Constant regret in Exchangeable Action Models:

Overbooking, Bin Packing & Beyond

Daniel Freund (MIT)

Wentao Weng
(MIT)

Kamessi Zhao
(MIT)

Chamsi Hssaine
(soon USC)

Elisabeth Paulson
(HBS)

Thodoris Lykouris
(MIT)

Brad Sturt
(UIC)

Sid Banerjee
(Cornell)







My fantastic collaborators

Agenda

- Define exchangeable action MDPs as a problem framework
- Characterize sets of assumptions under which O(1) loss is achievable
  - Types of objective/constraints
  - Types of information structures
  - Amount of adaptivity / computation required
- Algorithmic and analytical ideas
- Takeaways

Define *exchangeable actions*

- Exogenous arrivals $\theta_1, \ldots, \theta_T \in \Theta$, where $|\Theta| = k$
- In period t, observe θ_t & take an irrevocable action:
  $$a_{\theta_t j} \in \mathcal{A}(\theta_t), where\ |\mathcal{A}(\theta_t)| = \ell$$
- Denote by $x_{\theta j}$ the number of times we take $a_{\theta j}$
- Then our objective is to maximize (minimize)
  $$f(\vec{x})$$
  for some known function $f(\cdot)$

We'll make assumptions on $f(\cdot)$ and on the arrivals.

First, just consider what's captured

Bin packing

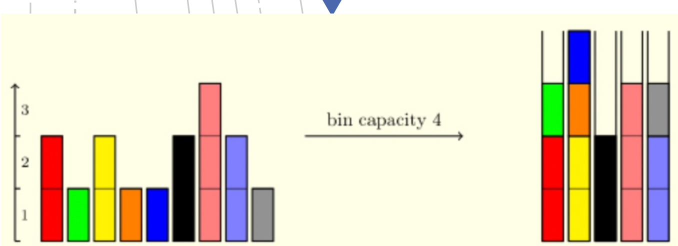

- We have bins of a given size

  e.g., all bins have size 10

- $\theta_1, \dots, \theta_T \in \Theta$, are items of sizes $w_1, \dots, w_k$

  e.g., suppose $w_\theta \in \{1,3,4,5,8\}$

- In period t, observe w_t & place it into bin of type j, where $b_{\theta,j} \geq 1$ denotes # of θ that fit into bin type j

e.g., j indexes the possible configurations of items in a bin: (1,1,8), (1,4,5), (5,5)… And $b_{5,(5,5)} = 2$

- Denote by $x_{\theta j}$ the number of times we take $a_{\theta j}$

e.g., how many size-1 items did we put in (1,1,8) bins

- Then our objective is to minimize

$$f(\vec{x}) = \sum_j \max_\theta \lceil x_{\theta j} / b_{\theta,j} \rceil$$

The ceiling is a boring technical detail

Network revenue management

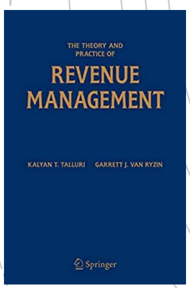

- We have some resources $\vec{B} \in \mathbb{N}^m$
- $\theta_1, \dots, \theta_T \in \Theta$, are types with values & resource reqs

  e.g., $v_1, \dots, v_k$ *and* $\vec{r_1}, \dots, \vec{r_k} \in \mathbb{N}^m$, R matrix of $\vec{r_j}$

- In period t, observe θ_t & either accept or reject

  i.e., $a_{\theta j} \in \{accept, reject\}$

- We want to maximize the accepted values w/o violating resource constraints
- Denote by x_θ the number of times we *accept* θ
- Then our objective is to maximize

For today, the negative part is a boring technical detail

$$f(\vec{x}) = \sum_\theta v_\theta x_\theta - v_{max} \left| \left((R \cdot \vec{x}) - B \right)^+ \right|_1$$

Models of job assignment

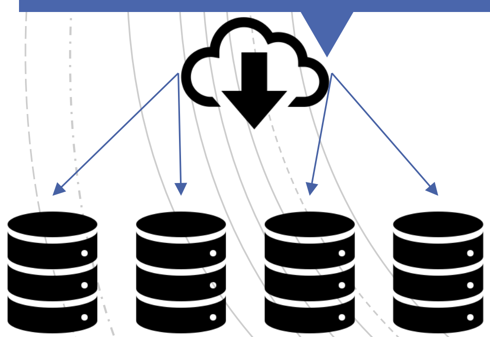

- Choice between m servers
- Arrivals $\theta_1, \dots, \theta_T \in \Theta$

Job θ to be processed

- In period t, observe θ_t & take an irrevocable action:

  $a_{\theta_t j} \in \{1, \dots, m\}$ to process job θ_t at j
- Denote by $x_{\theta j}$ the number of times we take $a_{\theta j}$
- Then our objective is:

Cost to process all jobs where servers have (i) a fixed cost per job $c_{\theta j}$ & (ii) a minimum *average cost per job* m_j

$$\min \sum_j \max\left\{\sum_\theta x_{\theta j} c_{\theta j}, \sum_\theta x_{\theta j} m_j\right\}$$

Models of refugee placement

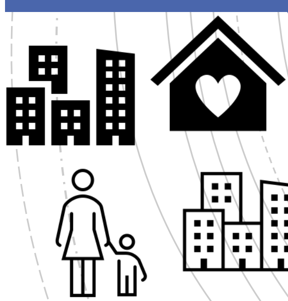

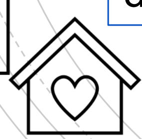

- m centers have room to absorb $b_1, \dots, b_m$ refugees

- Arrivals $\theta_1, \dots, \theta_T \in \Theta$

probability $v_{\theta j}$ for employment if placed at center j

- In period t, observe θ_t & take an irrevocable action:

  $a_{\theta_t j} \in \{1, \dots, m\}$ to place arrival θ_t at center j

- Denote by $x_{\theta j}$ the number of times we take $a_{\theta j}$

- Then our objective is:

For today, the negative part is a boring technical detail

$$\max \sum_j \sum_\theta x_{\theta j} v_{\theta j} - \left[b_j - \sum_\theta x_{\theta j} \right]^+$$

Not captured

- Unknown objective: unknown $f(\cdot)$

  (bandits / pricing)

- Time-sensitive actions: $f(\cdot)$ depends

  (Weina's talks!) not just on $\vec{x}$

- Overbooking: don't quite know $f(\cdot)$

Overbooking for a single resource

- High-level:
  - arrivals of type θ have value v_θ if accepted
  - arrivals of type θ are no-shows with prob. $1 - q_\theta$
  - no-shows pay but do not consume resources

(incentivizes us to admit more arrivals than there are resources for)

- If more than B (capacity) people show up, we pay a penalty of c per person we'll need to bump
- When we admit a type, we don't know whether they'll show up!
- So, we don't know $f(\cdot)$ ← it's random!

Expected objective

- If I knew all the arrivals, who should I accept?

