Pandora's Box Problem

Introduced by Weitzman (1979), models the *cost of information*

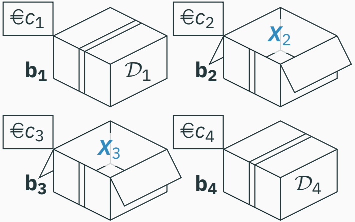

- n "boxes", $\mathbf{b_1}, \mathbf{b_2}, \dots, \mathbf{b_n}$, labeled with costs c_i and independent reward distributions $\mathcal{D}_i$.
- Pay c_i to open $\mathbf{b}_i$ and observe random reward value: $X_i \sim \mathcal{D}_i$.
- Only keep one reward! If opened set S, get $\max_{i \in S} X_i - \sum_{i \in S} c_i$.

Goal: Find the (*adaptive*) strategy achieving in expectation the largest net gain

The solution is a simple *threshold strategy*: The Pandora's Rule

1. For each box b_i Pre-compute *Reservation value* ζ_i such that $\mathbb{E}\left[(X_i - \zeta_i)_+\right] = c_i$
2. Open largest un-opened ζ_i, if have not seen larger value before
3. Repeat until none worth opening.

Note: Stopping time is *adaptive*, but order is not!

Pandora's Box with Order Constraints

This work

"Pandora's Box Problem with Order Constraints"

by Shant Boodaghians[1]*, Federico Fusco*[2]*, Philip Lazos*[2]*, Stefano Leonardi*[2]

Order Constraints: *Modeled by a directed acyclic graph G*

- Models *dependencies* in information, *e.g.* stages of medical trials.
- Can only open box *after* opening some parent
- Forces going through high-cost to get high-reward, risky

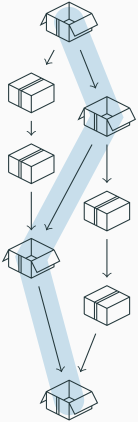

Our Results: general DAGs

- In general, no fixed-order strategy, hard to approximate*.
- We show how to build a 2-approximation* via "adaptivity gap" result

* Approximation model: Strategy π is β-approximation of the optimal solution π^* if

$$\mathbb{E}\left[\max_{i \in S(\pi^*)} X_i - \sum_{i \in S(\pi)} c_i\right] \geq \mathbb{E}\left[\beta^{-1} \cdot \max_{i \in S(\pi^*)} X_i - \sum_{i \in S(\pi^*)} c_i\right]$$

[1] Univerisy of Illinois at Urbana-Champagne, [2] Sapienza University of Rome

Pandora's Box on a tree

Special Case: The graph modeling the order constraint is a *rooted tree*

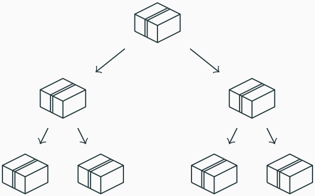

Challenges:

- Depth The value of a box depends from the *possibilities* its opening makes accessible.
- Breath *Distant directions of exploration* must be compared at every time step.

Our Results: Trees and Forests

- There exists a threshold strategy which is optimal
- The thresholds can be computed in polynomial time and space
- The threshold of a box depends only on the subtree of the descendants
- The same results hold for forests

Generalized Reservation Value

The thresholds for the optimal strategy on trees are defined similarly to the ζ of Weitzman original solution

Generalized reservation value: consider the subtree of the descendants as a *macro-box*, whose random cost and reward are given *implicitly* by an optimal strategy π^*

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

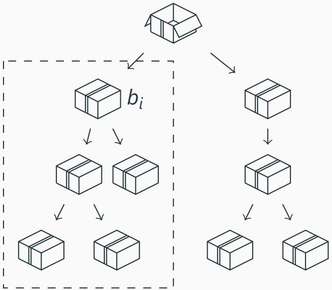

Algorithmic Idea

- Reduce trees to lines recursively: from leaves to root.
- Interleave progressively smallest linearized branches
- Simple dynamic programs to compute Generalized Reservation values.