ArXiv Version

Pandora's Box Problem Over Time

Artem Tsikiridis (CWI)

joint work with

Georgios Amanatidis
University of Essex

Federico Fusco
Sapienza University of Rome

Rebecca Reiffenhäuser
University of Amsterdam

Dutch Seminar on Optimization 31/10/24, to appear at WINE 2024

The Pandora's Box Problem [Weitzman 1979]

Models cost of information in search problems

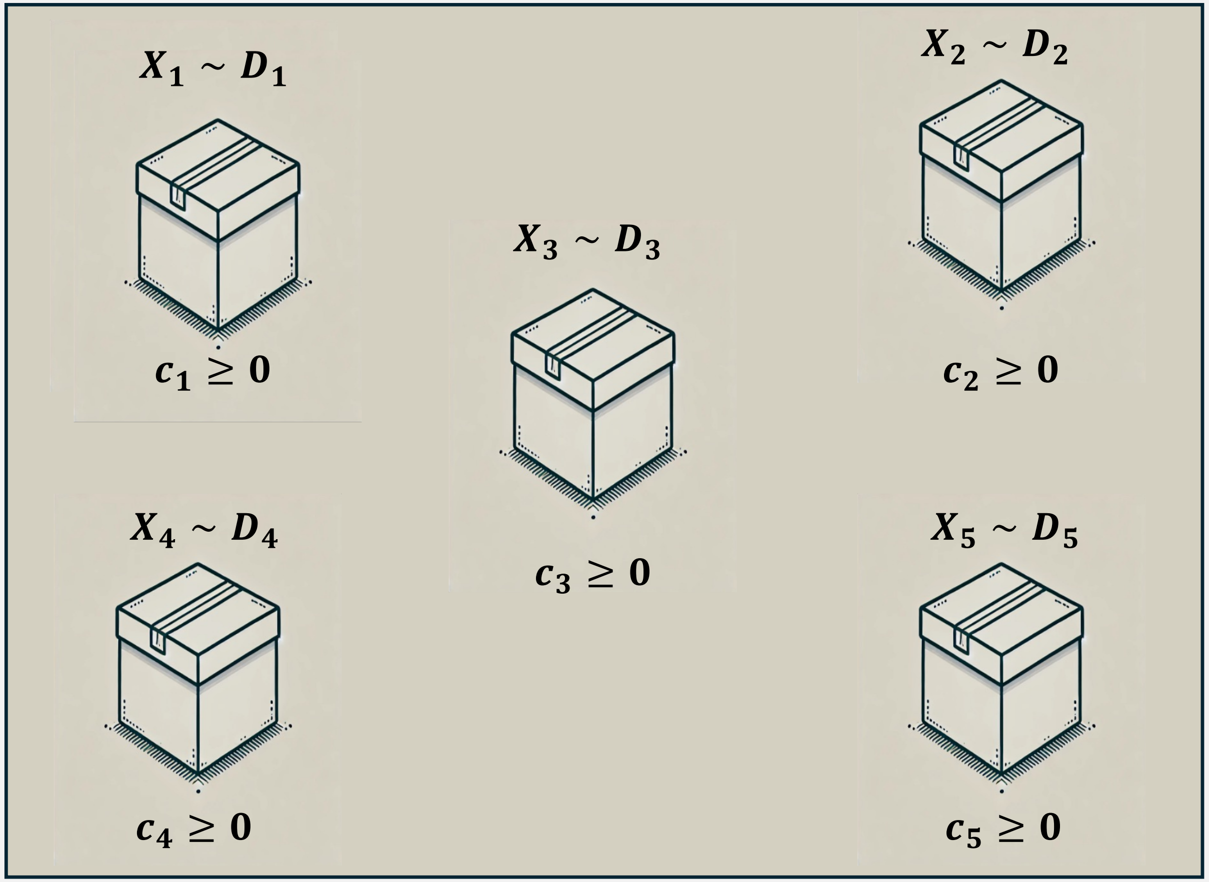

Set of boxes $[n] = \{1, \dots, n\}$

Each box $i \in [n]$ labeled with cost $c_i \geq 0$ and distribution of reward D_i . $(D_1, \dots, D_n)$ are independent and non-negative.

Strategy π: at each round t either:

Open a box $\boldsymbol{j}$ i.e., draw $X_j \sim D_j$ and pay c_j
or
Stop Searching.

