Dynamic Resource Allocation

Spectrum of Achievable Performances
& Algorithmic Design Principles

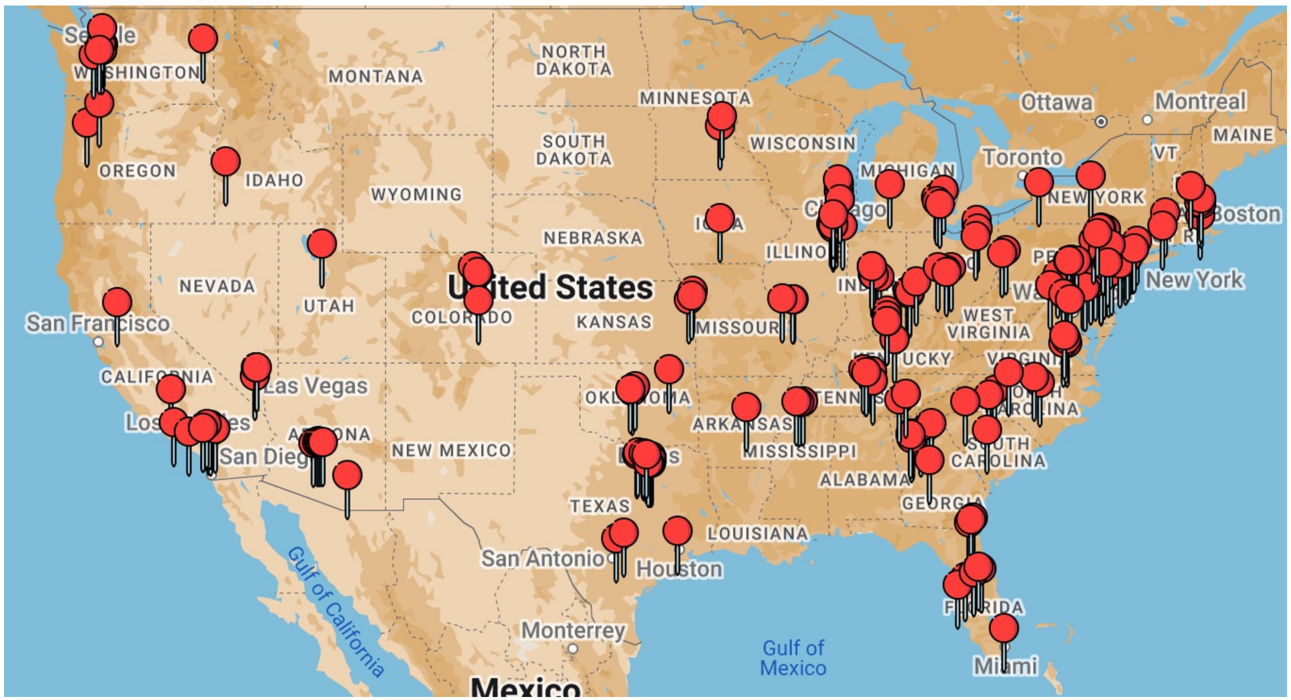

Akshit Kumar

Joint work with Omar Besbes and Yash Kanoria

Columbia Business School

Research Question

Research Question

A few (demand)
types are present

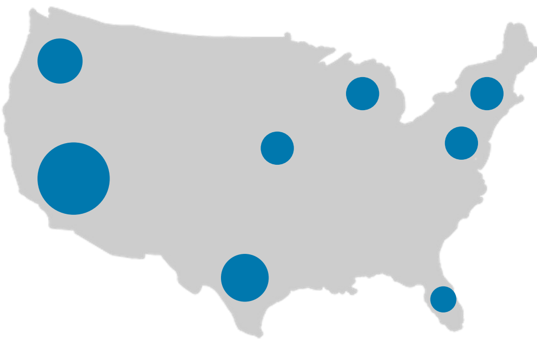

Theory

All (demand)
types are present

Theory

Research Question

A few (demand) types are present

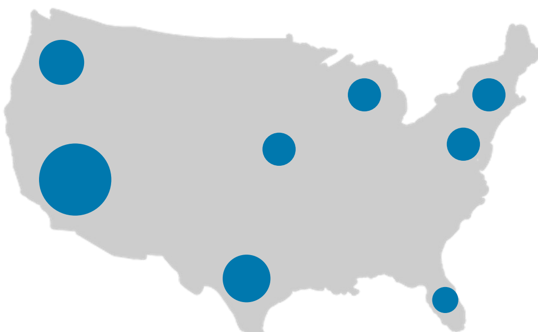

Theory

Many (demand) types are present

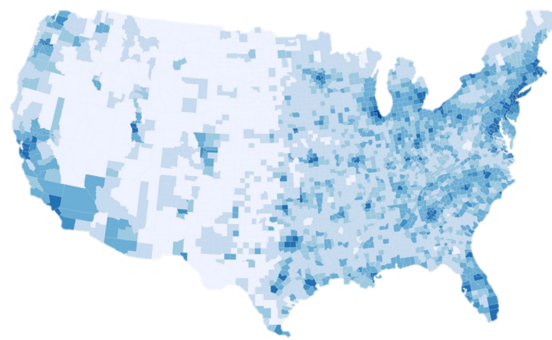

Practice

All (demand) types are present

Theory

Research Question

A few (demand) types are present | Many (demand) types are present | All (demand) types are present

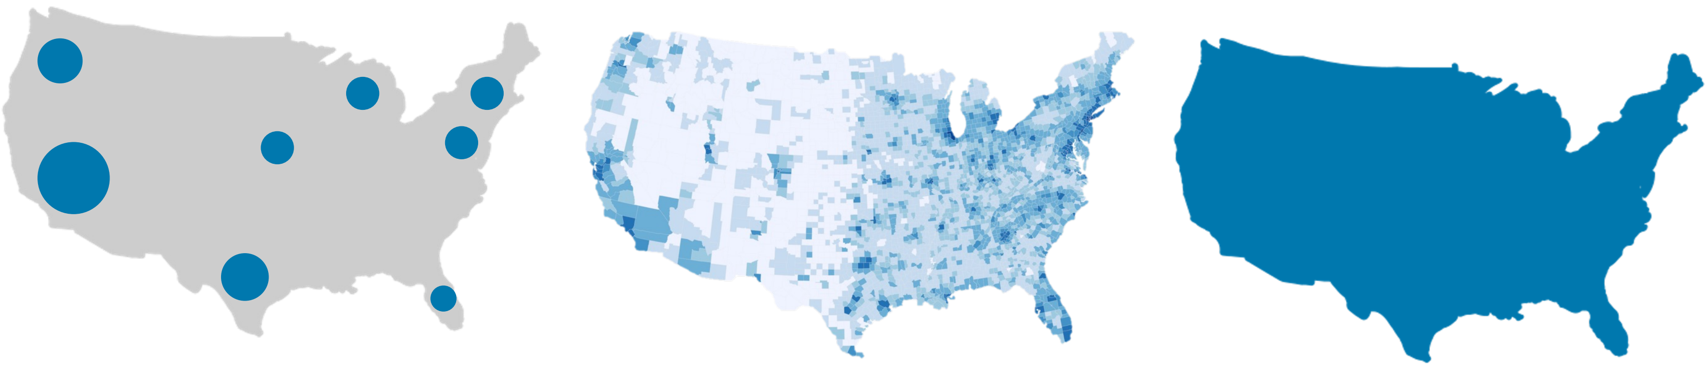

Theory | Practice | Theory

Entire spectrum of how to model (demand) types

Research Question

1. What is the interplay between the distribution of request types and achievable algorithmic performance?

Research Question

1. What is the interplay between the distribution of request types and achievable algorithmic performance?
2. Can we design a **unified**, **simple** and **near-optimal** algorithms which works for all type distributions?

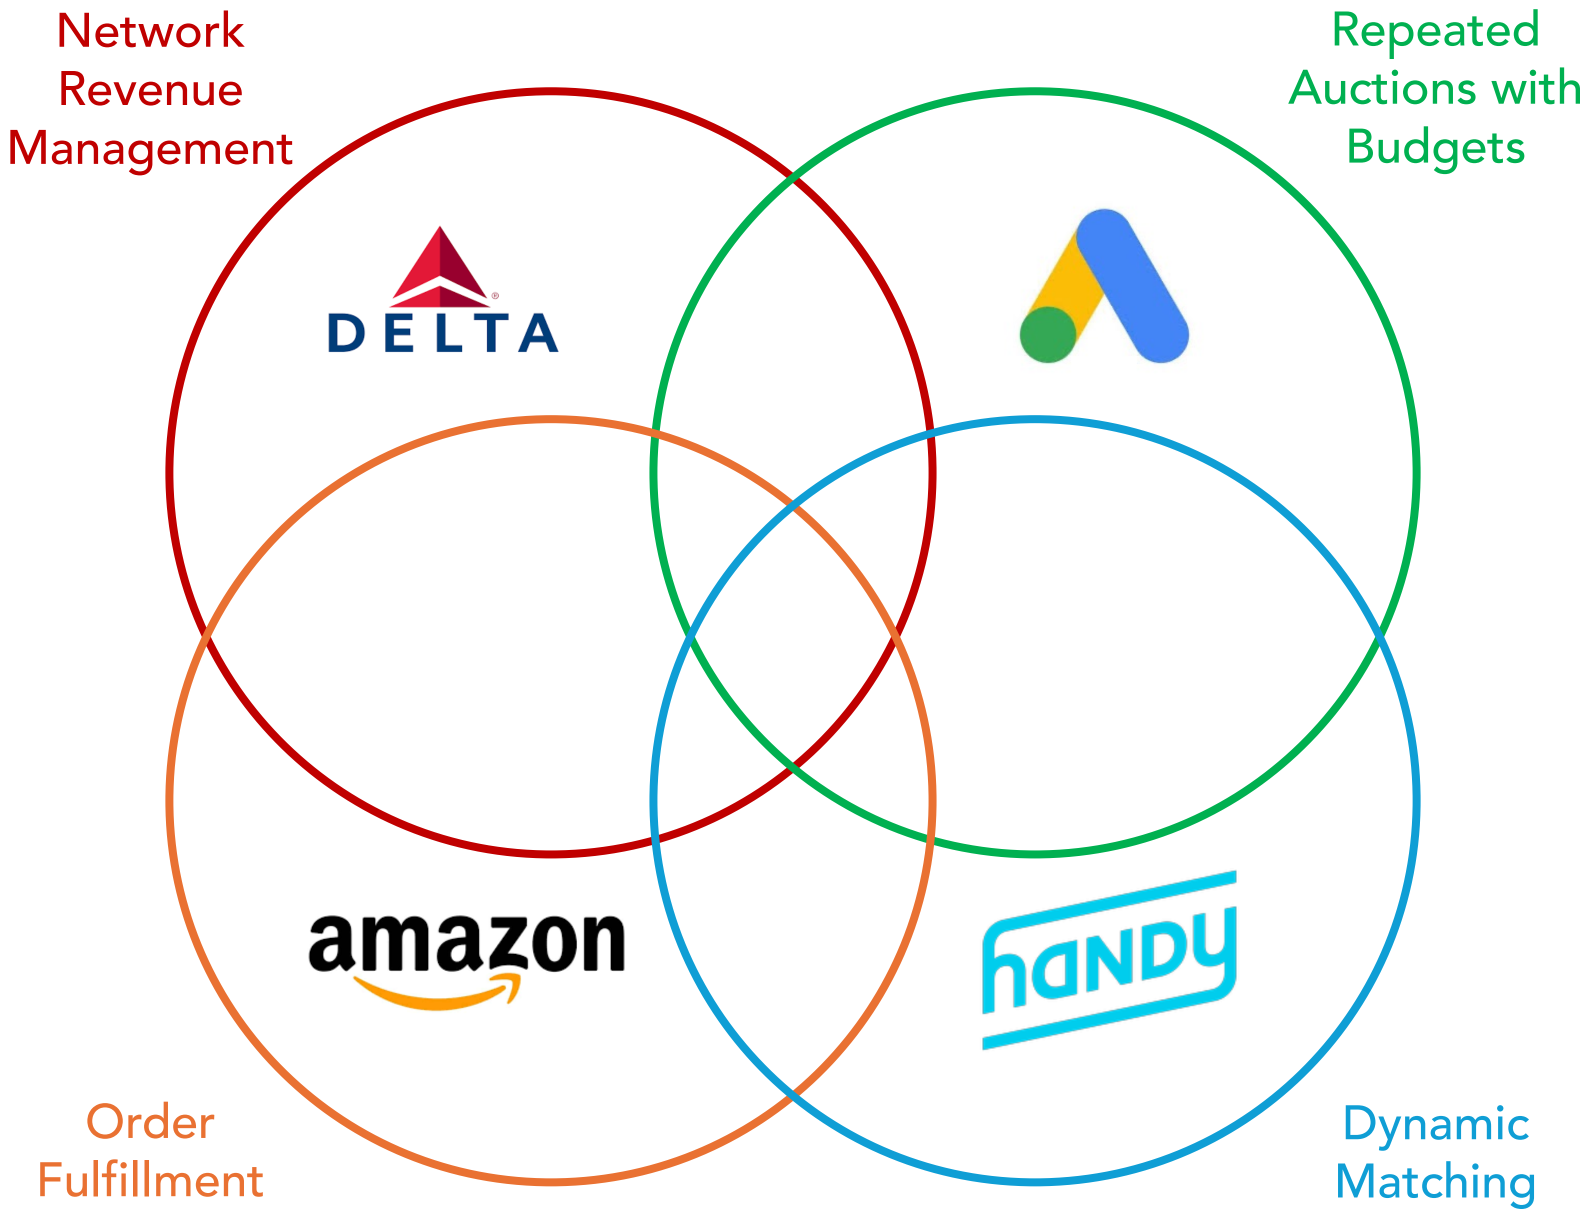
Network Revenue Management
Repeated Auctions with Budgets
Order Fulfillment
Dynamic Matching

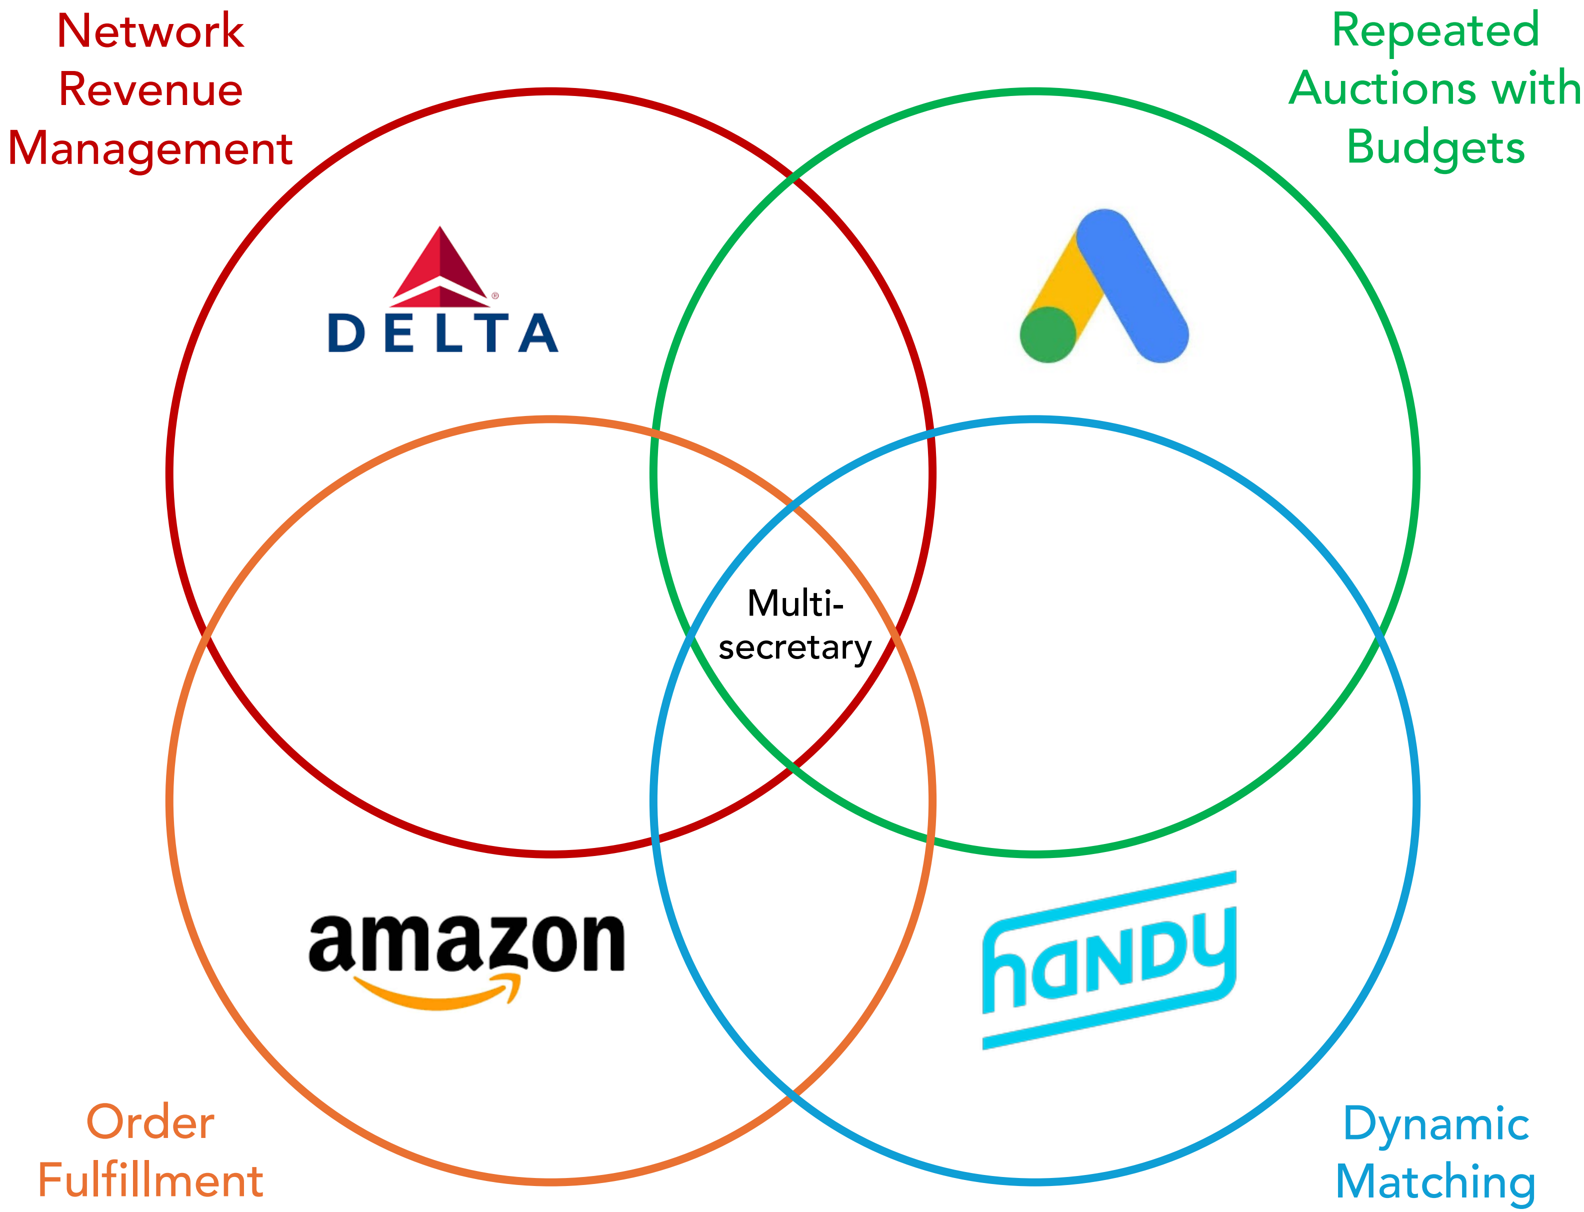
Network
Revenue
Management
Repeated
Auctions with
Budgets
DELTA
Multi-
secretary
amazon
handy
Order
Fulfillment
Dynamic
Matching

