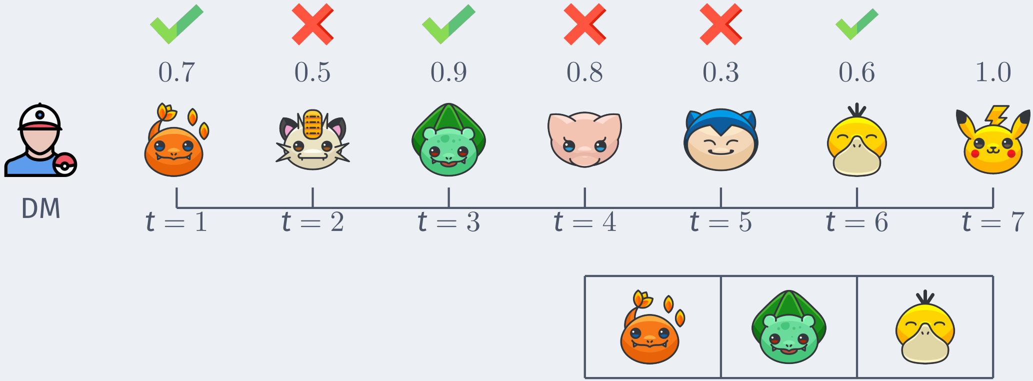

Gotta catch'em ~~all~~ some

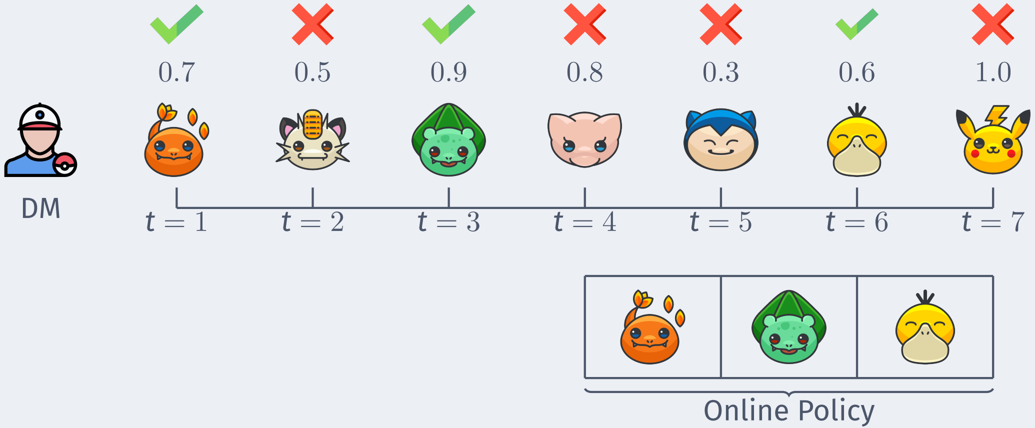

Gotta catch'em ~~all~~ some

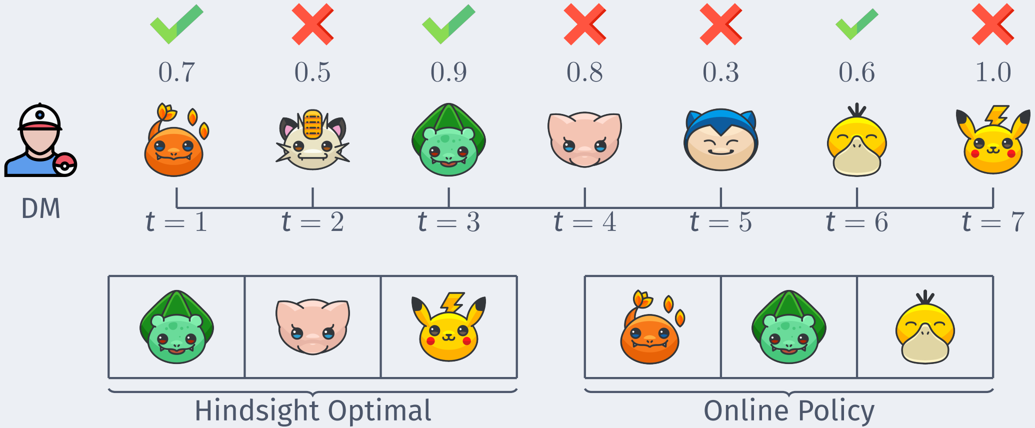

» Gotta catch'em ~~all~~ some

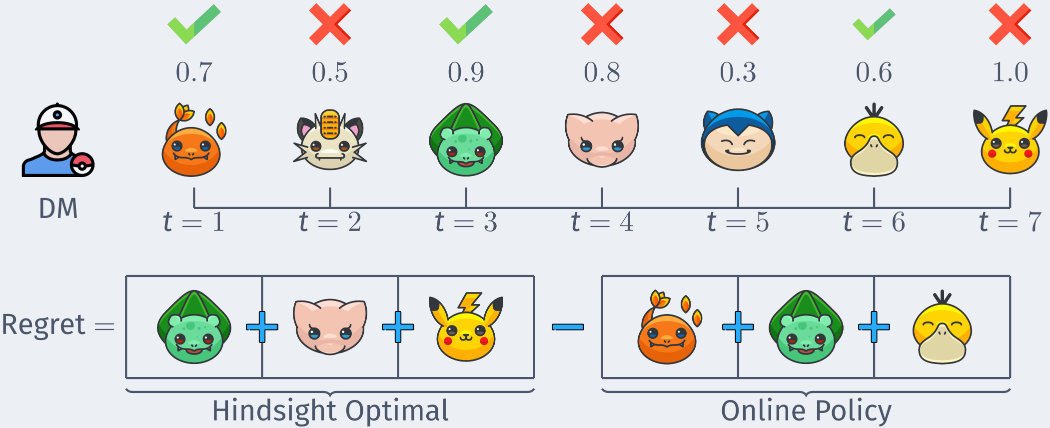

» Gotta catch'em ~~all~~ some

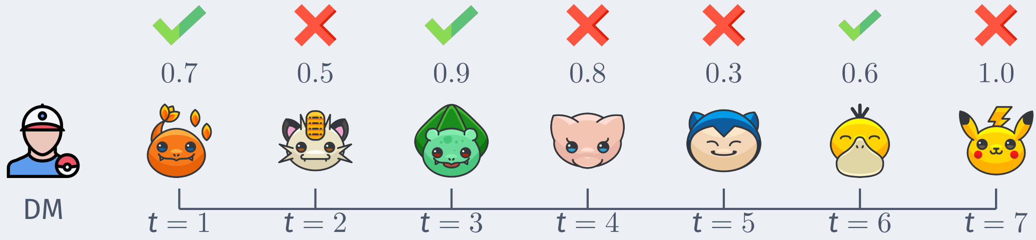

$$\text{Regret} = \underbrace{\boxed{0.9 + 0.8 + 1.0}}_{\text{Hindsight Optimal}} - \underbrace{\boxed{0.7 + 0.9 + 0.6}}_{\text{Online Policy}} = 0.5$$

» Gotta catch'em ~~all~~ some

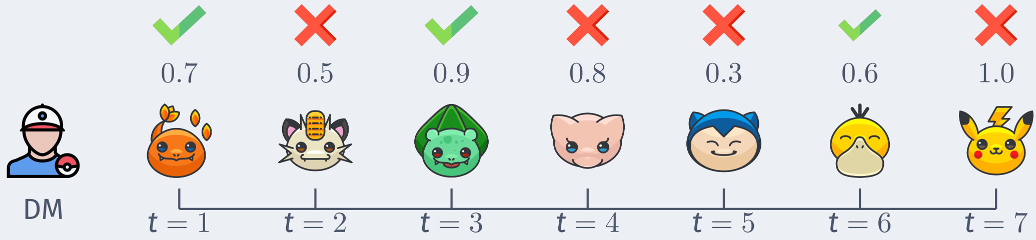

$$\text{Regret} = \underbrace{\boxed{0.9 + 0.8 + 1.0}}_{\text{Hindsight Optimal}} - \underbrace{\boxed{0.7 + 0.9 + 0.6}}_{\text{Online Policy}} = 0.5$$

We will focus on $\mathbb{E}\left[\text{Regret}\right]$

» Hindsight-based regret

Multisecretary is a 1d DSM problem, with an atomic "supply" distribution with B units at 1 and $T - B$ units at 0. $\Theta(\sqrt{T})$ optimal regret wrt fluid benchmark, which can be achieved by a trivial static policy.