The Pandora's Box Problem

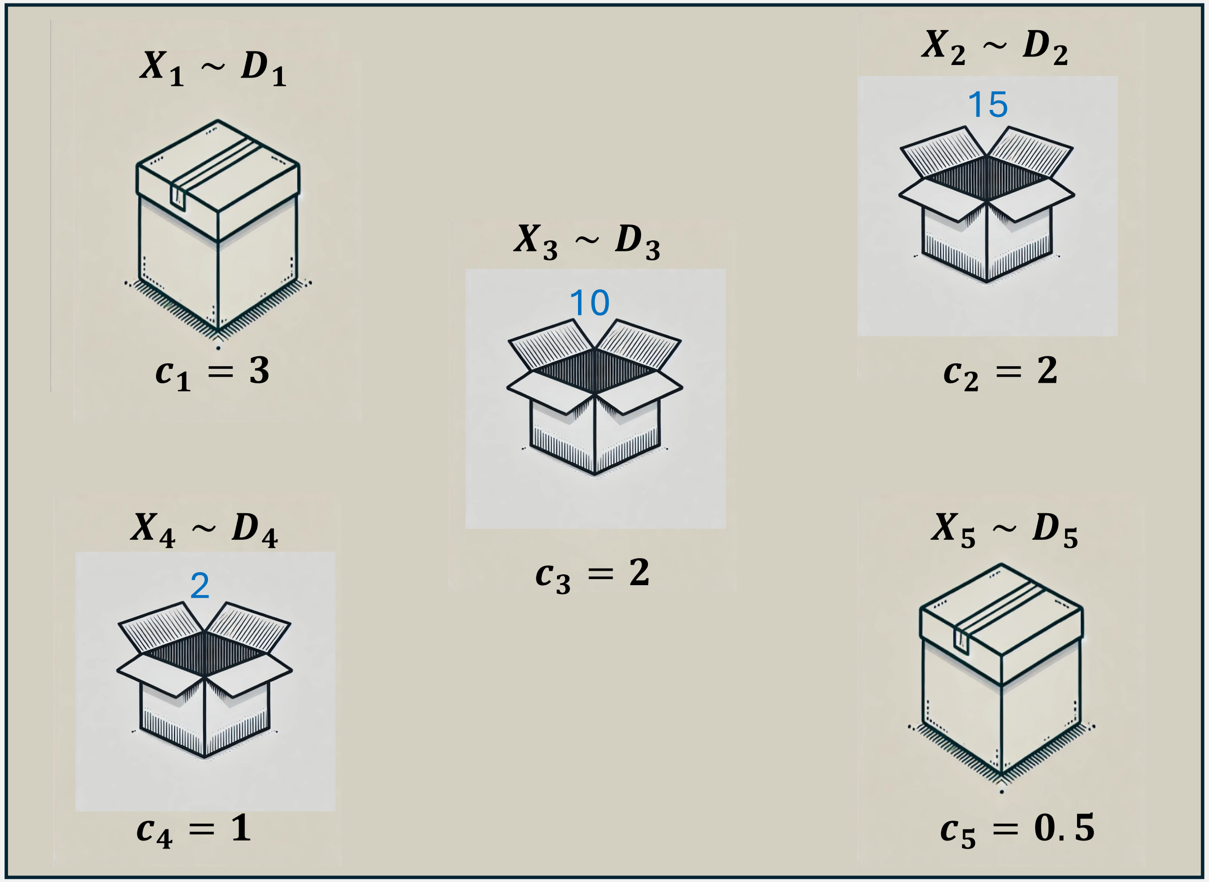

Let $S(\pi) \in [n]$ be the boxes opened by π.

Goal: Maximize expected net gain i.e., find

$$\pi^* \in \arg\max_{\pi} \mathbb{E}\left[\max_{j \in S(\pi)} X_j - \sum_{j \in S(\pi)} c_j\right]$$

Net gain $= 15 - (2 + 2 + 1) = 10$

Optimal strategy:

- Could be adaptive i.e., open different boxes depending on the history of the boxes that have already been opened.
- Surprisingly, Weitzman gave a very structured solution!

Weitzman's Rule

1. Before opening boxes, compute reservation value r_i s.t. $\mathbb{E}[\max\{X_i - r_i, 0\}] = c_i$ for every box $i \in [n]$.