(by arrivals I only mean their type, not whether they will show up; if I knew that, I'd accept everyone who won't show up… silly benchmark)

$$\max \sum_{\theta} v_\theta x_\theta - c \cdot \mathbb{E}\left[\left(\sum_{\theta} X_\theta - B\right)^{+}\right]$$

where $X_\theta \sim Bin(x_\theta, q_\theta)$

- Pretend our objective is $\mathbb{E}[f(\vec{x})]$ and we'll be able to compare ourselves with the best clairvoyant who knows the arrivals but not the no-show-realizations

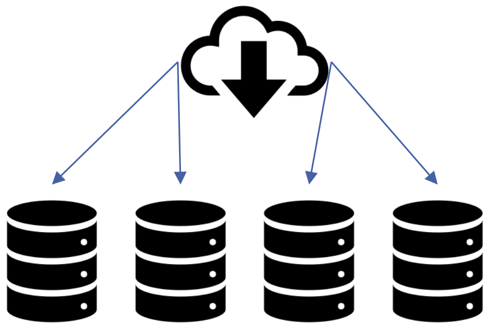

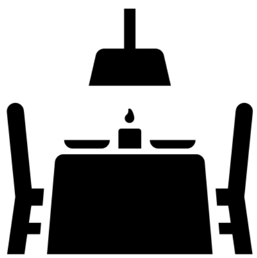

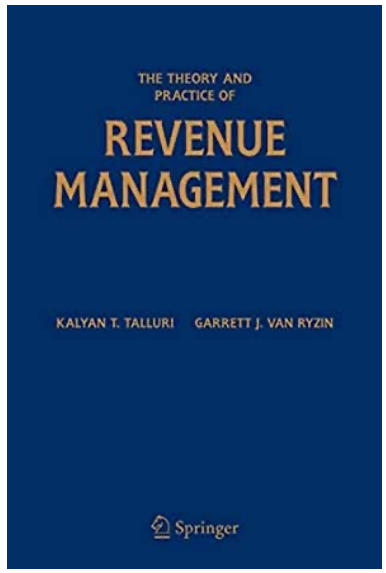

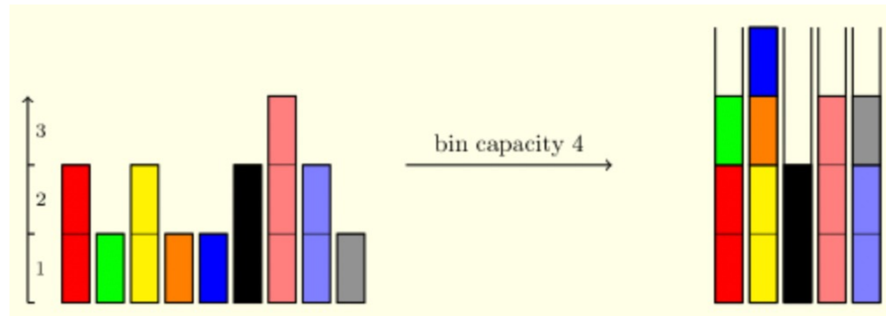

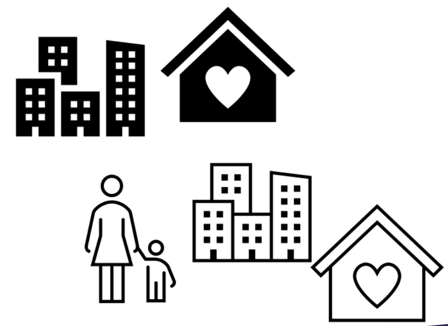

Many examples of exchangeable actions!

We'll keep it general!

Benchmark

Clairvoyant optimum:

$$OPT = \max f(\vec{x})$$

$$s.t. \forall\theta: \sum_j x_{\theta j} = N_\theta[T]$$

$$x_{\theta j} \geq 0$$

where $N_\theta[\tau] = \sum_{t=1,\ldots,\tau} \mathbb{I}_{\{\theta_t = \theta\}}$

Desired performance: Constant regret

- Denote an algorithm's objective by $ALG(\theta_1, \ldots, \theta_T)$

$$\mathbb{E}[\ OPT - ALG(\theta_1, \ldots, \theta_T)] \leq M \in O(1)$$

- Meaning we want to bound the performance loss of an algorithm independent of T

It's somewhat trivial in most/all our settings to achieve $\tilde{O}(\sqrt{T})$ loss; so the name of the game is to obtain something better/constant!

Possible assumptions on T

($T^{\frac{1}{2}+\epsilon}$ $for\ some\ \epsilon > 0$ works, but we don't want to carry the ϵ)

T1: Known time-horizon T

Fairly standard in many settings

T2: T is a priori unknown but revealed at $\hat{T}$ with

$$T - \hat{T} \in \Omega(T^{\frac{3}{4}})$$

Slight variation of an adversarial end point; unknown, but there's a heads-up when a few periods are left.

Example: we've been running an open-ended marketing campaign since mid-August and we're told today (10/10) that it will end on 10/15

Example 2: there's an unknown number of batches, with $\Omega(T^{\frac{3}{4}})$ arrivals, last one is announced as such.

Possible assumptions on arrivals

A1: iid with unknown $p_\theta \geq p_{min}\ \forall\theta$

A2: independent with known $p_\theta(t) \geq p_{min}\ \forall\theta, t$

A3: iid with known $p_\theta \geq p_{min}\ \forall\theta$

A4: We have a single sample of T arrivals & we know that it's drawn from a distribution with certain density/concentration properties

Possible assumptions on f

O1: $\frac{L}{2}$-Lipschitz-continuous

$$|f(\vec{x}) - f(\vec{y})| \leq |\vec{x} - \vec{y}|L/2$$

Genuinely innocent!

O2: Stability of optimal solution

Denote by $S(\vec{N})$ the set of optimal solutions under $\vec{N}$

$$\forall \vec{N}, \vec{N}': \forall \vec{x} \in S(\vec{N}) \exists \vec{y} \in S(\vec{N}'): |\vec{x} - \vec{y}| \leq \delta |\vec{N} - \vec{N}'|$$

Looks weird, but always fulfilled when $f(\cdot)$ is linear

(key challenge for overbooking is not having this)

O3: Homogeneous ($f(\lambda\vec{x}) = \lambda f(\vec{x})$)

Needed under T2! E.g., a marketing campaign with a fixed budget per customer

O4: Existence of unique opt

Only required in special cases or for being able to compute an offline optimal solution