Gap between fluid and hindsight benchmarks is already $\Omega(\sqrt{T})$.

As in the recent NRM literature, we adopt the tighter hindsight benchmark.

Multi-secretary Problem: What's known and research question?

Multi-secretary Problem: What's known and research question?

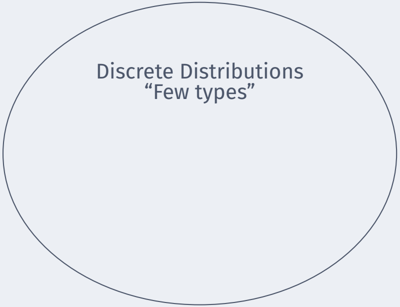

Multi-secretary Problem: What's known and research question?

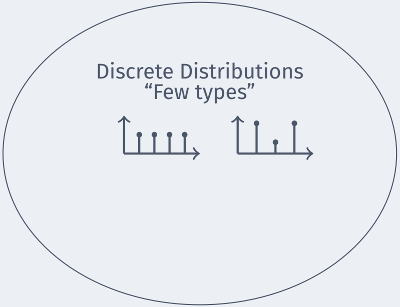

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]

Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]
Regret can scale with # of types

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$

[Arlotto and Gurvich'19]

Regret can scale with # of types

Continuous Distributions
"Many Types"

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]

Regret can scale with # of types

Continuous Distributions
"Many Types"

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]
Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$
[Leuker'98, Bray'22]

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]
Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$
[Leuker'98, Bray'22]
Need density to be bounded below

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]
Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$
[Leuker'98, Bray'22]
Need density to be bounded below

» Multi-secretary Problem: What's known and research question?

Discrete Distributions
"Few types"

Regret $= \Theta(1)$
[Arlotto and Gurvich'19]
Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$
[Leuker'98, Bray'22]
Need density to be bounded below
Do not admit distributions with gaps

Multi-secretary Problem: What's known and research question?

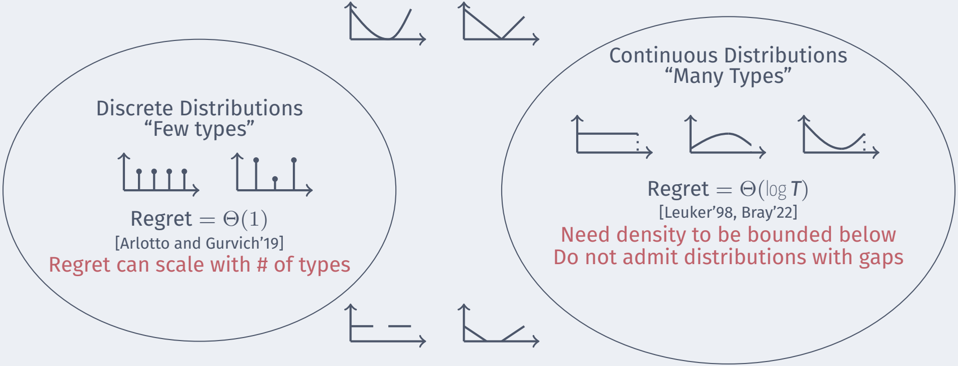

» Multi-secretary Problem: What's known and research question?

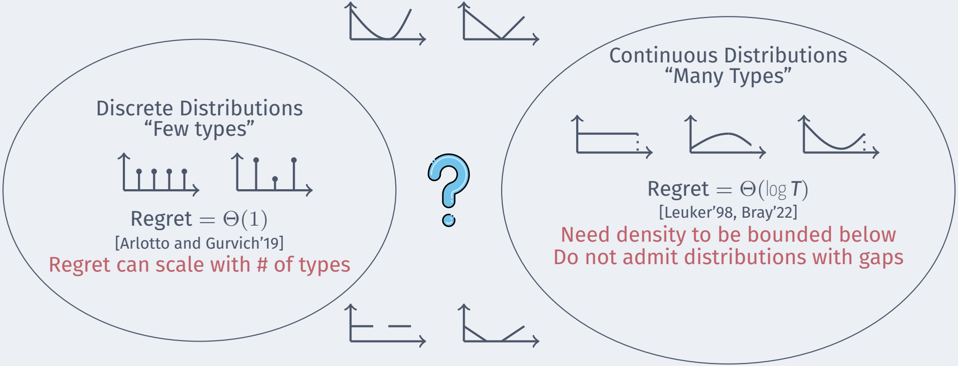

How to interpolate between these distribution classes?
How to deal with gaps? What are the possible regret scalings?