Multi-secretary Problem

Multi-secretary Problem

A few types are present

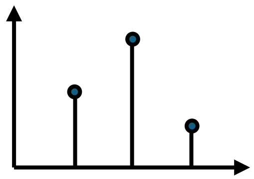

All types are present

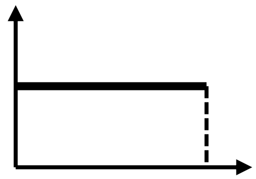

Multi-secretary Problem

A few types are present

Bounded Regret

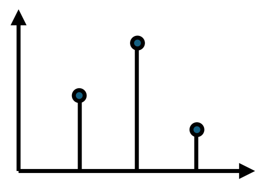

All types are present

Logarithmic Regret

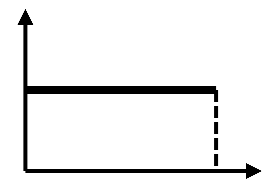

Regret is the additive gap b/w the value of hindsight opt. and value under some algorithm

Multi-secretary Problem

A few types are present

Bounded Regret

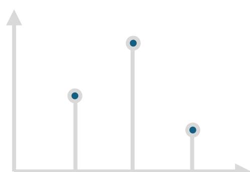

All types are present

Logarithmic Regret

β –clustered distributions

Regret is the additive gap b/w the value of hindsight opt. and value under some algorithm

Multi-secretary Problem

A few types are present
Bounded Regret

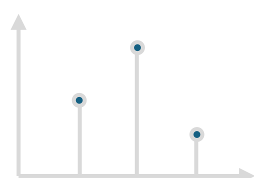

All types are present
Logarithmic Regret

Entire spectrum of regret scalings is possible

β –clustered distributions

Regret is the additive gap b/w the value of hindsight opt. and value under some algorithm

Multi-secretary Problem

one algorithm to solve them all

Repeatedly Act using Multiple Simulations (RAMS)

A few types
are present
Bounded Regret

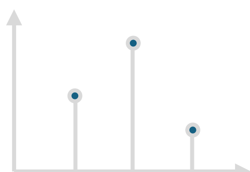

All types are
present
Logarithmic Regret

Entire spectrum of regret scalings is possible

β –clustered distributions

Regret is the additive gap b/w the value of hindsight opt. and value under some algorithm

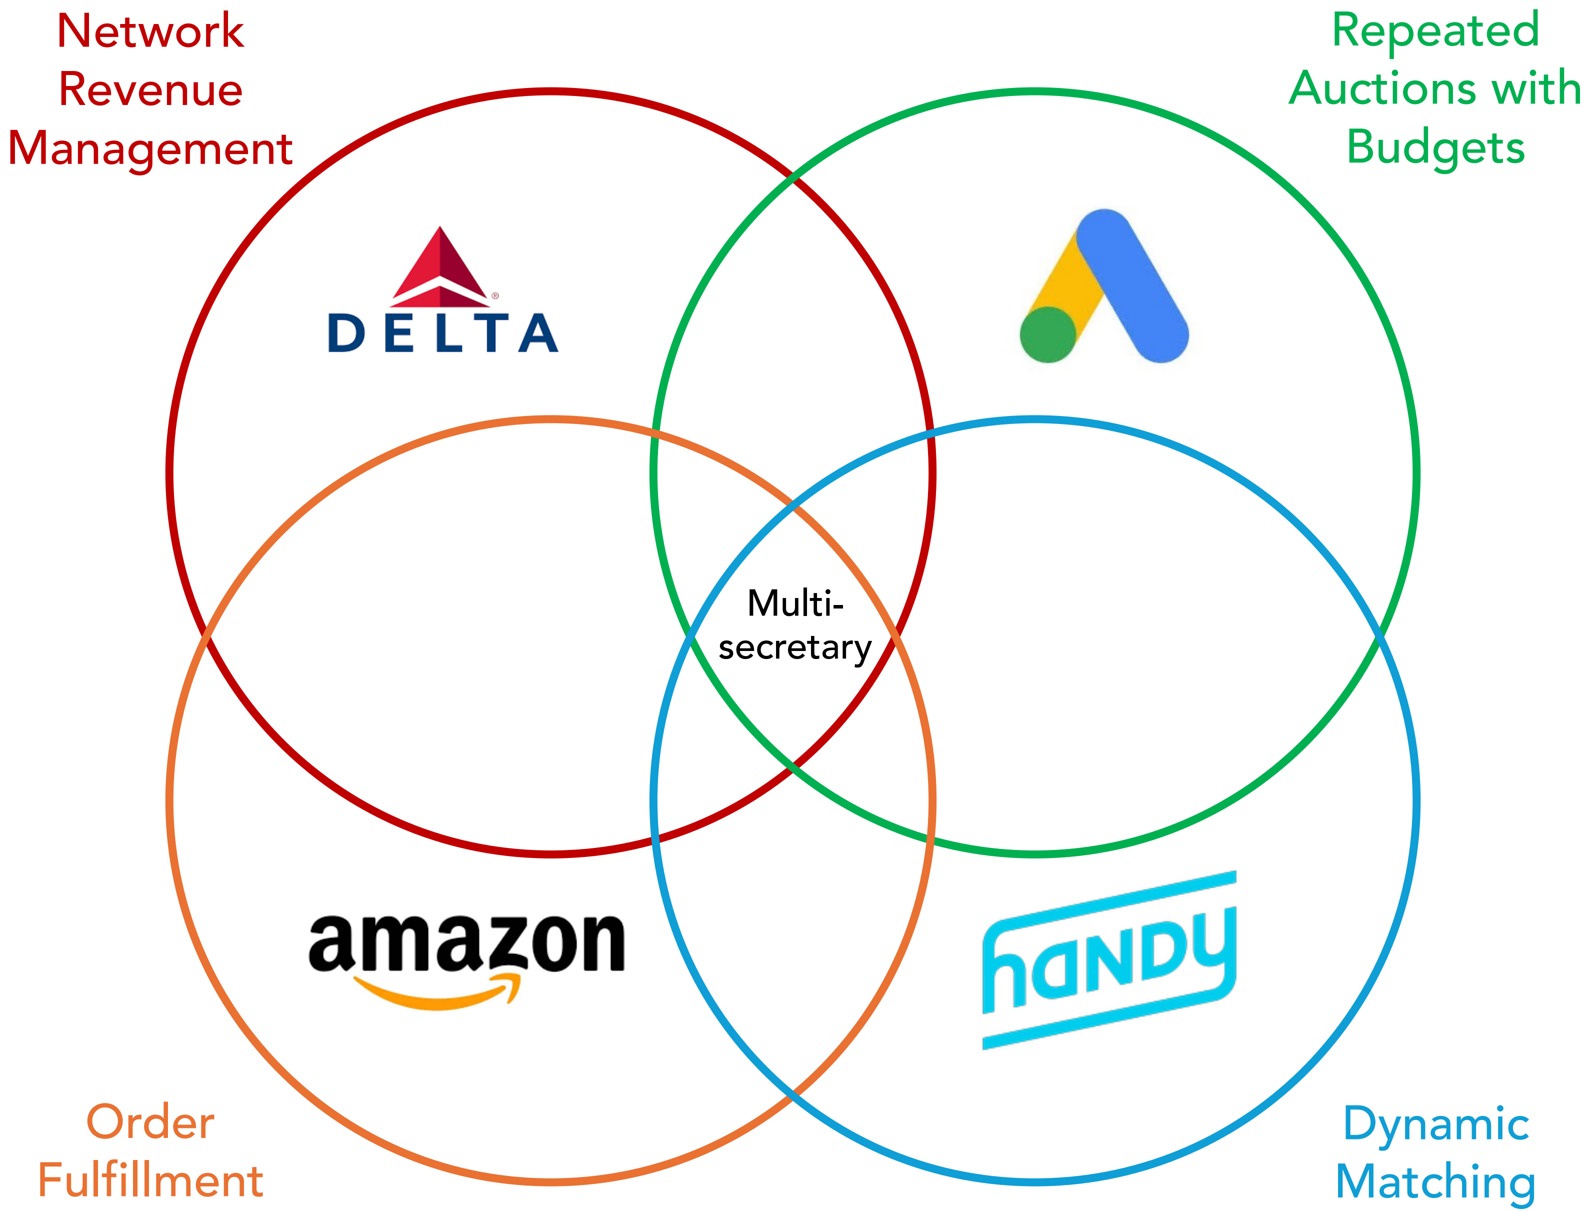
Network
Revenue
Management
Repeated
Auctions with
Budgets
DELTA
Multi-
secretary
amazon
handy
Order
Fulfillment
Dynamic
Matching

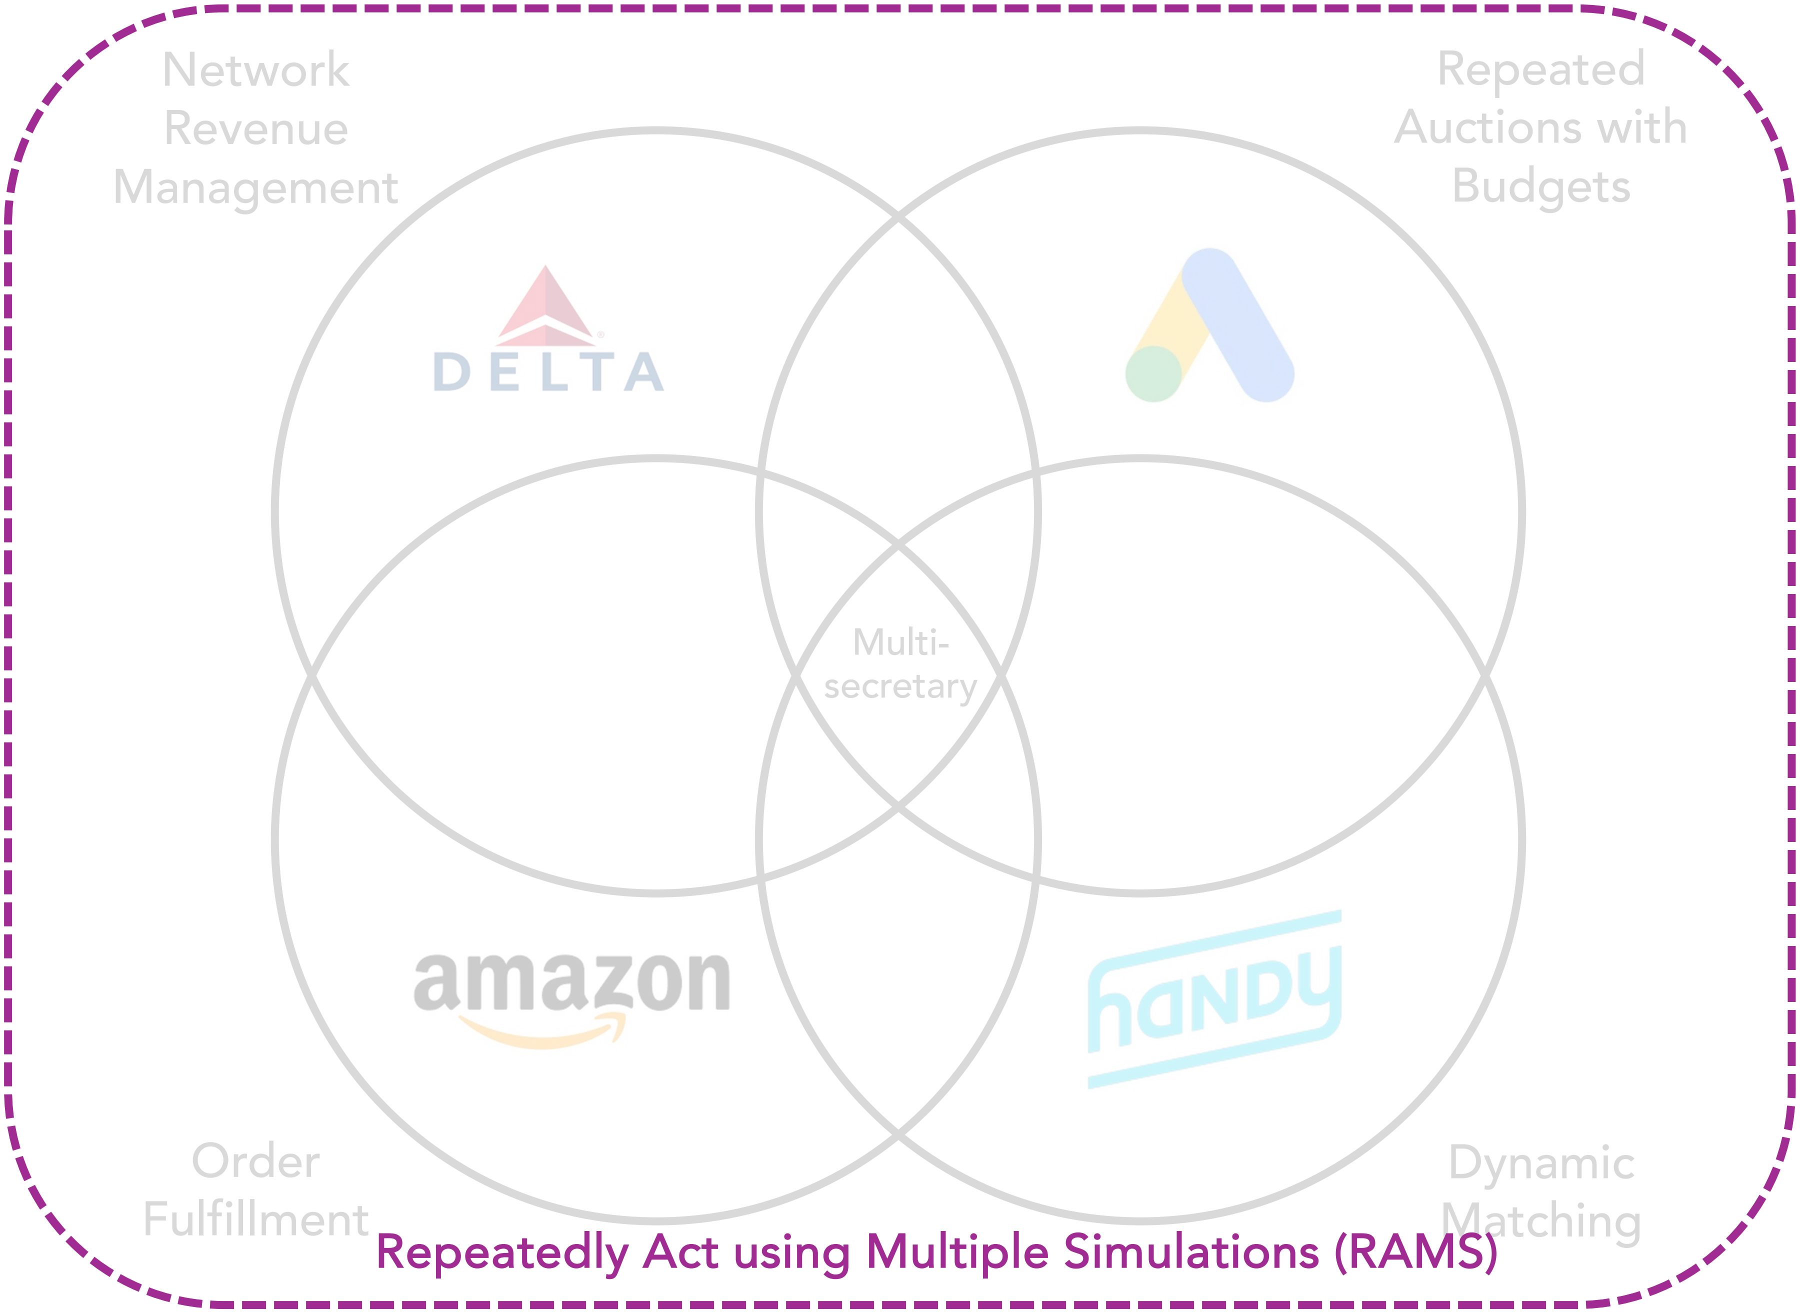
Network Revenue Management
Repeated Auctions with Budgets
Multi-secretary
Order Fulfillment
Dynamic Matching
Repeatedly Act using Multiple Simulations (RAMS)

Multi-secretary Problem

Online selection of top-B values

Online selection of top-B values

- Given a sequence of T values and budget B, the DM wants to select the top B values

Online selection of top-B values

- Given a sequence of T values and budget B, the DM wants to select the top B values
- The values arrive in an online fashion

Online selection of top-B values

- Given a sequence of T values and budget B, the DM wants to select the top B values
- The values arrive in an online fashion
- The DM must make irrevocable accept/reject decisions

Online selection of top-B values

- Given a sequence of T values and budget B, the DM wants to select the top B values
- The values arrive in an online fashion
- The DM must make irrevocable accept/reject decisions
- Assumption: The values are drawn i.i.d from a known distribution F

Online selection of top-B values

- Given a sequence of T values and budget B, the DM wants to select the top B values
- The values arrive in an online fashion
- The DM must make irrevocable accept/reject decisions
- Assumption: The values are drawn i.i.d from a known distribution F
- Performance Metric: Minimize Regret

  Regret(π) = (Expected Maximum Value in Hindsight) – (Expected Value under π)

Multi-secretary Problem

A few types are present
Bounded Regret

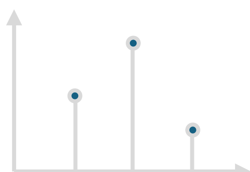

All types are present
Logarithmic Regret

Entire spectrum of regret scalings is possible

β –clustered distributions

Regret is the additive gap b/w the value of hindsight opt. and value under some algorithm