Informal results

ALGORITHMIC

Pick the right combination of the above & there exists an algorithm ALG such that

$$\mathbb{E}[OPT - ALG(\theta_1, \dots, \theta_T)] < M \in O(1)$$

for some constant M that depends on all above, except for T

IMPOSSIBILITY

Drop one from the right combination of the above & no algorithm achieves

$$\mathbb{E}[OPT - ALG(\theta_1, \dots, \theta_T)] < M \in O(1)$$

for any constant M independent of T

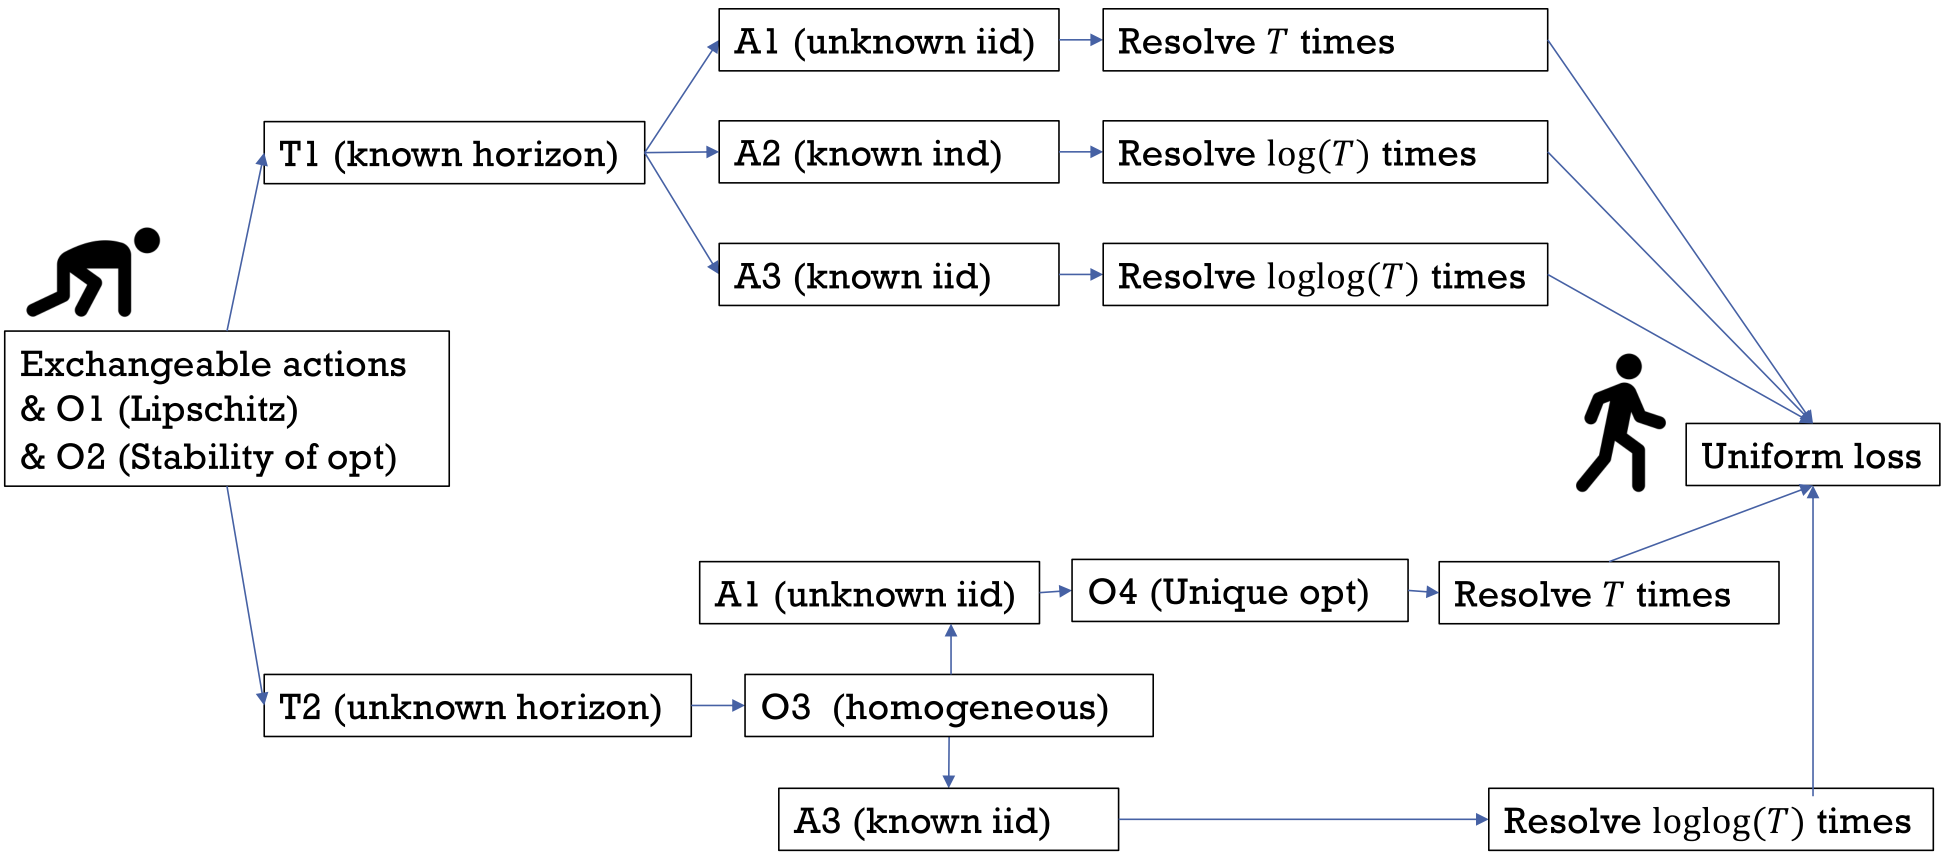

A1 (unknown iid)
Resolve T times
T1 (known horizon)
A2 (known ind)
Resolve $\log(T)$ times
A3 (known iid)
Resolve $\log\log(T)$ times
Exchangeable actions
& O1 (Lipschitz)
& O2 (Stability of opt)
Uniform loss
A1 (unknown iid)
O4 (Unique opt)
Resolve T times
T2 (unknown horizon)
O3 (homogeneous)
A3 (known iid)
Resolve $\log\log(T)$ times

Necessity of assumption O2 (stability of opt)

Suppose in each period we accept/reject an arrival

Each arrival has iid probability ½ to be type 1 or 2

Our objective is to maximize, over known horizon T

$$\max\{x_1, x_2\}$$

$$s.t.\ \ x_1 + x_2 \leq \frac{T}{2}$$

Lipschitz, exchangeable actions, iid… no O2!