2. Order boxes in non-increasing order by r_i.

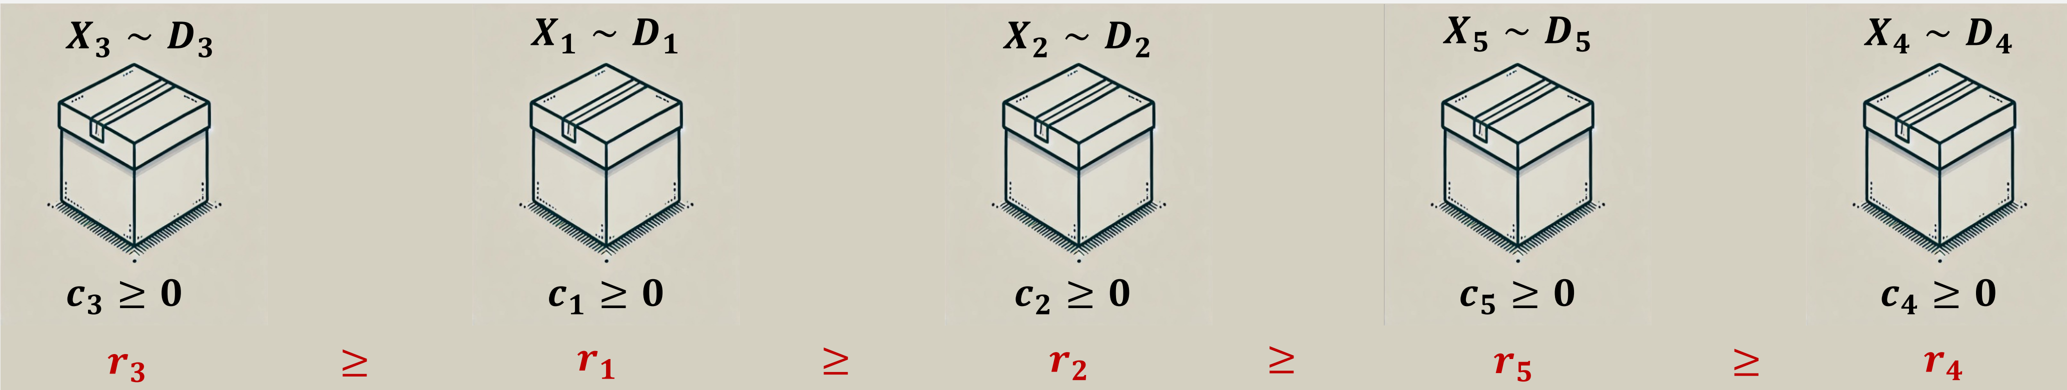

Greedy strategy: Open box with largest r_i if you have not observed a larger reward before. When no such box remains uninspected, stop searching.

Theorem [Weitzman 1979]: Weitzman's Rule maximizes expected net gain.

Order non-adaptive strategy:
- search order is predetermined!
- stopping rule is adaptive

Recent Work on the Pandora's Box Problem

[Kleinberg, Waggoner & Weyl 2016]: connection with mechanism design **and** new proof template

Survey Paper: [Beyhaghi and Cai 2023].

- non-obligatory inspection (may claim the expected reward without opening box): [Doval 2018; Beyhaghi and Kleinberg 2019; Fu, Li and Liu 2023; Beyhaghi and Cai 2023]
- "committed" variant (must claim last opened box): [Fu, Li and Xu, 2018; Esfandiari, HajiAghayi, Lucier and Mitzenmacher, 2019; Segev and Singla, 2021]
- combinatorial constraints: [Kleinberg, Waggoner & Weyl 2016; Singla 2018]
- interdependent valuations [Chawla, Gergatsouli, Teng, Tzamos, and Zhang 2020], combinatorial costs [Berger, Ezra, Feldman and Fusco 2023], contract design [Hoefer, Schecker and Schewior 2024], online learning [Guo, Huang, Tang and Zhang 2021; Gergatsouli and Tzamos, 2022; Gatmiry, Kesselheim, Singla and Wang, 2024]

Our Results

An extension of Pandora's Box to settings where time is a factor:

1. after opening a box, its value starts "degrading"
2. cost of opening a box changes over time
3. each box has a processing time

Setting	Approximation
(**1**)	1.37
(**1**) & (**2**)	8+ε
(**1**) & (**2**) & (**3**)	21.3

Technical Challenge: No Weitzman's rule, different exploration-exploitation dilemma. Problem is NP-Hard.

Our Results: $O(1)$-approximation for the "adaptive" case, using order-non-adaptive strategies and a new algorithm for a submodular maximization problem (to be defined).