Complete Characterization of the Multi-secretary Problem

Complete Characterization of the Multi-secretary Problem

Complete Characterization of the Multi-secretary Problem

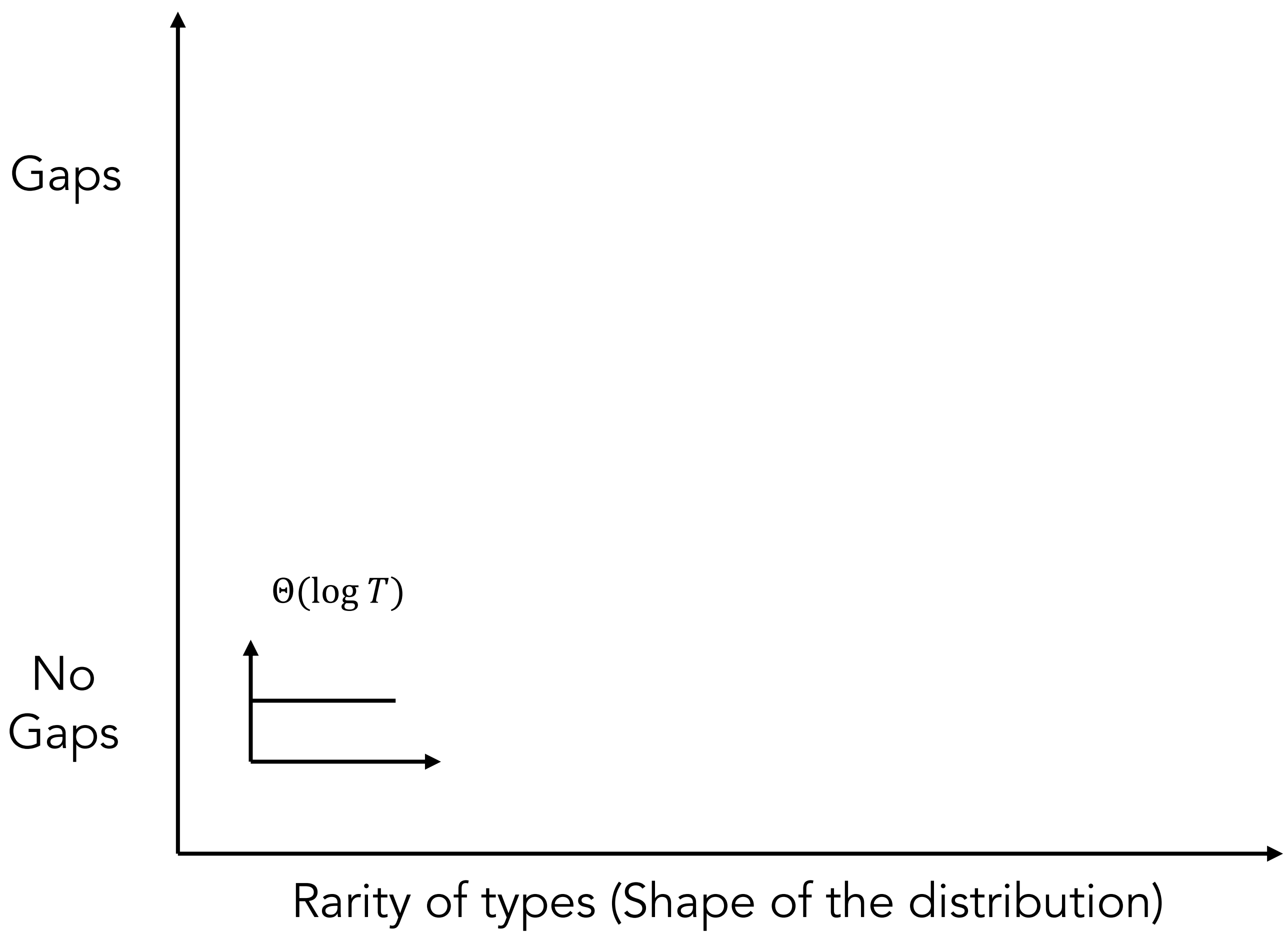

Complete Characterization of the Multi-secretary Problem

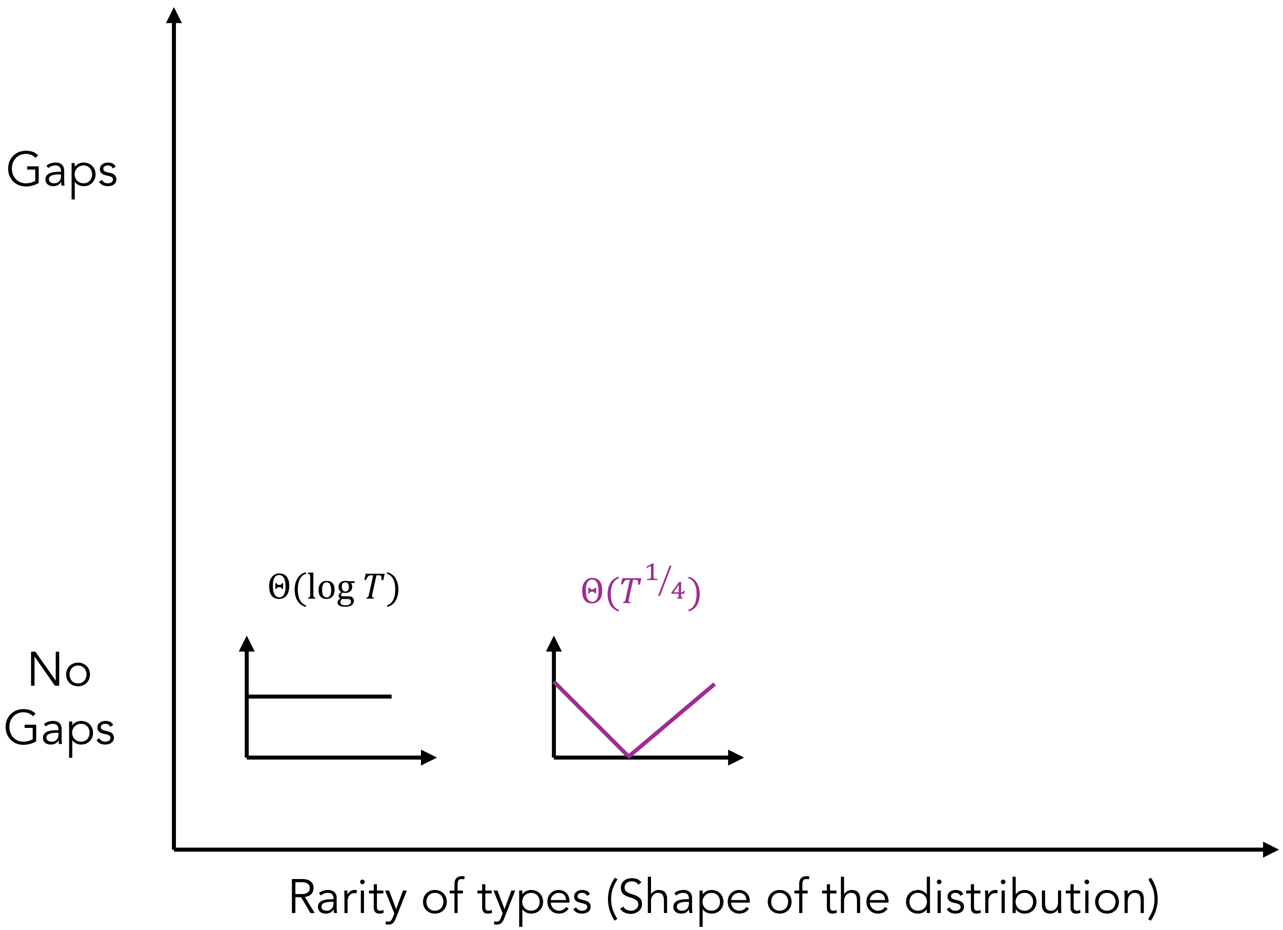

Complete Characterization of the Multi-secretary Problem

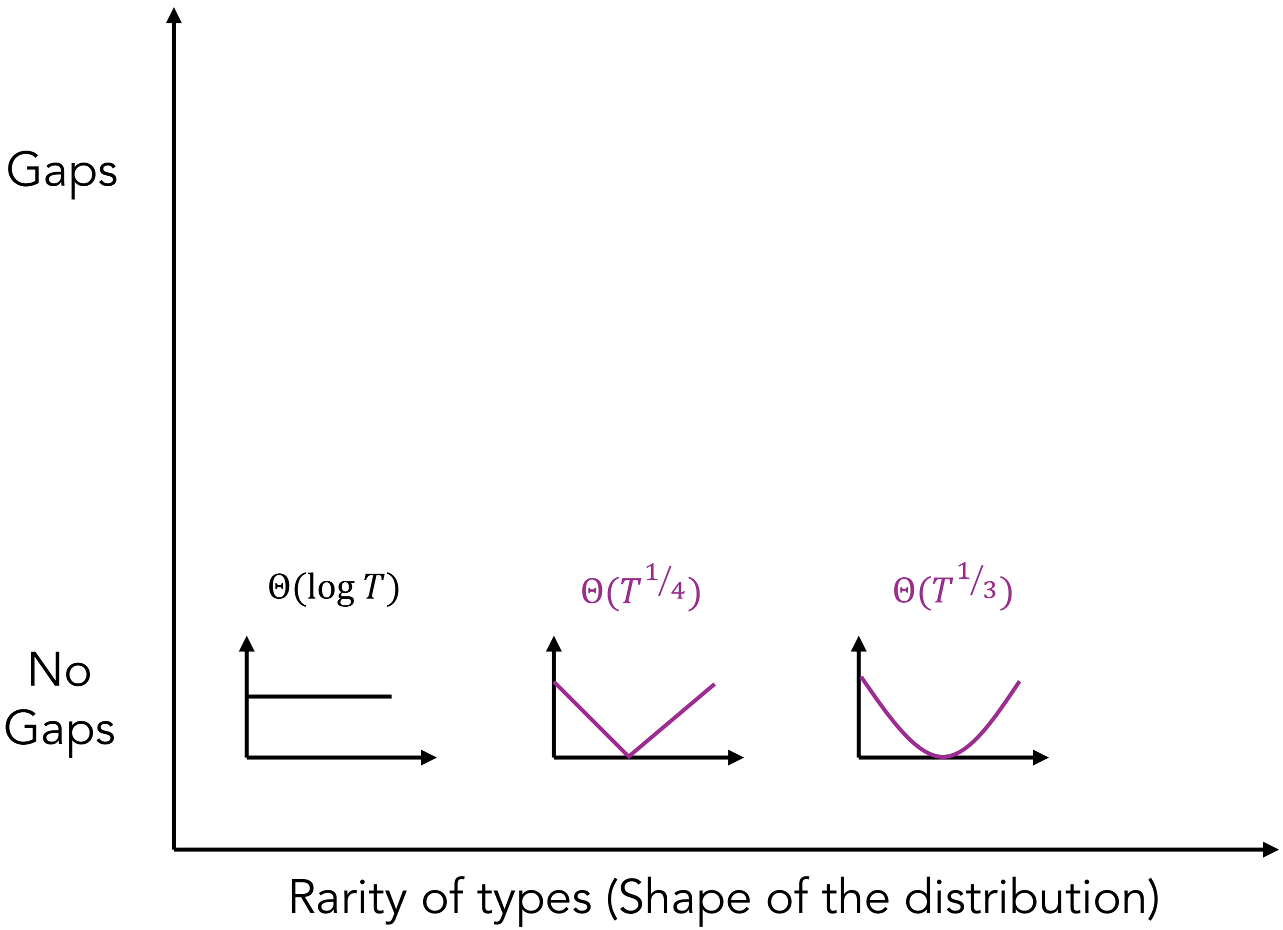

Complete Characterization of the Multi-secretary Problem

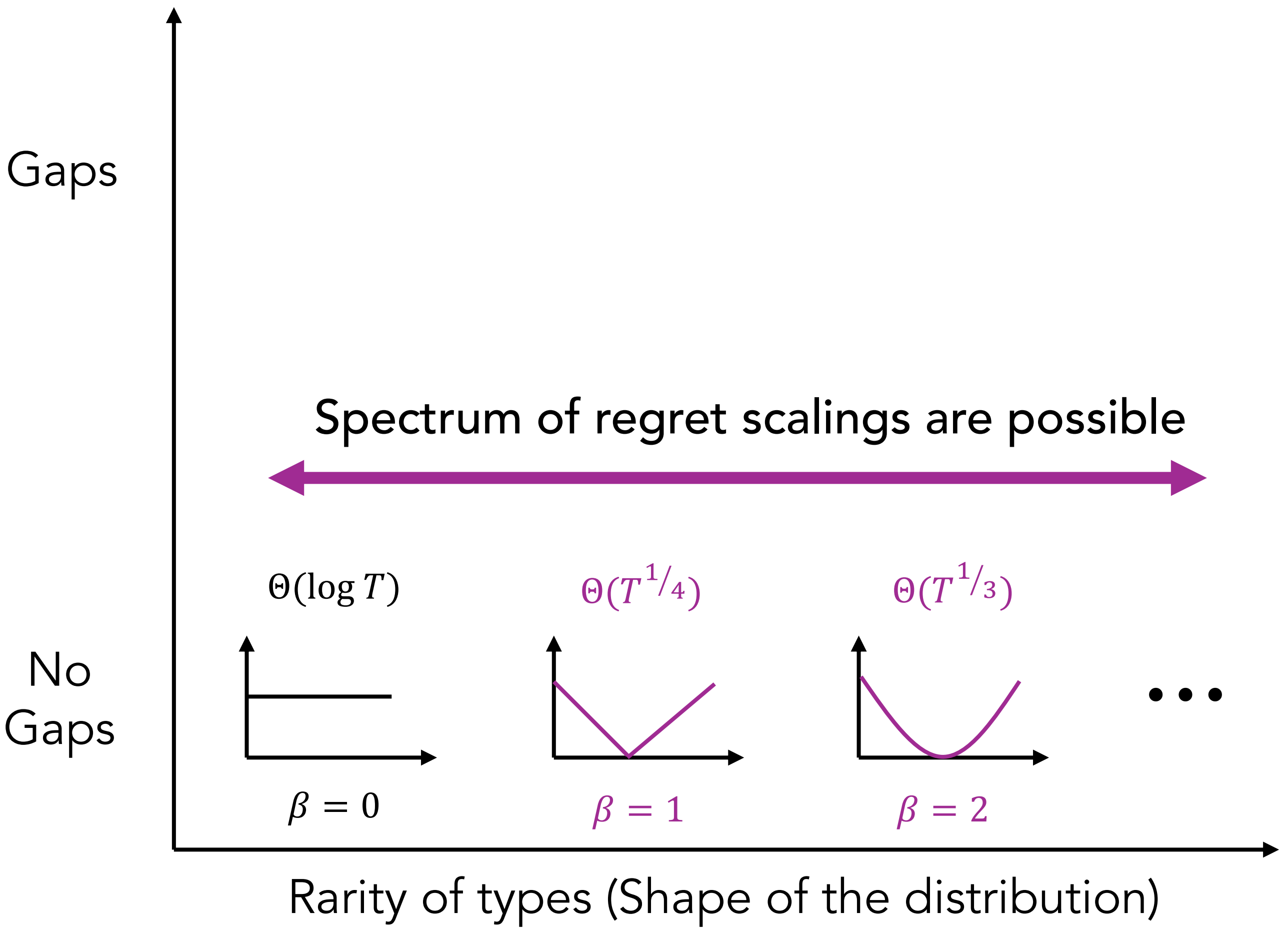

Complete Characterization of the Multi-secretary Problem

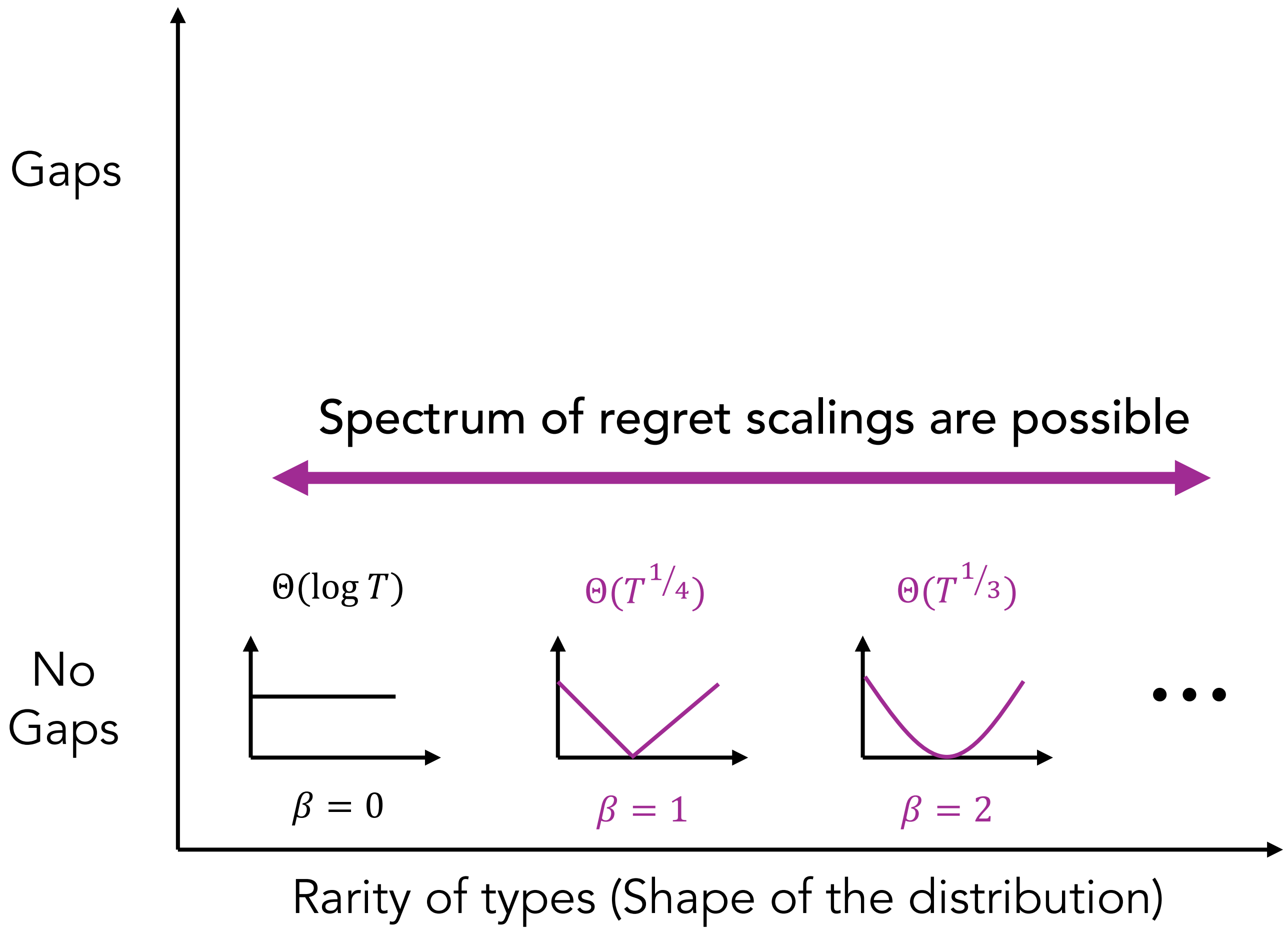

Distribution shape is a fundamental driver of performance

Complete Characterization of the Multi-secretary Problem

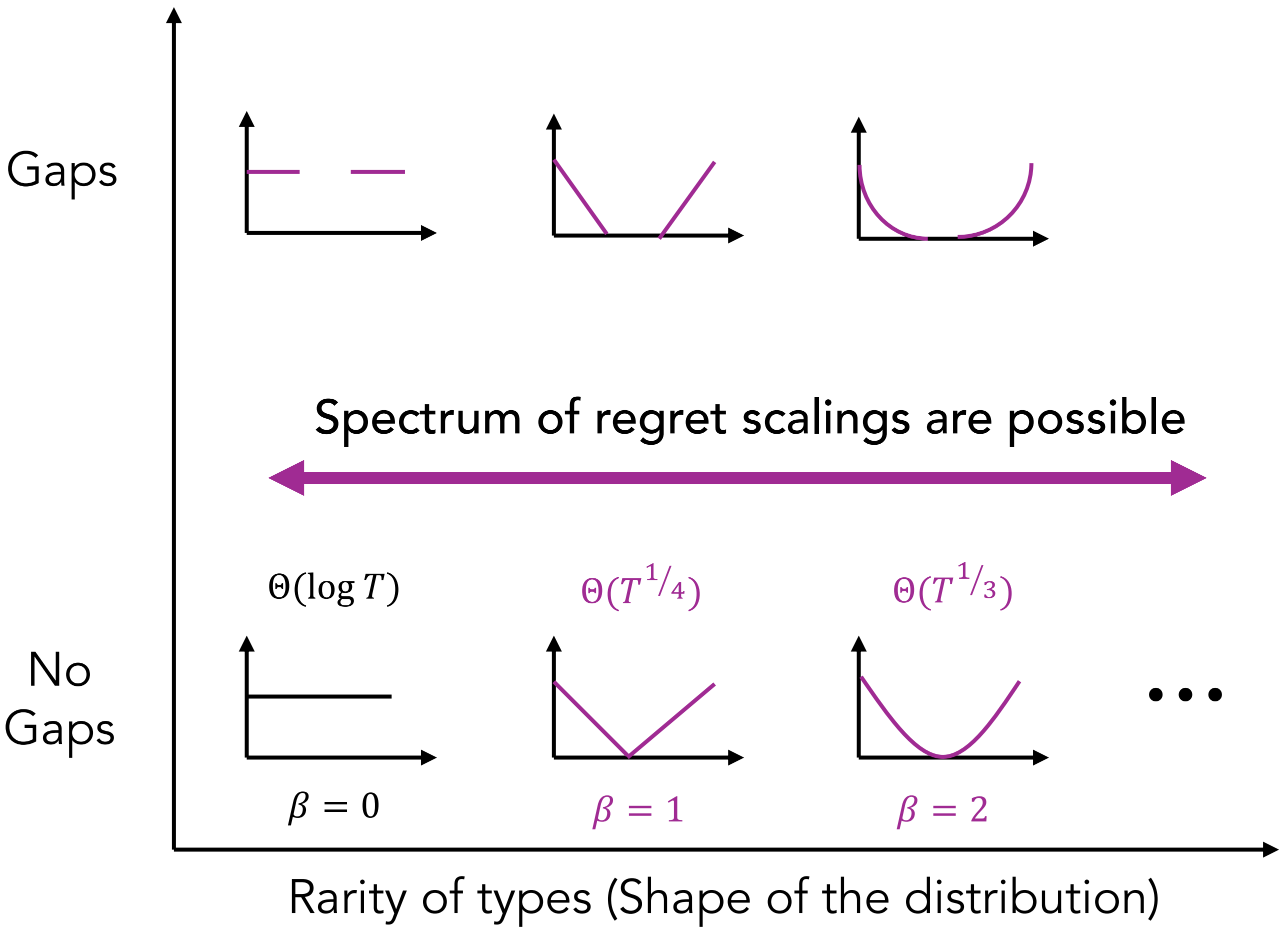

Distribution shape is a fundamental driver of performance

Complete Characterization of the Multi-secretary Problem

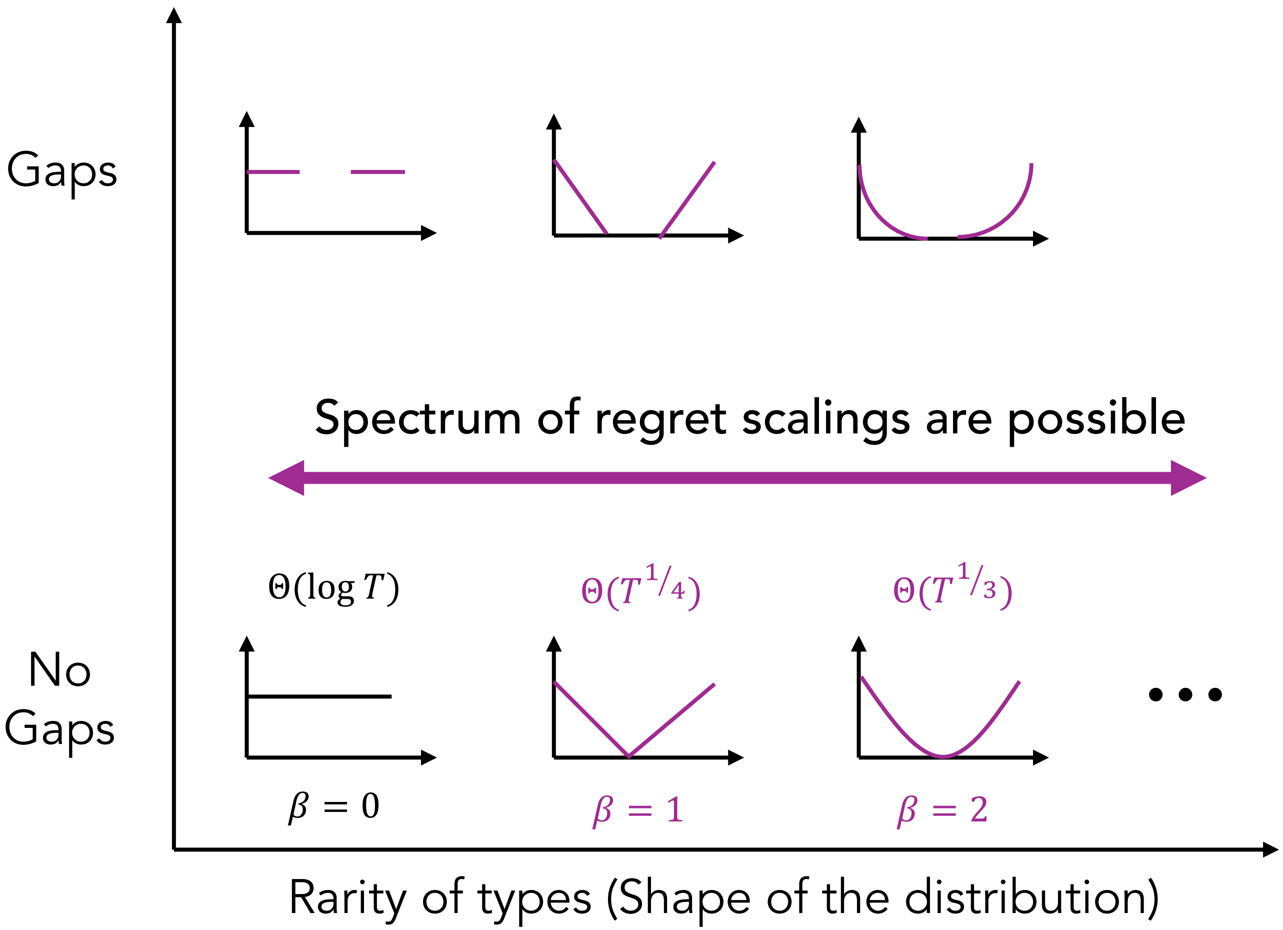

Distribution shape is a fundamental driver of performance

Dealing with gaps in an algorithmic challenge

Complete Characterization of the Multi-secretary Problem

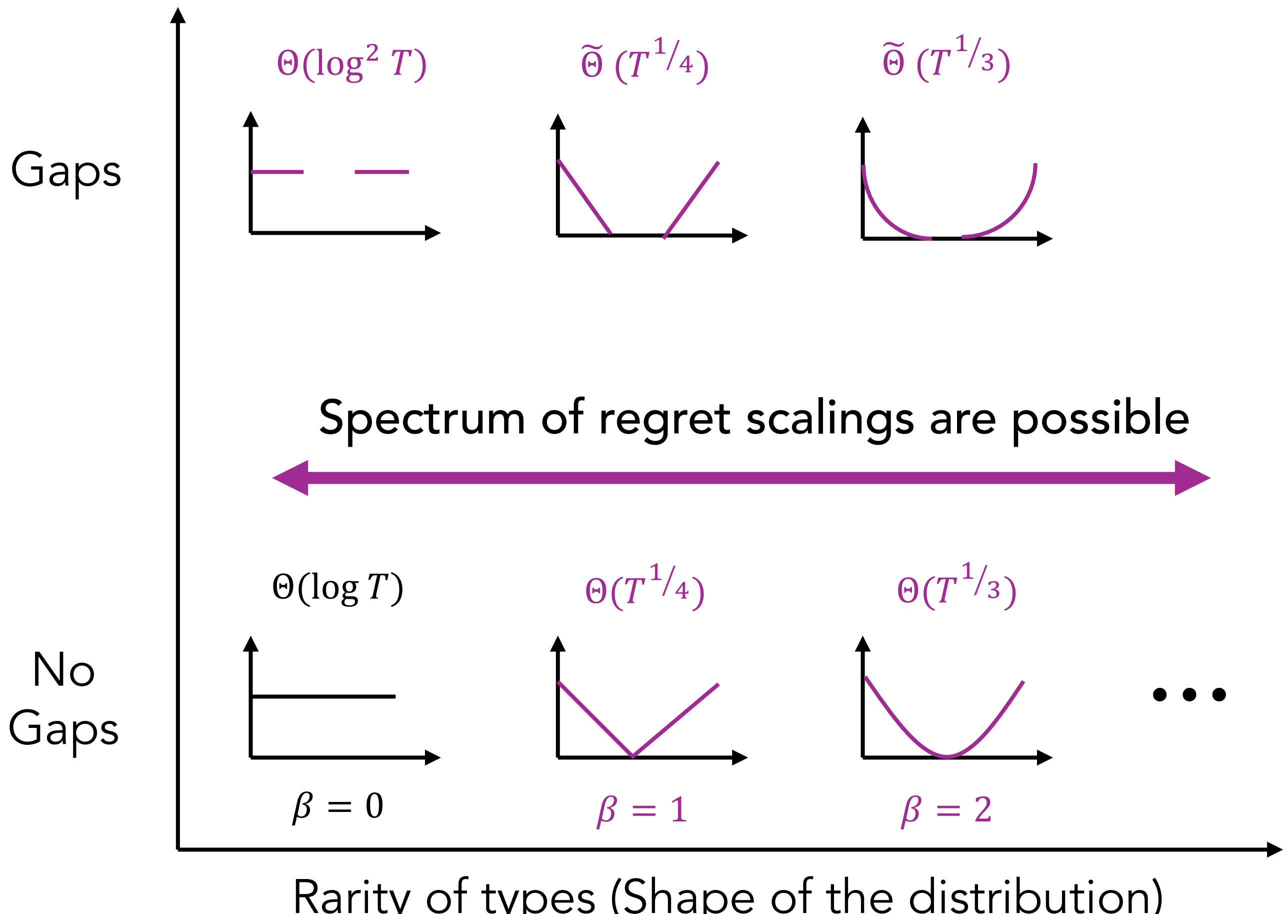

Distribution shape is a fundamental driver of performance

Dealing with gaps in an algorithmic challenge

Conservativeness with respect to gaps (CwG) principle enables near-optimal performance