Punchline for the Multi-secretary Problem

Drivers of Regret

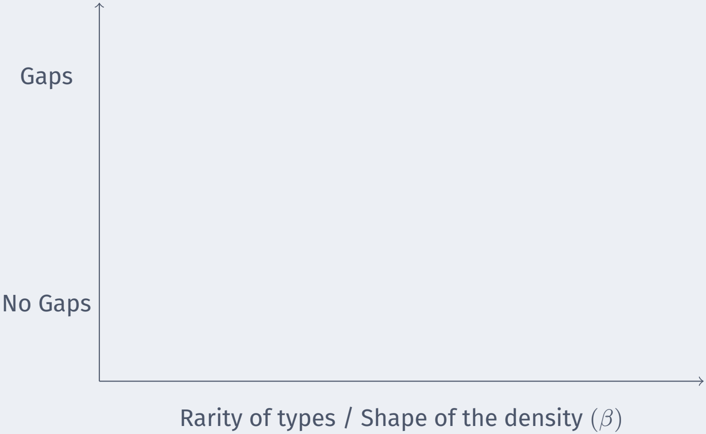

Punchline for the Multi-secretary Problem

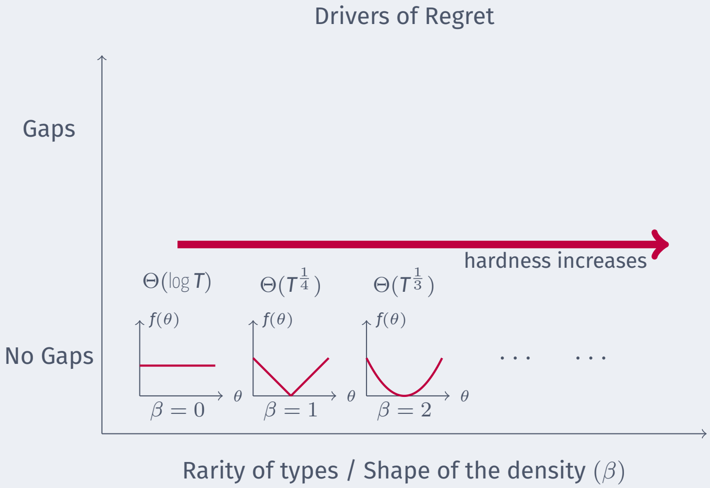

Punchline for the Multi-secretary Problem

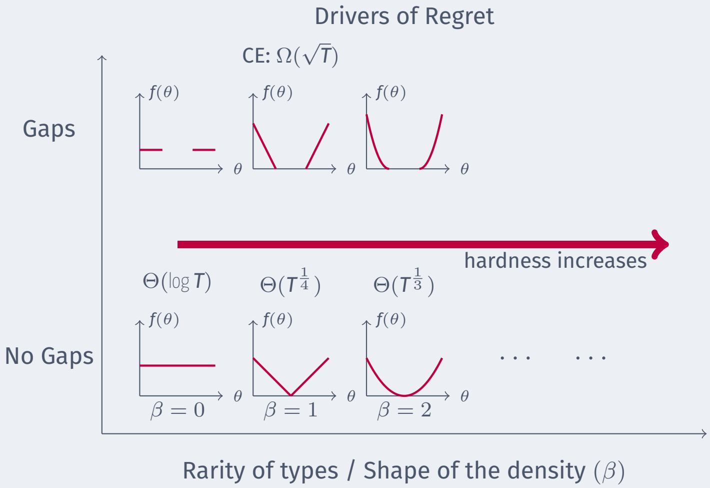

Punchline for the Multi-secretary Problem

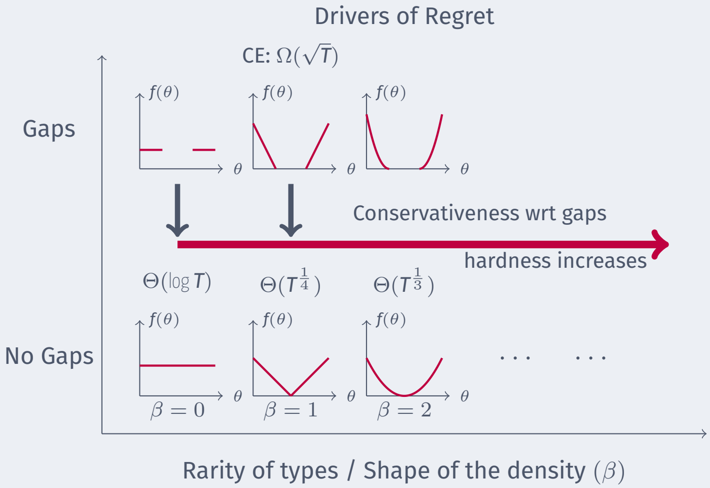

» Punchline for the Multi-secretary Problem

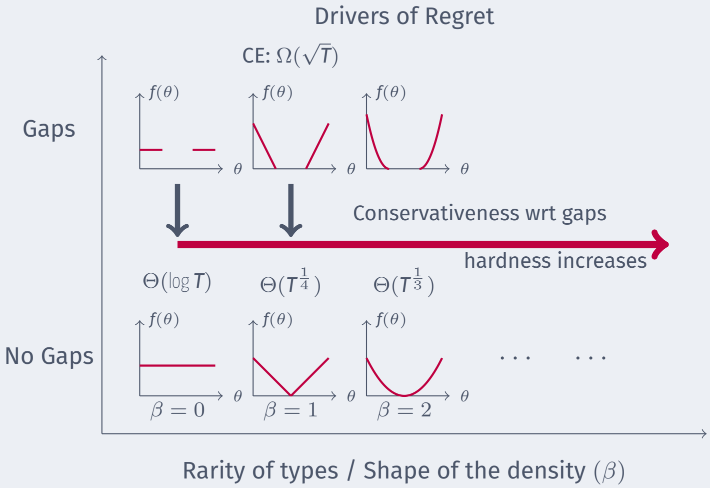

- Distribution shape is a **fundamental** driver of regret.
- Dealing with gaps is an algorithmic challenge.
- Novel Principle: Conservativeness wrt gaps (CwG)
- multi-sim SOAR variant automatically pursues CwG

Towards unification of the multi-secretary problem

Discrete Distributions
"Few types"

Regret $= \Theta(1)$

[Arlotto and Gurvich'19]

Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$

[Bray'22, Leuker'98]

Need density to be bounded below
Do not admit distributions with gaps

» Towards unification of the multi-secretary problem

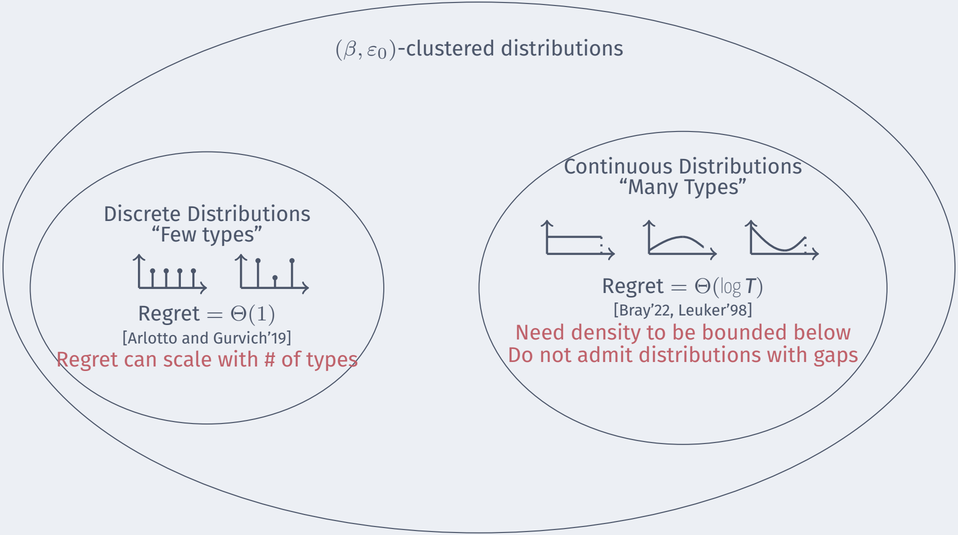

Fundamental Limits

Universal Lower Bound

For every $\beta \in [0, \infty)$, there exists a distribution F_β such that

$$\sup_{B \in [T]} \mathbb{E}_{F_\beta} \left[\text{Regret(DP)} \right] = \begin{cases} \Omega\left(\log T\right), & \beta = 0, \\ \Omega\left(T^{\frac{1}{2} - \frac{1}{2(1+\beta)}}\right), & \beta > 0. \end{cases}$$

» Fundamental Limits

Universal Lower Bound

For every $\beta \in [0, \infty)$, there exists a distribution F_β such that

$$\sup_{B \in [T]} \mathbb{E}_{F_\beta} \left[\text{Regret(DP)} \right] = \begin{cases} \Omega\left(\log T\right), & \beta = 0, \\ \Omega\left(T^{\frac{1}{2} - \frac{1}{2(1+\beta)}}\right), & \beta > 0. \end{cases}$$

$\Omega(\log T)$ $\Omega(T^{\frac{1}{4}})$ $\Omega(T^{\frac{1}{3}})$

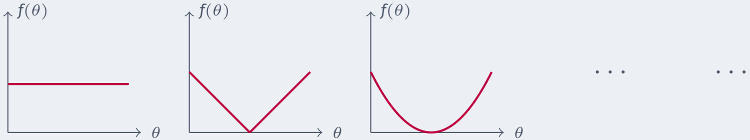

$\beta = 0$ $\beta = 1$ $\beta = 2$

» Certainty Equivalent Control

For $(\beta, 1)$-clustered distributions

* Let B_t be the remaining budget at time t
* Compute the budget ratio $br_t = \frac{\text{Remaining Budget}}{\text{Remaining Time}} = \frac{B_t}{T-t}$
* Define a quantile threshold $p_t^{ce} = 1 - br_t$
* Define a ability threshold $\gamma_t^{ce} = F^{-1}(p_t^{ce})$
* hire $\iff \theta_t \geq \gamma_t^{ce}$