Clairvoyant is guaranteed $\frac{T}{2}$; any ALG gets at most $\frac{T}{2} - \Omega(\sqrt{T})$ in exp

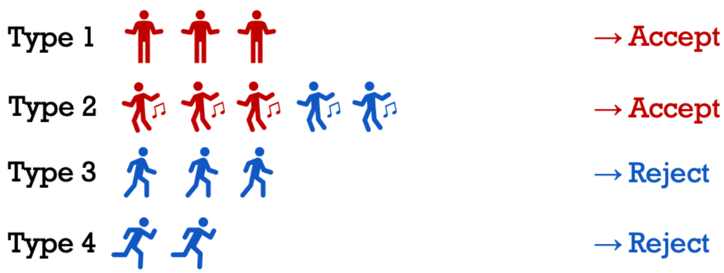

Alternative to O2: Overbooking problem

- Would want to maximize

$$\sum_{\theta} v_\theta x_\theta - c \cdot \mathbb{E}\left[\left(\sum_{\theta} X_\theta - B\right)^+\right]$$

where $X_\theta \sim Bin(x_\theta, q_\theta)$ subject to $x_\theta \leq N_\theta[T]$

- Change of optimal solution when perturbing $N_\theta[T]$

  (Bound for O2)

- Index solution: order types by $\frac{v_1}{q_1} > \frac{v_2}{q_2} > \cdots > \frac{v_k}{q_k}$
- Accept lower-indexed types first

Observe: Index solutions are suboptimal

- Index solutions are NOT optimal in general
- Asymptotically the clairvoyant general and the clairvoyant index solutions look “similar”

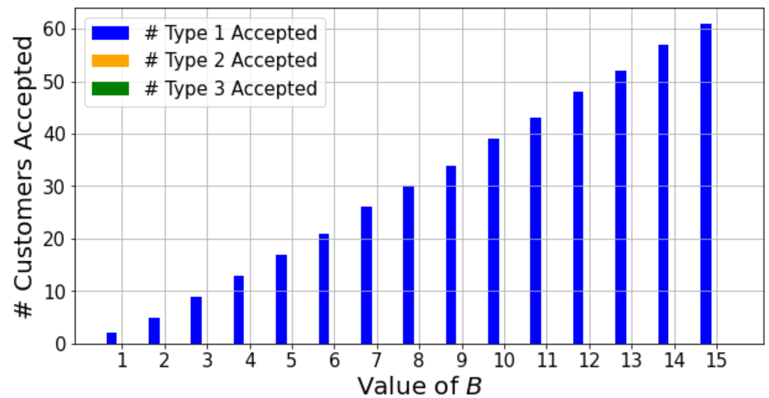

Index Solution

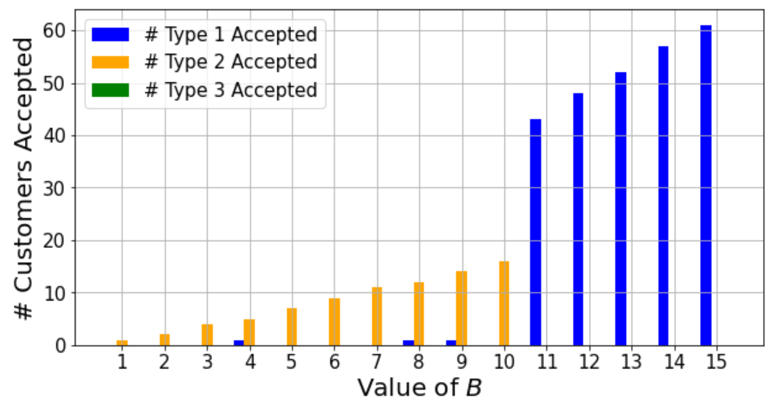

Optimal Solution

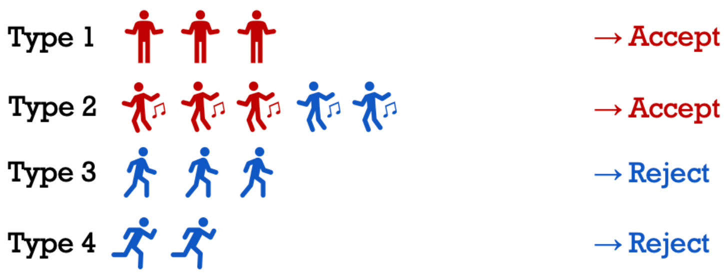

Alternative to O2: Overbooking problem

- Would want to maximize

$$\sum_{\theta} v_\theta x_\theta - c \cdot \mathbb{E}\left[\left(\sum_{\theta} X_\theta - B\right)^{+}\right]$$

where $X_\theta \sim Bin(x_\theta, q_\theta)$ subject to $x_\theta \le N_\theta[T]$

- Change of optimal solution when perturbing $N_\theta[T]$

  (Bound for O2)

- Index solution: order types by $\frac{v_1}{q_1} > \frac{v_2}{q_2} > \cdots > \frac{v_k}{q_k}$
- Accept lower-indexed types first

- Can bound as O(1)
  - loss of only considering index solutions
  - change of best index solution when perturbing $N_\theta[T]$
  - Effectively proves O2 for a restricted set of solutions

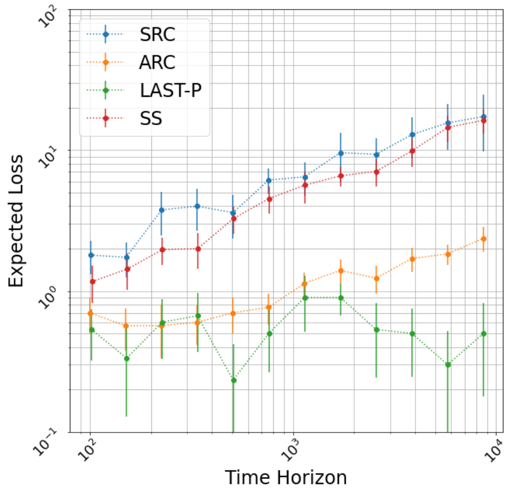

$$\Theta = \{2,3\}; B = 9$$
$$p_2 = \frac{3}{4}, p_3 = \frac{1}{4}$$

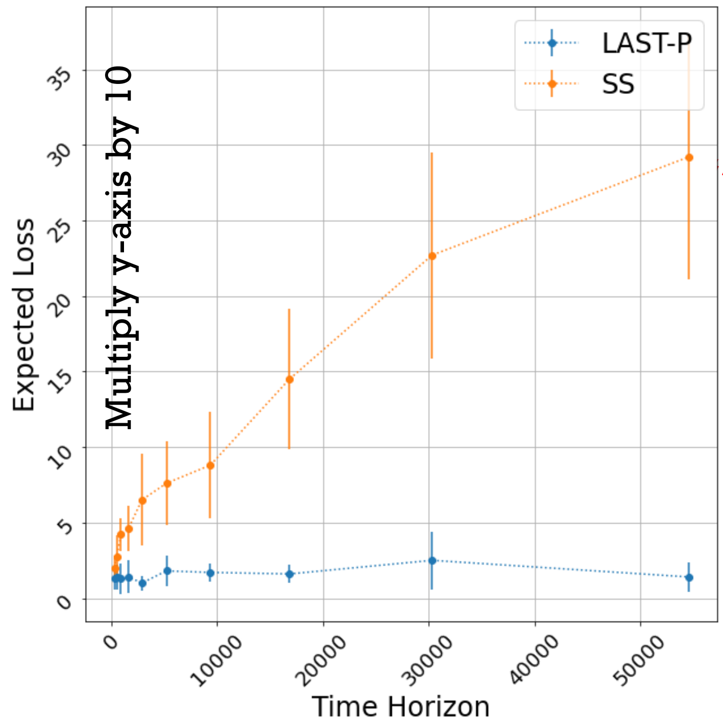

Comparison with SS from Gupta & Radovanovic, *OR'20*

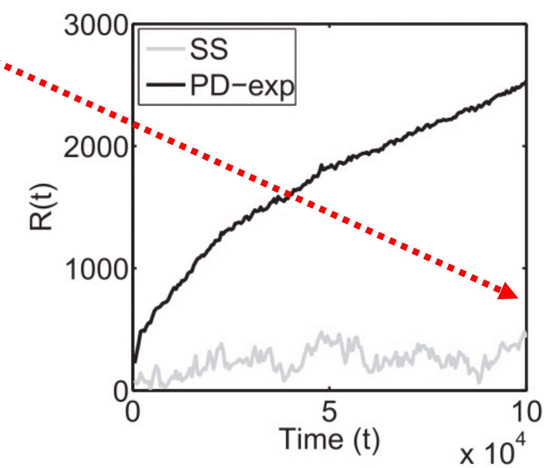

$$\Theta = \{1,3,4,5,8\}; B = 10$$
$$p_1 = \frac{1}{4}, p_3 = \frac{1}{4}, p_4 = \frac{1}{8}, p_5 = \frac{1}{4}, p_8 = \frac{1}{8}$$