Complete Characterization of the Multi-secretary Problem

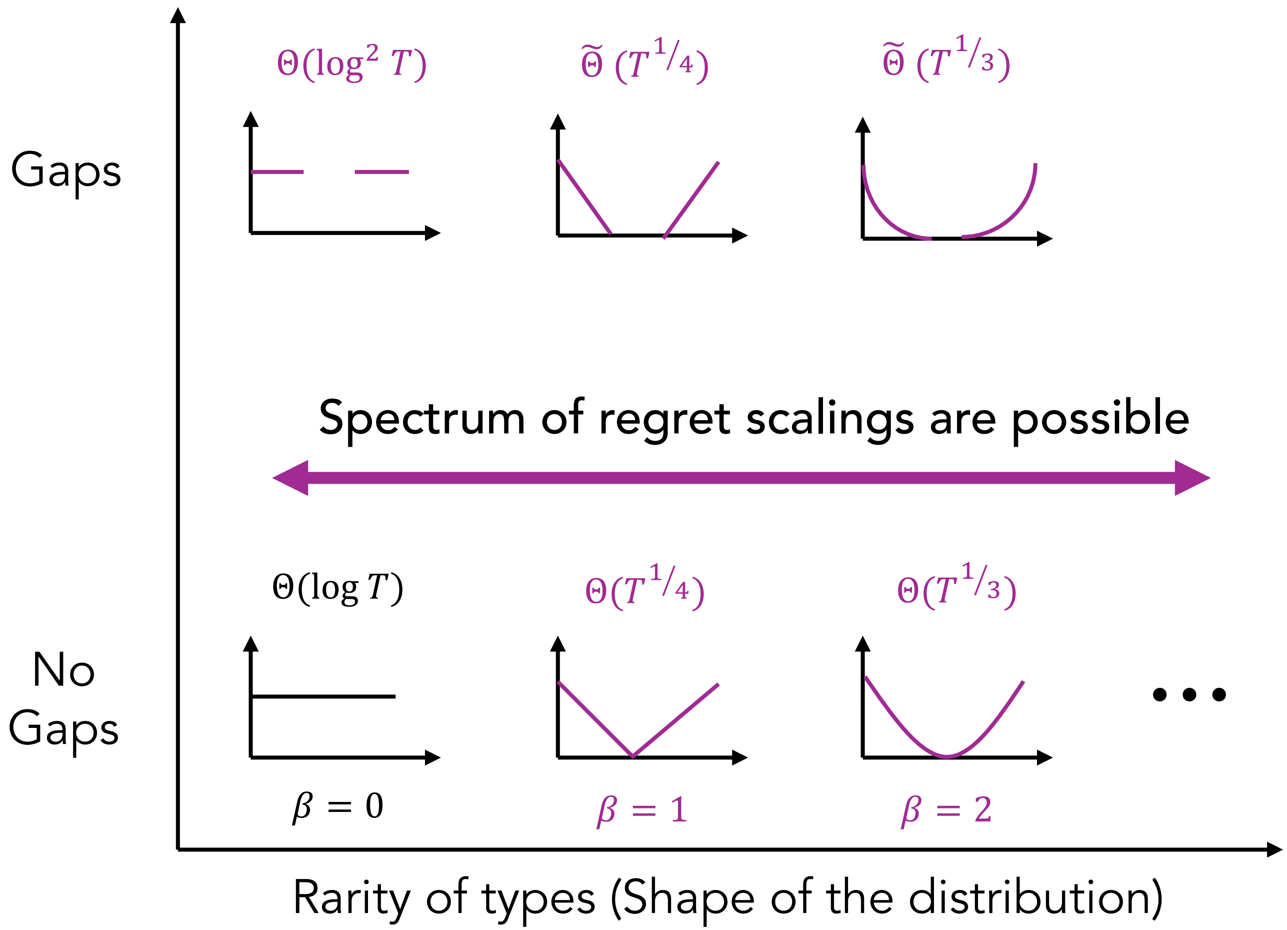

Distribution shape is a fundamental driver of performance

Dealing with gaps in an algorithmic challenge

Conservativeness with respect to gaps (CwG) principle enables near-optimal performance

Use RAMS to operationalize CwG

Multi-secretary Problem

one algorithm to solve them all

Repeatedly Act using Multiple Simulations (RAMS)

A few types are present

Bounded Regret

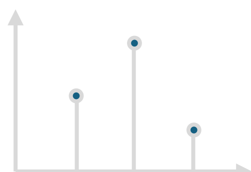

All types are present

Logarithmic Regret

Entire spectrum of regret scalings is possible

β –clustered distributions

Regret is the additive gap b/w the value of hindsight opt. and value under some algorithm

Gaps pose an algorithmic challenge

Gaps pose an algorithmic challenge

Gaps pose an algorithmic challenge

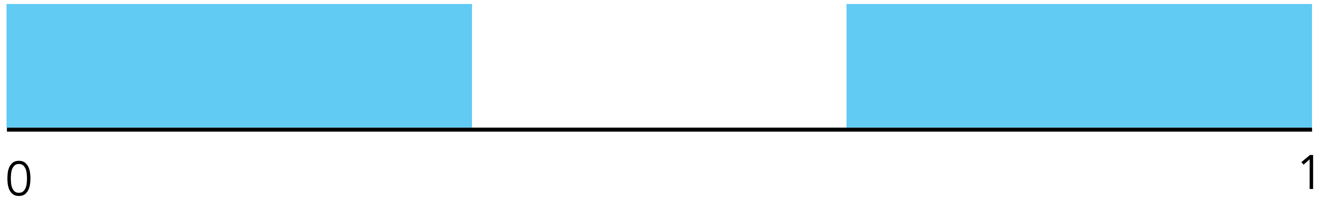

Certainty Equivalent Control computes the budget ratio

Gaps pose an algorithmic challenge

Certainty Equivalent Control computes the budget ratio

br = Budget Ratio = (Remaining Budget) / (Remaining Time)

Gaps pose an algorithmic challenge

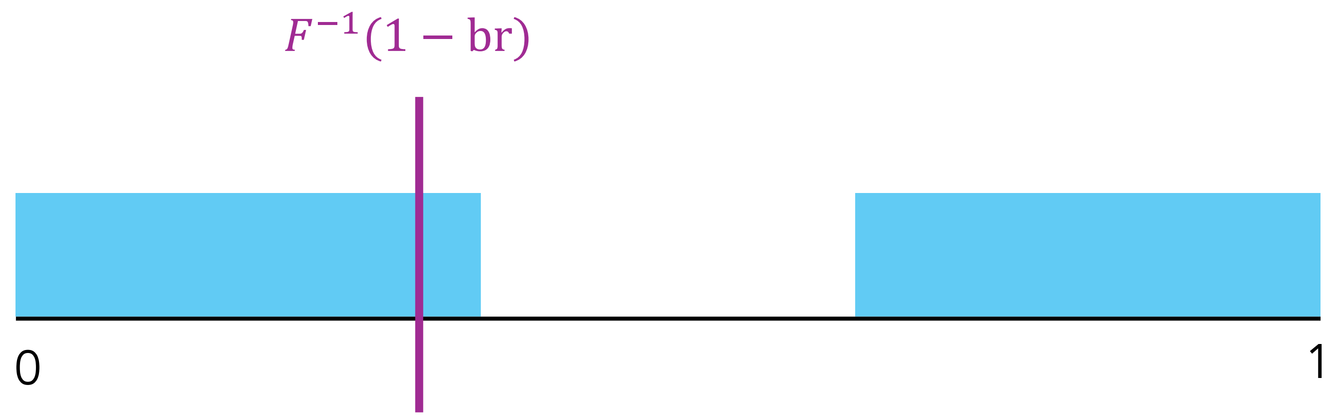

Certainty Equivalent Control computes the budget ratio

br = Budget Ratio = (Remaining Budget) / (Remaining Time)

Accept if the request type value is more than $F^{-1}(1 - \mathrm{br})$, else reject the request