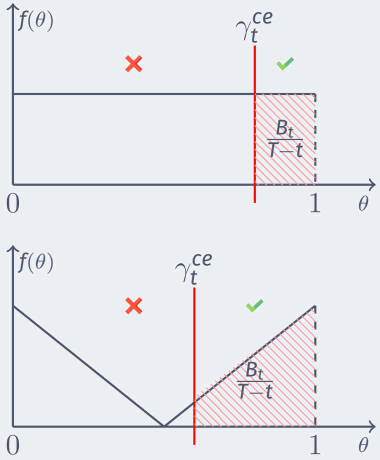

» Certainty Equivalent Control

For Bi-modal Uniform Distribution

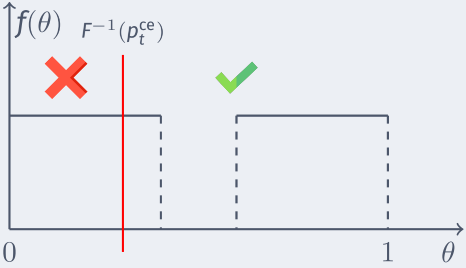

Let B_t be the remaining budget at time t

$$\text{Budget Ratio} = \frac{\text{Remaining Budget}}{\text{Remaining Time}} = \frac{B_t}{T-t}$$

$$\text{CE Quantile Threshold} = 1 - \frac{B_t}{T-t} \triangleq p_t^{ce}$$

$$\text{Decision: hire} \iff \theta_t \geq F^{-1}(p_t^{ce})$$

» Failure of Certainty Equivalent Control

Regret Lower Bound

Insufficiency of Certainty Equivalent Control

Assume that $F = \text{Unif}([0, \frac{1}{4}] \cup [\frac{3}{4}, 1])$, for $B = T/2$, we have

$$\mathbb{E}\left[\text{Regret}(CE)\right] = \Omega\left(\sqrt{T}\right)$$

» Failure of Certainty Equivalent Control

Regret Lower Bound

Insufficiency of Certainty Equivalent Control

Assume that $F = \text{Unif}([0, \frac{1}{4}] \cup [\frac{3}{4}, 1])$, for $B = T/2$, we have

$$\mathbb{E}\left[\text{Regret}(CE)\right] = \Omega\left(\sqrt{T}\right)$$

Remark

* Same scaling is achievable under a static threshold policy.

Why does CE fail?

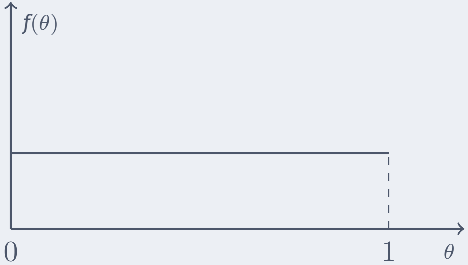

Why does CE fail?

$$B/\tau = \frac{1}{2} + \epsilon$$

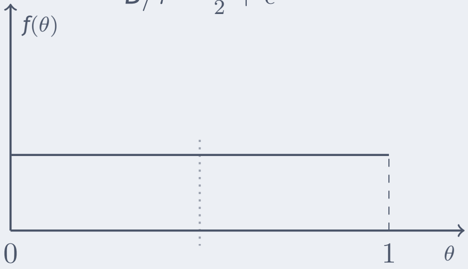

Why does CE fail?

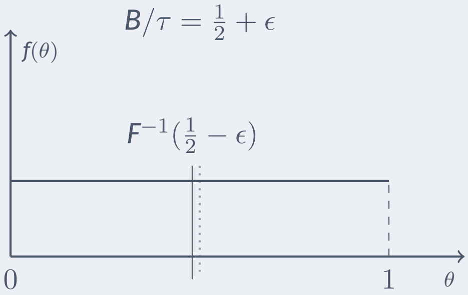

Why does CE fail?

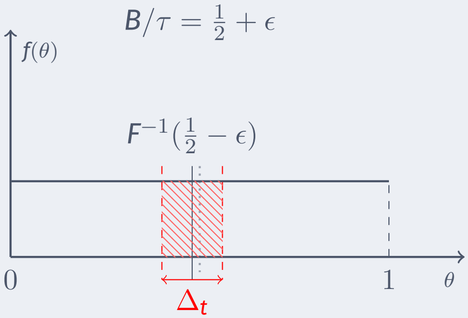

» Why does CE fail?

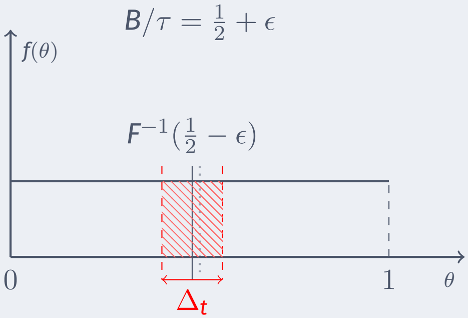

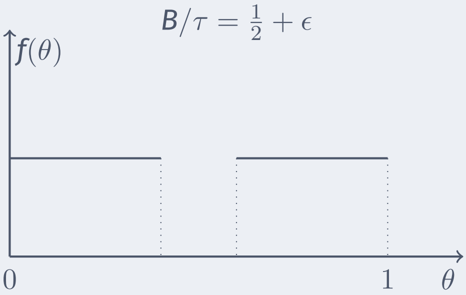

Why does CE fail?

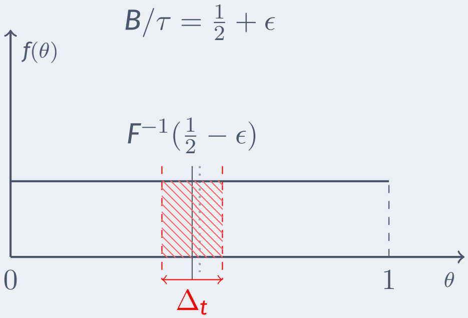

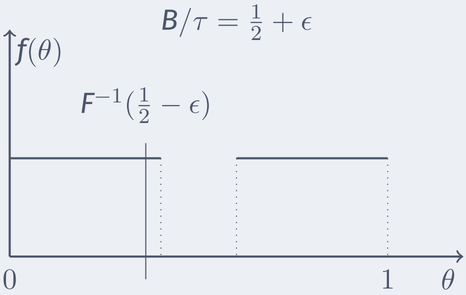

» Why does CE fail?

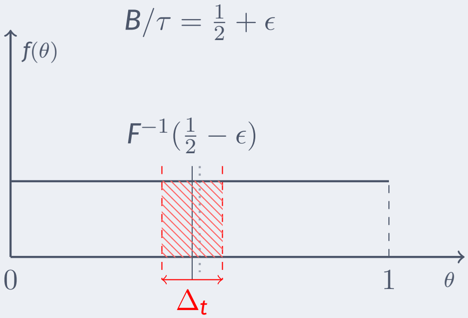

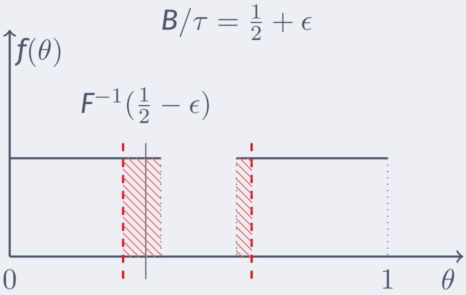

» Why does CE fail?