Numerical results (Bin packing)

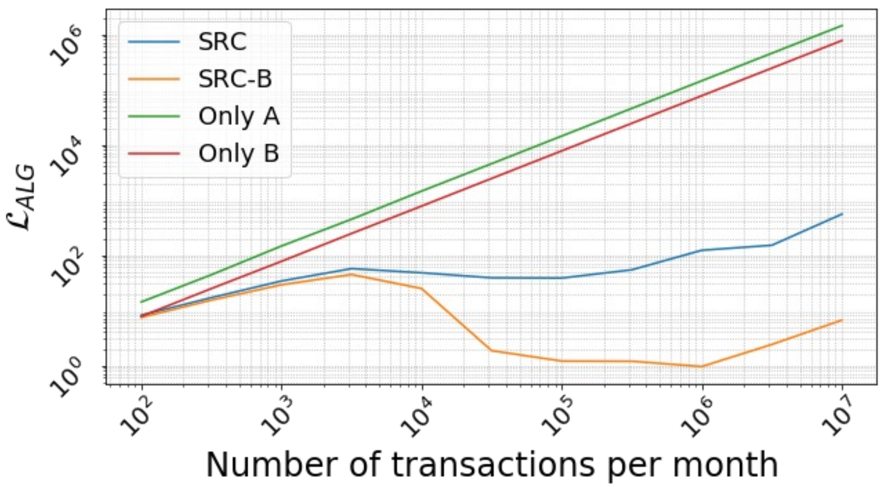

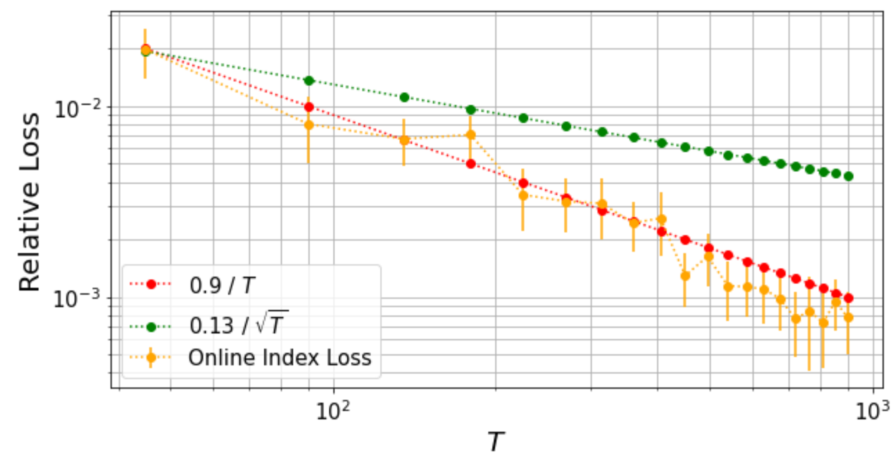

Numerical results (Load balancing & overbooking)