Gaps pose an algorithmic challenge

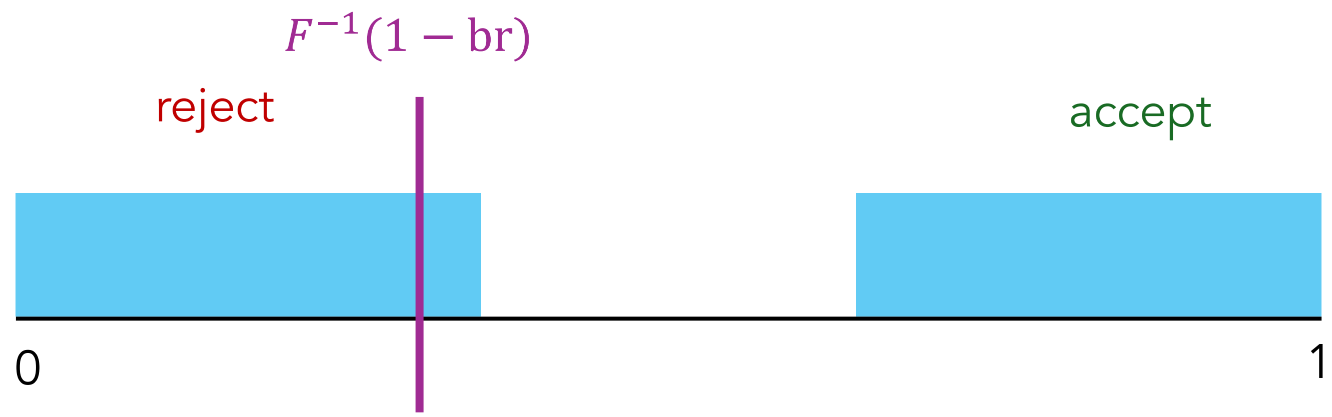

Certainty Equivalent Control computes the budget ratio

br = Budget Ratio = (Remaining Budget) / (Remaining Time)

Accept if the request type value is more than $F^{-1}(1-\mathrm{br})$, else reject the request

Gaps pose an algorithmic challenge

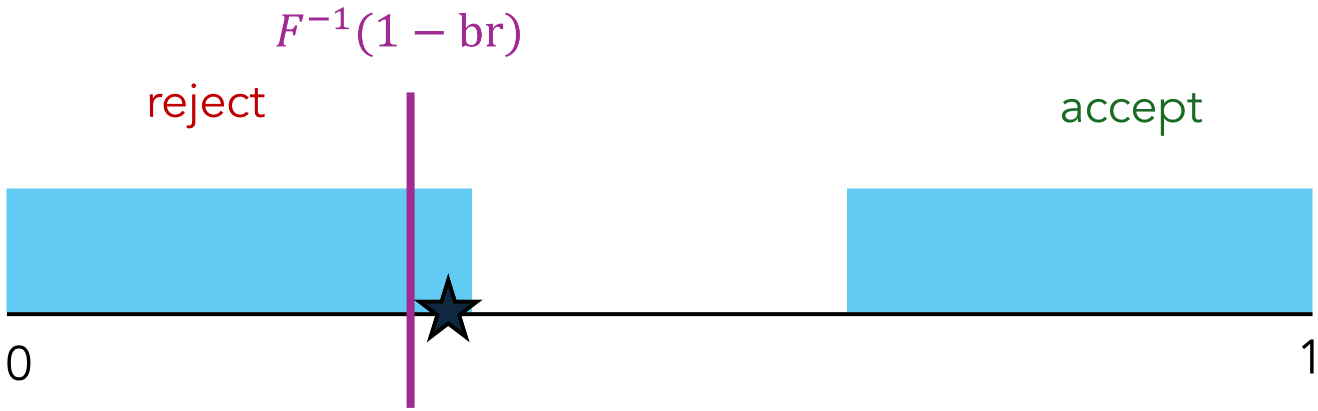

Certainty Equivalent Control computes the budget ratio

br = Budget Ratio = (Remaining Budget) / (Remaining Time)

Accept if the request type value is more than $F^{-1}(1 - \mathrm{br})$, else reject the request

Gaps pose an algorithmic challenge

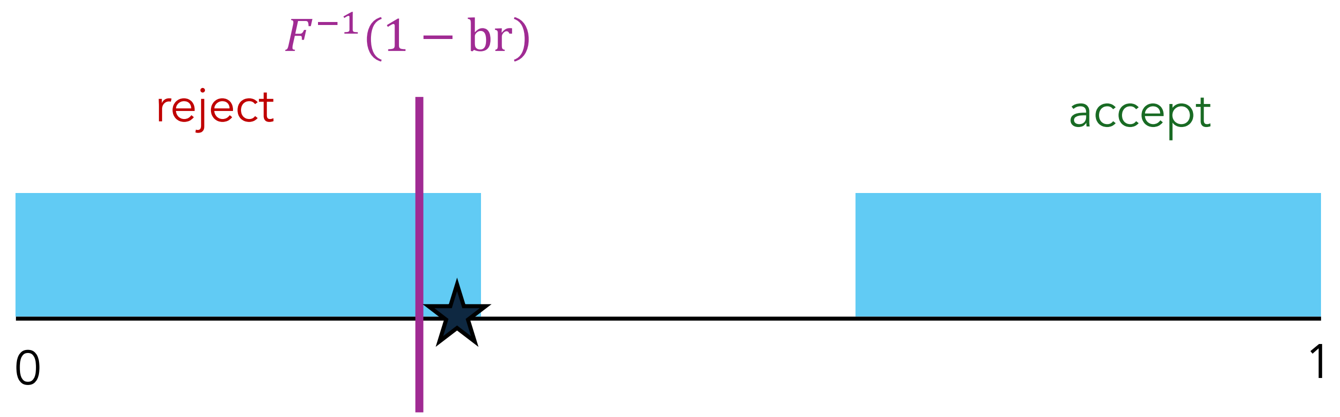

Certainty Equivalent Control computes the budget ratio

br = Budget Ratio = (Remaining Budget) / (Remaining Time)

Accept if the request type value is more than $F^{-1}(1-\mathrm{br})$, else reject the request

Regret(CE) = $\Omega(\sqrt{T})$ (highly sub-optimal regret scaling)