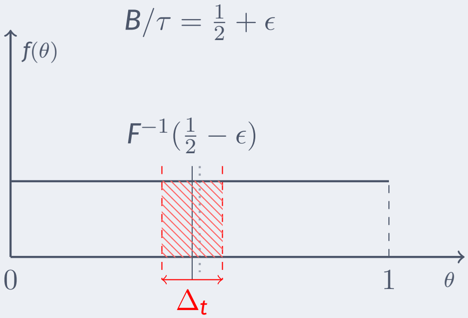

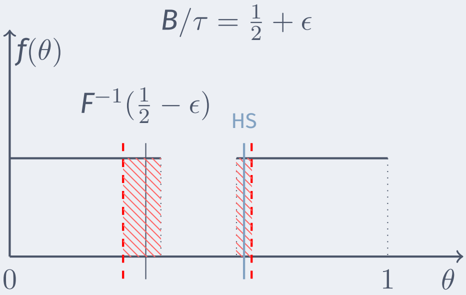

Why does CE fail?

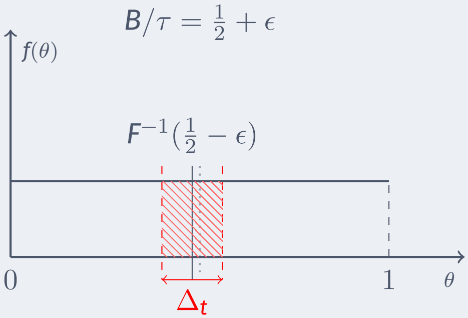

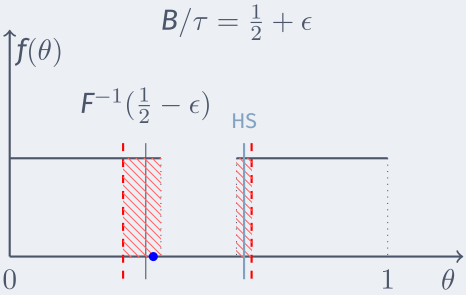

Why does CE fail?

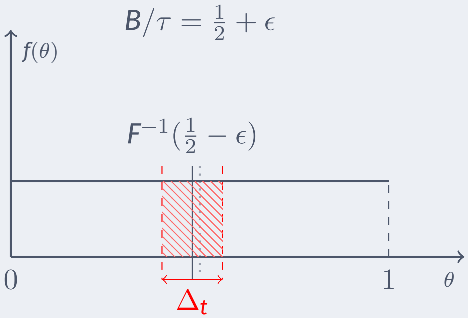

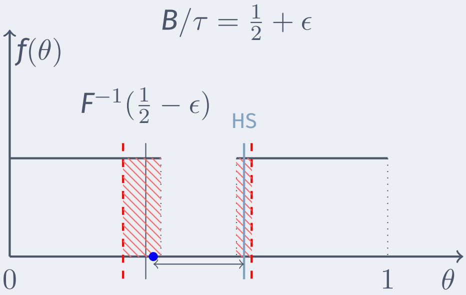

» Why does CE fail?

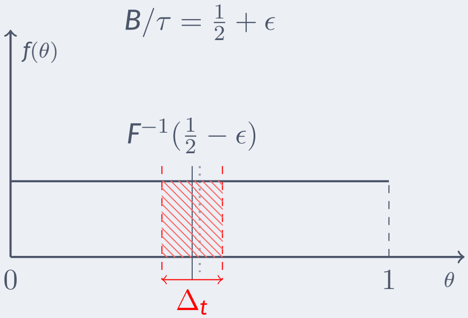

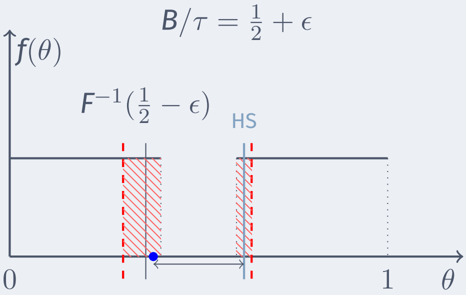

This is $\Omega(1/\sqrt{\tau})$ expected compensation

Why does CE fail?

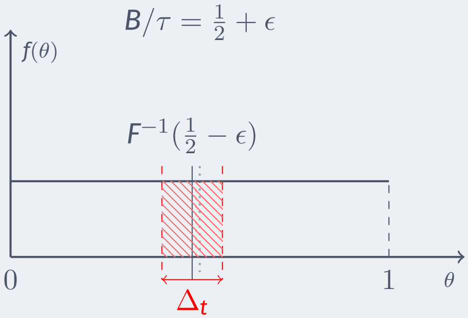

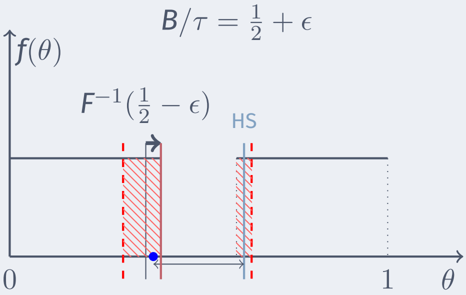

This is $\Omega(1/\sqrt{\tau})$ expected compensation

» Why does CE fail?

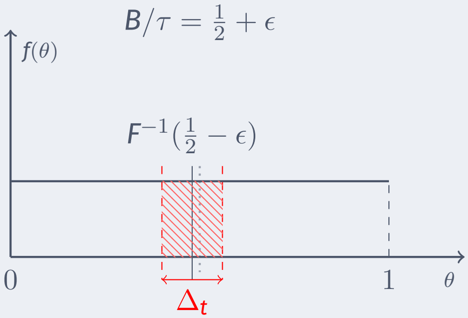

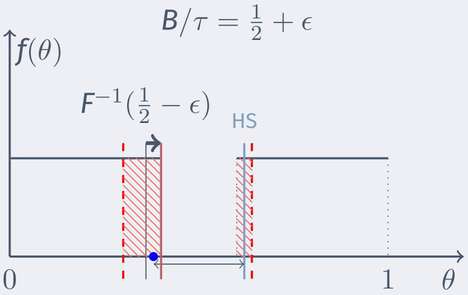

This is $\Omega(1/\sqrt{\tau})$ expected compensation

Conservativeness wrt gaps

Why does CE fail?

What if?

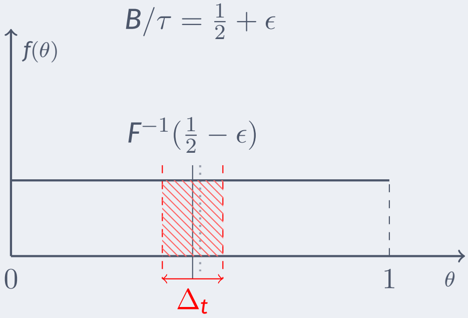

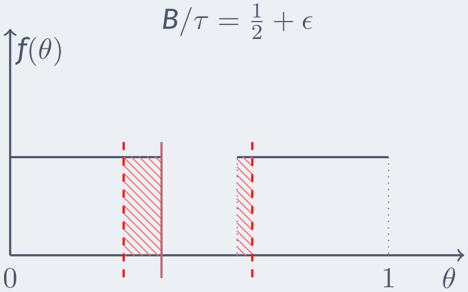

» Why does CE fail?

What if?