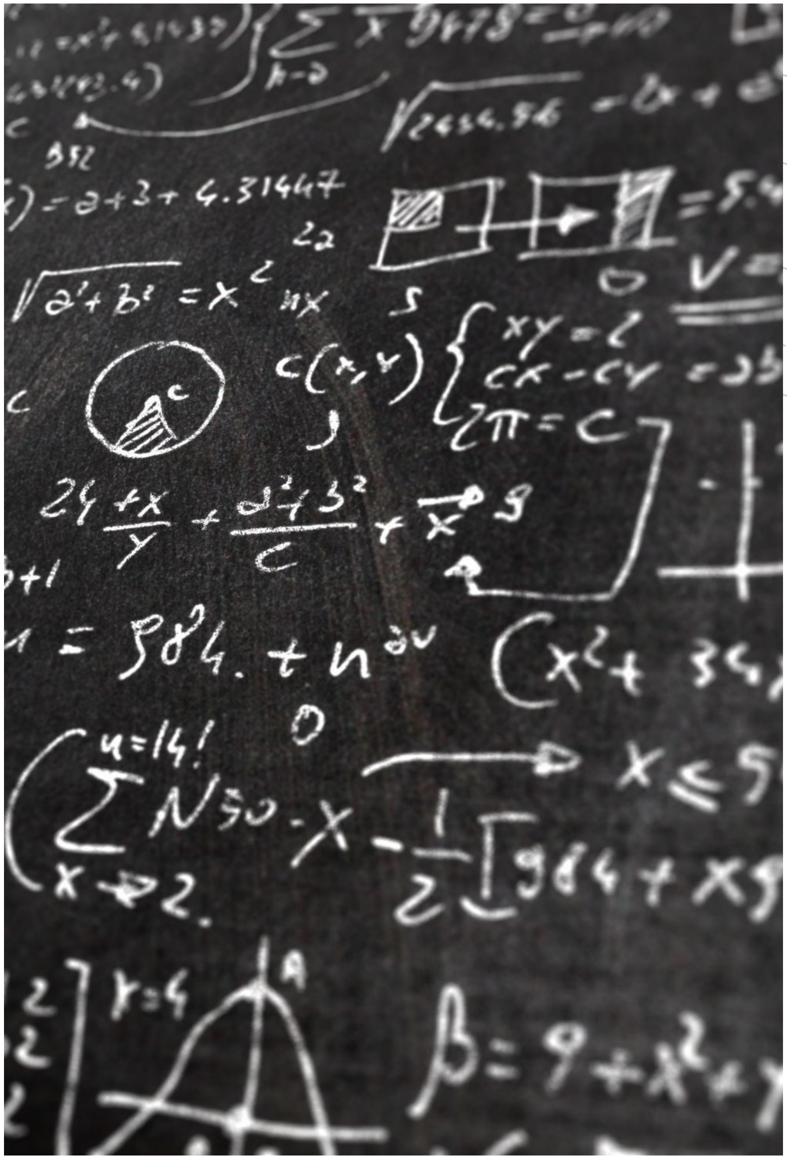

Algorithmic ideas

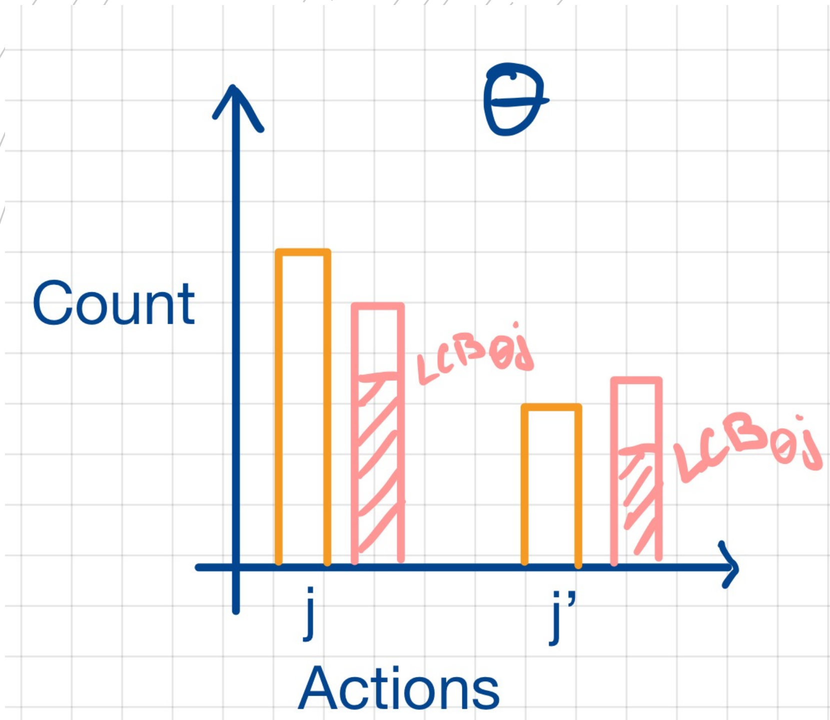

(by O2/stability of optimum)

- In period t, define semi-clairvoyant $OPT[t]$ that follows ALG until t-1, then is clairvoyant until T

$$OPT[t] = \max f(\vec{x})$$
$$\sum_j x_{\theta j} = N_\theta[T]$$
$$x_{\theta j} \geq x_{\theta j}[t-1]$$

Solve deterministic problem in which remaining arrivals are replaced by expectation (or proxy)

$$\text{DLP[t]} = \max f(\vec{x})$$
$$s.t. \forall\theta \sum_j x_{\theta j} = N_\theta[t] + (\underbrace{\bar{t}}_{= T - t} - 1)p_\theta$$
$$x_{\theta j} \geq x_{\theta j}[t-1]$$

Create a LCB on # of times an action is played by considering that whp $\left(\frac{1}{\bar{t}^3}\right)$

$$\forall\theta\colon\ |N_\theta[T] - N_\theta[t] - (T - t - 1)p_\theta| \leq \sqrt{\bar{t}\log(\bar{t})}$$

and consequently, whp, $OPT[t]$ uses $a_{\theta j}$ as often as DLP[t] does $-\delta\sqrt{\bar{t}\log(\bar{t})}$

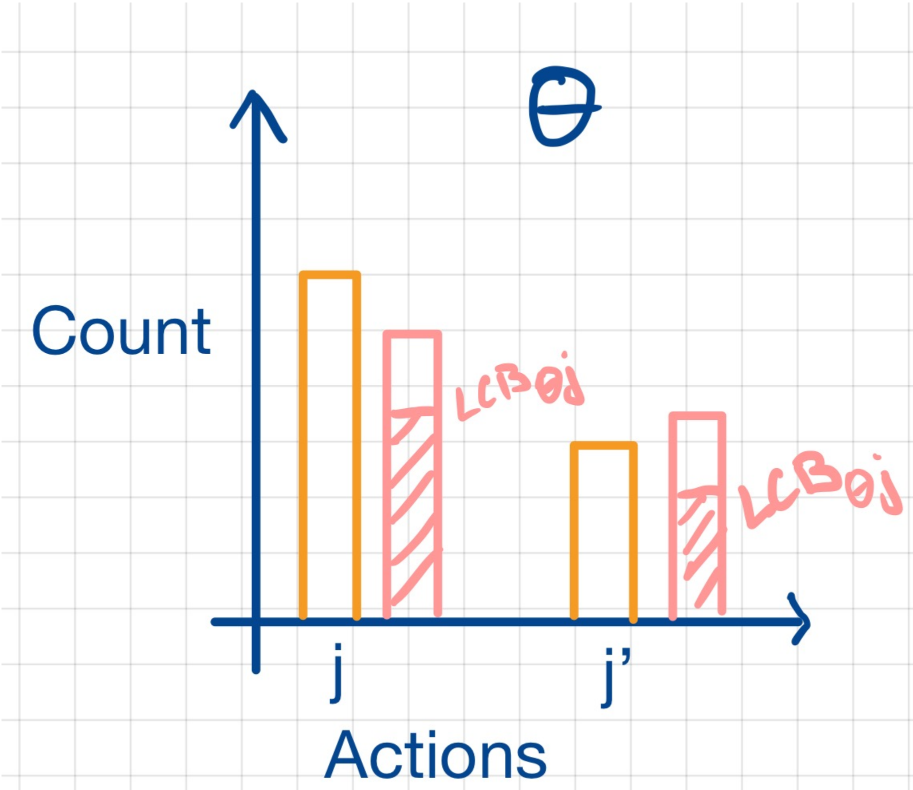

Suppose the lower confidence bounds hold true for every type and every action. If some periods later, each action $a_{\theta j}$ has been taken at most $LCB_{\theta j}$ times, then semi-clairvoyant achieves the same objective after these periods as it did before (old sol'n still feasible).

Mistakes

Loss bound

- Say at t we find LCBs that we use until t'
- In period t' we resolve to obtain new LCBs
- If we resolve in periods $t_1 = 1, \dots t_s = T$:

$$\mathbb{E}[\,OPT - ALG(\theta_1, \dots, \theta_T)] \\ = \mathbb{E}[OPT[1] - OPT[T]]$$

$$= \mathbb{E}[\sum_{\tau=1,..,s-1} OPT[t_\tau] - OPT[t_{\tau+1}]\ ]$$

$$\leq \sum_{\tau=1,..,s-1} L \cdot (t_{\tau+1} - t_\tau)\mathbb{P}[LCBs\ wrong\ at\ t_\tau]$$

$$\leq \sum_{\tau=1,..,s-1} L \cdot (t_{\tau+1} - t_\tau)\ {}^{1}/_{(T - t_\tau)^3}$$

$$\leq M$$

Exchangeable actions
& O1 (Lipschitz)
& O2 (Stability of opt)

T1 (known horizon)

A3 (known iid)

Resolve $\log\log(T)$ times

Uniform loss

First path to uniform loss

- Requires us to see that after each resolving we have actions until a sublinear number of periods is left
- Resolve with $\bar{t} := T - t$ periods left
- *Budget of actions* for type θ is equal to at least
$$\bar{t}p_\theta - \delta\ell\sqrt{\bar{t}\log(\bar{t})}$$
- Will need to resolve after that many type θ arrivals
- Whp we won't need to resolve until
$< \bar{t}^{3/4}$ periods left

Unknown horizon (known iid dist)

Clairvoyant optimum:

$$OPT = OPT[1] = \max f(\vec{x})$$

$$s.t. \forall\theta \sum_j x_{\theta j} = N_\theta[T]$$

$$x_{\theta j} \geq 0$$

Stochastic policy:

$$\max f(\vec{x})$$

$$s.t. \forall\theta \sum_j x_{\theta j} = \mathbb{E}[N_\theta[T]]$$

$$x_{\theta j} \geq 0$$

Denote solution by $y_{\theta j}$; take action $a_{\theta j}$ w.p. $y_{\theta j}/\mathbb{E}[N_\theta]$

Denote $p_{\theta j} = p_\theta y_{\theta j}/\mathbb{E}[N_\theta]$ (prob. of playing $a_{\theta j}$)

Observe: if $f(\cdot)$ is homogeneous (O3) we don't need to know T to obtain this policy!