CwG via multiple simulations

Remaining Budget = 5

Remaining Time = 9

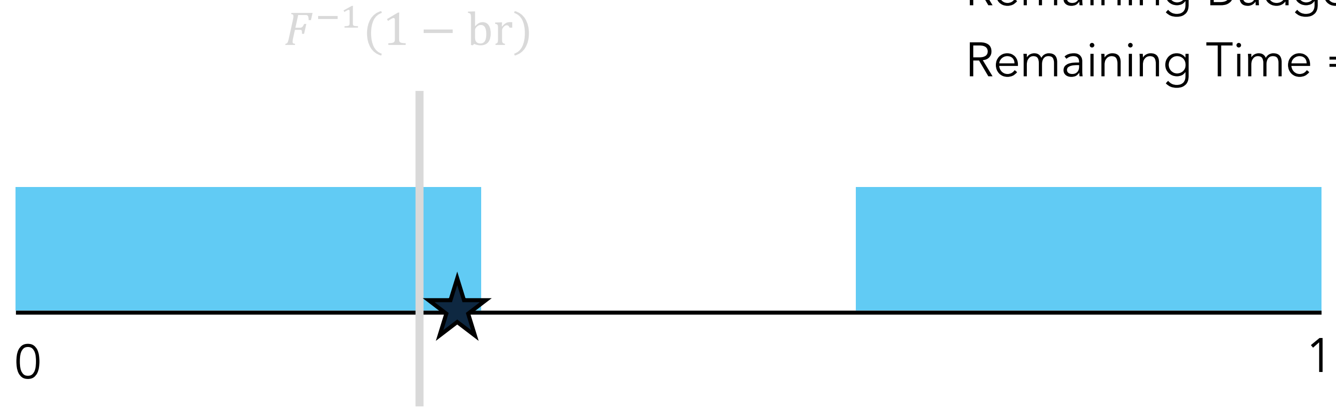

CwG via multiple simulations

Remaining Budget = 5

Remaining Time = 9

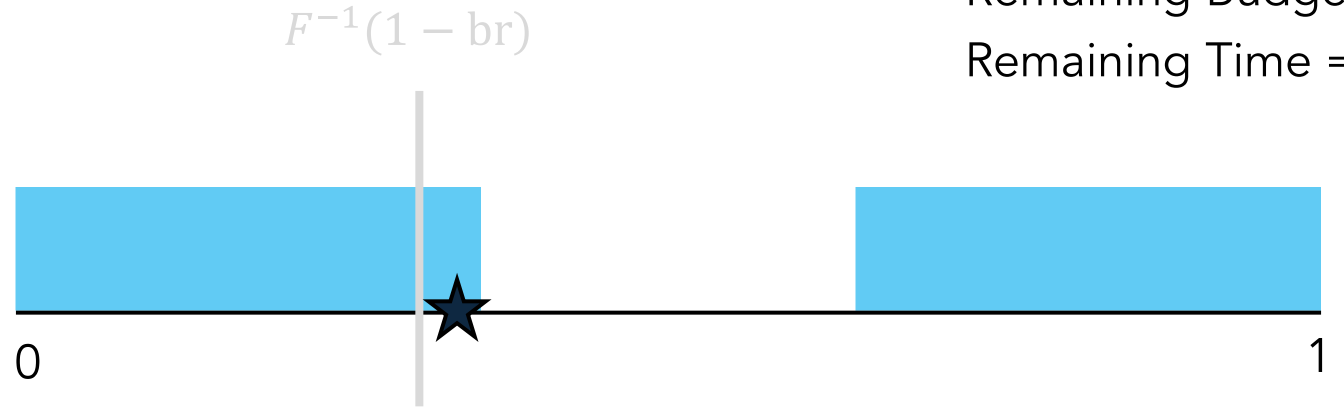

CwG via multiple simulations

Remaining Budget = 5
Remaining Time = 9

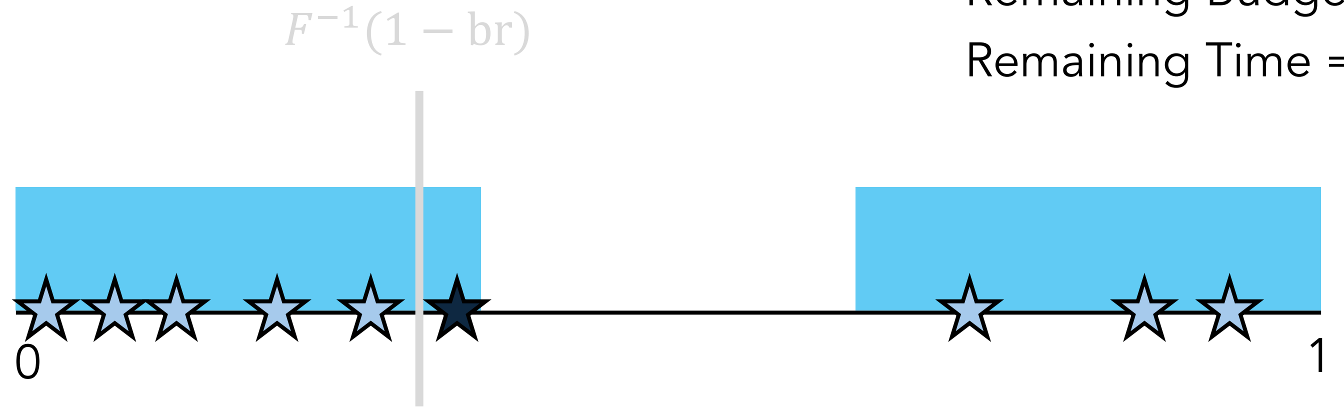

CwG via multiple simulations

Remaining Budget = 5

Remaining Time = 9

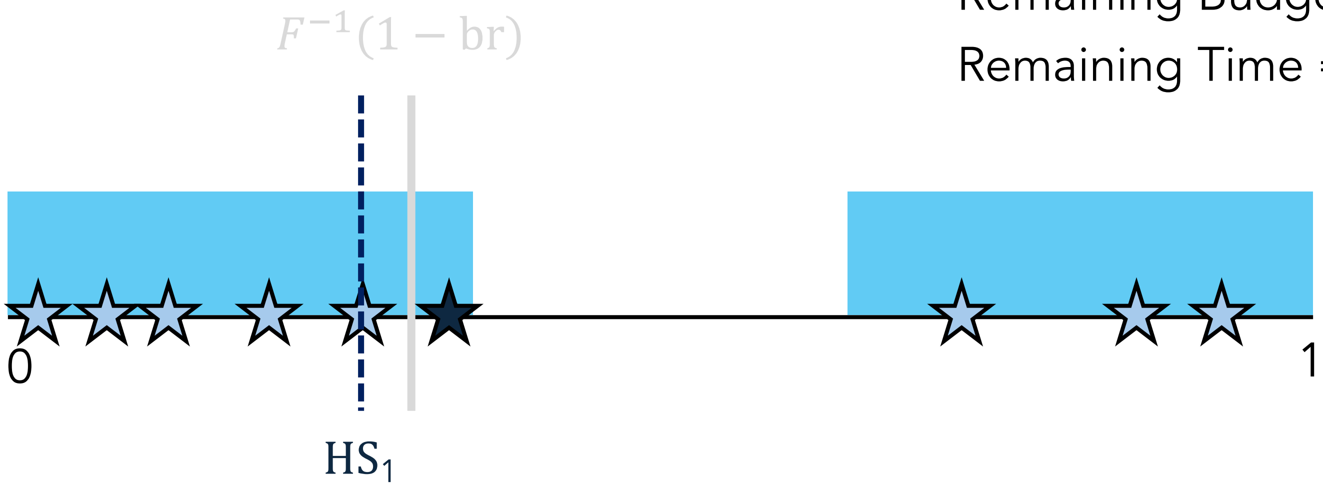

CwG via multiple simulations

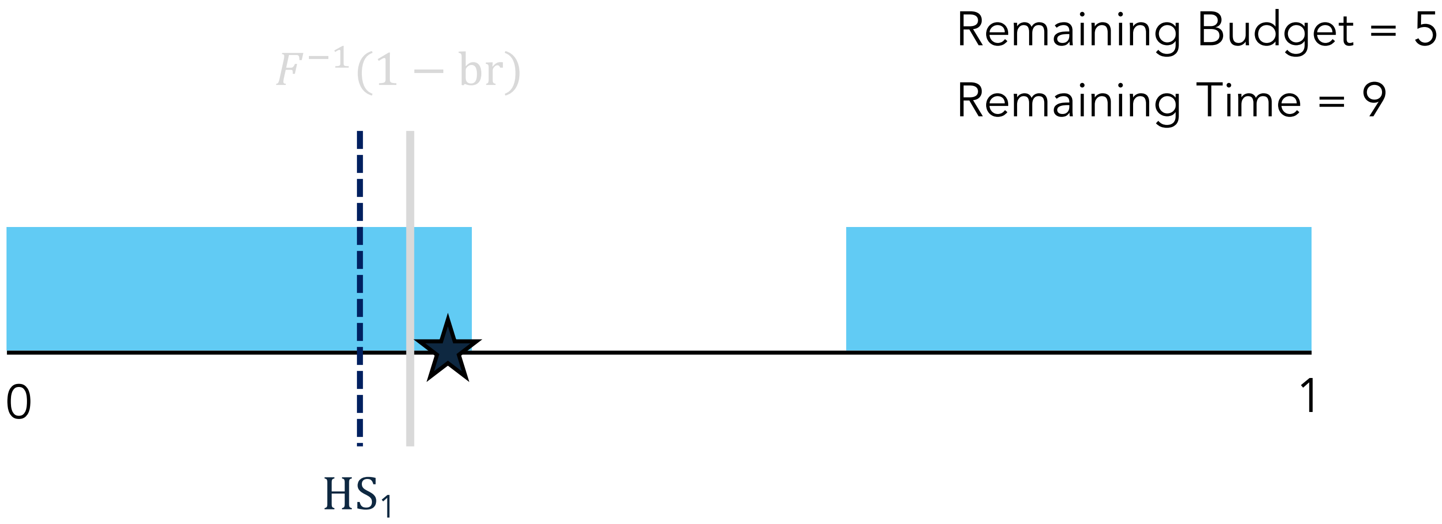

CwG via multiple simulations

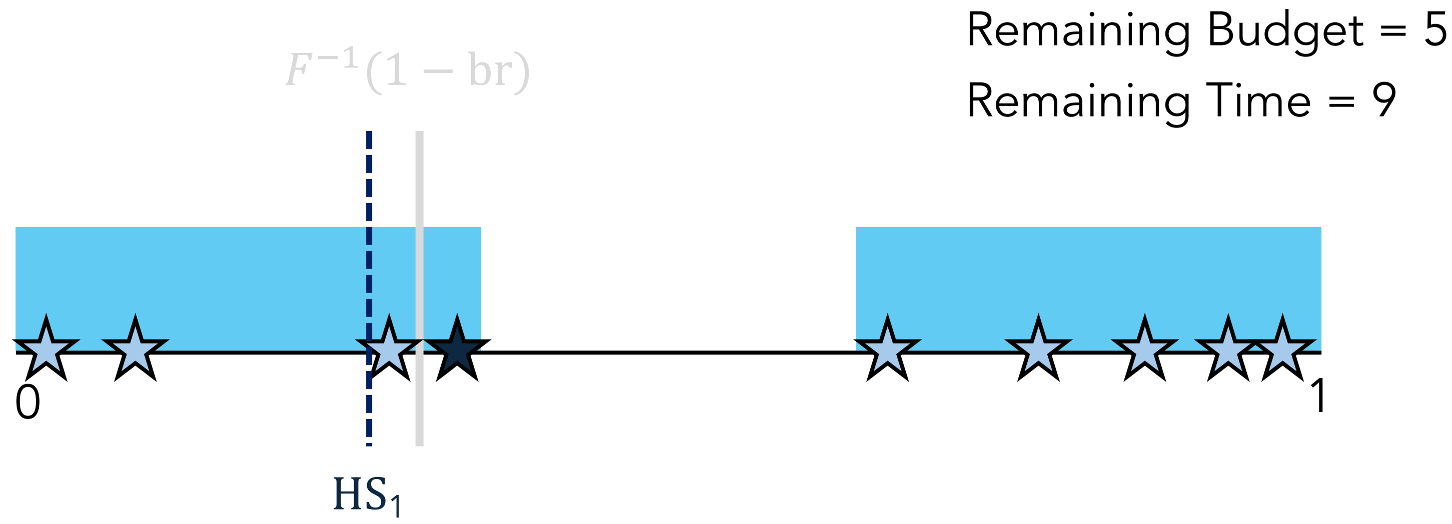

CwG via multiple simulations

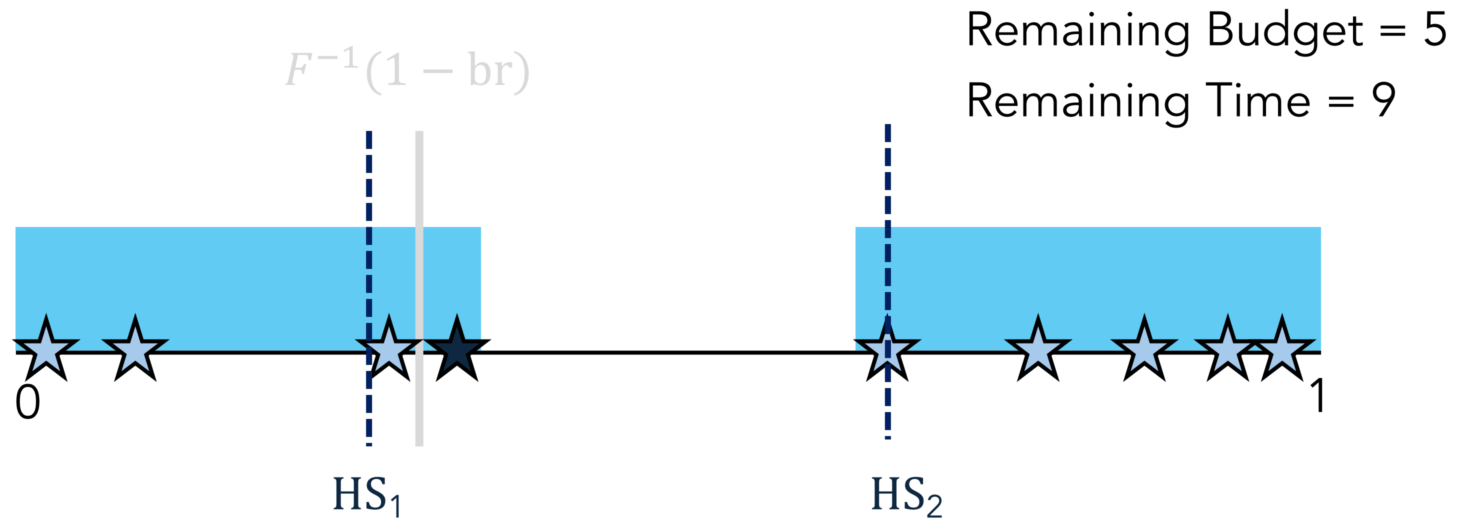

CwG via multiple simulations

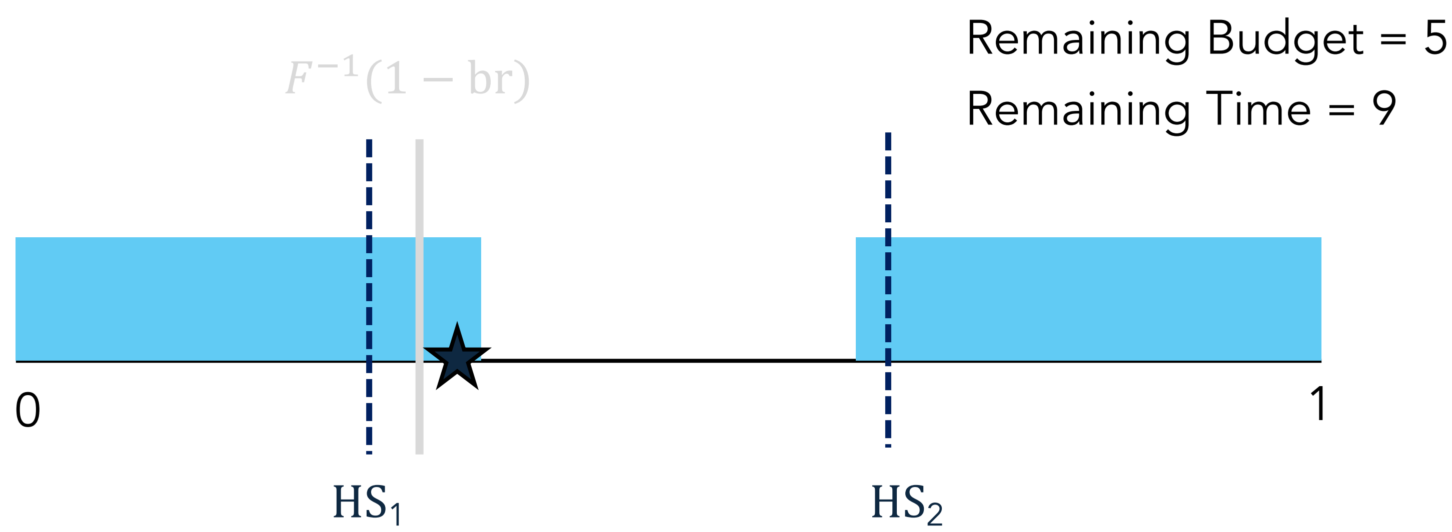

CwG via multiple simulations

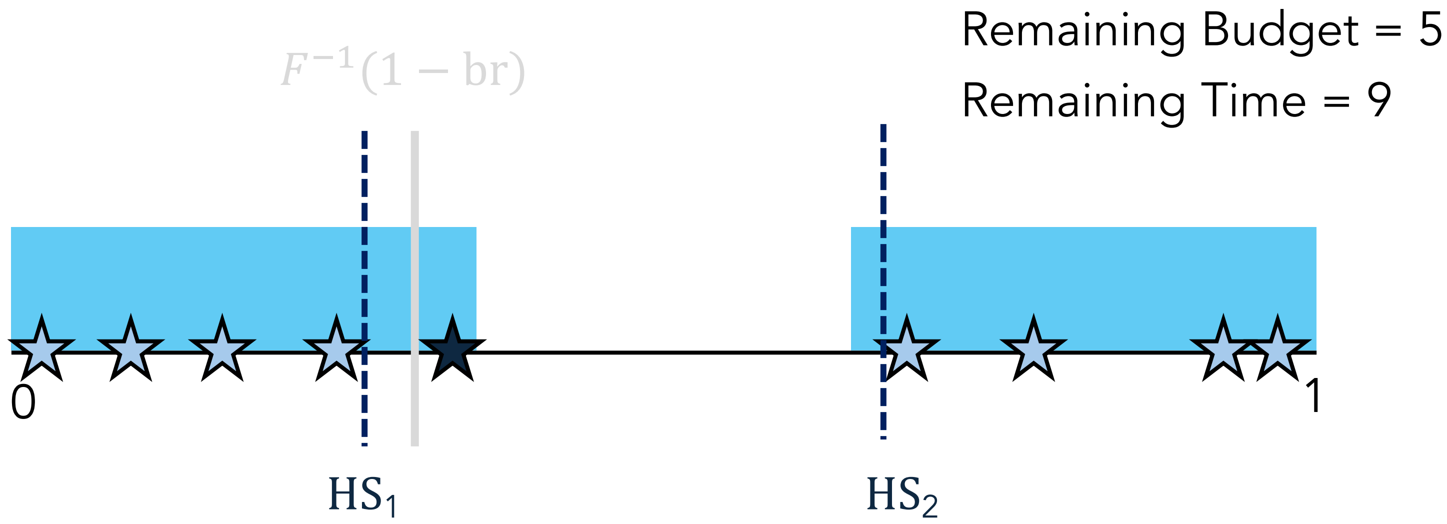

CwG via multiple simulations

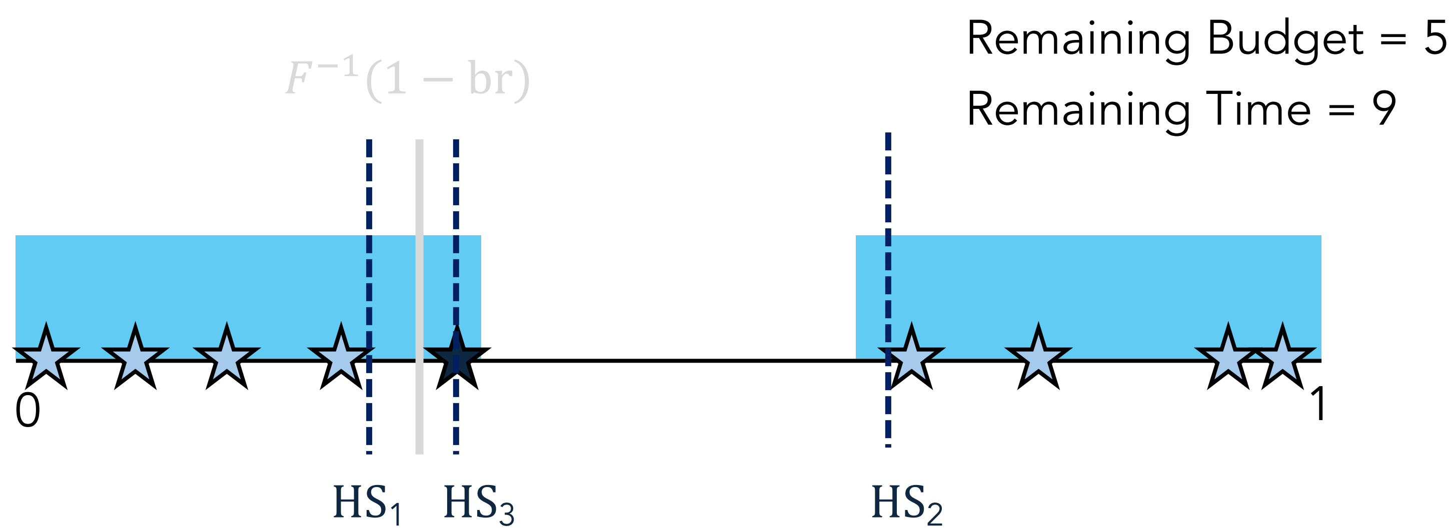

CwG via multiple simulations

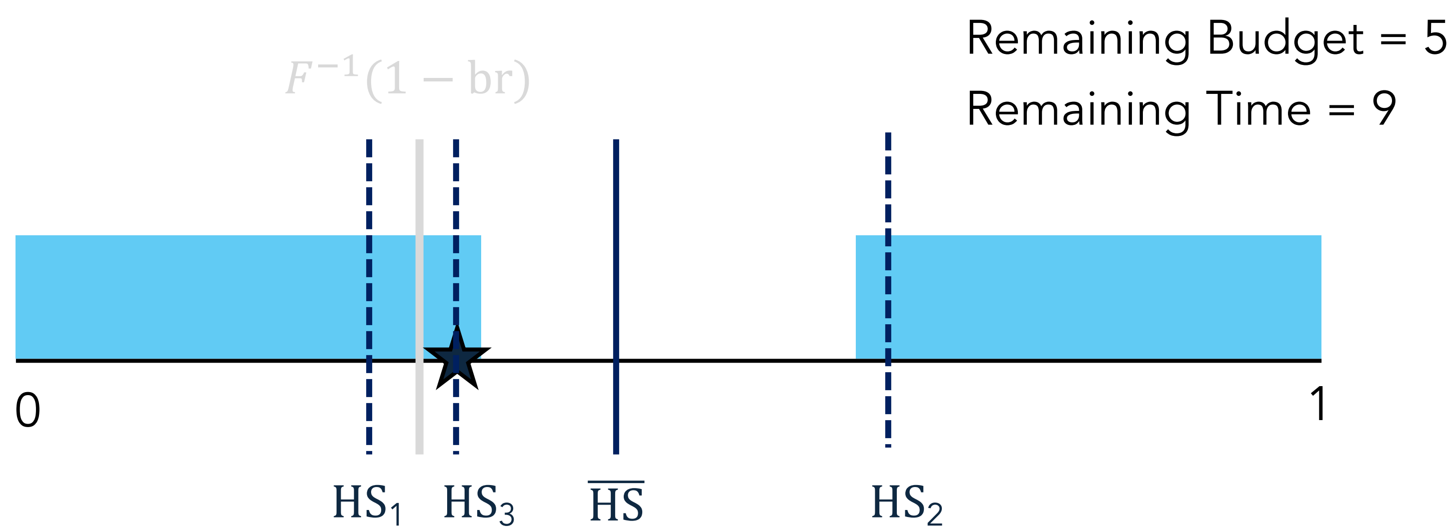

CwG via multiple simulations

Remaining Budget = 5

Remaining Time = 9

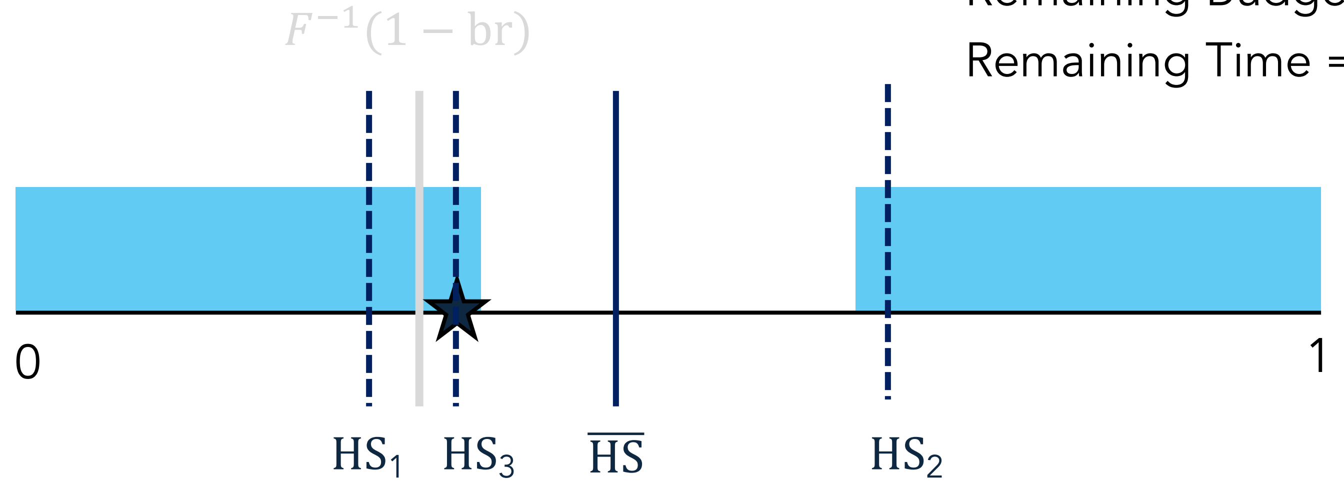

Accept if the request type value is more than $\overline{\mathrm{HS}}$, else reject the request

CwG via multiple simulations

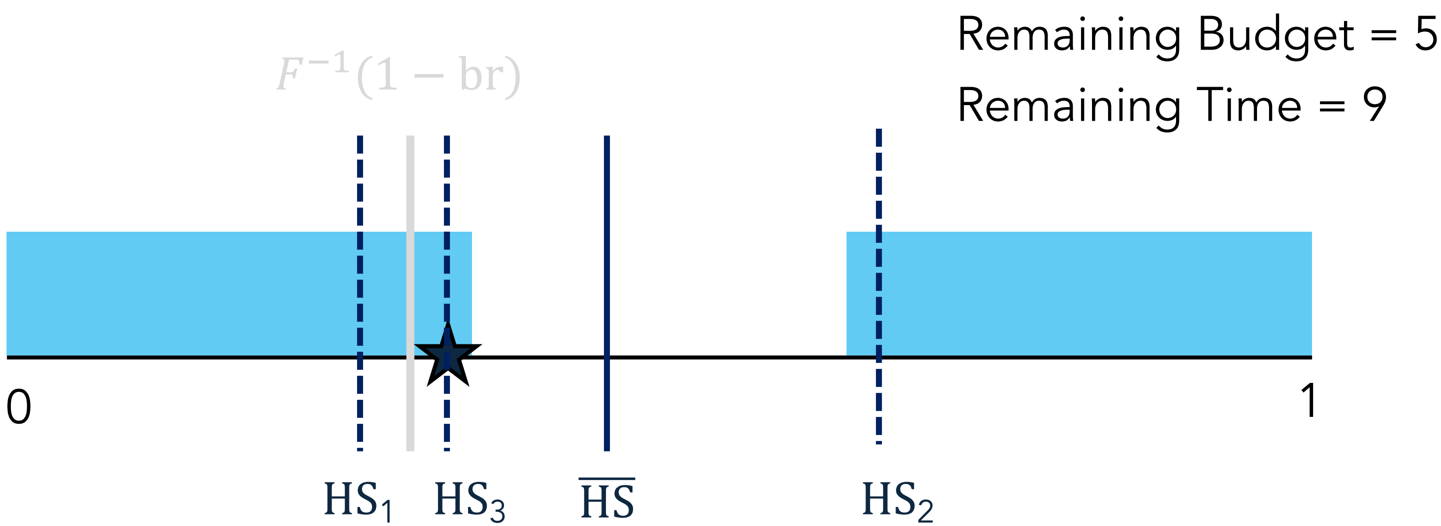

Accept if the request type value is more than $\overline{\mathrm{HS}}$, else reject the request

Regret(RAMS) $= O(\log^2 T)$

CwG via multiple simulations

Remaining Budget = 5

Remaining Time = 9

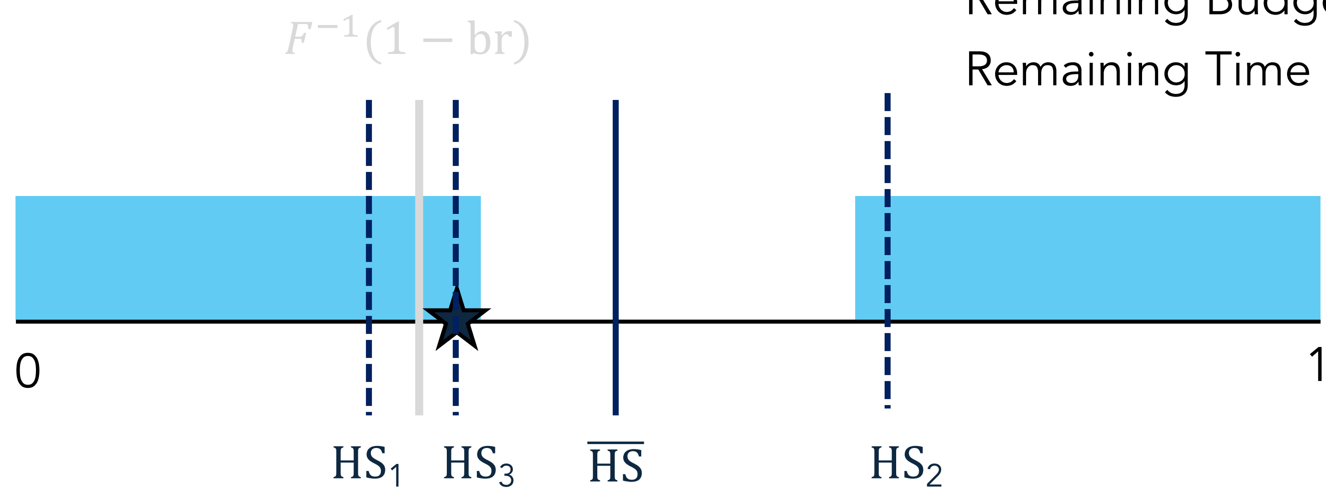

Accept if the request type value is more than $\overline{\mathrm{HS}}$, else reject the request

Regret(RAMS) = $O(\log^2 T)$

Regret(CE) = $\Omega(\sqrt{T})$ (highly sub-optimal regret scaling)

CwG via multiple simulations

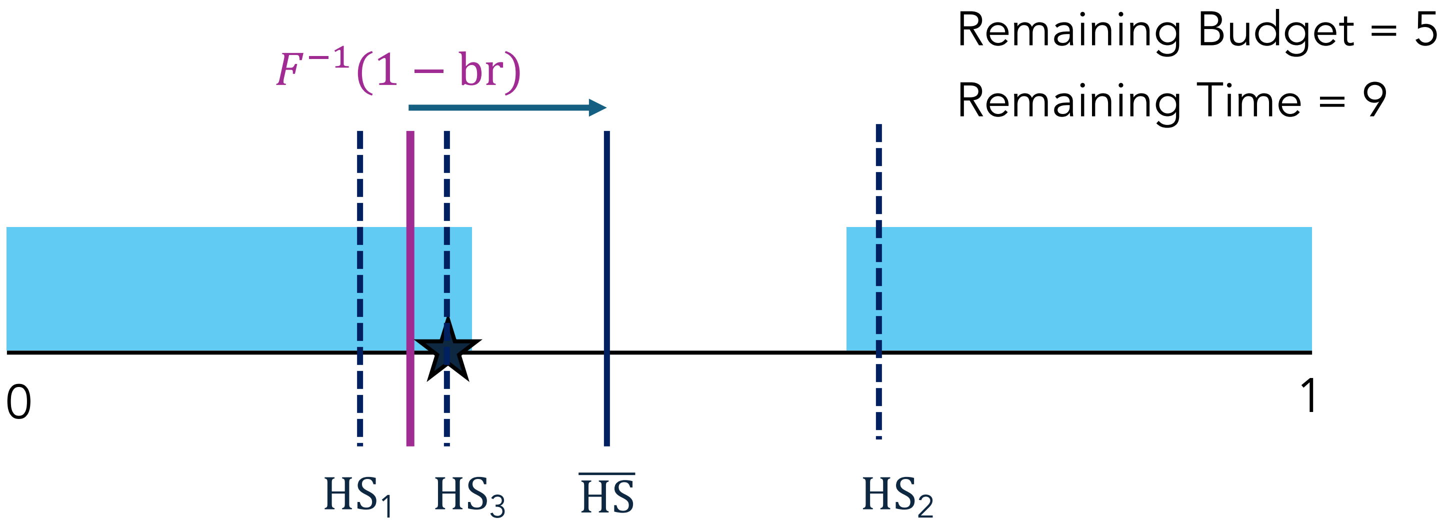

Accept if the request type value is more than $\overline{\mathrm{HS}}$, else reject the request

Conservativeness with respect to Gaps Principle
If the CE threshold $F^{-1}(1 - \mathrm{br})$ is close to a gap, use the gap as the threshold. Otherwise continue using the CE threshold.