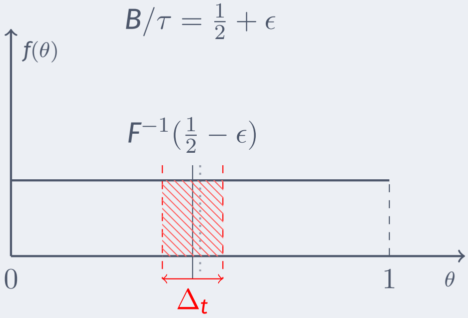

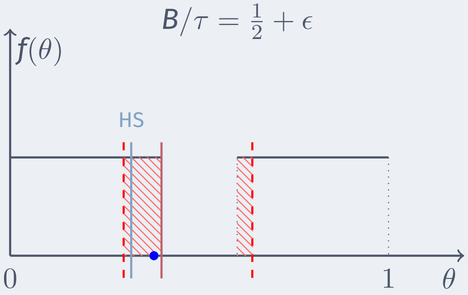

» Why does CE fail?

What if?

$B/\tau = \frac{1}{2} + \epsilon$

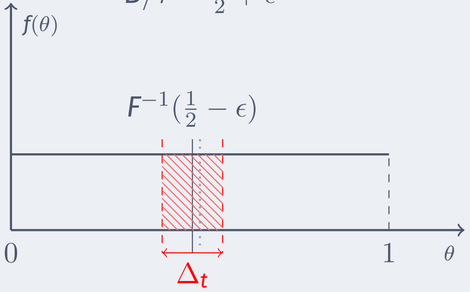

$B/\tau = \frac{1}{2} + \epsilon$

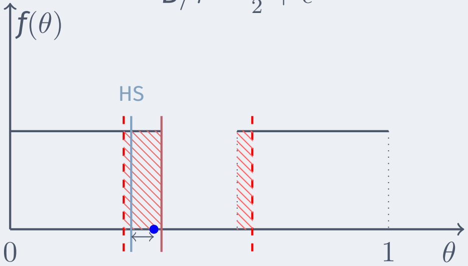

This is $\tilde{O}(1/\tau)$ expected compensation

Good in theory but practically infeasible

* What is the conservativeness parameter I should use?
* How to find where these gaps are? What happens if gaps shift?
* E.g., no chance of deploying for Amazon's fulfillment problem

Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

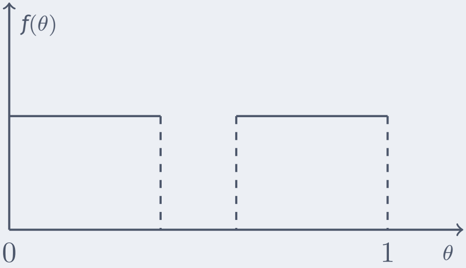

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

$$B/\tau = 1/4$$

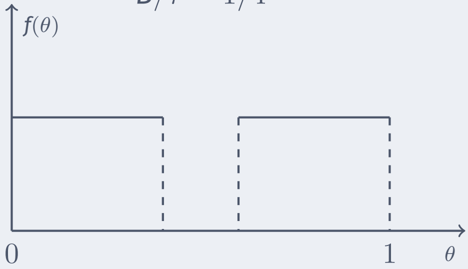

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

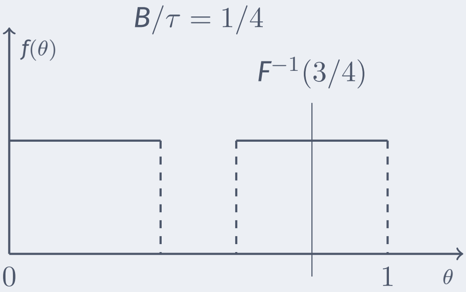

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

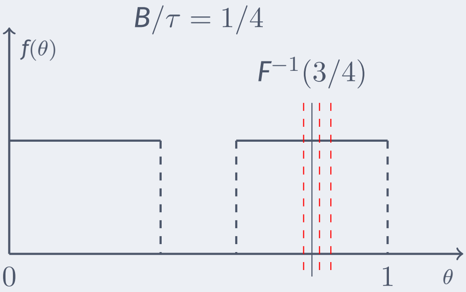

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

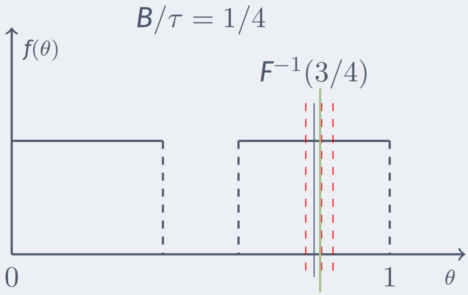

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

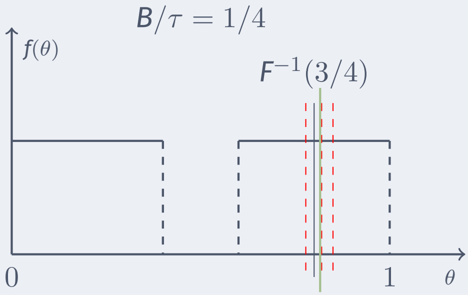

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

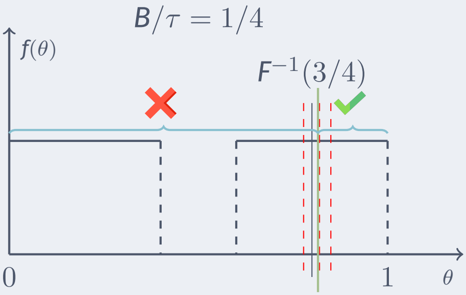

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

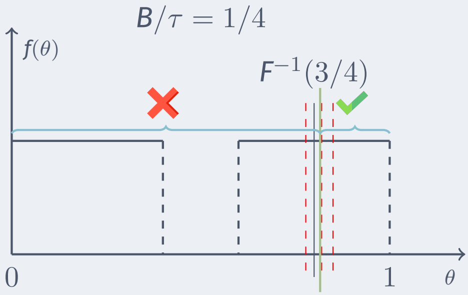

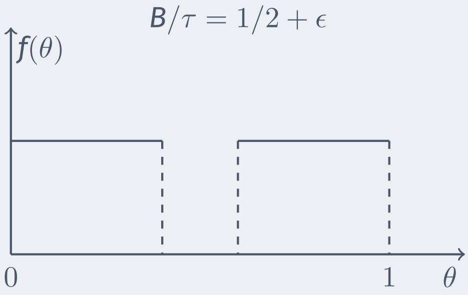

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

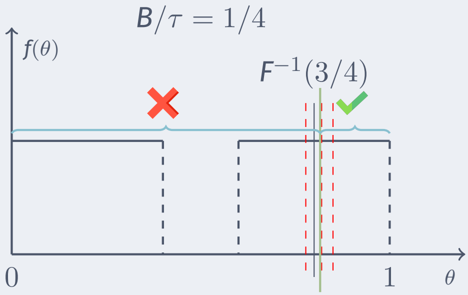

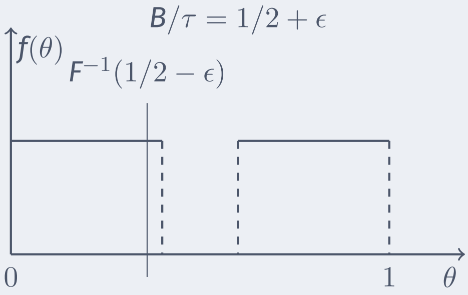

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

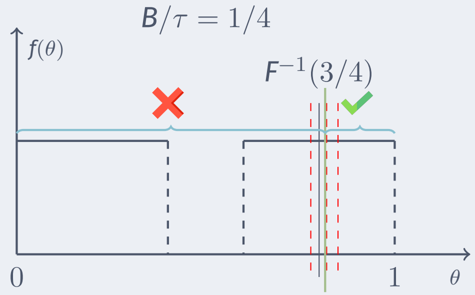

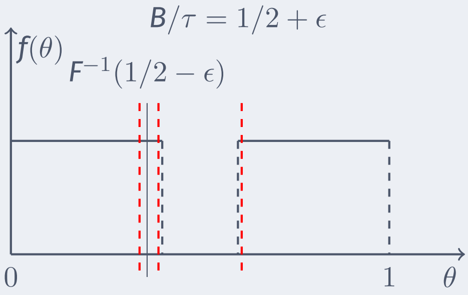

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

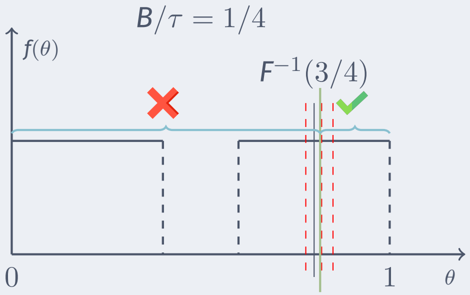

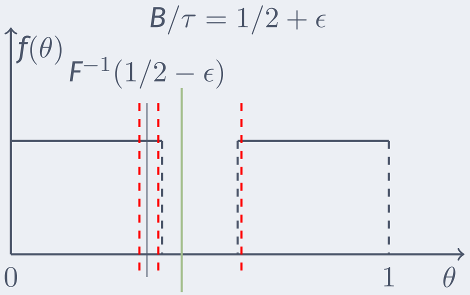

If far from a gap, use the CE threshold

» Conservativeness with respect to gaps

Algorithmic Idea: Simulate into the future

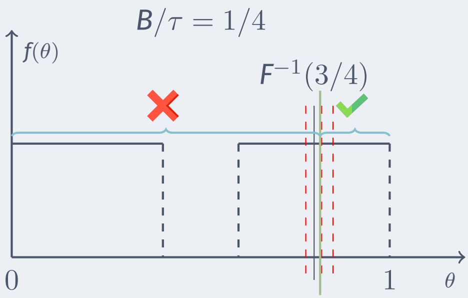

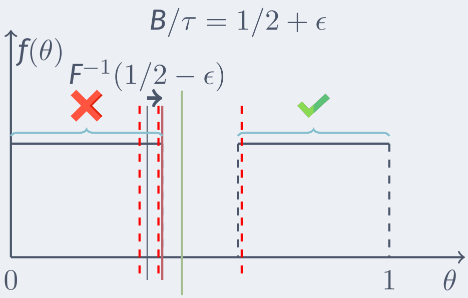

If far from a gap, use the CE threshold

If close to gap, use the gap as threshold

» Conservativeness with respect to gaps

Punchline

Regret of RAMS Policy

If F is a (β, ε_0)-clustered distribution, then

$$\mathbb{E}\left[\text{Regret(RAMS)}\right] = \begin{cases} \mathcal{O}\left((\log T)^2\right), & \beta = 0, \\ \mathcal{O}\left(\text{poly}(\log T) T^{\frac{1}{2} - \frac{1}{2(1+\beta)}}\right), & \beta > 0 \end{cases}$$

If F is a discrete distribution, $\mathbb{E}\left[\text{Regret (RAMS)}\right] = \mathcal{O}(1/\varepsilon_0)$

» Conservativeness with respect to gaps

Punchline

Regret of RAMS Policy

If F is a (β, ε_0)-clustered distribution, then

$$\mathbb{E}\left[\text{Regret(RAMS)}\right] = \begin{cases} \mathcal{O}\left((\log T)^2\right), & \beta = 0, \\ \mathcal{O}\left(\text{poly}(\log T) T^{\frac{1}{2} - \frac{1}{2(1+\beta)}}\right), & \beta > 0 \end{cases}$$

If F is a discrete distribution, $\mathbb{E}\left[\text{Regret (RAMS)}\right] = \mathcal{O}(1/\varepsilon_0)$

Remark

* $F = \text{Unif}([0, \frac{1}{4}] \cup [\frac{3}{4}, 1])$, RAMS ($\mathcal{O}((\log T)^2)$) outperforms CE ($\Omega(\sqrt{T})$).
* Matches the universal lower bound upto polylog factors $\Longrightarrow$ RAMS is near-optimal.

» Multi-secretary with general distributions

Brief Summary

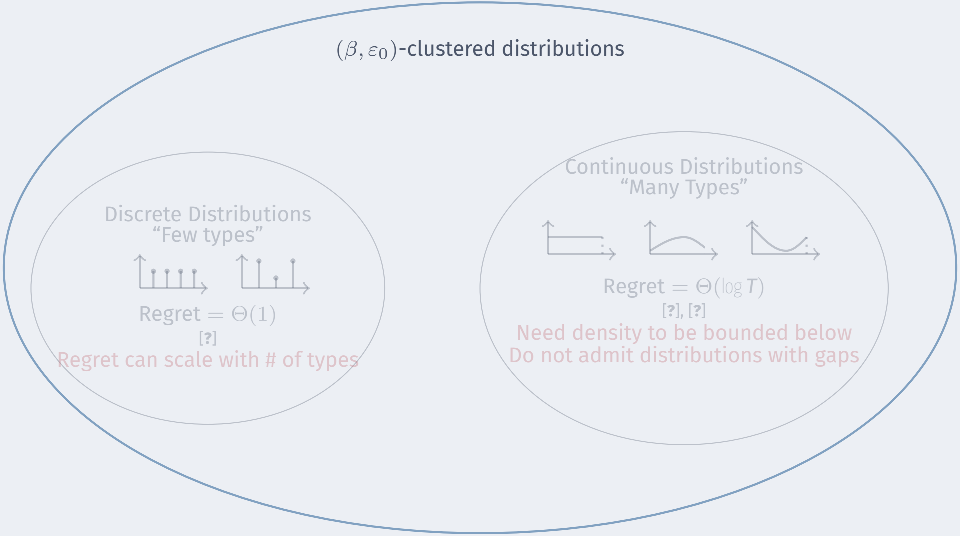

» Multi-secretary with general distributions

Brief Summary

(β, ε_0)-clustered distributions

$$\mathbb{E}\left[\text{Regret}(DP)\right] = \begin{cases} \Omega\left(\log T\right), & \beta = 0, \\ \Omega\left(T^{\frac{1}{2}-\frac{1}{2(1+\beta)}}\right), & \beta > 0 \end{cases}$$

Discrete Distributions
"Few types"

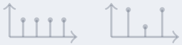

Regret $= \Theta(1)$
[?]
Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$
[?], [?]
Need density to be bounded below
Do not admit distributions with gaps

» Multi-secretary with general distributions

Brief Summary

(β, ε_0)-clustered distributions

$$\mathbb{E}\left[\text{Regret}(DP)\right] = \begin{cases} \Omega\left(\log T\right), & \beta = 0, \\ \Omega\left(T^{\frac{1}{2}-\frac{1}{2(1+\beta)}}\right), & \beta > 0 \end{cases}$$

Discrete Distributions
"Few types"

Regret $= \Theta(1)$

[?]

Regret can scale with # of types

Continuous Distributions
"Many Types"

Regret $= \Theta(\log T)$

[?], [?]