Pandora With Value Discounting

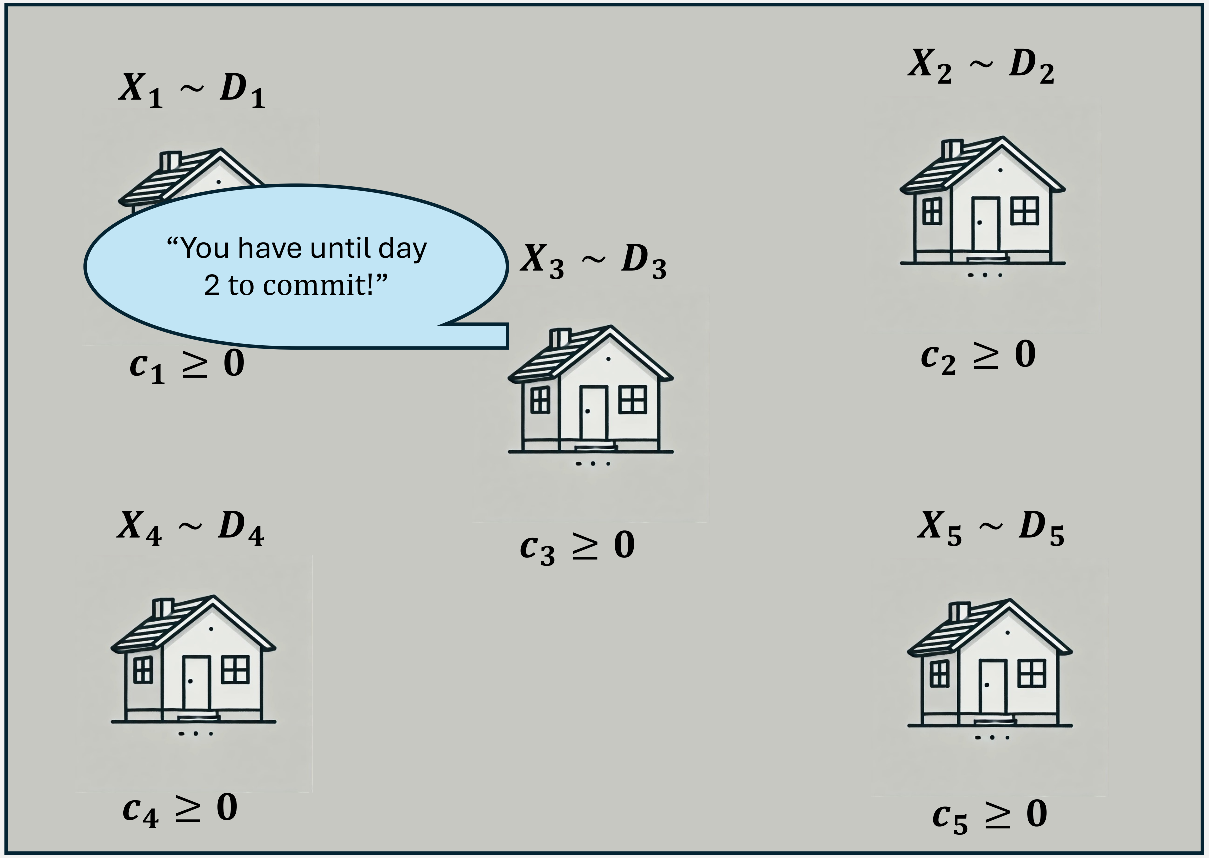

Search Problem Example: House Rental

- set of houses $[n] = \{1, \dots, n\}$
- cost c_i: cost of viewing house i
- draw $X_i \sim D_i$: viewing of house i

Suppose you do a viewing of house 3 on t = 1 (one viewing per day).

But then… X_i will become 0 at $t = 3$.

Pandora With Value Discounting:

If you draw $X_i \sim D_i$ and halt after τ rounds, you may only get

$\bar{v}_i(X_i, \tau)$ = discounted reward after τ rounds

For all $\tau \in \mathbb{Z}^+$, $\bar{v}_i(X_i, \tau) \leq \bar{v}_i(X_i, 0) = X_i$.

Pandora With Value Discounting (2)

For every π, let T_π be the round we stop searching and, for $i \in [n]$, let $t_i(\pi)$ be the (random) round we open box i .

Goal: Find strategy which achieves

$$OPT = \max_{\pi} \mathbb{E}\left[\max_{i \in S(\pi)} \bar{v}_i(X_i, T_\pi - t_i(\pi)) - \sum_{j \in S(\pi)} c_j\right]$$

Recall: reservation value r_i s.t. $\mathbb{E}[\max\{X_i - r_i, 0\}] = c_i$ for every box $i \in [n]$.