CwG via multiple simulations

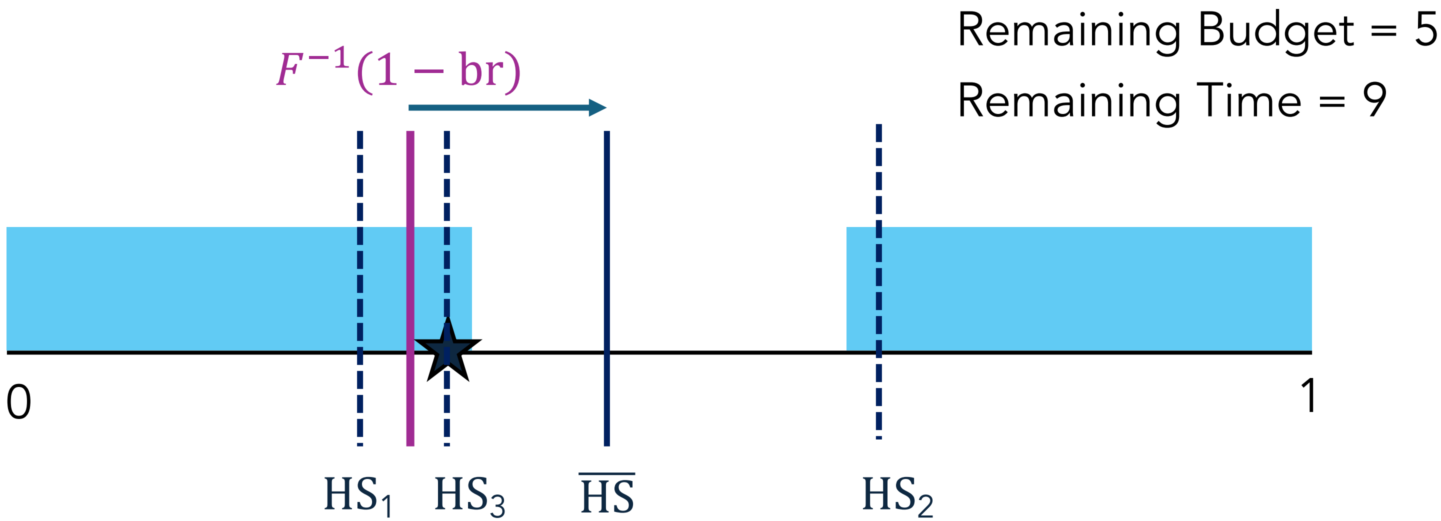

Accept if the request type value is more than $\overline{\mathrm{HS}}$, else reject the request

Connections to "Dual Averaging"
The different HS thresholds are the shadow prices of the budget for different scenarios, the bid price is computed by averaging the HS thresholds

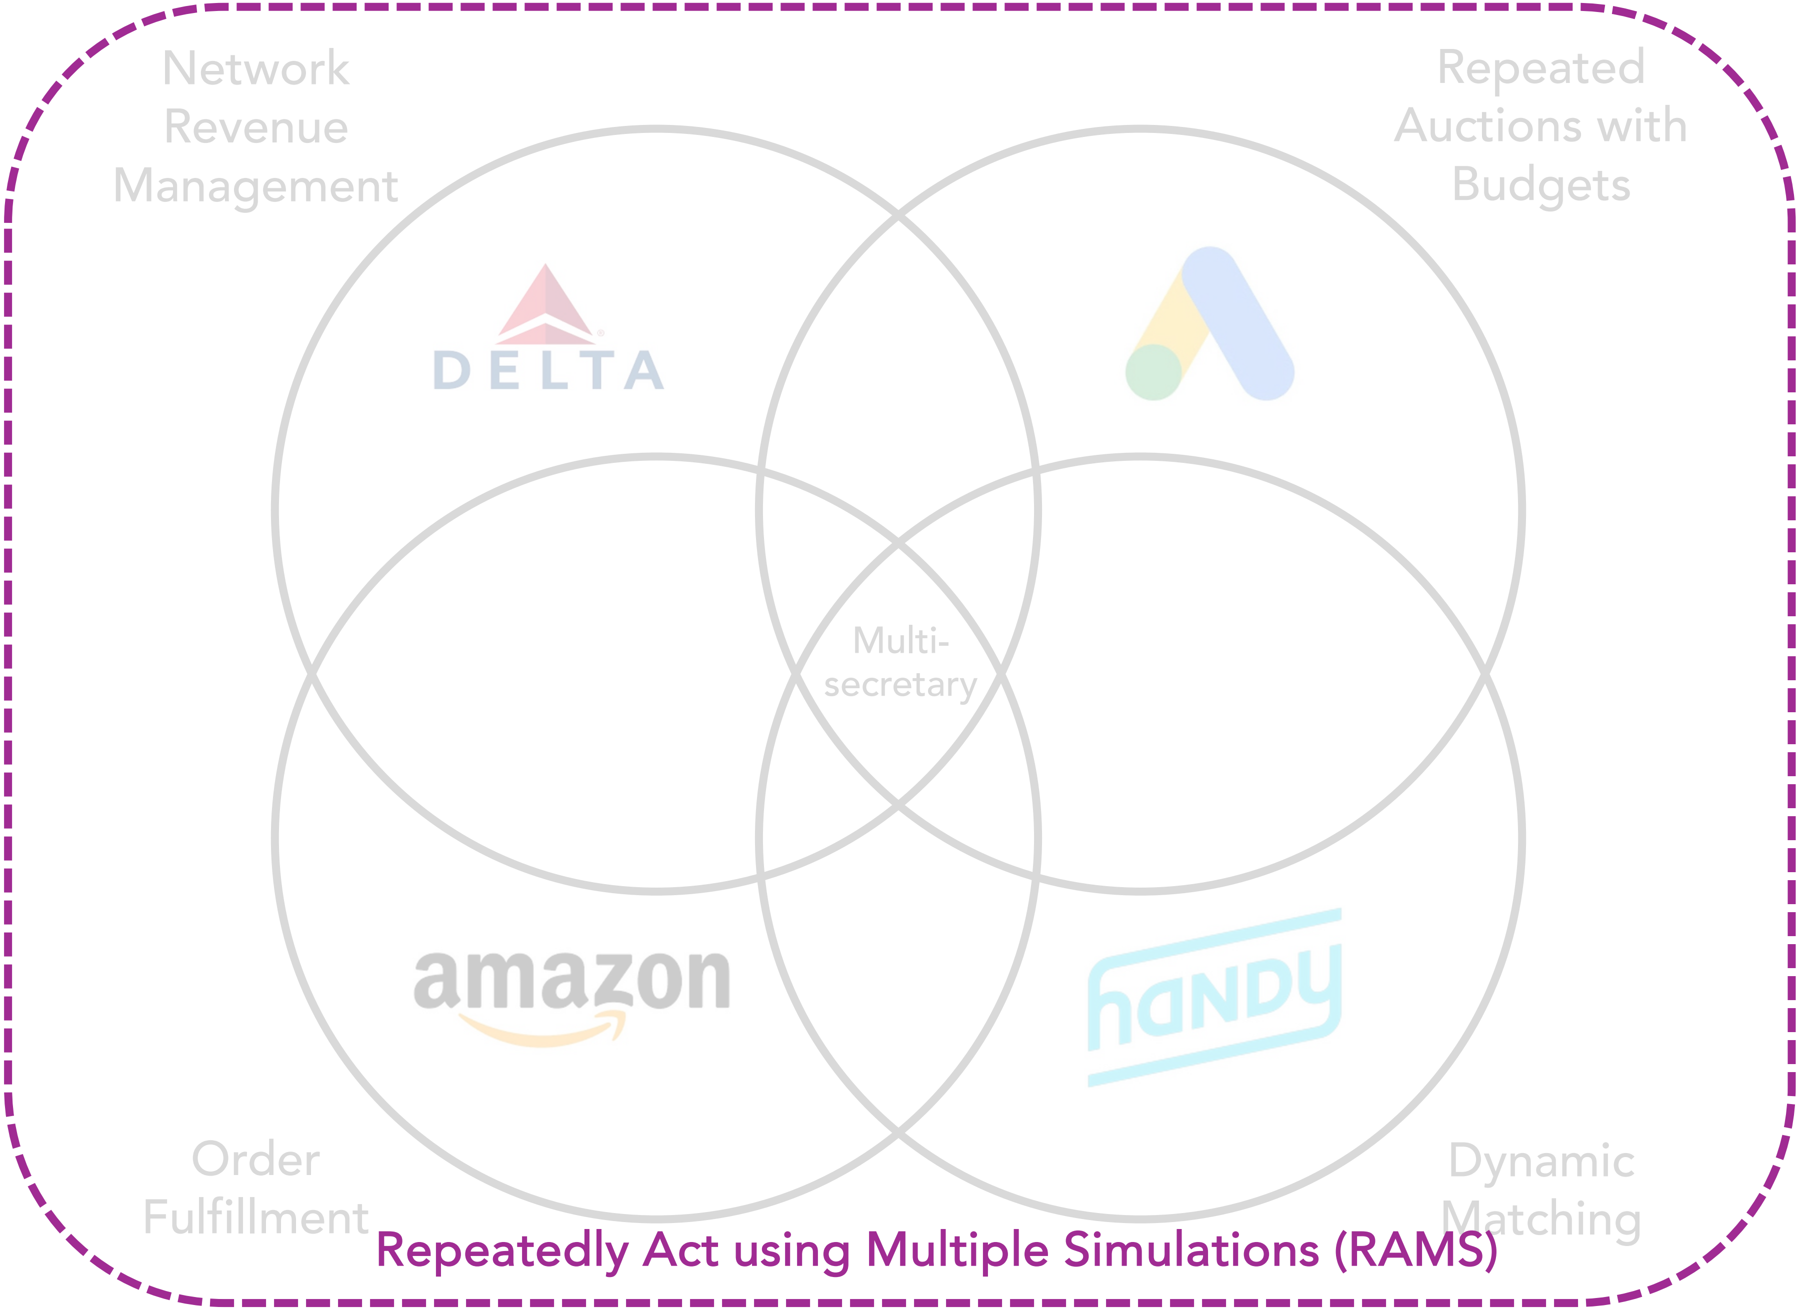
Network
Revenue
Management
Repeated
Auctions with
Budgets
DELTA
Multi-
secretary
amazon
handy
Order
Fulfillment
Dynamic
Matching
Repeatedly Act using Multiple Simulations (RAMS)

RAMS beyond multi-secretary

RAMS beyond multi-secretary

State (Budget) B_t and feasible set of actions A_t

RAMS beyond multi-secretary

State (Budget) B_t and feasible set of actions A_t

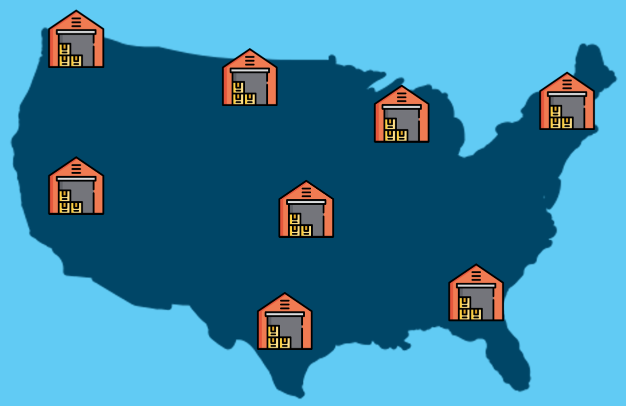

RAMS beyond multi-secretary

State (Budget) B_t and feasible set of actions A_t

Request $\theta_t = (r_t, c_t)$ arrives at time t

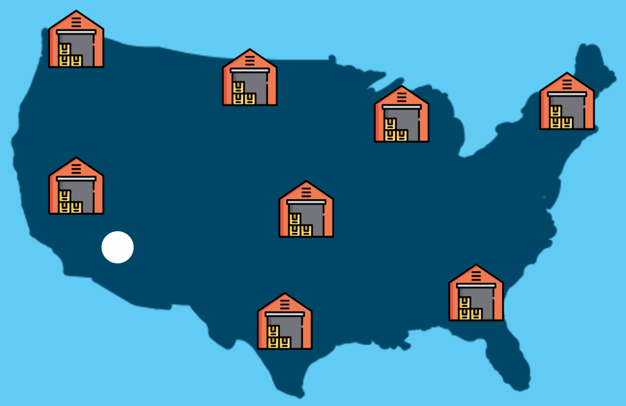

RAMS beyond multi-secretary

State (Budget) B_t and feasible set of actions A_t

Request $\theta_t = (r_t, c_t)$ arrives at time t

Simulate multiple request scenarios

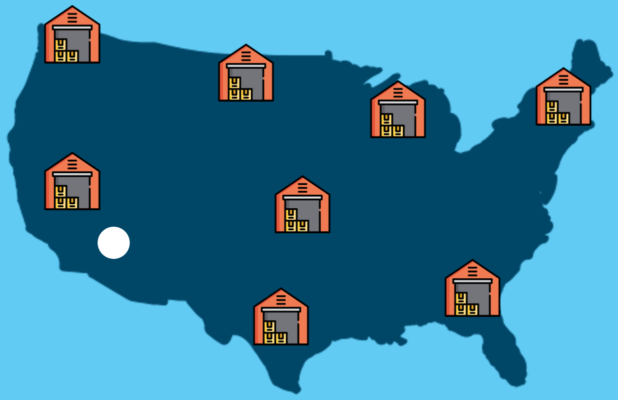

RAMS beyond multi-secretary

State (Budget) B_t and feasible set of actions A_t

Request $\theta_t = (r_t, c_t)$ arrives at time t

Simulate multiple request scenarios

$\theta_t, \theta_{t+1}^{(1)}, \theta_{t+2}^{(1)}, \dots, \theta_T^{(1)}$ Scenario 1

$\theta_t, \theta_{t+1}^{(2)}, \theta_{t+2}^{(2)}, \dots, \theta_T^{(2)}$ Scenario 2

$\vdots$

$\theta_t, \theta_{t+1}^{(m)}, \theta_{t+2}^{(m)}, \dots, \theta_T^{(m)}$ Scenario m

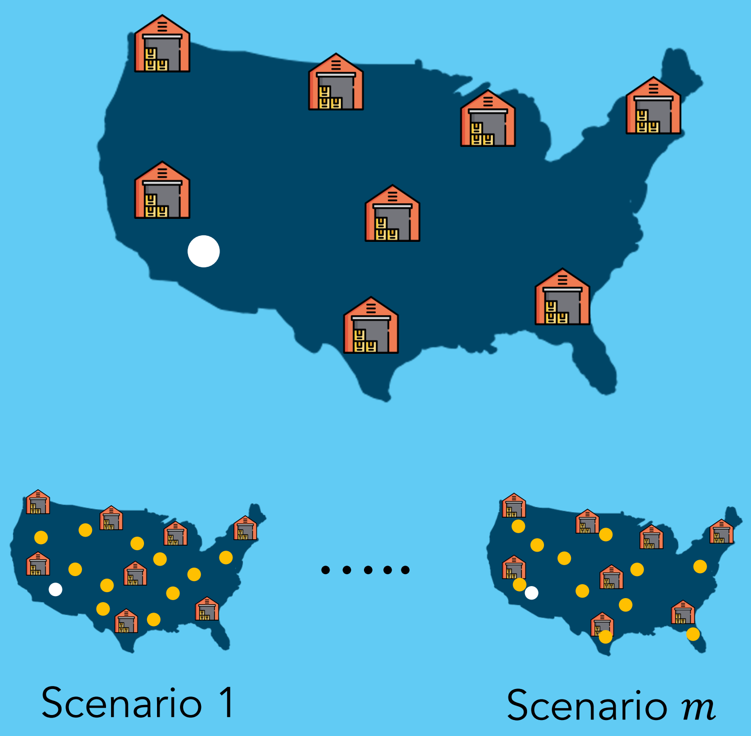

RAMS beyond multi-secretary

State (Budget) B_t and feasible set of actions A_t

Request $\theta_t = (r_t, c_t)$ arrives at time t

Simulate multiple request scenarios

$\theta_t, \theta_{t+1}^{(1)}, \theta_{t+2}^{(1)}, \dots, \theta_T^{(1)}$ Scenario 1

$\theta_t, \theta_{t+1}^{(2)}, \theta_{t+2}^{(2)}, \dots, \theta_T^{(2)}$ Scenario 2

$\vdots$

$\theta_t, \theta_{t+1}^{(m)}, \theta_{t+2}^{(m)}, \dots, \theta_T^{(m)}$ Scenario m