Need density to be bounded below
Do not admit distributions with gaps

$$\mathbb{E}\left[\text{Regret(RAMS)}\right] = \begin{cases} \mathcal{O}\left((\log T)^2\right), & \beta = 0, \\ \mathcal{O}\left(\text{poly}(\log T)T^{\frac{1}{2}-\frac{1}{2(1+\beta)}}\right), & \beta > 0 \end{cases}$$

» One Policy to solve them all?

* The multi-secretary problem is special but RAMS is general: in each period, simulate several futures and choose the action which minimizes the expected "compensation" in hindsight. Compensation $\equiv$ How to much we need to pay an agent who knows the future to take a particular action, for a given future.
* Can be applied to NRM and stochastic online matching problems to recover almost all known guarantees in the literature.

» One Policy to solve them all?

Proposition (RAMS is as good as any algorithm)

Given an NRM setting P, consider *any* algorithm A for P, such that with τ periods remaining, uniformly over the state, the expected compensation under A is bounded above by $\delta_\tau(\text{A})$. Then RAMS achieves an expected compensation bounded uniformly by $\delta_\tau + 1/\tau^{1.1}$. As a result the regret of RAMS is bounded above by a constant plus the regret guarantee for algorithm A,

$$\text{Regret(RAMS)} \leq \text{Constant} + \sum_{\tau=1}^{T} \delta_\tau(\text{A}) \,.$$

» What to take away from this talk?

Simple and practical simulation-based policy SOAR is broadly applicable:

* Recovers the guarantees for almost all settings in the NRM literature (e.g., constant regret for finite types, $\log^2 T$ for semi-infinite types)
* Establishes novel guarantees for dynamic spatial matching problems

Thank you!

APPENDIX on (β, ε_0)-clustered distributions

» (β, ε_0)-clustered distribution

Examples

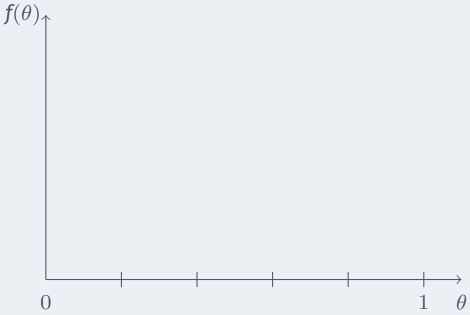

» (β, ε_0)-clustered distribution

Examples

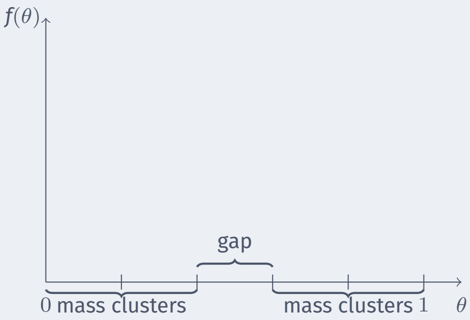

» (β, ε_0)-clustered distribution

Examples

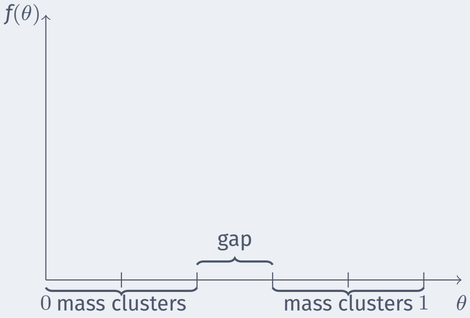

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

» (β, ε_0)-clustered distribution

Examples

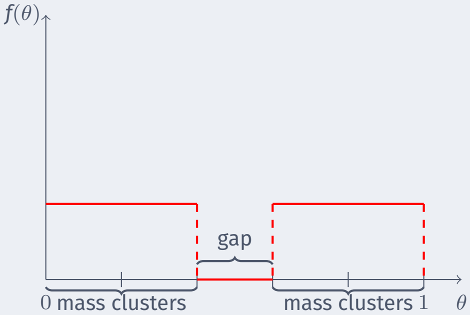

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 0$ (mass accumulation around gaps)

» (β, ε_0)-clustered distribution

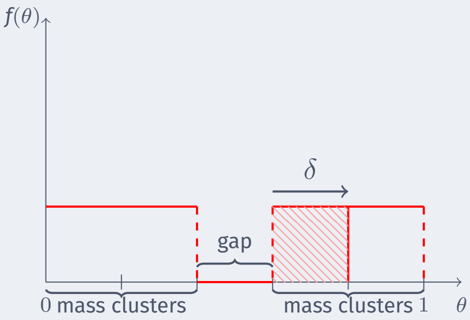

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 0$ (mass accumulation around gaps)

» (β, ε_0)-clustered distribution

Examples

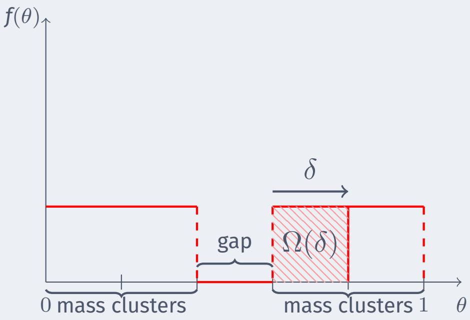

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 0$ (mass accumulation around gaps)

$|F(m+\delta) - F(m)| \geq \delta$ on the same mass cluster

» (β, ε_0)-clustered distribution

Examples

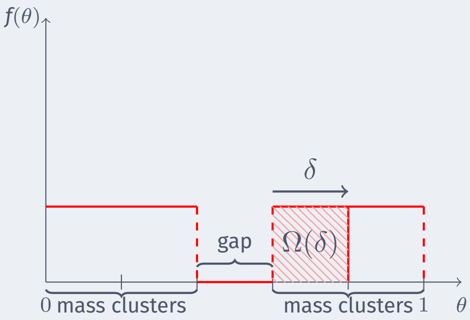

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 0$ (mass accumulation around gaps)

$|F(m+\delta) - F(m)| \geq \delta$ on the same mass cluster

μ(mass clusters) $\geq \varepsilon_0$

» (β, ε_0)-clustered distribution

Examples

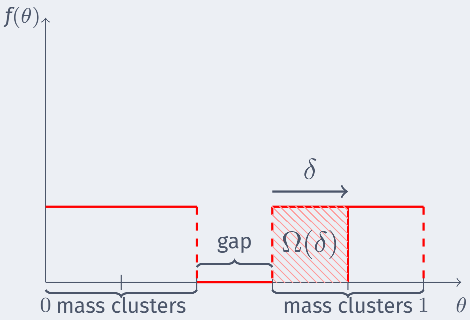

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 0$ (mass accumulation around gaps)

$|F(m+\delta) - F(m)| \geq \delta$ on the same mass cluster

μ(mass clusters) $\geq \varepsilon_0$

For discrete distrbutions, $\beta = 0, \varepsilon_0 = \min_j p_j$

» (β, ε_0)-clustered distribution

Examples

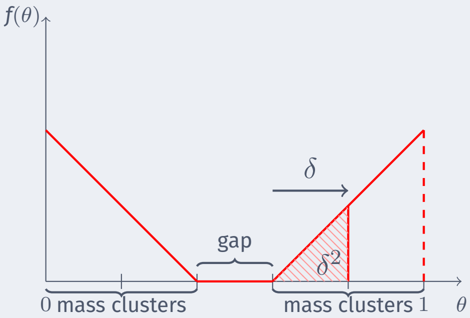

Gap $\equiv$ intervals of positive length with zero mass

mass cluster $\equiv$ interval with positive mass

$\beta = 1$ (mass accumulation around gaps)

$|F(m+\delta) - F(m)| \geq \delta^2$ on the same mass cluster

μ(mass clusters) $\geq \varepsilon_0$