For $i \in [n]$, define random variable $Y_i = \min(X_i, r_i)$.

Lemma 1: $OPT \leq \mathbb{E}[\max_{i \in [n]} Y_i]$.

Question: How to interpret the RHS?

Prophet Inequalities

Input: A set of random variables $Y_1, \dots, Y_n$, with $Y_i \sim D_i$. $(D_1, \dots, D_n)$ are independent and non-negative.

A permutation σ of $[Y_1, \dots, Y_n]$ is given to the gambler. At round i, the gambler samples $Y_{\sigma(i)} \sim D_{\sigma(i)}$ and chooses whether to stop and accept $Y_{\sigma(i)}$ or go to the next round.

The prophet knows the random realizations of $Y_1, \dots, Y_n$ (and thus $\arg\max Y_i$) beforehand.

Given a threshold τ, let i^* be the (random) index for which $Y_{\sigma(i^*)} \geq \tau$ and $Y_{\sigma(j)} < \tau$ for all $j < i^*$.

Classical Prophet Inequality [Samuel-Cahn 1984, Kleinberg and Weinberg 2012]**:**

Given $Y_1, \dots, Y_n$ and permutation σ, set $\tau = \frac{1}{2}\mathbb{E}[\max_{i \in [n]} Y_i]$. It holds that $\mathbb{E}\left[Y_{\sigma(i^*)}\right] \geq \frac{1}{2}\mathbb{E}[\max_{i \in [n]} Y_i]$.

Free Order Prophet Inequality [Bubna and Chiplunkar 2023] :

Given $Y_1, \dots, Y_n$, there exists a permutation σ and a threshold τ so that $\mathbb{E}\left[Y_{\sigma(i^*)}\right] \geq 0.7258\mathbb{E}[\max_{i \in [n]} Y_i]$.

A strategy π for the value discounting model

Phase A: Determine tentative schedule of inspection

1. Construct random variables $Y_1, \dots, Y_n$.
2. Let σ be the permutation of $Y_1, \dots, Y_n$ and $\tau > 0$ be the threshold of free-order prophet inequality algorithm of [Bubna and Chiplunkar 2023].

Phase B: Open the boxes using the permutation σ from Phase A using a stopping rule such that

$$\mathbb{E}\left[\max_{i \in S(\pi)} \bar{v}_i\left(X_i, T_\pi - t_i(\pi)\right) - \sum_{j \in S(\pi)} c_j\right] = \mathbb{E}\left[Y_{\sigma(i^*)}\right]$$

Lemma 1: $OPT \leq \mathbb{E}[\max_{i \in [n]} Y_i]$.

Theorem: Strategy π is a 1.37-approximation to Pandora with Value Discounting.

Proof Idea: $\mathbb{E}\left[\max_{i \in S(\pi)} \bar{v}_i(X_i, T_\pi - t_i(\pi)) - \sum_{j \in S(\pi)} c_j\right] = \mathbb{E}\left[Y_{\sigma(i^*)}\right] \geq 0.7258 \cdot \mathbb{E}[\max_{i \in [n]} Y_i] \geq 0.7258 \cdot OPT$

Free order prophet ineq.