For each scenario k, compute the **compensation** for each action in A_t

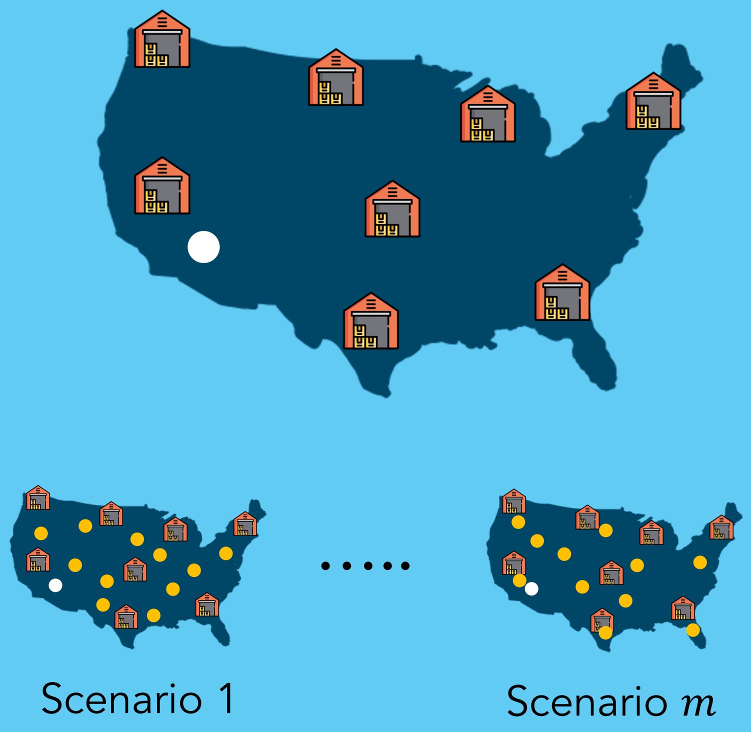

RAMS beyond multi-secretary

State (Budget) B_t and feasible set of actions A_t

Request $\theta_t = (r_t, c_t)$ arrives at time t

Simulate multiple request scenarios

$\theta_t, \theta_{t+1}^{(1)}, \theta_{t+2}^{(1)}, \dots, \theta_T^{(1)}$ Scenario 1

$\theta_t, \theta_{t+1}^{(2)}, \theta_{t+2}^{(2)}, \dots, \theta_T^{(2)}$ Scenario 2

$\vdots$

$\theta_t, \theta_{t+1}^{(m)}, \theta_{t+2}^{(m)}, \dots, \theta_T^{(m)}$ Scenario m

For each scenario k, compute the **compensation** for each action in A_t

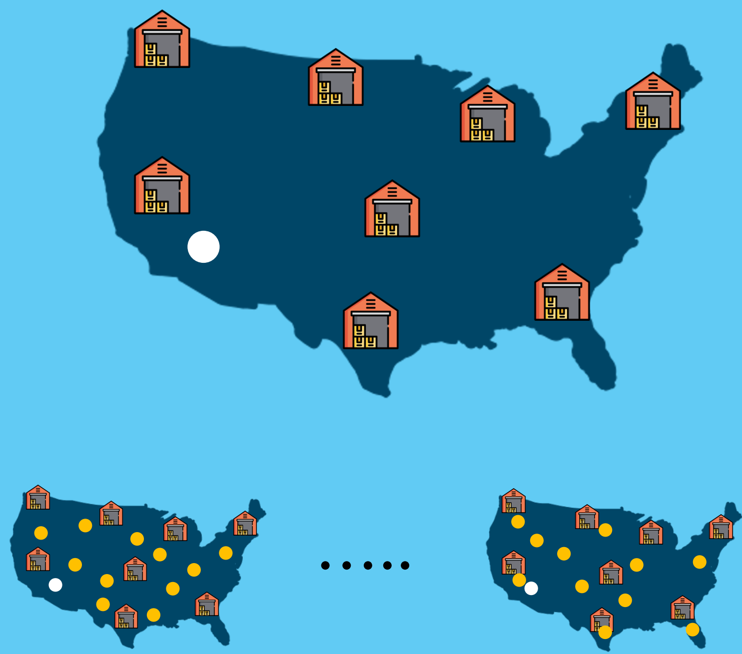

Scenario 1 Scenario m

Compensation(scenario k, a) =
(Max total reward in scenario k) –
(Max total reward in scenario k if action a is taken at time t)

RAMS beyond multi-secretary

State (Budget) B_t and feasible set of actions A_t

Request $\theta_t = (r_t, c_t)$ arrives at time t

Simulate multiple request scenarios

$\theta_t, \theta_{t+1}^{(1)}, \theta_{t+2}^{(1)}, \dots, \theta_T^{(1)}$ Scenario 1

$\theta_t, \theta_{t+1}^{(2)}, \theta_{t+2}^{(2)}, \dots, \theta_T^{(2)}$ Scenario 2

$\vdots$

$\theta_t, \theta_{t+1}^{(m)}, \theta_{t+2}^{(m)}, \dots, \theta_T^{(m)}$ Scenario m

For each scenario k, compute the **compensation** for each action in A_t

Take the **action** with the **minimum compensation** averaged over m scenarios

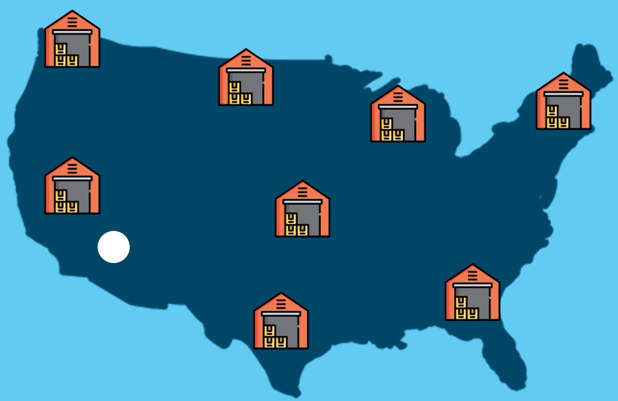

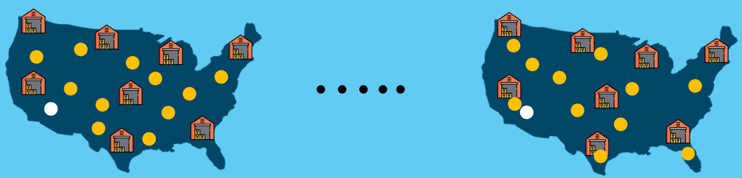

Scenario 1 ….. Scenario m

Compensation(scenario k, a) =
(Max total reward in scenario k) –
(Max total reward in scenario k if action a is taken at time t)

RAMS beyond multi-secretary

State (Budget) B_t and feasible set of actions A_t

Request $\theta_t = (r_t, c_t)$ arrives at time t

Simulate multiple request scenarios

$\theta_t, \theta_{t+1}^{(1)}, \theta_{t+2}^{(1)}, \dots, \theta_T^{(1)}$ Scenario 1

$\theta_t, \theta_{t+1}^{(2)}, \theta_{t+2}^{(2)}, \dots, \theta_T^{(2)}$ Scenario 2

$\vdots$

$\theta_t, \theta_{t+1}^{(m)}, \theta_{t+2}^{(m)}, \dots, \theta_T^{(m)}$ Scenario m

For each scenario k, compute the **compensation** for each action in A_t

Take the **action** with the **minimum compensation** averaged over m scenarios

Repeat the process

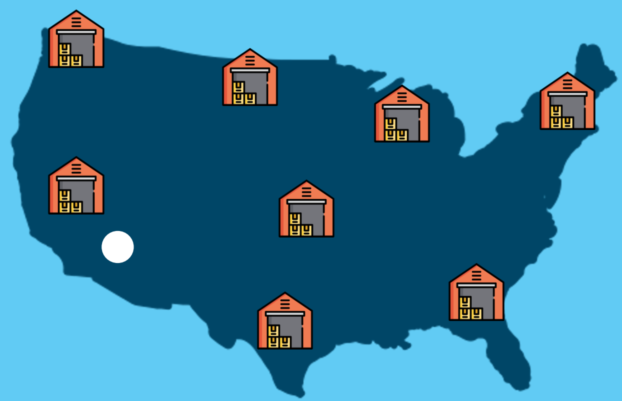

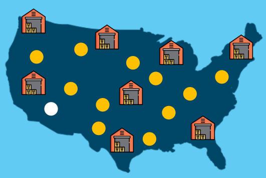

.....

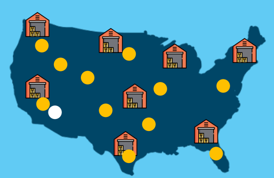

Scenario 1

Scenario m

Compensation(scenario k, a) =
(Max total reward in scenario k) –
(Max total reward in scenario k if action a is taken at time t)

RAMS minimizes hindsight-based regret

RAMS minimizes hindsight-based regret

Informal Meta Theorem [RAMS inherits guarantees of near-optimal algos].

Given a dynamic resource allocation setting, if there exists an algorithm **ALG** satisfying certain technical conditions, then

Regret(RAMS) ≤ Regret Upper Bound of **ALG** + Sampling Error

RAMS minimizes hindsight-based regret

Informal Meta Theorem [RAMS inherits guarantees of near-optimal algos].

Given a dynamic resource allocation setting, if there exists an algorithm **ALG** satisfying certain technical conditions, then

Regret(RAMS) $\leq$ Regret Upper Bound of **ALG** + Sampling Error

Proof of the Informal Meta Theorem.

$$\text{Regret(RAMS)} = \sum_{t=1}^{T} \mathbb{E}\left[\mathrm{Comp}_t\left(a_t^{\mathrm{RAMS}}\right)\right] \leq \sum_{t=1}^{T} \mathbb{E}\left[\mathrm{Comp}_t\left(a_t^{\mathrm{ALG}}\right)\right]$$

Compensated Coupling or Performance Diff. Lemma

RAMS chooses the action with the minimum compensation

RAMS is on-par with SOTA

RAMS is on-par with SOTA

Corollary of the Meta Theorem.

Polynomial regret for multi-secretary problem under different type distributions **[this work]**

Bounded regret for Network Revenue Management and Online Matching for a **few types** [Vera and Banerjee '21]

Logarithmic regret for Network Revenue Management with many types and non-degeneracy assumps. [Bray '23]

Log-Squared regret for Network Revenue Management with many types and w/o non-degeneracy assumps. [Jiang et. al '22]

RAMS is on-par with SOTA

Corollary of the Meta Theorem.

Polynomial regret for multi-secretary problem under different type distributions **[this work]**

Bounded regret for Network Revenue Management and Online Matching for a **few types** [Vera and Banerjee '21]

Logarithmic regret for Network Revenue Management with many types and non-degeneracy assumps. [Bray '23]

Log-Squared regret for Network Revenue Management with many types and w/o non-degeneracy assumps. [Jiang et. al '22]

Network Revenue Management problem with a few types

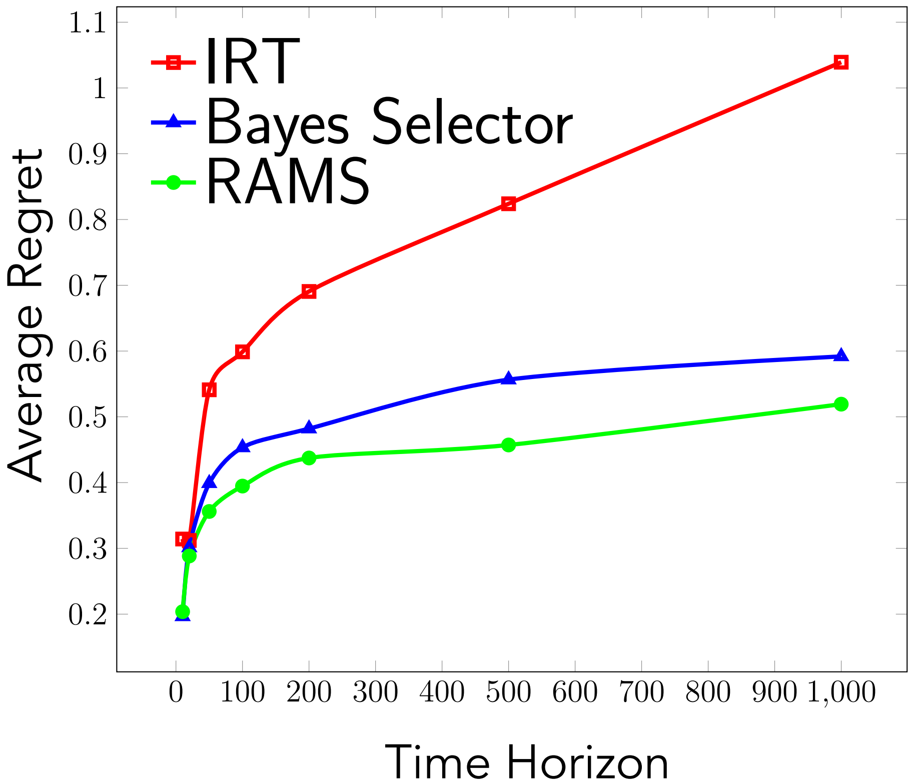

What is the interplay between the distribution of request types and achievable algorithmic performance?

A few types are present

Bounded Regret

All types are present

Logarithmic Regret

Multi-secretary Problem

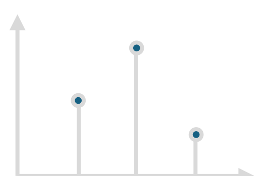

Entire spectrum of regret scalings are possible

β –clustered distributions

Can we design a **unified**, **simple** and **near-optimal** algorithms which works for all type distributions?

one policy to solve them all

Repeatedly Act using Multiple Simulations (RAMS)

A few types are present

Bounded Regret

All types are present

Logarithmic Regret

Multi-secretary Problem

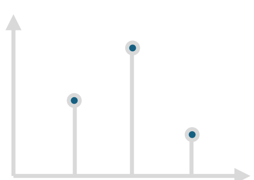

Entire spectrum of regret scalings are possible

β –clustered distributions

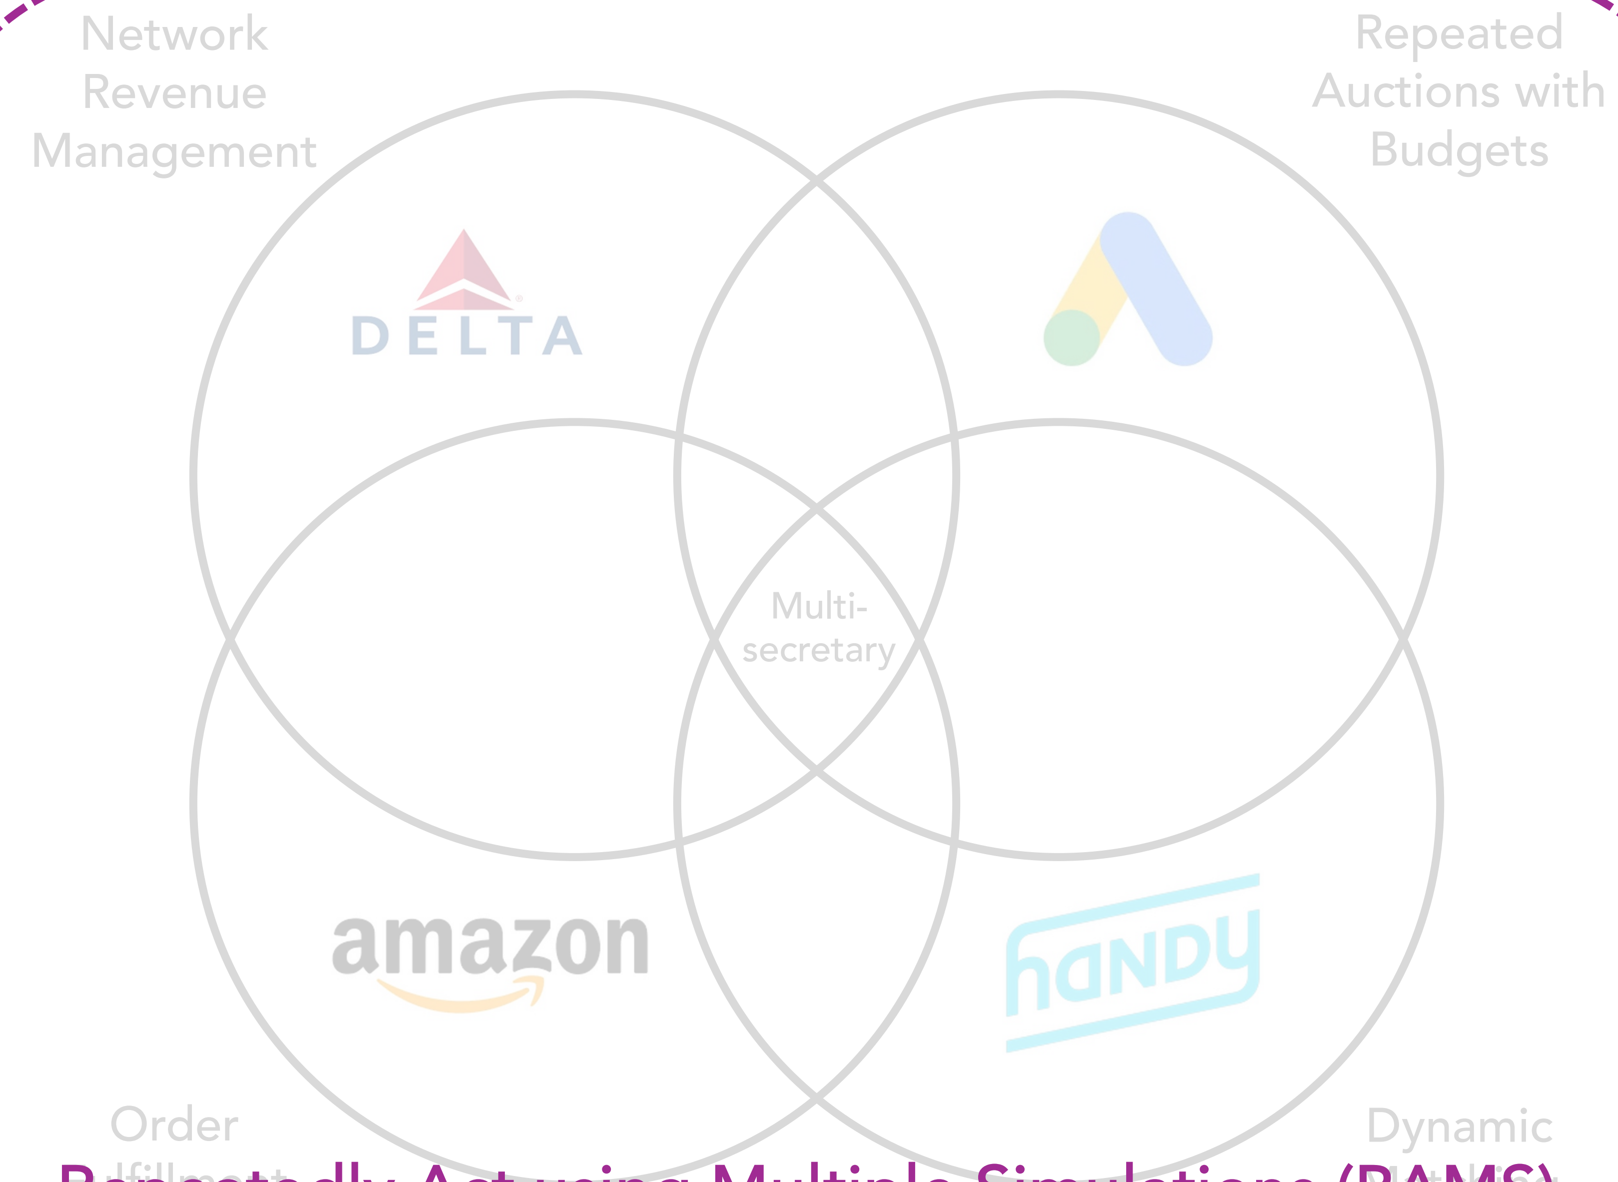

Repeatedly Act using Multiple Simulations (RAMS)

So long and **Thanks** for all the fish