Stochastic policy upper confidence bounds

- How often does $OPT[1]$ take action $a_{\theta j}$?
- The DLP uses action $a_{\theta j}$ exactly $Tp_{\theta j}$ times
- With high probability (whp)

$$\forall\theta |N_\theta - \mathbb{E}[N_\theta]| \leq \sqrt{Tlog(T)}$$

If so, then there exists $OPT[1]$ that uses action $a_{\theta j}$ at least

(LCB) $Tp_{\theta j} - \delta\sqrt{Tlog(T)}$ times (by O2/stability of optimum)

- The stochastic policy that follows DLP takes $a_{\theta j}$

$Bin(p_{\theta j}, T - T^{3/4})$ in the first $T - T^{3/4}$ periods

(UCB) $Bin(p_{\theta j}, T - T^{3/4}) \leq (T - T^{\frac{3}{4}})p_{\theta j} + \sqrt{Tlog(T)}$ whp

Large T, **constant** $p_{\theta j}, \delta$ $< Tp_{\theta j} - \delta\sqrt{Tlog(T)} =$ **(LCB)** for $OPT[1]$

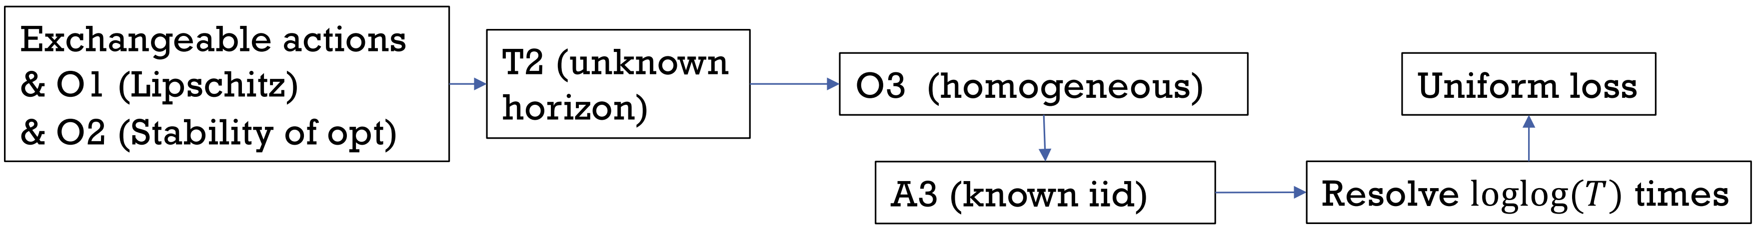

2nd path to uniform loss

Caveats for unknown distribution

- Want to just use empirical estimates so far

- Careful: We don't have good LCBs for actions!

$$Tp_{\theta j} - \delta\sqrt{Tlog(T)}$$

  - Especially true in initial periods
  - Especially true when we don't know T

- Advantage:
  - Stochastic policy initially makes no mistakes whp
  - ➢ may compare ourselves to stochastic policy instead

Algorithm for unknown distributions

Technical subtlety here requires O4 (unique solution for DLP):
Problem arises If the ``optimal'' offline solution varies too much across periods…

$$\widehat{DLP}[t] = \max f(\vec{x}) \quad s.t. \forall\theta \sum_j x_{\theta j} = N_\theta[t] \quad x_{\theta j} \geq x_{\theta j}[t]$$

$$DLP = \max f(\vec{x}) \quad s.t. \forall\theta \sum_j x_{\theta j} = \mathbb{E}[N_\theta[t]] \quad x_{\theta j} \geq 0$$

- Difference between solutions for $\widehat{DLP}[t]$ & DLP:
  - With probability $1 - 1/t^2$ we have (*good event*)

    $$|\hat{x}_{\theta j}[t] - x_{\theta j}[t]| \leq \frac{\delta}{\sqrt{t\log(t)}}$$

  - Threshold to avoid taking $a_{\theta j}$ with $x_{\theta j} = 0$:

    $$\hat{y}_{\theta j}[t] = 0 \ if \ \hat{x}_{\theta j}[t] < \frac{\delta}{\sqrt{t\log(t)}}$$ (& scale other actions up)

  - Randomize based on $\hat{y}$
- May make mistakes if either
  - *Good event not true* (errors are summable) or
  - We scale an action (that DLP takes) up by too much

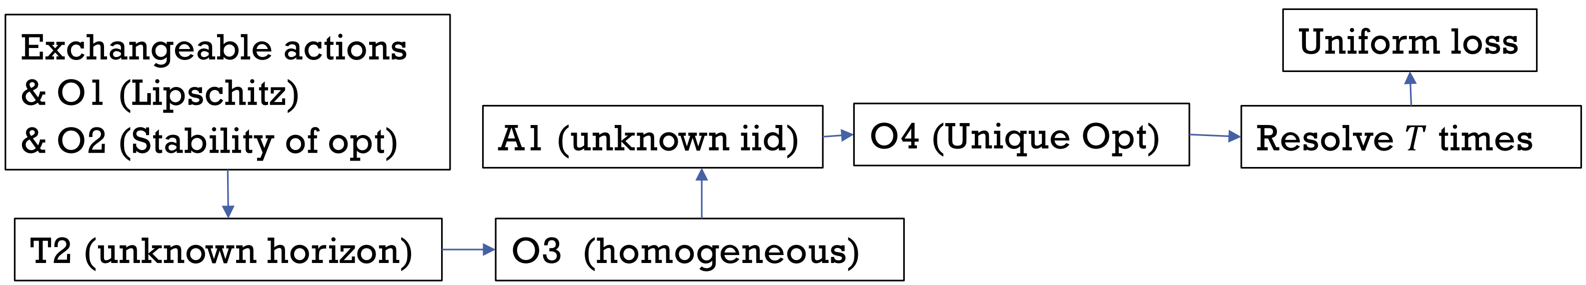

3rd path to uniform loss

Necessity of heads-up (T2)

Bin-packing with bins of size 3

Items are, with prob. ½ , of size 1 or 2

Possible configurations are (1,1,1) and (1,2)

Horizon of length T or $\frac{T}{2}$ (with no heads-up)