Pandora With Time-varying Costs

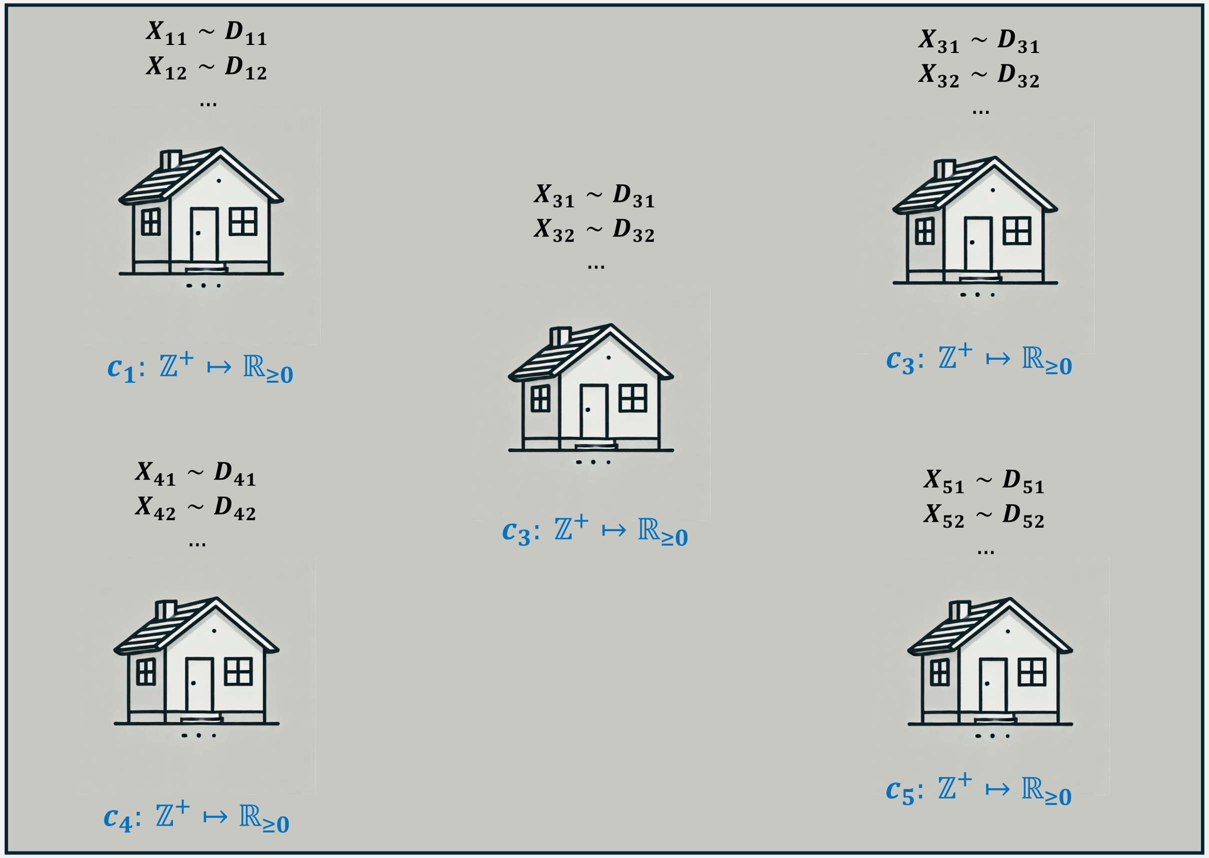

- cost function $c_i\colon \mathbb{Z}^+ \mapsto \mathbb{R}_{\geq 0}$ for each $i \in [n]$
- $c_i(t)$ = cost of opening box i at round t
- different reward distribution per round
- value discounting still there

