Opening Pandora's Box: the Correlated Case

Evangelia Gergatsouli
University of Wisconsin-Madison

Joint work with: Shuchi Chawla[1], Yifeng Teng[2], Jeremy McMahan[3], Christos Tzamos[3,4], Ruimin Zhang[2]
[1]UT-Austin, [2]Google Research, [3]UW-Madison, [4]University of Athens

INFORMS '23, Phoenix AZ, October 2023

A Search Problem

Find the best alternative, with costly information!

- Information is not free!
- Explore alternatives (open boxes)
- Stop anytime & take best so far

} **Strategy**

A Search Problem

Find the best alternative, with costly information!

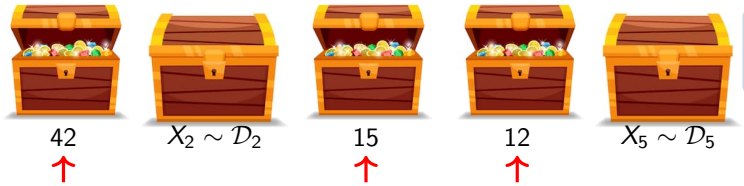

Opening cost: 3
Final option: box 4
Total cost: 12 + 3

- Information is not free!
- Explore alternatives (open boxes)
- Stop anytime & take best so far

} **Strategy**

Goal: find min cost strategy!

Previous Work

Weitzman's algorithm gives the optimal! [Weitz 1979]

- Calculate *reservation value* σ_i for every box[1]
- Search boxes in order of increasing index until:
  - Current min price seen smaller than index of next box

[1] (Gittins index)

Previous Work

Weitzman's algorithm gives the optimal! [Weitz 1979]

- Calculate *reservation value* σ_i for every box[1]
- Search boxes in order of increasing index until:
  - Current min price seen smaller than index of next box

Partially Adaptive:
fix order from beginning

[1] (Gittins index)

Previous Work

Weitzman's algorithm gives the optimal! [Weitz 1979]

- Calculate *reservation value* σ_i for every box[1]
- Search boxes in order of increasing index until:
  - Current min price seen smaller than index of next box

Crucial assumption: distributions are **independent**!

What about **correlation?**

[1] (Gittins index)

Previous Work

Weitzman's algorithm gives the optimal! [Weitz 1979]

- Calculate *reservation value* σ_i for every box[1]
- Search boxes in order of increasing index until:
  - Current min price seen smaller than index of next box

Crucial assumption: distributions are **independent**!

What about **correlation?**

Our work: Algorithms for Pandora's Box with Correlations

[1] (Gittins index)

Why Correlation?

Pandora's Box with Correlations: Setting

Given:

- m scenarios with probabilities p_i
- Matrix of values as below

Goal: Open boxes, stop & pick value
Minimize: Sum of
opening cost + value chosen

	Box 1	**Box 2**	**Box 3**	...	**Box n**
$p_1 \to$ **Scenario** 1:	42	13	15		24
$p_2 \to$ **Scenario** 2:	0	24	94	...	2
$\vdots$			$\vdots$		
$p_m \to$ **Scenario** m:	31	15	9	...	2

$\mathcal{D}$

Finding the Optimal

Sample access to $\mathcal{D}$

Can we approximate the optimal? $\rightarrow$ No: arbitrarily encode location of box

Cannot learn arbitrary mapping with finitely many queries

Finding the Optimal

Sample access to $\mathcal{D}$

Can we approximate the optimal? $\rightarrow$ No: arbitrarily encode location of box

Cannot learn arbitrary mapping with finitely many queries

Explicitly given $\mathcal{D}$

Can we approximate the optimal? $\rightarrow$ (probably) Not within constant in poly-time

Why? Pandora's Box equivalent to Uniform Decision Tree [Chawla, G., McMahan, Tzamos, APPROX '23]

- $\tilde{O}(\log m)$ in poly-time
- $\tilde{O}(1/\alpha)$ in time $n^{\tilde{O}(m^\alpha)}$ for $\alpha \in (0,1)$.

Strategies

Strategy: (1) what is the next box? (2) when do we stop?

- Fully Adaptive (FA): next box/stopping rule both **adaptive**
- Partially Adaptive (PA): fixed order, adaptive stopping time

Optimal for Independent!

Strategies

Strategy: (1) what is the next box? (2) when do we stop?

- Fully Adaptive (FA): next box/stopping rule both **adaptive**
- Partially Adaptive (PA): fixed order, adaptive stopping time

Optimal for Independent!

Our results:
We can approximate PA within **constant**

- 9.22-approx [Chawla, G, Teng, Tzamos, Zhang FOCS '20]
- 4.42-approx [G, Tzamos, NeurIPS 2023]

Initial Approach

Recall Strategy: (1) decide order (2) decide stopping time
Idea: separately decide (1) and (2)

Algorithm in 2 parts:

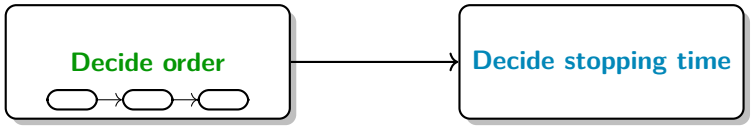

Main Algorithm:
Reduce to Min Sum Set Cover

Myopic stopping lemma: Reduce to generalized Ski-Rental

This Algorithm: 9.22-approx to the Partially Adaptive optimal

A Simpler Approach

- Calculate a reservation value σ_i for every box i for every round
- Open box b with min σ_i, stop if $v_b \leq \sigma_b$
- Box is now free $c_b = 0$
- Update Prior with $v_b > \sigma_b$

This Algorithm: 4.42-approx to the Partially Adaptive optimal

A Simpler Approach

- Calculate a reservation value σ_i for every box i for every round
- Open box b with min σ_i, stop if $v_b \leq \sigma_b$
- Box is now free $c_b = 0$
- Update Prior with $v_b > \sigma_b$

Almost tight: **4** is best possible!

This Algorithm: 4.42-approx to the Partially Adaptive optimal

Summary & Similar settings

Correlated Distributions

- **vs Fully Adaptive**: equivalent to Uniform Decision Tree [Chawla, G., McMahan, Tzamos: APPROX '23]
- **vs Partially Adaptive**: constant approximation, results even for combinatorial settings [Chawla, G., Teng, Tzamos, Zhang : FOCS '20, G. & Tzamos: NeurIPS '23]
  - k boxes (constant approx), matroid of basis k ($\log k$-approx)
  - all algorithms learnable from samples
  - maximization: cannot approximate within constant unless P=NP
- **Online setting**: play a new instance of the game for T rounds:
  - approx. no regret algorithms for adversarial instances [G. & Tzamos: ICML '22]
  - No regret algorithms when context is available (using independent boxes) [Atsidakou, Caramanis, G., Papadigenopoulos, Tzamos: ArXiv '22]

Similar Settings

Many more work on variants of original

- Non-obligatory inspection [STOC '23: Beyhaghi, Cai & Fu, Li, Liu] (**next talk**!)
- Order-constrained [Boodaghians et al.: EC '20]
- Commited Pandora's Box [Fu, Li, Xu: ICALP '18, Segev, Singla: EC '21]
- Combinatorial costs [Berger, Ezra, Feldman, Fusco: EC '23]
- many more...

Survey: Recent Developments in Pandora's Box Problem: Variants and Applications [Beyhaghi, Cai, 2023]

Thank you!