With constant probability the following **both** occur

$N_1[T/2] \geq \frac{T}{4} + \sqrt{T}$

$N_1[T] \leq \frac{T}{2} - \sqrt{T}$

$o(\sqrt{T})$ loss at time T/2 requires creating $\Omega(\sqrt{T})$ bins of configuration (1,1,1) whereas $o(\sqrt{T})$ loss at time T requires having created $o(\sqrt{T})$ such bins

Similar result applies to geometric horizon length

Necessity of p_{min}

(A1/A2/A3)

Multi-secretary with budget $\frac{T}{2}$ iid arrival types

$v_3 = 3$ has probability $\frac{1}{2} - \frac{1}{T^{\frac{3}{4}}}$ (mean $\frac{T}{2} - T^{1/4}$)

$v_2 = 2$ has probability $\frac{1}{T^{\frac{3}{4}}}$ (mean $T^{1/4}$)

$v_1 = 1$ has probability $\frac{1}{2}$ (mean $\frac{T}{2}$)

After $\frac{T}{2}$ (whp) one has either

accepted at least $T^{\frac{1}{4}}/8$ arrivals of type 2

or rejected most $T^{\frac{1}{4}}/8$ of type 2

Berry-Esseen: constant probability to have

at least $\frac{T}{2}$ type-3 over entire horizon

at most $\frac{T}{2} - T^{\frac{1}{2}}$ type-3 over entire horizon

Even with full knowledge of the first $\frac{T}{2}$ arrivals do not know, whether to accept 0 or all type-2 arrivals

Takeaways (Overbooking)

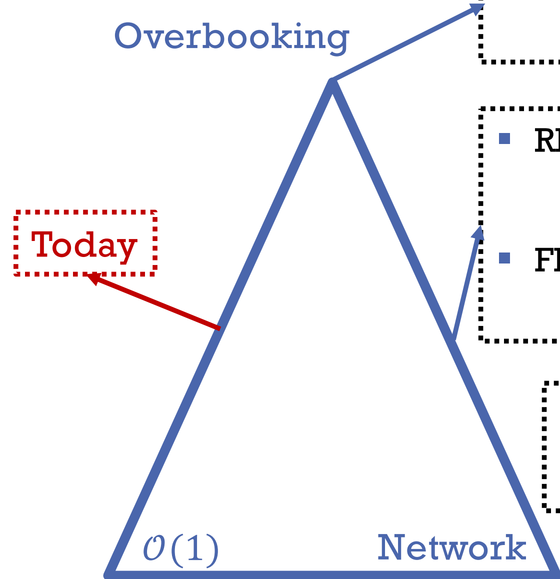

- DPD algorithm (Erdelyi and Topaloglu, 2010)
  - w/ no-shows: $\Omega(T)$ loss

- RLP algorithm (Kunnumkal et al., 2012)
  - w/ no-shows: $\Theta(\sqrt{T})$ loss
- Fluid policy (Dai et al., 2019)
  - w/ cancellations and no-shows: $\Theta(\sqrt{T})$ loss

- RLP Estimator (Talluri and Van Ryzin, 1999)
- T^2 Policy (Reiman and Wang, 2008)

- Fluid Bayes Selector (Vera and Banerjee, 2020)
- Resolving Heuristics (Jasin & Kumar, 2012; Bumpensanti and Wang (2020))

Budget-Ratio Policy (Arlotto and Gurvich, 2019)

➢ Based on *Overbooking with bounded Loss* with Kamessi Zhao (EC'21, MOR'22)

Takeaways (Bin packing)

	Regret	Distr.	Algorithm & Remarks
Shor (1986)	$\Omega(\sqrt{T}\log T)$	Unif$[0,1]$	Lower bound
Shor (1986); Asgeirsson (2002)	$\Theta(\sqrt{T})$	Unif$[0,1]$	Best Fit; Known T
Shor (1991)	$O(\sqrt{T}\log T)$	Unif$[0,1]$	Best Fit
Rhee and Talagrand (1993a,b)	$K\sqrt{T}\log^{3/4} T$	General	Double-overflow; unspecified constant K
Csirik et al. (2006)	$B\sqrt{T}$	Int. supp.	Sum-of-squares; bin size B
Gupta and Radovanović (2020)	$B\sqrt{T}$	Int. supp.	Lagrangian-based; bin size B
Banerjee and Freund (2020)	M	Int. supp.	Re-solving; Known T; problem-dependent M
Liu & Li (2021)	$C\sqrt{T}$	General	Adaptive; Known T; $C \leq 11$
Liu & Li (2021)	$C\sqrt{T}$	Ran. Perm.	Adaptive; Known T; $C \leq 13$

Table from *Online Bin Packing with Known T, Liu & Li, '21*

- Based on *Good prophets know when the end is near* with Sid Banerjee (SIGMETRICS'20, ??'??)

Takeaways (Modeling)

- Heads-up for horizon end
  - In-between adversarial, stochastic, and known
- Positive results are comparable to known horizon
  - Provable improvements vs. geometric/adversarial
- In many applications it may be the most realistic(?)

Takeaways
(exchangeable actions lens)

- Captures wide set of problems, but precludes
  - Many inventory problems (arrivals & departures)
  - Resource allocation with (traditional) cancellations
- Instance-dependent for the most part
  - In some cases (overbooking): provably unavoidable
  - Though: numerically, the constants don't kick in!
- Prove O2 (stability) for nonlinear objectives
  - Potential alternative: near-optimal alternate solution
  - Requires ad hoc machinery (as for overbooking)

Summary

- Algorithmic/analytical framework
- Different sets of assumptions for O(1) loss
- New guarantees
  - Bin packing
  - Single-leg RM with overbooking
- (Almost) minimal set of assumptions

T: time horizon
B: capacity
v_j: revenue of type j
p_j: show up probability of type j

Appendix

Instance-independent Bound

$\widehat{OPT}_{\vec{A}}$: clairvoyant general obj.
$OPT_{\vec{A}}[1]$: clairvoyant index obj.
$OPT_{\vec{A}}[t]$: semi-clairvoyant index obj. at t
$OBJ_{\vec{A}}$: online index obj.

- Instance-independent: v, p allowed to change with T
- Any online policy incurs a loss of $\Omega(\sqrt{T})$ due to the inherent uncertainty in arrivals
- E.g. Suppose $\mathrm{B} = \frac{T}{6}$. Moreover,

$$\lambda_1 = \frac{1}{6}, v_1 = \frac{1}{2}, p_1 = 1$$
$$\lambda_2 = \frac{1}{3}, v_2 = \frac{1}{\sqrt{T}}, p_2 = \frac{3}{\sqrt{T}}$$
$$\lambda_3 = \frac{1}{2}, v_3 = 0, p_3 = 1$$

- Do not know how many type 1 customers arrive (error $\sim \Theta(\sqrt{T})$) and are thus likely to make mistakes in type 2
  - $N_1 \geq \frac{T}{6}$: no type 2 customer should be accepted
  - $N_1 \leq \frac{T}{6} - \sqrt{T}$: "almost" all type 2 customer should be accepted