Reduction to a Constrained Pandora's Box Instance

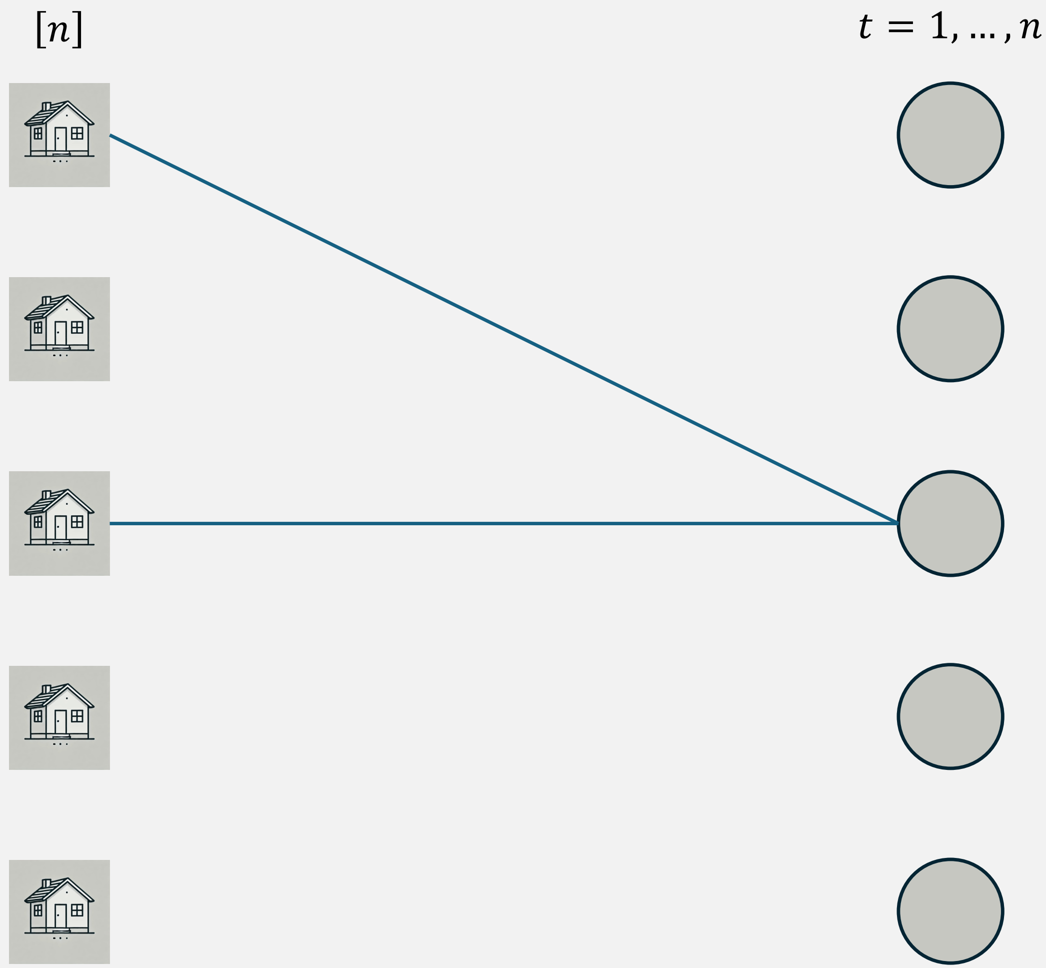

Idea 1: Construct a complete bipartite graph with:
- boxes $[n]$ as left nodes
- timeslots as right nodes

Idea 2: Construct a Pandora's Box instance with:

- a box for every edge (i, t) with:
  - cost $c_i(t)$
  - reward distribution D_{it}

- Feasibility Constraint: a set of opened boxes must be a **matching**

Constrained instance captures inspection schedules in Pandora Over Time instance!

An Upper Bound on OPT

$Y_{(i,t)} = \min\left(X_{(i,t)}, r_{(i,t)}\right)$ for every "box"/edge (i,t)

OPT = the optimal net gain in a given Pandora Over Time instance.

$\mathcal{M}$ = the set of bipartite matchings on $[n]\times[n]$.

Lemma 2: It holds that:
$OPT \leq 2 \cdot \max_{M\in\mathcal{M}} \mathbb{E}\left[\max_{(i,t)\in M} Y_{(i,t)}\right].$

maximize $f(M) = \mathbb{E}[\max_{e\in M} Y_e]$

s.t.

$M \in \mathcal{M}$

Observation: $f(M)$ is monotone submodular:

f is monotone and $\forall\, S, T$ it holds that $f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$.

Our Order Non-Adaptive Strategy

Phase A: Use the $(2+\varepsilon)$-approximation algorithm of [Lee, Sviridenko and Vondrak, 2010] to solve the constrained submodular maximization problem and obtain a tentative schedule i.e., a bipartite matching $\widehat{M}$.

Phase B: Open the boxes following the permutation from Phase A using a stopping rule such that:

expected net gain= $\mathbb{E}\left[Y_{\sigma(i^*)}\right]$ (for the classical prophet inequality)

Theorem: Strategy is an $(8+\varepsilon)$-approximation.

Idea: exp. net gain $= \mathbb{E}\left[Y_{\sigma(i^*)}\right] \underbrace{\geq}_{\text{Prophet inequality}} \frac{1}{2}\mathbb{E}[\max_{(i,t)\in\widehat{M}} Y_{(i,t)}] \underbrace{\geq}_{\text{Submodular maximization.}} \frac{1}{4+\varepsilon}\cdot\max_{M\in\mathcal{M}}\mathbb{E}[\max_{(i,t)\in M} Y_{(i,t)}] \underbrace{\geq}_{\text{Lemma 2}} \frac{OPT}{8+\varepsilon}$.

Processing Time

Suppose that each box $i \in [n]$ has:

- a cost function $c_i\colon \mathbb{Z}^+ \mapsto \mathbb{R}_{\geq 0}$ of the time it is opened

- a processing time $p_i \in \mathbb{Z}_0^+$ i.e., number of rounds you must wait to inspect box $i \in [n]$

Example: *If you open box i at time t, you may only open another box at time $t + p_i + 1$.*

Exploration/Exploitation Dilemma: If you open box i, you may "miss out" on "good" boxes being cheap for the next p_i rounds!

Block Bipartite Hypergraphs

$p_i = 0$ for all $i \in [n]$

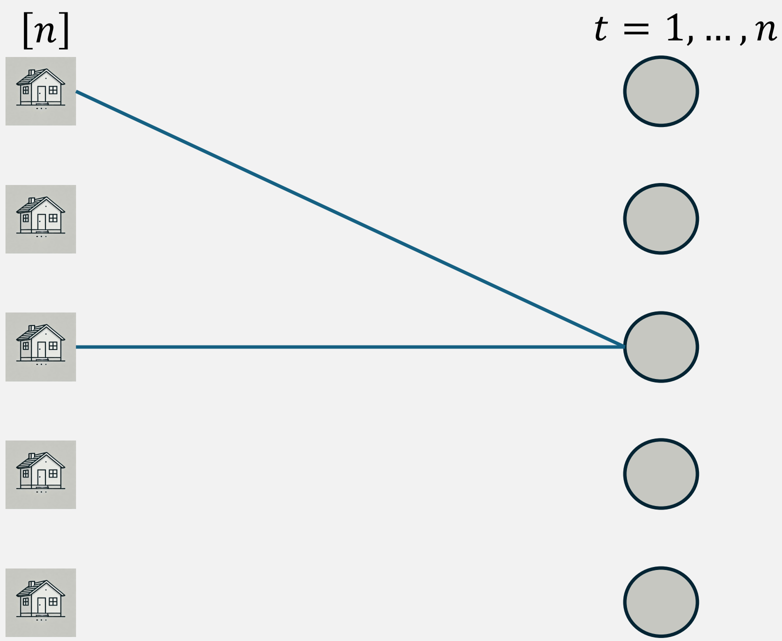

Matching in Bipartite Graph

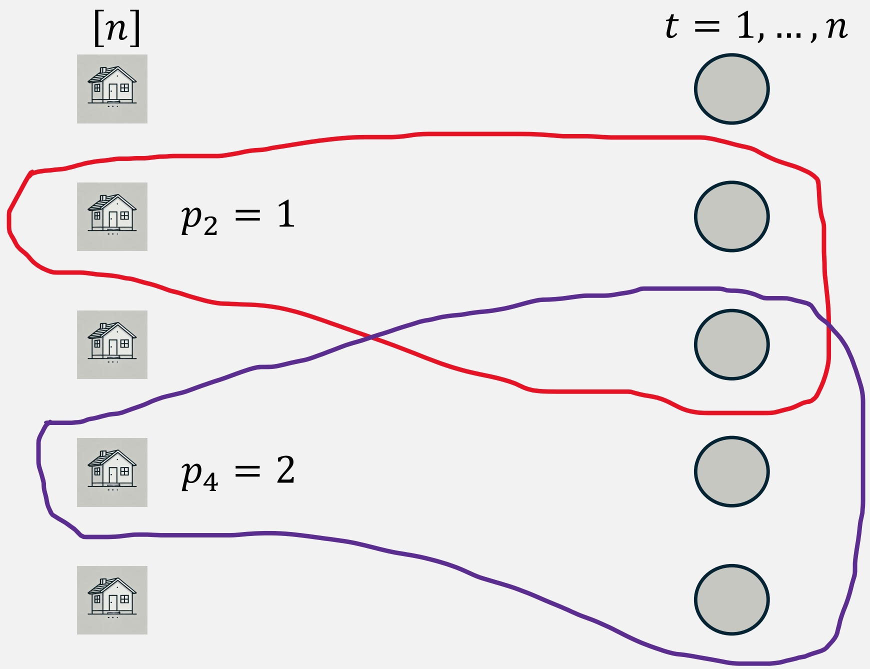

Matching in Block Bipartite Hypergraph

Another Optimization Problem

Maximize monotone submodular f(M)

s.t.

M is a matching in given Block Bipartite Hypergraph

Theorem: There is a 5.32-approximation algorithm for this problem.

Notes:
1) Each hyperedge has exactly one node on the left and a consecutive block of nodes on the right.
2) Contention resolution scheme [Feige and Vondrak, 2006] for the measured continuous greedy algorithm [Buchbinder and Feldman, 2018].

Concluding Remarks

New class of stochastic optimization problems which extend the Pandora's Box Problem.

Not the first mention of time in Pandora-related literature!

- [Weitzman, 1979] proposed a particular form of "exponential" discounting
- [Berger et al., 2024] considered a model with deadlines
- [Fu et al., 2018] study the committed version
- [Singla, 2018] studies the problem with a knapsack constraint (fixed time horizon).

Our model captures all the above.

Open Problems:

1) Improve approximation guarantees for Pandora Over Time
2) Hardness of approximation (e.g. APX-Hardness) ?

Thank you for your attention!